

## **DX NetOps Performance Management - 23.3**

---

# Table of Contents

<b>Getting Started.....</b>	<b>16</b>
Product Architecture.....	17
Videos.....	19
Network Discovery and Monitoring.....	22
Get Started as a New User.....	24
<b>Release Notes.....</b>	<b>27</b>
Features and Enhancements 23.3.6.....	27
Features and Enhancements 23.3.5.....	27
Features and Enhancements 23.3.4.....	28
Features and Enhancements 23.3.3.....	30
Features and Enhancements 23.3.2.....	32
Features and Enhancements 23.3.1.....	32
Fixed Issues.....	33
Known Issues.....	47
Data Source Compatibility.....	52
Language Support.....	52
Third-Party Software Acknowledgments.....	55
<b>Installing.....</b>	<b>56</b>
Installation Requirements and Considerations.....	57
Review Cloud Sizing Guidelines.....	67
Fault Tolerance.....	71
Prepare to Install NetOps Portal.....	73
Install NetOps Portal.....	78
Prepare to Install the Data Repository.....	82
Prepare to Install the Data Repository for 23.3.2 and Lower.....	90
Install the Data Repository.....	99
Install the Data Repository for 23.3.2 and Lower.....	105
Prepare to Install the Data Aggregator.....	116
Install the Data Aggregator.....	119
Install Fault-Tolerant Data Aggregators.....	120
Install a Non-Fault-Tolerant Data Aggregator.....	123
Prepare to Install the Data Collectors.....	126
Install the Data Collectors.....	131
Complete the Post-Installation Configuration.....	135
Determine Monitoring Requirements.....	145
Configure Monitoring in a New Environment.....	148

Configure Reporting in a New Environment.....	149
<b>Install a Disaster Recovery System.....</b>	<b>150</b>
<b>Install DX NetOps Mediation Manager.....</b>	<b>162</b>
<b>Install NetOps Kafka.....</b>	<b>163</b>
Verify the Prerequisites for a NetOps Kafka Installation.....	163
Install NetOps Kafka from the Command Line.....	165
Install NetOps Kafka in Silent Mode.....	171
<b>Install and Configure Log Analytics for Insights.....</b>	<b>174</b>
<b>Deploy NetOps Flow.....</b>	<b>177</b>
<b>Enable the NetOps Business Reports.....</b>	<b>195</b>
Verify the Prerequisites for Installing the NetOps Report Manager Service.....	196
Install the NetOps Report Manager Service from the Command Line.....	201
Install the NetOps Report Manager Service in Silent Mode.....	205
Set up to Run NetOps Business Reports.....	208
<b>Install a Low-Scale System.....</b>	<b>211</b>
Install the Data Collector on the Same Host as the Data Aggregator.....	212
<b>Uninstall Performance Management.....</b>	<b>215</b>
Uninstall NetOps Portal.....	216
Uninstall the Data Aggregator.....	216
Uninstall the Data Collectors.....	218
Uninstall the Data Repository.....	218
<b>Uninstall NetOps Kafka.....</b>	<b>219</b>
<b>Uninstall the NetOps Report Manager Service.....</b>	<b>220</b>
<b>Upgrading.....</b>	<b>222</b>
<b>Migrating.....</b>	<b>223</b>
Migrate NetOps Portal.....	224
Migrate the Data Repository.....	225
Reinstall or Migrate the Data Aggregator.....	234
Reinstall or Migrate the Data Collectors.....	238
Rehydrate Data in a Cloud Environment.....	242
Migrate Virtual Disk Usage Data for SDN Devices.....	248
<b>Building.....</b>	<b>251</b>
<b>Self-Certification.....</b>	<b>251</b>
Manage Components.....	253
Create or Extend Metric Families.....	254
Create or Extend Vendor Certifications.....	257
Edit Custom Vendor Certifications Using NetOps Portal.....	260
Manage Vendor Certification Priorities.....	261
Create or Edit Vendor Certification Expressions.....	263

Functions and Global Variables.....	264
Vendor Certification Expression Operators.....	281
Self-Certification XML.....	282
Certification Schema Files and Examples.....	282
Restricted XML Tags.....	282
Vendor Certification XML Structure.....	283
Multi-MIB Table Support.....	293
Metric Family XML Structure.....	298
Component XML Structure.....	314
Self-Certification Workflows.....	318
Add Metrics to Existing Metric Families.....	329
Add a Filter to a Vendor Certification.....	332
Change the Calculation Method for an Existing Metric.....	333
Deploy a New Metric Family or Vendor Certification on the Production System.....	335
<b>Manage SNMP Profiles.....</b>	<b>335</b>
Show Secure SNMP Data in Clear Text.....	340
<b>IP Domains.....</b>	<b>340</b>
Manage IP Domains.....	341
Assign Network Flow Analysis Items to an IP Domain.....	343
Assign Application Delivery Analysis Items to an IP Domain.....	344
Assign CA Unified Communications Monitor Items to an IP Domain.....	344
<b>Discovery.....</b>	<b>345</b>
Manage Discovery Profiles.....	346
Run Device Discovery.....	352
Quickly Discover SNMP Devices.....	355
Run Device Rediscovery.....	357
Discover Logical Systems Through SNMP Context.....	359
Discovery From Other Data Sources.....	360
Discovery and Polling in VMware Environments.....	361
SystemEDGE System Response Path Test Metrics.....	362
<b>Groups.....</b>	<b>363</b>
Device Collections.....	366
Manage Groups.....	369
Identify Empty and Unused Groups.....	375
Manage Subgroups.....	377
Manage Group Rules.....	377
View Groups Change Log.....	380
Organize Group Items Geographically.....	381
Organize Devices and Component Items Using Groups and Group Rules.....	383
<b>Manage Monitoring Profiles.....</b>	<b>386</b>



<b>Manage Custom Attributes.....</b>	<b>393</b>
<b>Configure Threshold Profiles.....</b>	<b>396</b>
Manage Threshold Profiles and Event Rules.....	404
Threshold Assessment Logic for Window and Duration.....	408
View Event Rules and Threshold Violation Events.....	410
<b>Configure Business Hours Filtering.....</b>	<b>413</b>
<b>Schedule Maintenance Indicators.....</b>	<b>415</b>
<b>Manage Monitored Devices.....</b>	<b>417</b>
Change the Primary IP Address for a Device.....	421
Device Reconfiguration.....	422
Manage Device Life Cycles.....	425
Manage Hostname Changes.....	428
Override Device Types.....	428
Device Deduplication.....	433
Delete Components That Are Not Present.....	435
Delete Devices.....	441
<b>Manage Metric Families.....</b>	<b>441</b>
Configure Metric Filtering.....	442
Edit a Metric.....	443
Populate Components List for Response Path Metric Family.....	444
Rediscover Metric Families.....	444
<b>Manage Aggregated Components.....</b>	<b>445</b>
<b>Manage Application Mappings.....</b>	<b>448</b>
<b>Manage Interfaces.....</b>	<b>449</b>
Poll Critical Interfaces Faster than Non-critical Interfaces.....	449
Interface Components Naming Convention.....	453
Override Speed In and Speed Out Values on Interfaces.....	453
Configure Counter Behavior.....	454
Manage Interface Polling Behavior.....	456
Manage Network Flow Processing.....	457
Manage Interface Aggregation for Flow.....	460
<b>Configure Round Trip Time (RTT) Tests.....</b>	<b>461</b>
RTT Configuration Details.....	464
RTT Configuration Examples.....	472
IPSLA Polling.....	476
<b>Set Alias Names Using a Script.....</b>	<b>477</b>
<b>Using.....</b>	<b>480</b>
Launch NetOps Portal.....	480
Search and Filter in NetOps Portal.....	480
Customize Your User Settings.....	483

<b>Share Data with Other Users.....</b>	<b>487</b>
Download Dashboard Data as Reports.....	487
Manage Scheduled Reports.....	488
Generate a URL for a View.....	493
Download View Data.....	494
Print View Data.....	494
<b>Inventory Pages and Views.....</b>	<b>495</b>
<b>Dashboards.....</b>	<b>500</b>
Manage Dashboards.....	500
Organize Dashboards in Menus.....	506
Configure the Dashboard Settings.....	507
Out-of-the-Box Dashboards.....	508
Technology-Specific Dashboards.....	512
Vendor-Specific Dashboards.....	513
<b>Context Pages.....</b>	<b>516</b>
Manage Context Pages.....	516
<b>Views.....</b>	<b>520</b>
Customize Views.....	527
Set the Resolution for Reported Data.....	532
Alarms View.....	535
Browser Views.....	542
Bar Chart Views.....	544
Calendar Heat Chart Views.....	547
Card Views.....	549
Dynamic Trend Views.....	551
Dynamic Table Views.....	556
Gauge Views.....	559
Group Scorecard Trend Views.....	561
Group Scorecard Table Views.....	565
Inventory Hierarchy Views.....	569
Map Views.....	571
Pie Chart Views.....	572
Table Views.....	572
Time Bar Chart Views.....	580
Trend Views.....	581
<b>NetOps Business Reports.....</b>	<b>595</b>
<b>Run Business Reports.....</b>	<b>597</b>
Manage NetOps Business Reports.....	597
<b>Manage On-Demand Reports.....</b>	<b>601</b>
<b>Manage On-Demand Report Templates.....</b>	<b>613</b>

<b>Performance Metrics.....</b>	<b>617</b>
Baseline Calculations.....	617
Rate Metrics.....	621
Interface Reporting.....	621
CPU Utilization.....	623
Memory Utilization.....	623
Device Availability and Reachability.....	623
Reachability Status and Contact Status.....	624
Scorecard Projections.....	626
Percentiles.....	626
Metric Projection.....	627
Total, Average, Minimum, and Maximum Values.....	630
<b>Events.....</b>	<b>631</b>
Event Types.....	632
Change the Event Manager Properties.....	636
Threshold Monitoring and Threshold Limiter Behavior.....	641
Threshold Event Processing Self-Monitoring Metrics.....	645
<b>Insights.....</b>	<b>646</b>
Use Log Analytics for Insights.....	646
<b>NetOps Flow.....</b>	<b>647</b>
Flow Dashboards.....	647
View Network Flow.....	648
Update the NetOps Flow Configuration.....	649
<b>Modern Network Monitoring.....</b>	<b>650</b>
Monitor SDN/NFV Virtual Inventory.....	650
Monitor SDN/NFV Virtual Resource Usage.....	651
Monitor SDN/NFV Physical Host Resource Usage.....	651
Monitor SDN/NFV vSwitch Performance.....	652
Monitor Service Chains.....	654
Monitor AWS.....	656
Monitor Cisco ACI.....	657
Monitor Cisco DNA Center.....	660
Monitor SD-WAN Devices.....	663
Monitor VMware NSX-T.....	669
Monitor Wi-Fi Device Inventory.....	678
Monitor Monitoring Point Devices.....	688
Monitor Client Device Inventory.....	689
Monitor VNA System Health.....	692
Virtual Network Assurance Health / Polling Dashboard.....	692
<b>Fault Monitoring.....</b>	<b>693</b>

Customize Event Integration.....	693
Device Configuration View.....	699
Monitor Device Inventory with Alarm States.....	700
<b>Configure Notifications.....</b>	<b>702</b>
Traps Usage.....	708
<b>Identify Volatility in Network Performance.....</b>	<b>709</b>
<b>Generate a REST API Token.....</b>	<b>711</b>
<b>Log Out of NetOps Portal.....</b>	<b>712</b>
<b>Administrating.....</b>	<b>713</b>
<b>Onboard a New Product Operator.....</b>	<b>713</b>
<b>Manage Data Sources.....</b>	<b>714</b>
Configure a Data Source.....	715
Synchronize Data Sources.....	719
Update the Configuration of the Alarm Service.....	722
<b>Manage Roles and User Accounts.....</b>	<b>724</b>
Role Rights.....	726
Data Source Role Rights.....	732
Manage User Account Roles.....	734
Product Privilege.....	735
Data Source Product Privileges.....	736
Manage Product Access.....	738
Manage User Accounts.....	738
Enable or Disable Users.....	744
<b>Proxy Users and Tenants.....</b>	<b>745</b>
<b>Multi-tenancy.....</b>	<b>745</b>
Multi-tenancy Deployment Considerations.....	746
Configure a Tenant Environment.....	749
Manage Tenants.....	751
Administer Tenants.....	755
Automate Tenant Configuration with REST Web Services.....	757
Configure Tenant-Agnostic Data Collectors.....	768
Multi-tenancy and Application Delivery Analysis.....	773
Multi-tenancy and Network Flow Analysis.....	774
Multi-tenancy and CA Unified Communications Monitor.....	774
Multi-tenancy and DX NetOps Spectrum.....	775
<b>NetOps Portal Administration.....</b>	<b>775</b>
Back Up NetOps Portal.....	775
Manage the Themes for NetOps Portal.....	779
Update the NetOps Portal IP Address and Hostname.....	783
Move the NetOps Portal Database to a Separate Node.....	785

Modify Maximum Memory Usage for NetOps Portal.....	786
Restore NetOps Portal.....	787
Configure the Email Server.....	794
Configure Users for Internal Communications.....	797
Configure the Proxy Server on NetOps Portal.....	800
Configure a Reverse Proxy.....	802
Configure Java Options for NetOps Portal Services.....	803
NetOps Portal Scripts.....	804
<b>Data Aggregator Administration.....</b>	<b>805</b>
Automate Device Inventory Synchronization.....	805
Back Up the Data Aggregator.....	812
Choose Another Host in a Cluster When Selected Host Fails.....	816
Update the Data Aggregator IP Address and Hostname.....	817
Configure the Data Collectors When the Data Aggregator IP Address Changes.....	819
Update the Data Collector IP Address and Hostname.....	821
Data Aggregator Configuration Changes During Network Disconnects to a Data Collector Host.....	822
Assign Data Collectors to Tenants and IP Domains.....	823
Configure Data Collectors for Fault Tolerance.....	823
Modify Maximum Memory Usage for the Data Aggregator and Data Collector.....	827
Modify the External ActiveMQ Memory Limit.....	829
Configure Java Options for the Data Aggregator and the Data Collectors.....	831
Configure the Data Aggregator Cleanup.....	831
Monitor Data Aggregator System Health.....	832
Data Aggregator / Data Collector Health Dashboard.....	832
Data Aggregator Polling Dashboard.....	833
Data Collector Polling Dashboard.....	833
Data Aggregator Queries Dashboard.....	834
Data Aggregator General Processing Dashboard.....	834
Data Aggregator Event Processing Dashboard.....	834
Rebalance the Load on the Data Collectors.....	835
Restore the Data Aggregator.....	836
View Data Aggregator Details.....	839
Configure the Failover Settings for Fault Tolerance.....	839
Install or Uninstall the Proxy Server.....	840
Install Nginx.....	842
Configure the Connection from NetOps Portal to the Proxy Server Using Nginx as Reverse Proxy.....	843
Update the Proxy Server IP Address or Hostname.....	850
Data Aggregator Scripts.....	851
<b>Data Repository Administration.....</b>	<b>853</b>
Configure Data Retention Rates.....	853

Back Up the Data Repository.....	855
Configure Passwordless SSH.....	861
Configure the Data Repository Host for a Local Backup.....	862
Update the Data Repository IP Address and Hostname.....	862
Restore the Data Repository.....	864
Add a Node to the Data Repository Cluster.....	867
Data Repository Heartbeat Monitor Process.....	872
Data Repository Audit Process.....	872
Run Data Repository Diagnostic Utilities.....	872
Segment Database Tables.....	875
Move the Data Repository Data Directory.....	881
Data Repository Scripts.....	884
<b>Flow Administration.....</b>	<b>885</b>
<b>Data Extraction.....</b>	<b>894</b>
Configure Streaming Metric Export to a Kafka Cluster.....	894
Bulk Data Export.....	902
<b>View System Status.....</b>	<b>907</b>
<b>View Health Monitoring Information.....</b>	<b>909</b>
<b>Restart Performance Management Component Services.....</b>	<b>912</b>
Restart the Data Aggregator.....	912
Restart the Data Collector.....	914
Restart the Data Repository.....	915
Restart the ActiveMQ Broker.....	917
Restart NetOps Portal.....	918
<b>Syslog Integration.....</b>	<b>920</b>
Configure NetOps Portal to Send Messages to Syslog.....	920
Configure the Data Aggregator to Send Messages to Syslog.....	922
Watch for Syslog Messages.....	923
<b>Logs.....</b>	<b>924</b>
Data Aggregator Logs.....	924
NetOps Portal Logs.....	924
SSO Audit Log.....	925
<b>Activate a Disaster Recovery System.....</b>	<b>926</b>
<b>Securing Performance Monitoring.....</b>	<b>929</b>
<b>Integrating.....</b>	<b>930</b>
<b>Integrate with DX NetOps Spectrum for Fault Management.....</b>	<b>931</b>
Integrate with DX NetOps Spectrum.....	936
<b>Integrate with DX Operational Intelligence.....</b>	<b>938</b>
OI Connector (DX Operational Intelligence) Release Notes.....	940

Install and Upgrade the OI Connector.....	941
Secure Connections for the OI Connector.....	950
Configure the OI Connector.....	951
Uninstall the OI Connector.....	959
<b>Integrate with Application Delivery Analysis.....</b>	<b>959</b>
ADA Metrics.....	959
ADA Dashboards.....	963
Application Performance Dashboard (ADA).....	963
Network Overview Dashboard (ADA).....	964
Network Performance Dashboard (ADA).....	964
Performance Events Dashboard (ADA).....	965
Server Overview Dashboard (ADA).....	965
Server Performance Dashboard (ADA).....	965
ADA Views.....	966
Engineering Trend.....	966
Incident Count by Application.....	966
Incident Count By Network.....	968
Incident Count by Server.....	969
Incident Counts.....	970
Incident List by Network.....	973
Incident List by Server.....	974
Incident Lists.....	975
Performance by Application.....	977
Performance by Network.....	977
Performance by Server.....	978
Performance Map by Server.....	979
Performance Maps.....	980
Performance Scorecard.....	981
Performance Views.....	982
Top Performance by Network.....	985
Top Performance Map by Network.....	986
<b>Integrate with DX NetOps Network Flow Analysis.....</b>	<b>986</b>
Configure Network Flow Analysis in NetOps Portal.....	986
Configure Flow Collection.....	987
Configure Traps.....	991
Results of Unregistering Network Flow Analysis Data Sources.....	991
Set Up User Accounts.....	991
Set Up Groups.....	992
Test Data Source Connections (Register and Configure NFA Use Case).....	992
Verify IP Domains.....	992

Verify SNMP Profiles.....	995
Verify That Data Is Received.....	995
Network Flow Analysis Views in NetOps Portal.....	996
Enterprise-Level Views.....	998
Interface Stacked Trend View.....	1003
Interface ToS Summaries.....	1008
Interface Top Conversations.....	1011
Interface Top Hosts.....	1015
Interface Top Protocols.....	1019
Calendar Chart (Flow).....	1022
<b>Integrate with CA Unified Communications Monitor.....</b>	<b>1023</b>
Call Quality Breakdown.....	1023
Call Quality Service Level Agreement.....	1024
Call Quality Trend.....	1024
Performance Overview Dashboard.....	1025
Top Volume and Utilization Dashboard.....	1026
Worst Performance Dashboard.....	1029
<b>Monitor Server Performance with DX Application Performance Management.....</b>	<b>1031</b>
<b>Generate Mediation Manager Device Packs.....</b>	<b>1033</b>
<b>Integrate with Virtual Network Assurance.....</b>	<b>1033</b>
Manage Connections to Virtual Network Assurance.....	1033
Secure Connections to Virtual Network Assurance.....	1050
<b>Integrate with Splunk.....</b>	<b>1050</b>
<b>Troubleshooting.....</b>	<b>1063</b>
Access Denied to MySQL Utilities.....	1066
Automatic Rediscovery Does Not Run After Updating Vendor Group Priority.....	1067
Browser Shows Error when Logging In.....	1067
Cannot Create a Vendor Certification.....	1068
Cannot Remove a Custom Vendor Certification.....	1068
Cannot Find the Data Aggregator RIB Document.....	1068
Cannot Remove a Metric Family.....	1069
Cannot View More than 5000 Device Components in Inventory List.....	1069
Clean Up After a Failed NetOps Portal Installation.....	1069
Data Aggregator Disk Space is Decreasing.....	1070
Data Aggregator Fails to Synchronize.....	1070
Data Aggregator or Data Collector Does Not Initialize.....	1071
Data Collector Dropped Polling Event Message.....	1074
Data Collector Installs But Does Not Appear in the Data Collector List Menu.....	1074
Data Collector Shows Polling Status Not Connected.....	1076
Data Is Missing from Views.....	1076



Data Source Registration Fails.....	1076
Data Source Synchronization Fails.....	1077
Data Source Test Fails.....	1078
Discovery Does Not Start.....	1078
Gaps Appear in Reports or Views.....	1079
Gaps in Data Appear during Throttling.....	1079
Group Membership Is Not Updated During Synchronization.....	1079
Insecure Connection Message in Firefox.....	1080
Inventory is Empty After a Data Source is Registered.....	1080
Low Data Aggregator Disk Space.....	1081
Metric Family is Incomplete.....	1082
Metric Family is Not Supported.....	1083
Metric Values Do Not Appear in Table in OpenAPI.....	1084
MIB Fails to Compile.....	1084
Multiple SNMP Devices Trigger Intrusions Alarms.....	1084
No Charts or Images are Visible in IE with HTTPS.....	1085
'No Data to Display' Message in Views.....	1085
No Output is Generated After Running the Device Pack Generator.....	1086
No Performance Data for a Device Pack.....	1086
OpenAPI Query Results in Empty Table.....	1086
NetOps Portal Cannot Contact the Data Aggregator.....	1086
Polling Does Not Complete for My Sensitive Device.....	1087
Polling Has Stopped on Discovered Metric Family.....	1087
Polling Safety Valve Event Message.....	1088
Polling Stopped Event Message.....	1088
PrimaryIPAddress ATTRIBUTE_VALUE_NOT_ALLOWED Error in Karaf Log.....	1088
QueryBuilder Certificate Warning.....	1089
Report on All Pages Times Out.....	1089
Services Do Not Start After Installing the Data Collector and VNA.....	1089
Spectrum Alarms Are Missing from the Alarms View.....	1090
Unable to Back Up the Data Repository.....	1091
Unable to Resolve Issue.....	1092
Unexpected Data Aggregator Shutdown.....	1094
Vendor Certification Expression is Erroneous.....	1095
Vertica Fails to Install in a Cluster Environment.....	1095
Vertica Fails to Start.....	1095
Vertica Fails to Install due to 'iptables' Error.....	1096
View Shows Invalid RIB Query Syntax Error.....	1096
<b>APIs.....</b>	<b>1097</b>
NetOps Portal REST Web Services.....	1097

Use NetOps Portal Web Services.....	1098
Dashboards Web Service.....	1100
Data Sources Web Service.....	1102
Devices Web Service.....	1105
Set Device Alias Names Using the Devices Web Service.....	1106
Set Interface Alias Names Using the Devices Web Service.....	1107
Set Component Alias Names Using the Devices Web Service.....	1108
Manage Device Life Cycles Using the Devices Web Service.....	1110
Domains Web Service.....	1110
Groups Web Service.....	1113
Manage Groups Using the Groups Web Service.....	1113
Roles Web Service.....	1133
Users Web Service.....	1136
Tenants Web Service.....	1143
Console Info Web Service.....	1146
Automate Provisioning and Configuration using the tenants Web Service.....	1147
Business Hours Web Service.....	1150
Maintenance Indicators Web Service.....	1155
Alarm Attributes Web Service.....	1158
<b>Data Aggregator REST Web Services.....</b>	<b>1161</b>
Change When Same Day, Same Hour Baseline Averages Are Calculated.....	1166
Manage Polling Behavior for Components.....	1167
Manage Default Polling Behavior.....	1168
Poll Sensitive and Critical Devices Without a Performance Impact.....	1171
Schedule Data Purges.....	1174
Schedule Rollup Processing and Baseline Calculations.....	1175
Manage Discovery Using REST.....	1176
Automate Creating Discovery Profiles.....	1177
Automate Running Discovery.....	1179
Automate Configuring Aggregated Components.....	1186
<b>OpenAPI.....</b>	<b>1191</b>
Use the OpenAPI QueryBuilder.....	1191
OpenAPI QueryBuilder Examples.....	1195
Advanced OpenAPI Query Examples.....	1209
Configure OpenAPI Defaults and Limits.....	1212
OpenAPI Apps.....	1214
Audit OpenAPI Usage.....	1218
<b>Product Accessibility Features.....</b>	<b>1221</b>
<b>Product References and Abbreviations.....</b>	<b>1225</b>

<b>Documentation Legal Notice.....</b>	<b>1226</b>
--	-------------

# Getting Started

---

Learn about the basic features and use cases of DX NetOps Performance Management.

DX NetOps Performance Management monitors, stores, analyzes, and displays information for assuring service quality across large, complex, multi-technology, multi-vendor network infrastructure. The solution helps the largest networks successfully monetize service offerings while lowering the cost and complexity of service.

*Communications service providers* can improve network monitoring and delivery of revenue-generating services, such as 4G LTE, Voice over LTE, Mobile Backhaul, Metro Ethernet, and more, using DX NetOps Performance Management.

*Enterprises* can assure underlying network services for applications that drive their internal business processes and revenue-generating customer interactions using DX NetOps Performance Management.

## Key Features

DX NetOps Performance Management includes the following key features:

- **High-scale monitoring architecture on a platform that scales efficiently.**

The system architecture provides scale that supports the largest networks.

### TIP

You can get the system requirements to support your scale using the [DX NetOps Sizing Tool](#).

- **Unified multi-technology, multi-vendor device monitoring. Certifications for classic network devices and specialized carrier Ethernet, Wi-Fi offloading, and mobile wireless equipment.**

DX NetOps Performance Management supports the common vendors, metrics, and components in your network infrastructure. You can customize communication with monitored devices by way of the extensible certification model.

For more information about how to support new vendor devices and technology types, see [Self-Certification](#).

- **Intelligent analytics, high-scale visualization, and fast processing for instant reporting. Flexible, easily customizable dashboards and reports.**

The customizable dashboards and views provide flexible visualization. For example, a regional manager uses a dashboard that pins views to each site group in that region and systems administrator uses a dashboard to monitor all servers. Dashboards show information that is scoped to customizable groups. Context pages show information that is related to a specific item in the system. Views provide visualization options.

For more information, see [Dashboards](#), [Views](#), and [Context Pages](#).

- **Extensible architecture for easy integration and automation.**

DX NetOps Performance Management supports downstream integration through the customer-facing OpenAPI.

For more information, see [OpenAPI](#).

- **Predictive analytics to give a complete, unencumbered view of the network, and business key performance indicators.**

Configurable and dynamic projections provide insight into capacity planning and situations to watch.

For more information, see [Metric Projection](#).

- **Modern network monitoring for software-defined architectures and hybrid cloud platforms.**

Dashboards and views designed for modern network monitoring.

For more information, see [Modern Network Monitoring](#).

## Integrated Solutions

DX NetOps Performance Management includes the following integrated solutions:

- **DX NetOps Mediation Manager (DX NetOps MM)**

Integrate non-SNMP monitoring in your DX NetOps Performance Management environment.

For more information:

- About this integrated solution, see [Install DX NetOps Mediation Manager](#).
- About DX NetOps Mediation Manager, see the [DX NetOps Mediation Manager documentation](#).

- **DX NetOps Network Flow Analysis (NFA)**

Monitor interface performance and bandwidth utilization with insight into the traffic going through your network.

For more information:

- About this integrated solution, see [Integrate with DX NetOps Network Flow Analysis](#).
- About DX NetOps Network Flow Analysis, see [the DX NetOps Network Flow Analysis documentation](#).

- **DX NetOps Spectrum (Spectrum)**

Map network topology and create alarms from performance events.

For more information:

- About this integrated solution, see [Integrate with DX NetOps Spectrum for Fault Management](#).
- About Spectrum, see [the DX NetOps Spectrum documentation](#).

- **DX NetOps Virtual Network Assurance (VNA)**

Extend traditional network monitoring to the virtual network, provide modern network monitoring for software-defined architectures and hybrid cloud platforms, and correlate logical network entities with physical resources.

For more information:


- See [Integrate with Virtual Network Assurance](#).
- See [the DX NetOps Virtual Network Assurance documentation](#).

## Product Architecture

DX NetOps Performance Management uses an extensible multiserver architecture to support monitoring the largest networks. This article covers the basic system architecture.

DX NetOps Performance Management collects performance data and integrates with other data sources, such as CA Application Delivery Analysis and DX NetOps Network Flow Analysis, and provides more information about your network.

The following diagram shows the basic system architecture for DX NetOps:

**Figure 1: Basic System Architecture****NOTE**

For DX NetOps Performance Management to work properly in a firewall-protected environment, certain ports must be open.

For more information about these ports, see [Review Installation Requirements and Considerations](#).

**Components**

DX NetOps Performance Management includes the following primary components:

- **NetOps Portal**
  - Front-end client for the console, dashboards, and reporting
  - Administration and event management
  - Integrates with other data sources
- **Data Aggregator**
  - Data loading and normalization
  - Threshold monitoring
  - OpenAPI
- **Data Repository**
  - A database that stores performance data for the data aggregator
  - A multi-node (3-5) database cluster for large networks, or a single database node for small networks
- **Data Collector**
  - Simple Network Management Protocol (SNMP) and Internet Control Message Protocol (ICMP) discovery and data collection
  - Integration point for DX NetOps Virtual Network Assurance and DX NetOps Mediation Manager
  - Each supports monitoring up to 500,000 monitored items
  - Multiple data collectors support large or geographically dispersed networks
- **DX NetOps Mediation Manager**
  - Collects non-SNMP inventory and performance data
  - Loads data to the data aggregator through a dedicated data collector
- **DX NetOps Virtual Network Assurance**
  - Collects inventory and performance data from software-defined networking (SDN) and network functions virtualization (NFV) environments
  - Loads data to the data aggregator through a data collector

### Other Architectures

The following architectures reduce data loss:

- **Disaster Recovery**  
If a large-scale disaster occurs, disaster recovery enables a switchover to a recovery system.  
For more information about disaster recovery, see [Disaster Recovery](#).
- **Fault Tolerance**  
Fault tolerance enables your DX NetOps Performance Management environment to continue operating properly when a hardware failure or network issue occurs.  
For more information about fault tolerance, see [Fault Tolerance](#).

A low-scale architecture is also available. A low-scale architecture matches the basic architecture, but the data aggregator and data collector are on a single node. If the sizing tool indicates a low-scale deployment, see [Install a Low-Scale System](#).

## Videos

Use the quick overview videos to get started with DX NetOps Performance Management.

This article includes videos for the following:

- [Integrations](#)
- [Context Pages and Dashboards](#)
- [Discovery](#)
- [Events and Notifications](#)
- [Groups and Collections](#)
- [Install and Upgrade](#)
- [Device Management](#)
- [Monitoring](#)
- [Administering](#)

#### **NOTE**

- These videos are also embedded throughout the articles.
- References to the `karaf` directory in the videos apply for current version.

### **Access Video Transcripts**

You can access video transcripts, for example, for accessibility reasons, from YouTube.

#### **Follow these steps:**

1. Click the video title to open it in YouTube.
2. Click the **More (...)** icon.
3. Click **Open transcript**.

The transcript appears in the Transcript pane.

### **Integrations**

The following videos examine integrations.

#### **DX Operational Intelligence Integrations**

The following video examines a DX Operational Intelligence integration:

For more information, see [Integrate with DX Operational Intelligence](#).

#### **DX NetOps Spectrum Integrations**

The following video examines a DX NetOps Performance Management-DX NetOps Spectrum (Spectrum) integration.

For more information, see [Integrate with DX NetOps Spectrum for Fault Management](#).

### **Context Pages and Dashboards**

The following video examines how to customize dashboards in NetOps Portal with views and context pages to display meaningful information and visualizations for specific network monitoring requirements:

For more information, see [Dashboards](#).

The following video shows how to drill down from a dashboard in the NetOps Portal into context pages, as well as how to configure context pages to display meaningful views, including information about device components and interfaces:

For more information, see [Manage Context Pages](#).

The following video shows how to generate digital and printed reports from the content of dashboards and context pages in NetOps Portal:



For more information, see [Share Data with Other Users](#).

### **Discovery**

The following videos examines how to create, run, and view the results of a discovery profile in NetOps Portal, using an IP address or IP address ranges, to define which devices to discover on the network for performance monitoring:

For more information, see [Discovery](#).

### **Events and Notifications**

The following video examines how to set up notifications for threshold violation events in NetOps Portal to notify teams when threshold violation events occur or change:

For more information, see [Configure Notifications](#).

### **Groups and Collections**

The following video examines how to logically organize network items in NetOps Portal by creating and managing groups, and using groups as data set filters for dashboards, views, and context pages:

For more information, see [Groups](#).

The following video examines how to create a custom collection in NetOps Portal to organize network items with common monitoring configurations:

For more information, see [Manage Groups](#).

The following video examines how to verify DX NetOps Virtual Network Assurance (VNA) data integration with NetOps Portal to confirm VNA plug-in data is populating dashboards as required:

For more information, see [Groups](#).

### **Install and Upgrade**

The following video examines the NetOps Portal installation:

For more information, see [Install NetOps Portal](#).

The following video examines the data repository installation:

For more information, see [Install the Data Repository](#).

The following video examines the data aggregator installation:

For more information, see [Install the Data Aggregator](#).

The following video examines the data collector installation:

For more information, see [Install the Data Collectors](#).

The following video shows how to bind the data aggregator to NetOps Portal and confirm that the data collector is properly connected to the data aggregator:

For more information, see [Configure a Data Source](#) and [Synchronize Data Sources](#).

## **Device Management**

The following video examines how to create and prioritize an SNMP profile in NetOps Portal to provide access credentials for monitored devices:

For more information, see [Manage SNMP Profiles](#).

The following video examines new device status verification:

The following video examines viewing metric family and vendor certification details:

For more information, see [Manage Monitored Devices](#).

## **Monitoring**

The following video examines discovery and monitoring by way of monitoring profiles:

For more information, see [Manage Monitoring Profiles](#).

The following video examines controlling what metrics to monitor for each metric family within a monitoring profile by way of *metrics filters*:

For more information, see [Configure Metric Filtering](#).

The following video examines how to create, edit, and add *component filters* to determine which components to discover for a metric family on a monitoring profile:

For more information, see [Manage Monitoring Profiles](#).

The following video examines how to create a threshold profile in NetOps Portal to generate threshold violation events when major changes happen to baseline performance conditions:

For more information, see [Configure Threshold Profiles](#).

## **Administrating**

The following video examines how to configure the settings for user roles to determine the menu rights, NetOps Portal functionality rights, and data source rights for users assigned to that specific role:

The following video examines how to create a user in NetOps Portal by assigning a role, defining access permissions, assigning groups, defining group administration rights, and defining product privileges. Then, once a user is configured, proxy the user, viewing NetOps Portal with their permissions to confirm the user is configured properly:

For more information, see [Manage User Accounts](#) and [Proxy Users and Tenants](#).

## **Network Discovery and Monitoring**

NetOps Portal discovers your network, and then collects performance data.

Monitoring and discovery are made through the Internet Control Message Protocol (ICMP) and Simple Network Management Protocol (SNMP) protocols. The following are the items that control discovery and monitoring and that you can configure:

- **Collections** are groups of devices that share monitoring behavior through association with monitoring profiles.
- **Discovery profiles** attempt to discover devices by identifying Internet Protocol (IP) address ranges for the data collectors. Discovery profiles define which SNMP profiles CAPM uses to contact devices in the defined IP ranges.
- **Metric families** are related sets of metrics that NetOps Portal collects across several technologies. The metric definitions determine how to report the values for the metrics. Metric families normalize performance data from different devices and device types.
- **Monitoring profiles** define the monitoring behavior, poll cycle duration and which metrics to collect, for the associated device collections.
- **SNMP profiles** include access credentials to use for discovery and monitoring.
- **Vendor certifications** map the MIB attributes for a particular vendor device to the metrics in supported metric families.

## **Discovery**

During discovery, NetOps Portal identifies devices in your network:

1. The *data aggregator* assigns IP addresses in the discovery profile IP ranges to the data collectors.
2. The *data collectors* determine whether devices respond to ICMP or SNMP.
3. For devices that respond to SNMP, the data collector associates the device with a vendor certification and determines configuration data, which includes the following information:
  - The classification of the device, such as router or switch.
  - The device vendor, such as Cisco or Juniper.
  - The device type, such as 7700 or 8200.
4. *Group rules* add the devices to collections.
5. The *data collectors* identify components for each device using the information in the vendor certifications.

## **Monitoring**

During operation, the data collectors query devices using SNMP MIB requests:

1. For each device, the data collector consolidates information from each monitoring profile that is associated with each collection to which the device belongs.
2. The data collectors:
  - a. Request the supported metrics in the monitoring profiles at the fastest rate among all monitoring profiles for the device.
  - b. Stagger SNMP requests within the time of the poll cycle.
  - c. Batch poll responses and send the messages to the data aggregator.
3. The data aggregator loads the performance data to the data repository.

## **Non-SNMP Inventory and Performance Data**

DX NetOps Mediation Manager is the standard non-SNMP data collection tool. DX NetOps Virtual Network Assurance (VNA) is the data collection and normalization tool for SDN and NFV controllers and orchestrators. DX NetOps Mediation Manager and VNA collect inventory and performance data and load that data to DX NetOps Performance Management through the data collectors. If DX NetOps Performance Management discovers and monitors the same devices through SNMP, the data aggregator deduplicates devices with the inventory collected from DX NetOps Mediation Manager and VNA.

For more information, see [the DX NetOps Mediation Manager documentation](#) and [the DX NetOps Virtual Network Assurance documentation](#).

## Get Started as a New User

You can manage your physical and virtual networks, applications, and devices using the NetOps Portal web-based user interface (UI).

The NetOps Portal context pages, dashboards, and reports show performance data from your network and systems-monitoring products (data sources). Compare large amounts of statistical data from multiple sources in a single web page. NetOps Portal takes a performance-first approach to application service delivery by placing end users in the primary role. Capture and analyze data from applications, devices, and the network using NetOps Portal, and get an understanding of how well an IT organization supports application delivery.

NetOps Portal includes role-specific views of application response times, traffic composition, infrastructure health, and flow-based diagnostics.

DX NetOps Performance Management integrated with DX NetOps Virtual Network Assurance (VNA) includes modern network monitoring. The integration enables comprehensive coverage, with monitoring that is scalable and heterogeneous across the greatest number of technology stacks in the following architectures:

- Traditional
- Software-defined networking (SDN)
- Software-defined data center (SDDC)
- Software-defined wide area network (SD-WAN)
- Network functions virtualization (NFV)
- Hybrid-cloud

The following video highlights several key features of NetOps Portal:

In this article:

- [Customize Your User Settings](#)
- [Explore Managed Items](#)
- [View Performance Data and Customize Your Experience](#)
- [Access and Share Reports](#)
- [Organize Managed Items](#)
- [Manage Events](#)

### **Customize Your User Settings**

You can customize each user account with your personal settings, such as your preferred language for NetOps Portal.

For more information about how to customize your user account settings, see [Customize Your User Settings](#).

### **Explore Managed Items**

Data sources discover and monitor your managed items (for example, applications, devices, or interfaces). After you have configured monitoring, you can [explore your managed items](#) using the Inventory pages or [search as a launch point](#).

For more information about how to explore managed items using the Inventory pages, see [Inventory Pages and Views](#).

### **Navigate the Inventory**

All the managed items to which you have permission to view are available from the Inventory. From the Inventory, you can navigate by item type category (for example, Devices) to lists of those managed items. You can also drill down to the context page of an individual item for more details.

For more information, see [Inventory Pages and Views](#).

---

## **Search for Managed Items**

You can search for text that is contained in an item string using the global-level search box. The search returns inventory lists of all the managed items that match your search, which are sorted by item type category.

For more information about how to search for managed items, see [Search and Filter in NetOps Portal](#).

## **View Performance Data and Customize Your Experience**

You can view performance data on dashboards and context pages, and customize your display settings, dashboards, and views. For example, you can edit dashboards to add, remove, rearrange, or customize views:

- **Dashboards**

Provide performance and status data that is scoped to a group. For example, a dashboard page can provide the average performance of monitored items in a group. Dashboards often provide a drill-down path to more detailed, related pages from a selected context.

The following video examines how to customize dashboards in NetOps Portal with views and context pages to display meaningful information and visualizations for specific network monitoring requirements:

For more information, see [Dashboards](#).

- **Context Pages**

Provide focused performance and status data that is scoped to a specific managed item, such as a single router or server. These pages are available as drill-down links or tabs from dashboard pages.

The following video shows how to drill down from a dashboard in the NetOps Portal into context pages, as well as how to configure context pages to display meaningful views, including information about device components and interfaces:

For more information, see [Context Pages](#).

Dashboards and context pages render views, which report collected data in a chart or a table format. Depending on the view, the data comes from the various registered data sources. Views that show data for a group contain collated and aggregated data from data sources. Views that show data for individual items provide a drill-down path to the context page for the item.

For more information, see [Views](#).

## **Access and Share Reports**

You can create or access reusable On-Demand report templates, which dynamically retrieve the most recent data sets from specific sets of items or groups. You can then download a report for sharing. You can also access, download, and share dashboards and views.

For more information, see [On-Demand Reports](#) and [Share Data with Other Users](#).

## **Organize Managed Items**

A group is a filter definition that functions as a container for managed items. Groups let you logically organize managed items in a hierarchical tree structure, with each group containing subgroups or managed items. The structure is propagated to the data sources, where it enables drill down from top-level groups into data from an increasingly narrow but related context.

While creating custom groups, you can define relationships, policies, and dependencies among services, devices, applications, locations, and users within your organization using the Groups tree. You can then organize your managed items by creating group rules. The groups that appear in your **My Custom Groups** area are visible to only you.

The following video examines how to logically organize network items in NetOps Portal by creating and managing groups, and using groups as data set filters for dashboards, views, and context pages:

For more information about how to organize managed items, see [Manage Groups](#) and [Manage Group Rules](#).

## **Manage Events**

An event is a message that provides information about what is happening in DX NetOps Performance Management. Events provide information for monitoring the health and status of your system and your environment. All events include basic information, such as related devices and the time of the occurrence that triggered the event.

To view events, access or add one of the following views:

- **Events View**

This view displays all the events that occurred in the selected time range for the dashboard. This view can be filtered for a specific group. This view is the default view in the Events Display dashboard.

- **Filtered Event Views**

This view includes filters for data source, severity, event type, event subtype, and threshold profile.

You can configure notifications for events that come from a data source to the Event Manager. The incoming events are evaluated against the conditions that you configure for the notification criteria. Only when the criteria are met does Event Manager take a notification action. If an event does not trigger a notification, the event can still be displayed in the Event List.

The following video examines events and notifications:

For more information, see [Events](#) and [Configure Notifications](#).

## Release Notes

---

The Release Notes section outlines the new features and current and fixed issues for DX NetOps Performance Management. In addition, this section contains the Third-Party Software Acknowledgements, which detail the terms and conditions of using third-party software in the creation of DX NetOps Performance Management.

### Features and Enhancements 23.3.6

View a summary of features and enhancements for DX NetOps Performance Management 23.3.6.

This release includes the following new features and enhancements:

- [Support for MySQL 8.0.35](#)
- [Updates to the NetOps Report Manager Service Installation and Upgrade](#)
- [Enhanced Support for SNMPv3 Device Discovery](#)
- [Manage the Notification Filters for a VNA Gateway](#)

For feature and enhancement descriptions for previous releases, use the **Version** drop-down.

#### **Support for MySQL 8.0.35**

NetOps Portal now supports MySQL 8.0.35 to leverage the latest enhancements and vulnerability fixes.

#### **Updates to the NetOps Report Manager Service Installation and Upgrade**

The NetOps Report Manager Service installer now:

- Removes the `InstallAnywhere` files.
- Uses Unix shell commands.

The NetOps Report Manager Service kit no longer includes Java. Certain commands, such as when enabling HTTPS for the NetOps Portal Report Manager Service, require `keytool`, which is part of JRE. This update also affects the path to `keytool`.

Prior to installing or upgrading the NetOps Report Manager Service, ensure that you have the required Java version installed and configured.

For more information, see [Verify the Prerequisites for Installing the NetOps Report Manager Service](#).

#### **Enhanced Support for SNMPv3 Device Discovery**

When DX NetOps Performance Management discovers SNMPv3-supported devices, it no longer requires a unique engineID.

#### **Manage the Notification Filters for a VNA Gateway**

You can now manage (add, edit, view details, and delete) the notification filters for a VNA Gateway from NetOps Portal.

For more information, see [Manage Connections to Virtual Network Assurance](#).

### Features and Enhancements 23.3.5

View a summary of features and enhancements for DX NetOps Performance Management 23.3.5.

This release includes the following new features and enhancements:

- [Identify Empty and Unused Groups for Removal](#)
- [Import New or Updated CA-Signed Certificates using SslConfig](#)

For feature and enhancement descriptions for previous releases, use the **Version** drop-down.

#### **Identify Empty and Unused Groups for Removal**

As part of regular maintenance, you remove groups that are empty and are unused. Now, prior to removing them, you can generate a list of those groups.

For more information, see [Identify Empty and Unused Groups](#).

#### **Import New or Updated CA-Signed Certificates using SslConfig**

Before certificate authority (CA)-signed certificates expire, you must import the new or updated CA-signed certificates. You can now import these certificates using the SSL Configuration tool (SslConfig) or the `keytool` command.

For more information, see [Update NetOps Portal HTTPS Certificates](#).

## **Features and Enhancements 23.3.4**

View a summary of features and enhancements for DX NetOps Performance Management 23.3.4.

This release includes the following new features and enhancements:

- [Hide Authentication Details During Login](#)
- [Print and Magnify Views](#)
- [Upgrades to the Proxy Server](#)
- [Manage the Scheduling Options for NetOps Business Reports](#)
- [Show Device/Interface Threshold Lines in Trend Views](#)
- [Provide Ingress TLS Certificate and Private Key for NetOps Flow as Files](#)
- [Simplification of the NetOps Flow Deployment Process](#)
- [Integrate with Splunk](#)
- [Support for ActiveMQ 5.18.3](#)
- [Support for AdoptOpenJDK 11.0.21+9](#)
- [Manage User Domains in a VNA Gateway](#)
- [Updates on Access Point Context Page and Client Context Page](#)

For feature and enhancement descriptions for previous releases, use the **Version** drop-down.

#### **Hide Authentication Details During Login**

On the **Log In** page when specifying your NetOps Portal username and password and authentication fails, you can now specify whether extra details about the failure are shown.

For more information, see [Update the Single Sign-On Website Settings](#).

#### **Print and Magnify Views**

The **Export/Print Dashboard Views** role right now replaces the **Export to CSV** role right. This new role rights allows users to export view data as comma-separated values (CSV) files or to print view data as Portable Document Format (PDF) files.

For more information:

- About these role rights, see [Role Rights](#).
- About how to print view data, see [Print View Data](#).



## **Upgrades to the Proxy Server**

You can now upgrade the proxy server without having to reinstall. The upgrade process backs up and restores the HTTPS settings. No further steps are required post install to restore these settings.

There is now a separate process for upgrading the proxy server depending on the version from which you are upgrading.

For more information:

- About how to upgrade the proxy server to 23.3.4 and higher, see [Upgrade the Proxy Server](#).
- About how to upgrade the proxy server to 23.3.1 through 23.3.3, see [Upgrade the Proxy Server to 23.3.3](#).

## **Manage the Scheduling Options for NetOps Business Reports**

You can now define whether NetOps Portal creates a copy of a NetOps business report when you schedule the report so that you can:

- Schedule the report with different settings.
- View the settings of the report.
- Edit the settings of the report.

For more information, see [Manage NetOps Business Reports](#).

## **Show Device/Interface Threshold Lines in Trend Views**

You can now set to show a device/interface threshold (line) in Multi-View and Multi-Trend customizable views, trend views, dynamic trend views, and on-demand reports. If a threshold is applied, NetOps Portal displays threshold line trend lines in pink color and shaded in the region above/below the threshold line. You can also now view the assigned thresholds, and then download this data as a report.

For more information, see [Customize Views](#) and [Manage On-Demand Reports](#).

## **Provide Ingress TLS Certificate and Private Key for NetOps Flow as Files**

You now provide the Ingress TLS certificate and private key securely as files to the NetOps Helm chart when deploying NetOps Flow instead of as clear text in the command line.

For more information, see [Deploy NetOps Flow](#).

## **Simplification of the NetOps Flow Deployment Process**

The following changes have been made in the deployment process for NetOps Flow:

- The NetOps Flow deployment no longer includes or requires the `netops-flow-override.yaml` file.
- The `global.imagePullPolicy` Helm chart value now defaults to **Always**.
- The NetOps Flow deployment no longer requires the `<IP_DOMAIN_ID>-flow-aggregator-windowed-data-store-changelog` and `<IP_DOMAIN_ID>-flow-aggregator-windowed-data-suppress-store-changelog` Kafka topics.

For more information, see [Deploy NetOps Flow](#).

## **Integrate with Splunk**

### **IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make

Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

You can now integrate with Splunk to retrieve Syslogs in the context of alarms and devices using the out-of-the-box Syslog connector. The Syslog connector retrieves Syslogs from your Splunk instance, and then constructs the retrieved data in a format that NetOps Portal can consume.

For more information, see [Integrate with Splunk](#).

### **Support for ActiveMQ 5.18.3**

The data aggregator and data collector now use the apache-activemq-5.18.3 version of Apache ActiveMQ (AMQ).

### **Support for AdoptOpenJDK 11.0.21+9**

DX NetOps Performance Management now supports AdoptOpenJDK 11.0.21+9 to leverage the latest enhancements and vulnerability fixes that the 11.0.21+9 version offers.

### **Manage User Domains in a VNA Gateway**

You can now perform Swagger functionality on the **VNA Administration** page in NetOps Portal. From this page, you can now:

- Create a user domain
- Update the name for a user domain in a VNA Gateway
- Get a list of user domains

For more information, see [Manage Connections to Virtual Network Assurance](#).

### **Updates to the Access Point Context Page and Client Context Page**

The following updates have been made to the pages:

- **Client Context Page**

The **Clients** table now includes the following columns:

- **Signal Strength**
- **SNR-Average**
- **Bytes-In Total**
- **Bytes-Out Total**

- **Wireless Access Point Devices Page**

The **Wireless Access Point Devices** table now includes the **Clients Count-Average** column.

For more information, see [Monitor Wi-Fi Device Inventory](#).

## **Features and Enhancements 23.3.3**

View a summary of features and enhancements for DX NetOps Performance Management 23.3.3.

This release includes the following new features and enhancements:

- [NetOps Business Reports](#)
- [View Device Contact Status from NetOps Portal](#)
- [Updates to the Data Repository Installation and Upgrade Process](#)
- [Support for MySQL 8.0.34 and MySQL Connector 8.1.0](#)
- [Updates to Upgrading NetOps Flow from 22.2.11 - 23.3.2 to 23.3.3 or higher](#)
- [Map Groups of VNA Inventory into IP Domains](#)
- [Update to Installing and Upgrading the Data Repository](#)

## **NetOps Business Reports**

### **IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

NetOps Business Reports consists of the NetOps Report Manager Service standalone application and the new business reports in NetOps Portal. If you have integrated DX NetOps Spectrum (Spectrum) with CA Business Intelligence (CABI), you can now obtain alarms, availability, and asset reports using the NetOps business reports.

For more information about how to install the NetOps Report Manager Service and enable the NetOps business reports, see [NetOps Business Reports](#).

## **View Device Contact Status from NetOps Portal**

You can now view the whether the data collector can contact a device from the **Devices** page. You can also show/display the date and time that the device contact status changed on this page.

For more information, see [Inventory Pages and Views](#).

## **Updates to the Data Repository Installation and Upgrade Process**

The installation and upgrade includes changes to the scripts.

For more information, see [Install the Data Repository](#) and [Upgrade the Data Repository](#).

## **Support for MySQL 8.0.34 and MySQL Connector 8.1.0**

NetOps Portal now supports MySQL 8.0.34 and MySQL Connector 8.1.0 to leverage the latest enhancements and vulnerability fixes.

## **Updates to Upgrading NetOps Flow from 22.2.11 - 23.3.2 to 23.3.3 or higher**

When upgrading NetOps Flow from 22.2.11 - 23.3.2 to 23.3.3 or higher, *prior to upgrading NetOps Flow*, complete the prerequisite steps.

For more information about these steps, see [Upgrade NetOps Flow](#).

## **Map Groups of VNA Inventory into IP Domains**

Administrators can now map groups of VNA inventory into a specific IP domain. Based on the import rule, if the data aggregator moves a device from one IP domain to another, it also generates an event noting that. You can customize the groups that are available to be mapped. By default, the group is the VNA user domain.

For more information, see [Manage Connections to Virtual Network Assurance](#).

### **Update to Installing and Upgrading the Data Repository**

The process to prepare for installing the data repository, to install the data repository, and to upgrade the data repository has changed.

For more information, see [Prepare to Install the Data Repository](#), [Install the Data Repository](#), and [Upgrade the Database](#).

## **Features and Enhancements 23.3.2**

View a summary of features and enhancements for DX NetOps Performance Management 23.3.2.

This release includes the following new feature and enhancement:

- [NetOps Business Reports](#)

### **NetOps Business Reports**

#### **IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

NetOps Business Reports consists of the NetOps Report Manager Service standalone application and the new business reports in NetOps Portal. If you have integrated DX NetOps Spectrum (Spectrum) with CA Business Intelligence (CABI), you can now obtain alarms reports using the NetOps business reports.

For more information about how to install the NetOps Report Manager Service and enable the NetOps business reports, see [NetOps Business Reports](#).

## **Features and Enhancements 23.3.1**

View a summary of features and enhancements for DX NetOps Performance Management 23.3.1.

This release includes the following new features and enhancements:

- [Quick Discovery of Devices](#)
- [Debug Metric Family-Based Rediscovery](#)
- [Unified Securing Section in DX NetOps Documentation](#)
- [Identify Volatility in Network Performance](#)
- [Support for ActiveMQ 5.18.2](#)

### **Quick Discovery of Devices**

You can now have NetOps Portal quickly discover SNMP devices using the SNMPv3 protocol. NetOps Portal can no longer quickly discover virtual network devices.

For more information, see [Quickly Discover SNMP Devices](#).

### **Debug Metric Family-Based Rediscovery**

You can now optionally generate a discovery log for debugging purposes when rediscovering how a metric family is applied to a device.

For more information, see [Rediscover Metric Families](#).

## **Unified Securing Section in DX NetOps Documentation**

To provide cohesive and prescriptive end-to-end securing instructions, the secure DX NetOps documentation has been unified, and is now part of the [DX NetOps](#) documentation under the **Secure DX NetOps** section.

The **Secure DX NetOps** section includes high-level categories, each containing capability-specific information.

### **NOTE**

The previous securing URLs now redirect to the corresponding section in the new structure. If you have bookmarked these URLs, consider creating new ones to the new site.

## **Identify Volatility in Network Performance**

### **NOTE**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We fully intend to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per capability basis.

You can now determine where the performance of your interfaces specific to their volatility is variable and unpredictable from the **Interface Volatility** dashboard. Use this dashboard to accurately evaluate, predict, and plan resource availability and performance.

For more information, see [Identify Volatility in Network Performance](#).

## **Support for ActiveMQ 5.18.2**

The data aggregator and data collector now use the apache-activemq-5.18.2 version of Apache ActiveMQ (AMQ).

## **Fixed Issues**

This release provides fixes and enhancements to pre-existing DX NetOps Performance Management functionality.

In this article:

- [23.3.6 Fixes](#)
- [23.3.5 Fixes](#)
- [23.3.4 Fixes](#)
- [23.3.3 Fixes](#)
- [23.3.2 Fixes](#)
- [23.3.1 Fixes](#)

### **23.3.6 Fixes**

This release includes the following fixed issues:

- **Symptom:** When opening the view settings with report configurations, the processor load increases, causing the browser to lag.  
**Resolution:** With this fix, the dialog for selecting alarm types and causes now displays options divided into pages. (23.3.6, DE579027)
- **Symptom:** The **Inventory Devices** list might display an out-of-date **Device Contact Status**.  
**Resolution:** With this fix, the **Inventory Devices** list now displays the most current **Device Contact Status**. (23.3.6, DE582612)
- **Symptom:** The quick filter functionality does not work on a number of trend charts within the **System Health** dashboard pages.

- Resolution:** With this fix, some trend views within the **System Health** dashboard pages have been repurposed to support the quick filter functionality.  
(23.3.6, DE583000, 33580839)
- **Symptom:** The MUI DataGrid does not include explicit labels assigned for every grid. This causes VoiceOver to read every grid as a generic "Grid" rather than reading it as a specifically named grid.  
**Resolution:** With this fix, all grids now include an aria label that renders along side the role="grid". VoiceOver can now detect the name of the grid when reading the grid, and calls out the entire name of the grid.  
(23.3.6, DE585313)
  - **Symptom:** Baseline trends do not display when a language other than English is selected as default.  
**Resolution:** With this fix, baselines now display for non-english speaking users.  
(23.3.6, DE585480)
  - **Symptom:** The date format in Japanese locale is incorrect. It should be formatted as 1972年1月 instead of 1月 1972.  
**Resolution:** With this fix, the date format for Japanese is now formatted as YYYY MM (ex. 1972年1月).  
(23.3.6, DE585847)
  - **Symptom:** The quick filter and clearing of syncable groups on the VNA import rule page do not work properly.  
**Resolution:** With this fix, the quick filter and the clearing of issues in VNA import rule page now function properly.  
(23.3.6, DE586643)
  - **Symptom:** After the push synchronization stage has sent delete item relationships to all data sources, it cleans the `deleted_item_members` table of those relationships. It does not clean all relationships which can lead to a very large table after some time.  
**Resolution:** With this fix, the delete SQL command now includes all relationship types except for those marked with Excluded Item type. You can view excluded Items from the **Group Administration** page.  
(23.3.6, DE586696, 33608103)
  - **Symptom:** Tool tips are missing on trend charts with sparse (gap) data for percentile metric types.  
**Resolution:** With this fix, percentile metrics now correctly provide tool tips on trend charts.  
(23.3.6, DE587809)
  - **Symptom:** OpenAPI is not encoding colon (:), semicolon (;), and comma (,) characters in group paths. Since these characters are used as special characters to form the group path, OpenAPI might not parse the group path correctly.  
**Resolution:** With this fix, Open API now encodes colon (:), semicolon (;), and comma (,) characters as HTML characters, and now parses the group path correctly.  
(23.3.6, DE587835)
  - **Symptom:** In the data aggregator Administration pages, CTRL+Clicking an element in the left-hand nav (for example, Ctrl+Click **Monitored Items**), the page opens in the same tab (window) instead of in a new tab (window).  
**Resolution:** With this fix, on the data aggregator Administration pages, CTRL+Clicking an element in the left-hand nav (for example, Ctrl+Click on **Monitored Items**) now opens the page in a new tab (window).  
(23.3.6, DE588052)
  - **Symptom:** When upgrading VNA with the DNAC plug-in, a FileNotFoundException is thrown when the Filter.json file does not exist.  
**Resolution:** With this fix, the VNA upgrade now checks for the Filter.json file before upgrading the DNAC engine.  
(23.3.6, DE588137)
  - **Symptom:** Inventory discovery cannot discover some vendor's devices because the devices do not return proper sysUpTime value for the SNMP GetNext query.  
**Resolution:** With this fix, the new System Statistics (Alt) vendor cert uses SNMP Get query to retrieve the sysUpTime value, and inventory discovery can now discover all vendor's devices properly.  
(23.3.6, DE588321, 33629254)
  - **Symptom:** The Telemetry API endpoint takes 12 min to respond.  
**Resolution:** With this fix, Spectrum now optimizes the underlying mechanism to significantly reduce the response time of the Telemetry API (GET/spectrum/restful/telemetry/usageTelemetry). It now uses multi-threading to fetch the required data for the REST API. Users can now experience faster and more responsive interactions with the Telemetry API, leading to an improved overall experience.  
(23.3.6, DE588582, 33608984, 33608984)

### 23.3.5 Fixes

This release includes the following fixed issues:

- **Symptom:** Device type changes in the data aggregator are not always reflected in NetOps Portal.  
**Resolution:** With this fix, device type changes in the data aggregator are now synchronized and reflected in NetOps Portal.  
(23.3.5, DE508178, 32759736)
- **Symptom:** The plug-in global parameters that you remove still persist in the database as stale entries.  
**Resolution:** With this fix, during an upgrade, the parameters that are no longer part of plug-in global config are removed from the database.  
(23.3.5, DE569738)
- **Symptom:** NetOps Portal does not calculate and show Aruba bytes in and bytes out, as these are being calculated from the OC side and the data is not sent correctly sometimes due to timing issue.  
**Resolution:** With this fix, instead of calculating bytes in and bytes out from OC side, the logic has been shifted to NetOps Portal with new panel for the same.  
(23.3.5, DE576451)
- **Symptom:** The `PERFORMANCE_DELTA_TIME` parameter for the Cisco ACI, Viptela, and Versa plug-in have a hard-coded value from the Controller itself. The configure option has been removed for the `PERFORMANCE_DELTA_TIME` parameters for these plug-ins.  
**Resolution:** With this fix, the `PLUGIN_CONFIG` JSON file for the ACI, Versa, and Viptela plug-ins no longer include the `PERFORMANCE_DELTA_TIME` parameter since the controller supports a constant `DELTA_TIME`. The default for `PERFORMANCE_DELTA_TIME` for the ACI plug-in is 5 minutes (300 seconds), for Versa - 15 minutes (900 seconds), for Viptela - 2 minutes (120 seconds).  
(23.3.5, DE576521)
- **Symptom:** For all plug-ins that run subscription, the plug-in subscription state is not shown as part of the plug-in poll status.  
**Resolution:** With this fix, the subscription state is now part of the plug-in poll status.  
(23.3.5, DE577484)
- **Symptom:** The DNAC context page missing the graph for Health Score/Issue Count.  
**Resolution:** With this fix, a trend chart has been added to the DNAC context page to show Health Score/Issue Count.  
(23.3.5, DE577896)
- **Symptom:** On the **Alarm Details** drill-down Business Report, the **Alarm Duration** field does not display correct time format.  
**Resolution:** With this fix, on the Alarm Details drill-down Business Report, the **Alarm Duration** field now shows correct time format.  
(23.3.5, DE579141)
- **Symptom:** Valid flow data without corresponding monitored in/out interfaces are dropped from the query result.  
**Resolution:** Use local identifiers in select queries and resolve to global IDs in the model so rows with unresolved IDs are not dropped.  
(23.3.5, DE580035)
- **Symptom:** Interface Bandwidth Utilization is reported over 100%.  
**Resolution:** To resolve the data inconsistency, the vendor recommends to replace the `realTime` API with the `/aggregate/` API.  
(23.3.5, DE580224, 33561821, 33541773)
- **Symptom:** When running the NetOps Portal service using a port below 1024, like 443, and as a non-root user, NetOps Portal would not be able to bind to the port.  
**Resolution:** Updated the NetOps Portal service `systemd` config file to allow the the service to bind to ports under 1024 when running as a non-root user.  
(23.3.5, DE581769, 33540283)
- **Symptom:** DX NetOps Virtual Network Assurance (VNA) collects max power and power consumption metrics for access point devices, but not visible in NetOps Portal reports.  
**Resolution:** With this fix, the max power and power consumption metrics are now available to NetOps Portal reports.



(23.3.5, DE582786)

- Symptom:** The **View Properties of a Filtered Event** view do not correctly display after changes have made to which event types to display are saved and the dialog is re-opened.

**Resolution:** The **View Properties of a Filtered Event** view display changes made to which events are filtered correctly after the changes are saved and the properties dialog re-opened.

(23.3.5, DE583381)
- Symptom:** The inventory not getting completed for DNAC as License API requests are running continuously.

**Resolution:** There is issue in Query Params of License URL formation with Pagination logic for License retrieval. With this fix, pages are now requested using API properly.

(23.3.5, DE583905, 33576465)
- Symptom:** User gets 500 http status code error after VNA userdomain creation attempt with spaces in naming.

**Resolution:** With this fix, user can create VNA userdomain with spaces in naming (eg. "Test domain 1").

(23.3.5, DE584158)
- Symptom:** OpenAPI would not work correctly if the search string in the filter contains apostrophe.

**Resolution:** With this fix, OpenAPI will correctly if the search string in the filter contains apostrophe.

(23.3.5, DE584172)
- Symptom:** NetOps Portal might log an exception when an item name/alias contains an apostrophe.

**Resolution:** With this fix, NetOps Portal should work as expected when an item name/alias contains an apostrophe.

(23.3.5, DE584178)
- Symptom:** Some chart table views in the ACI and NSX-T dashboards do not print individually (only print as part of an entire page).

**Resolution:** All chart table views in the ACI and NSX-T dashboards print individually.

(23.3.5, DE584217)
- Symptom:** Properties that are prefixed with BROKER:, ZOOKEEPER:, PRODUCER:, and CONSUMER: in the answers file are only effective during the initial installation. After subsequent upgrades, the installer does not retain these custom values in the final properties file. As a result, you can lose your custom set values, and only the default values are present after an upgrade.

**Resolution:** Custom overrides for broker, zookeeper, producer, and consumer properties now persist across future upgrades. You can now specify custom properties in the `<INSTALL_DIR>/config/answers.properties` file, and during an upgrade, the installer writes these custom properties to the respective customization and final properties files. Regardless of other entries in the `answers.properties` file that might generate different values, the custom override always takes precedence.

(23.3.5, DE584371)
- Symptom:** In a secure NetOps Kafka environment, if the user provides custom keystore and truststore files with non-default file extensions, the netops-kafka installation or upgrade process may succeed. However, the SSL health check may fail due to the connection failure between broker, producer and consumer.

**Resolution:** The netops-kafka installation and upgrade process now prompts the user to specify the keystore and truststore types. The input is validated, ensuring that the provided value is one of JKS, PKCS12, or PEM. If no value is provided, the default type considered is PKCS12. Also during a silent install, keystoreType and truststoreType properties are validated from the answers file provided by the user.

(23.3.5, DE584374)
- Symptom:** Using GET `http://{{VNA_HOST}}:8080/vna/rest/v1/inventory/stats` returns "Default\_Domain-1" (domain\_name-domain\_id). Using GET `http://{{VNA_HOST}}:8080/vna/rest/v1/inventory/userdomains` returns "Default Domain" (domain name).

**Resolution:** The logic of filtering inventory stats response from VNA has been changed. Filtering only by Domain ID instead of Domain ID and Domain Name.

(23.3.5, DE585113)
- Symptom:** A `javax.crypto.BadPaddingException` might appear in the `PCService.log` file on startup of the NetOps Portal services.

**Resolution:** With this fix, the `javax.crypto.BadPaddingException` no longer appears in the `PCService.log` file on startup of NetOps Portal.



- (23.3.5, DE585304)
- **Symptom:** The password for the WildFly user is hashed using MD5 algorithm, and FIPS consider it as weak.  
**Resolution:** With this fix, the default hashing that is used for the WildFly password is now based on the policies set for the OS.  
(23.3.5, DE585396)
  - **Symptom:** The VNA interfaces and SNMP interfaces might not be reconciled and fail to receive the VNA performance data on the VNA interfaces.  
**Resolution:** With this fix, the reconciliation process now properly reconciles the VNA and SNMP interfaces and stores the VNA performance data on the reconciled interfaces.  
(23.3.5, DE586016, 33610019)
  - **Symptom:** For interface API responses, a label might not have been specified in the `INTERFACE` parameter for some interfaces. Also, due to API errors, the interfaceMap in the acquisition is not always correct.  
**Resolution:** With this fix, the issue of parsing of interface name in `INTERFACE` parameter for aggregate API interfaces is fixed. The acquisition now includes checks to handle API errors.  
(23.3.5, DE586050)
  - **Symptom:** In AppNeta, you can set multiple tags on a network path, specifically you can set multiple values for the same key/category. But when this data comes over to PM from VNA, only the last value that is associated with a key is displayed.  
**Resolution:** With this fix, the VNA and data collector code now allow multiple values for same AppNeta key/category.  
(23.3.5, DE586181, 10120039, 33610566)
  - **Symptom:** The **Network Paths (Time Bar)** view plots data utilization (inbound capacity utilization) and data loss (outbound and inbound data packet loss) metrics incorrectly.  
**Resolution:** With this fix, the **Network Paths (Time Bar)** view now displays all metrics correctly.  
(23.3.5, DE586555, 33612750, 33617839)
  - **Symptom:** The formula to compute the value for Bits in the F5 LTM Virtual Server Entity vendor cert is missing convert byte to bit, so the report for Bits is incorrect.  
**Resolution:** With this fix, the report for Bits in the F5 LTM Virtual Server Entity vendor cert now shows a correct value.  
(23.3.5, DE586625, 33613697)
  - **Symptom:** For the AppNeta plug-in, Vertica stores one-minute resolution data as 5-minute data with the same 5-minute timestamp, which can result in incorrect threshold evaluations.  
**Resolution:** With this fix, the AppNeta plug-in delta time now uses 1-minute (60 seconds) resolution.  
(23.3.5, DE586866)
  - **Symptom:** Missing data due to data not being available in AppNeta at the time of polling.  
**Resolution:** With this fix, the poll window now retrieves data for the previous 5-minute window instead of the most recent 5-minute window.  
(23.3.5, DE586869)
  - **Symptom:** The AppNeta REST API can deliver multiple records (data points) per network path within a single-minute timeframe. This can result in incorrect threshold evaluation.  
**Resolution:** With this fix, the AppNeta plug-in now only uses the earliest sample within any 1-minute window.  
(23.3.5, DE587097)
  - **Symptom:** Exceptions are thrown while handling poll responses, which results in missing inventory.  
**Resolution:** With this fix, the IP address parser can now handle missing network prefix information.  
(23.3.5, DE587315)
  - **Symptom:** RateLimit Parameters for inventory and performance are not retained post upgrade and issue enrichment API getting 429 errors.  
**Resolution:** With this fix, the Ratelimit parameter values for inventory and performance are now retained post upgrade, and the issue enrichment API is no longer getting errors.  
(23.3.5, DE587457)
  - **Symptom:** On the NetOps Flow dashboard, the Sankey diagram does not display even though there is data coming from the server.

**Resolution:** With this fix, the top conversations are now consolidated according to their hashKeys. This ensure that no more than 20 unique top conversations display in NetOps Portal.  
(23.3.5, DE587831)

- **Symptom:** After a data aggregator restart, the data aggregator sends new/updated VNA items without their VNA gateway UUID, so the consolidation in NetOps Portal does not work. It might appear that there were duplicates in the device, tunnel, slapath, and network path inventories.

**Resolution:** With this fix, data aggregator now initializes the VNA gateway to UUID map on restart, so all VNA items get a VNA gateway UUID synchronized to NetOps Portal and consolidation works correctly.  
(23.3.5, DE588207)

### 23.3.4 Fixes

This release includes the following fixed issues:

- **Symptom:** You cannot create a reference of a group from a non-My Custom Groups folder to a group under My Custom Groups.  
**Resolution:** With this fix, you cannot delete references to a non-My Custom Group folder which were referenced in a group under My Custom Groups.  
(23.3.4, DE369033)
- **Symptom:** If you do not have the Administrator role, but you have the Administer Groups Owned by You and Others or the Administer Groups Owned by You role right, you cannot copy and paste groups.  
**Resolution:** With this fix, if you do not have the Administrator role, but you have the Administer Groups Owned by You and Others or the Administer Groups Owned by You role right, you can now copy and paste groups.  
(23.3.4, DE537719, 33140307)
- **Symptom:** When NetOps Portal monitors VMware NSX-T Data Center 3.2.x environments, poll duration can exceed 4 hours and performance statistics are then missing in NetOps Portal.  
**Resolution:** With this fix, the NSX-T plug-in now correctly handles the NSX-T performance statistics poll ending and performance statistics is now displayed in NetOps Portal.  
(23.3.4, DE561509, 33389728)
- **Symptom:** The Groups API for the Aruba Central plug-in does not support pagination.  
**Resolution:** With this fix, the Groups API now supports pagination.  
(23.3.4, DE567769)
- **Symptom:** When installing the proxy server (DAProxy), if you choose a non-default location for the installation, the installer does not honor the location for the InstallLogs folder, and it is stored in the `/opt/CA/InstallLogs` location.  
**Resolution:** With this fix, the installer now honors the specified installation location for all installed files.  
(23.3.4, DE569942)
- **Symptom:** VNA does not include DB Backup/Restore tracking.  
**Resolution:** With this fix, VNA now includes DB Backup/Restore logs.  
(23.3.4, DE570077)
- **Symptom:** When creating a self-signed certificate by way of the SSL Configuration tool (SslConfig) to enable HTTPS for NetOps Portal and the SSL Configuration tool (sslConfig) to enable HTTPS for the data aggregator, the certificate uses 10 years as the expiration.  
**Resolution:** With this fix, to match the current industry standards, the self-signed certificate is now only valid for 1 year (372 days with 1 week padding).  
(23.3.4, DE572056, 33492919)
- **Symptom:** The file permissions for the `/etc/my.cnf` file are limited to the owner of the file.  
**Resolution:** With this fix, the file now includes explicit read permission through the script.  
(23.3.4, DE572632, 33496643)
- **Symptom:** You cannot access the `irepquery.html` file when using an HTTPS-enabled data aggregator.  
**Resolution:** With this fix, you can now access the `irepquery.html` file when the data aggregator is HTTPS-enabled.

- (23.3.4, DE575852, 33525675)
- **Symptom:** NetOps Portal does not honor the menu orientation when proxying between users using a different theme.  
**Resolution:** With this fix, NetOps Portal shows the menu orientation set for the theme for the active tenant.  
(23.3.4, DE576544)
  - **Symptom:** From the threshold profile **Link Event Rules** dialog, you can link to an event rule that is already linked to it, creating a bi-directional association.  
**Resolution:** With this fix, from the threshold profile **Link Event Rules** dialog, you cannot link to an event rule that is already linked to it.  
(23.3.4, DE577169)
  - **Symptom:** On the **Engine Stats** context tab, the Total Messages and Total Subscription Messages in the VNA Engine Table has formatting issues while displaying.  
**Resolution:** With this fix, the Total Messages and Total Subscription Messages now properly display on the **Engine Stats** context tab.  
(23.3.4, DE577488)
  - **Symptom:** You can create groups (subgroups) under groups to which you do not have access to administer. The pasted groups continue to be owned by the copied group's owner. You cannot copy groups from a non-custom group folder to custom groups.  
**Resolution:** With this fix, you can create groups (subgroups) only under those groups to which you have access to administer. The owner for the pasted groups now changes to the current user. You can now copy groups from a non-custom group folder to custom groups.  
(23.3.4, DE578098)
  - **Symptom:** With more and more users enabling HTTPS and the upgrade to Jetty 10, there are more instances of SNI errors being seen after upgrades. This is due to the value of the **Web Site Host** and **Web Service Host** properties not being in the HTTPS certificates. One reason is that the migration and disaster recovery script uses IPs instead of FQHNs.  
**Resolution:** With this fix, the `update_pc_da_database_references.sh` script now requires the FQHN hostname instead of an IP address. In addition, it now updates the four service properties files that include the hostname for NetOps Portal and the event manager to use the newly-provided FQHN.  
(23.3.4, DE578921, 33550366)
  - **Symptom:** An Internal Server error is seen on Inventory/Stats endpoint in VNA.  
**Resolution:** With this fix, you can now delete the USER DOMAIN in VNA.  
(23.3.4, DE579894)
  - **Symptom:** The Aruba Central plug-in does not process/publish access point (AP) alarms.  
**Resolution:** Streaming is not enabled in Aruba Central for specific alerts. After enabling streaming, Aruba Central can now receive AP Alerts through streaming. With this fix, the Aruba Central plug-in now processes/publishes AP alarms.  
(23.3.4, DE580128)
  - **Symptom:** The logs (`server.log`) are showing warnings with regard to Hibernate validator when VNA is running.  
**Resolution:** With this fix, these warnings are now avoided.  
(23.3.4, DE580259)
  - **Symptom:** Under some conditions, NetOps Portal might not sync Event Manager inventory.  
**Resolution:** With this fix, NetOps Portal now syncs Event Manager inventory as expected.  
(23.3.4, DE580439, 33562375)
  - **Symptom:** After running for a long period of time, the `netops-da-graphql` server might become unresponsive.  
**Resolution:** With this fix, the Spring Boot's default logging management, Logback, has been replaced with Log4j2.  
(23.3.4, DE581437)
  - **Symptom:** The **Manage Report Resolution Settings** page includes two lists of fields, one for users without the Run Dashboards At Higher Resolution role right and the other for users with the Run Dashboards At Higher Resolution role right. The same label is displayed over each list.  
**Resolution:** With this fix, the label for the list of fields on the right side of the page now correctly displays that the fields are for users with the Run Dashboards At Higher Resolution role right.

(23.3.4, DE581868, 33574177)

- **Symptom:** The **Aggregated Component** page member's table does not sort correctly by the name alias.  
**Resolution:** With this fix, the **Aggregated Component** page member's table now sorts correctly by the name alias.  
(23.3.4, DE582234)
- **Symptom:** When creating import rules that define to which IP domain the data aggregator maps groups of VNA inventory, the full group hierarchy is not displayed in the **Source Data** section, resulting in a flattened list of groups in some cases.  
**Resolution:** With this fix, the full group hierarchy is displayed.  
(23.3.4, DE582236)
- **Symptom:** The upgrade handles only if the proxy server and the data aggregators are all HTTPS or HTTPS-enabled.  
**Resolution:** With this fix, the upgrade now handles if the proxy server is HTTPS-enabled and the data aggregators are HTTP.  
(23.3.4, DE582396)
- **Symptom:** On the **Details** tab for a manageable device (**Administration, Monitored Items Management, Monitored Devices**), the **SNMP Set Profile** drop-down does not display SNMP profiles that can be used for a set operation, and you cannot select one to assign to the device.  
**Resolution:** With this fix, you can now select an SNMP profile from the drop-down.  
(23.3.4, DE582747, 33582317)
- **Symptom:** The `daproxy.toml` file address is empty and the URL is incorrect.  
**Resolution:** With this fix, the `daproxy.toml` file address and URL look as expected.  
(23.3.4, DE582862)
- **Symptom:** SDN devices might lose relation to parent groups, and subsequent changes to their SDN domain rule mappings are not applied to the device.  
**Resolution:** With this fix, deleted relations that are sent from NetOps Portal that have the `SDN_RELATION` facet are now ignored.  
(23.3.4, DE583291)
- **Symptom:** The AWS plug-in has compilation errors after an upgrade to 23.3.3.  
**Resolution:** With this fix, the AWS plug-in does not include this import statement.  
(23.3.4, DE583504)
- **Symptom:** The title for the Access point table on the **Wi-Fi Health** page is overridden with a page title. The **Clients** table on the **Clients** context tab (on the AP context page) does not come up when view suppression is on for the client table model.  
**Resolution:** The title for the Access point table on the **Wi-Fi Health** page is no longer overridden. View suppression is now off for the client table model, and is now always displayed.  
(23.3.4, DE583640)

### 23.3.3 Fixes

This release includes the following fixed issues:

- **Symptom:** The `hosteval` script that runs the installers depends on the `bc` package being installed and fails on systems without the package.  
**Resolution:** With this fix, the `hosteval` script now no longer depends on the `bc` package. In addition, the config files for `hosteval` now includes the required packages.  
(23.3.3, DE563214)
- **Symptom:** The NSX-T Manager nodes are not displayed in the list of devices when an NSX-T domain group is selected on the **NSX-T Inventory** page, and you click a device name (the link) in NetOps Portal.  
**Resolution:** With this fix, the NSX-T Manager nodes are now displayed under the NSX-T domain group.  
**Workaround:** After upgrading to 23.3.3 or higher, do a full data aggregator re-sync on NetOps Portal side to solve the issue.

- (23.3.3, DE566815)
- Symptom:** CA Remote Engineer has been a Python 3.6+ script since 22.2.8, but `hosteval` does not check for python 3.6+ is installed.  
**Resolution:** The `hosteval` config files now look for Python 3.6/3.8//3.9/3.11 depending on the base operating system availability.  
 (23.3.3, DE568140, 33457953)
  - Symptom:** The Cisco ACI severity mappings are not correct in VNA and Spectrum.  
**Resolution:** With this fix, based on the LC value from Notification attributes, the severity and state are now set as suggested.  
 (23.3.3, DE571060, 33484175)
  - Symptom:** When the data aggregator is very slow to respond, the system status background processes used to monitor different aspects of the data aggregator hang. This also causes the menu bar generation to hang when new users log into the system, as it determines the overall system status icon state.  
**Resolution:** With this fix, the system status icon color determination now uses the current cached value for new logins. Subsequent page loads now reflect any changes to system status overall state.  
 (23.3.3, DE572541, 33496885)
  - Symptom:** Event rules that use the **Aggregate Components by Device** option for the **Aggregation** field can incorrectly define the conditions that raise a threshold violation.  
**Resolution:** With this fix, the threshold processing logic now takes into account all applicable device component data when processing event rules that use the **Aggregate Components by Device** option for the **Aggregation** field.  
 (23.3.3, DE573898, 33443671, 33513387)
  - Symptom:** The OData v4 interface responds to all queries with Invalid tenant id: 0.  
**Resolution:** With this fix, the OData v4 interface now works as expected.  
 (23.3.3, DE576562)
  - Symptom:** Changing the **Flow Direction** or **SNMP Polling** field values on the **Edit Router(s)** page (to edit the device) changes the **SNMP Profile** field value to the default value (public), and the device stops collecting data.  
**Resolution:** With this fix, changing the field values no longer affects the **SNMP Profile** field value.  
 (23.3.3, DE576592, 33531237)
  - Symptom:** In large SNMPv3 environments, polling can slow down and back up due to how fast SNMP4j MPv3 cache is used. It is using an inefficient search algorithm.  
**Resolution:** With this fix, the MPv3 cache now includes an additional hash map that uses a different key, and the lookups are now faster. MPv3 cache no longer slows down polling.  
 (23.3.3, DE577202, 33529643)
  - Symptom:** The client count for DNAC Access Points (APs) is showing incorrectly, which is actually the radio count for an AP.  
**Resolution:** With this fix, the code now calculates the Clients per AP and report the same.  
 (23.3.3, DE577253)
  - Symptom:** The **ISSUE\_PRIORITY\_LIST** parameter has a spacing issue and this causes the URL invocation to fail.  
**Resolution:** With this fix, the **ISSUE\_PRIORITY\_LIST** parameter no longer has a spacing issue and the URL does not have the spacing and 505 issues.  
 (23.3.3, DE577398, 33538288)
  - Symptom:** VNA pulls only 500 devices for DNAC.  
**Resolution:** With this fix, pagination has been implemented for the Cisco DNAC plug-in as there is a limit for each API. Post pagination issues is resolved.  
 (23.3.3, DE578028, 33540054)
  - Symptom:** Expected switch port index value is Long but they are string values, and it fails to persist.  
**Resolution:** With this fix, there is validation for switch port index before persistence.  
 (23.3.3, DE578583, 33542344)
  - Symptom:** There is an error in the `gateway.log` with every poll (every 15 min):

```
2023-09-25 12:50:02,982 ERROR (Camel (SilverPeak Plugin-22.2.10-RELEASE) thread #56 -
vm://global/plugin/silverpeak/INVENTORY) [PLUGIN_SYSTEM] CamelLogger 205 Failed delivery
for (MessageId: queue_InventoryUpdateQueue_ID_91872f10-5bc3-11ee-b014-005056a7afad on
ExchangeId: B586B0A65F8B304-00000000000000003). Exhausted after delivery attempt: 1 caught:
java.lang.ArrayIndexOutOfBoundsException: Index 1 out of bounds for length 1 Message History (source
location and message history is disabled)
```

**Resolution:** With this fix, the parsing functionality now works for tunnel alias names.

(23.3.3, DE579087, 33541773)

- **Symptom:** The **Select Discovery Instance** list in the **Discovery History** dialog does not display anything if the user preferences for the date format was set to Short, 24 hour clock.

**Resolution:** With this fix, the list now shows all instances regardless of the date format in force.

(23.3.3, DE579747)

- **Symptom:** In the **Report Name** column on the **Manage Scheduled Reports** page, hyperlinks to reports that are not available appear. When you click the hyperlink, you get an error attempting to traverse to page:

```
"You do not have sufficient rights to perform this action"
```

**Resolution:** With this fix, if the reports are not available, a hyperlinks no longer appears in that column.

(23.3.3, DE579765)

- **Symptom:** Multiple notification updates for a time filter are not forwarded due to an exception in VNA.

**Resolution:** With this fix, the exception no longer occurs. The updates are forwarded and are no longer dropped.

(23.3.3, DE580015)

- **Symptom:** For the default configuration of the Aruba Central plug-in, the following parameters are false by default: **MONITORING\_STREAMING\_ENABLED**, **STREAM\_AP\_STATE\_EVENTS**, **STREAM\_CLIENT\_STATE\_EVENTS**, **POLL\_CLIENTS**.

**Resolution:** With this fix, these parameters are now true by default out of the box for new plug-in configuration.

(23.3.3, DE580125)

- **Symptom:** The `events`, `event_properties`, `unresolved_event_items`, `events_l10n`, and/or `event_identifiers` tables in the NetOps Portal Event Manager (em) database might contain events that occurred on a date older than the configured event retention period (default 30 days). This might be due to a failure in the nightly call to the `em.rotate_event_partitions` stored procedure and the failure will not be visible in the `EMService.log` or the DB error log. If the `em.rotate_event_partitions` stored procedure is manually executed, the following error is seen:

```
ERROR 1267 (HY000) at line 65: Illegal mix of collations (utf8_general_ci,IMPLICIT) and
(utf8_tolower_ci,IMPLICIT) for operation '='
```

**Resolution:** With this fix, The `em.rotate_event_partitions` stored procedure now forces a matching collation in the offending statement.

(23.3.3, DE580613, 33536611)

### 23.3.2 Fixes

This release includes the following fixed issues:

- **Symptom:** If NetOps Portal is performing a full synchronization of one data source and is performing an incremental synchronization of the other data sources, some Single Sign-On Configuration tool (SsoConfig)-modified settings might not be sent to the data sources for which NetOps Portal is performing incremental synchronization. The full synchronization sets the last updated timestamp on the setting to 2, and the incremental synchronization does not pick up that the settings must be synced.

**Resolution:** With this fix, the last updated timestamp when performing a full synchronization now uses the current timestamp. Incremental synchronizations now see the changes, and synchronizes them.

(23.3.2, DE570703, 33476070)

- **Symptom:** VNA sends the CIDR-formatted device IP address for Meraki events instead of the network.

**Resolution:** With this fix, VNA sends the CIDR-formatted network for Meraki events.



- (23.3.2, DE571803, 33491134)
- Symptom:** During a fresh installation and when creating a user domain in 22.2.11 and higher, the user-domain-creation logic sets the `datasource` context value for the user domain as a single digit.  
**Resolution:** With this fix, the logic now sets the context value using the `USERDOMAIN:<ID>` naming pattern. The upgrade script now updates the naming of existing user domains to `USERDOMAIN:<ID>`.  
 (23.3.2, DE575341)
  - Symptom:** IM Table, IM Filtered Table, and charts with table legends do not display their column sort indicators on initial display.  
**Resolution:** With this fix, IM Table, IM Filtered Table, and charts with table legends now display their column sort indicators.  
 (23.3.2, DE575389)
  - Symptom:** If SNMP-vendor information is not available, the vendor information for plug-ins that monitor Wi-Fi devices, such as the Cisco Meraki and Aruba Central plug-ins, and for the Cisco DNAC plug-in does not display in the **Wi-Fi Health** dashboard in NetOps Portal.  
**Resolution:** With this fix, if all the vendor information now displays in the **Wi-Fi Health** dashboard in NetOps Portal.  
 (23.3.2, DE575483)
  - Symptom:** The rate at which the Aruba Central plug-in collects inventory data (the `INVENTORY_POLL_RATE` parameter) is set to every 12 hours, instead of every hour.  
**Resolution:** The Aruba AP interface inventory collection API does not support bulk API and no longer supports AP interfaces. With this fix, the parameter is set to every hour.  
 (23.3.2, DE575586)
  - Symptom:** Radio Trends charts need to show Baseline-Average metrics trends as well.  
**Resolution:** With this fix, the charts now use the Multitrendchart model with custom properties and OnDemandChart properties. This model shows multiple charts in a single model and also shows baseline metrics.  
 (23.3.2, DE575594)
  - Symptom:** The fields on the **Manage Email Server Settings** page display incorrectly.  
**Resolution:** With this fix, the fields now display with the correct layout.  
 (23.3.2, DE575842)
  - Symptom:** The data collector Karaf logs fill up and rollover with `No SNMP profile was found for device` ERROR messages. This message appears for every component and every poll.  
**Resolution:** With this fix, the data collector now logs a single message for each IP indicating that the data collector could not find an SNMP profile for the device.  
 (23.3.2, DE575943)
  - Symptom:** The notification poller updates the last sync time for notifications even after API request failure, and it uses the same file for issues and events.  
**Resolution:** With this fix, the notification poller now writes the last sync time only after a successful request. A separate lastsynctime file has been introduced for issues and events.  
 (23.3.2, DE575981)
  - Symptom:** NetOps Portal does not properly apply seconds scale factor to goal trend lines with milliseconds values for those System Health trend chart views that show goal trend lines.  
**Resolution:** With this fix, NetOps Portal now properly applies the seconds scale factor to goal trend lines with milliseconds values.  
 (23.3.2, DE576277)
  - Symptom:** Clicking a threshold profile pollutes the NetOps Portal log with `Unable to find event profile` `CRUD dilaog` error messages.  
**Resolution:** With this fix, the log no longer contains these error messages.  
 (23.3.2, DE576381)
  - Symptom:** The Cisco Meraki plug-in groups Meraki inventory and performance polls into multiple poll categories. There is a mismatch of poll IDs, and self monitoring does not pick up the poll statuses.  
**Resolution:** With this fix, the `EngineStartupConfig.xml` file now handles poll name matching, and the poll-xml file now handles the poll IDs.

(23.3.2, DE576658)

- **Symptom:** If you create a monitoring profile filter using the `AdminStatus` and `OperStatus` attributes in the Alternate Interface or Integrated Adaptive Rate DSL metric families, the data aggregator `karaf.log` can fill with `Invalid enumeration value` errors, which can result in the data aggregator not polling the components of the two metric families.

**Resolution:** With this fix, the enumerations in the metric families are now correct, the monitoring profile filter now processes properly, and the data aggregator now polls the components accordingly.

(23.3.2, DE576887, 33531873, 33531873)

- **Symptom:** When creating or editing user settings and scheduling reports, the email address validator restricts potentially-supported special characters.

**Resolution:** With this fix, when creating or editing user settings and scheduling reports, you can include special characters in the email address.

(23.3.2, DE577191, 33536392)

- **Symptom:** VNA upgrades can fail with permission issues to the `keystore.jks` file.

**Resolution:** With this fix, the upgrade can now handle permissions. It ensures that the WildFly user has access to the `jks` file.

(23.3.2, DE577525, 33529299)

- **Symptom:** The properties in the **Event Details** log in NetOps Portal for linked event rules are missing or are empty.

**Resolution:** With this fix, content now displays in the properties in the **Event Details** dialog.

(23.3.2, DE577733)

- **Symptom:** When an SNMP profile is no longer valid for a device, the device's status becomes "Management Lost". A few minutes after you manually change the device's SNMP profile, the profile is reverted to the previous one.

**Resolution:** With this fix, the SNMP profile re-discovery process now does not replace user-manually-changed SNMP profiles with the previous SNMP profile.

(23.3.2, DE577874, 33542076)

- **Symptom:** When selecting group filters in a context page, then go to a different context page, then selecting other group filters from the group filter, the group filter breadcrumbs are not updated (unless the page was refreshed and the group filters were changed thereafter).

**Resolution:** With this fix, the page client code now listens to all server updates. Now, group breadcrumb updates to reflect the currently-selected group.

(23.3.2, DE578041)

- **Symptom:** Sometimes when you first click **Next Page** on table grids, the view reloads but does not properly set page navigation footer.

**Resolution:** With this fix, when you click **Next Page** on table grids, the view reloads and the page navigation footer is properly set.

(23.3.2, DE578419)

- **Symptom:** Duplicate alarms in the `alarm_clear` table can cause orphan processing to not clean up event/alarms for deleted items, deleted rules, and threshold profiles.

**Resolution:** With this fix, orphan processing now cleans up events/alarms for deleted items, deleted rules, and threshold profiles.

(23.3.2, DE578582)

- **Symptom:** With configured linked event rules, orphan processing might not remove ongoing events for deleted items, profiles, and rules, and logs a `NullPointerException` (NPE) in the log file.

**Resolution:** With this fix, orphan processing now removes ongoing events for deleted items, profiles, and rules.

(23.3.2, DE578618)

### 23.3.1 Fixes

This release includes the following fixed issues:

- **Symptom:** The tokenstore file contains a refresh token, and this file takes precedence for token generation. This causes the update plug-in (updating tokens) through REST APIs to not use the updated tokens.



- Resolution:** With this fix, the tokenstore file in the `work` directory is now deleted whenever the Aruba Central plug-in configuration is updated through the REST call.  
(23.3.1, DE558386)
- **Symptom:** There is a vulnerability in ActiveMQ.  
**Resolution:** ActiveMQ has been upgraded to 5.18.0. This is the latest version to date.  
(23.3.1, DE569011, 33464312)
  - **Symptom:** A Null Pointer exception comes in gateway logs sometimes for Meraki radio metrics.  
**Resolution:** The Meraki Response itself does not have the metric entry, which causes the issue. With this fix, the code has been updated.  
(23.3.1, DE569072)
  - **Symptom:** For some Meraki IPDevices, the vendor name is not populated as 'Meraki'.  
**Resolution:** With this fix, the vendor name is now populated as 'Meraki' for Meraki IPDevices.  
(23.3.1, DE569606)
  - **Symptom:** VNA self-monitoring events do not have unique notification IDs and names.  
**Resolution:** With this fix, VNA self-monitoring events now have unique notification IDs.  
(23.3.1, DE570069)
  - **Symptom:** You cannot edit rules where a Spectrum global collection has been deleted and only a group item ID is displayed in the rule editor table.  
**Resolution:** With this fix, you can now edit rules where a Spectrum global collection has been deleted. The group is now set to None in the rule editor dialog.  
(23.3.1, DE570685, 33474998, 33516298)
  - **Symptom:** Changing the vendor certification priority for a metric family can trigger massive metric family discoveries on many devices and then generate massive database operations for the changed components. At some point when the database cannot handle so many operations, a ROS exception can be thrown and cause the metric family to hang on some devices. The hanging metric family discovery stops the change detection of the metric family on the devices, and the components are not updated.  
**Resolution:** With this fix, to ensure that you can run new metric family discovery to have NetOps Portal update the components properly, metric family discovery now aborts the discovery when a discovery has been in running status for more than six hours.  
(23.3.1, DE571032, 33477350)
  - **Symptom:** The events that the Event Manager exports to Kafka do not contain the Severity value from the events table.  
**Resolution:** With this fix, the events that the Event Manager exports to Kafka now contain the Severity value from the events table.  
(23.3.1, DE571421)
  - **Symptom:** In some configurations, OpenAPI query cannot get availability and reachability metric family data for devices, for example, `/odata/api/devices?$expand=availabilitymfs,reachabilitymfs`.  
**Resolution:** With this fix, OpenAPI cannot get availability and reachability metric family data for devices in any configuration.  
(23.3.1, DE571700, 33483806)
  - **Symptom:** WildFly CPU Utilization for VNA self-monitoring is incorrect.  
**Resolution:** With this fix, WildFly CPU Utilization for VNA self-monitoring is now correct.  
(23.3.1, DE572445)
  - **Symptom:** Goal line trend lines might be missing from charts based on trend processing when scaling axis values.  
**Resolution:** With this fix, goal line trend lines are now on charts based on trend processing when scaling axis values.  
(23.3.1, DE573132, 33499889, 33487400)
  - **Symptom:** The value displayed for Mode in the **Radio Details** table in NetOps Portal is not readable.  
**Resolution:** With this fix, the value displayed for Mode in the **Radio Details** table in NetOps Portal is now Access Point, Air Monitor, or Spectrum Monitor.

- (23.3.1, DE573185)
- **Symptom:** For Viptela plug-in configuration for DXI onboarding, quickly discovering virtual network devices does not configure the Viptela plug-in.  
**Resolution:** With this fix, you can now use DXI onboarding to configure the Viptela plug-in.  
(23.3.1, DE573259)
  - **Symptom:** If you access discovery history for a discovery profile while a discovery is in progress, the discovery history dialog "hangs" loading data.  
**Resolution:** With this fix, you can no longer access discovery history while a discovery is in progress for a discovery profile instance.  
(23.3.1, DE573348)
  - **Symptom:** Data aggregator REST calls hang while executing large deletes.  
**Resolution:** With this fix, data aggregator REST calls now time out after five minutes if the item being read is in the process of being deleted.  
(23.3.1, DE573648, 33504902)
  - **Symptom:** Count of processed linked rule evaluations on self monitoring event page shows the number of DB queries instead of number of linked rule evaluations.  
**Resolution:** With this fix, the count of processed linked rule evaluations on self monitoring event page will show number of linked rule evaluations correctly.  
(23.3.1, DE573654)
  - **Symptom:** Customizable card views do not work when configured with aggregated components at the item level.  
**Resolution:** With this fix, customizable card views now work when configured with aggregated components at the item level.  
(23.3.1, DE573777, 33502421)
  - **Symptom:** The footer on most pages in NetOps Portal does not display in the correct position or z-order. This is especially pronounced in the CA-White UI theme.  
**Resolution:** With this fix, the footer displays in the correct position and z-order in all pages.  
(23.3.1, DE573942)
  - **Symptom:** Several instances of NetOps are incorrectly cased within NetOps Portal as Netops.  
**Resolution:** With this fix, all references to NetOps now have the proper casing.  
(23.3.1, DE573975)
  - **Symptom:** If the AppNeta plug-in is configured, you might see `FileNotFoundException` exception errors in the `oc.log` file.  
**Resolution:** With this fix, the AppNeta plug-in no longer generates `FileNotFoundException` exceptions.  
(23.3.1, DE574060)
  - **Symptom:** Customizable card views do not work when configured with aggregated components at the item level.  
**Resolution:** With this fix, customizable card views now work when configured with aggregated components at the item level.  
(23.3.1, DE574720)
  - **Symptom:** Quick Filter does not work with AppNeta devices when searching the **Monitoring Point** or **Monitored Target** fields.  
**Resolution:** With this fix, you can now Quick Filter AppNeta devices by searching the **Monitoring Point** or **Monitored Target** fields.  
(23.3.1, DE575033)
  - **Symptom:** When global synchronization runs for devices, it fails to break apart devices into their own itemid if they come from the same data source and have the same IP address.  
**Resolution:** With this fix, the global synchronization for devices now breaks apart existing devices that come from one data source and have the same IP address. All the devices are now visible in the Inventory page for devices.  
(23.3.1, DE575156, 33511955)
  - **Symptom:** Data collector-specific system health charts might not contain data for one or more data collectors.  
**Resolution:** With this fix, the modified data aggregator self-monitoring polling configuration now fully completes before it allows another thread to modify the same polling configurations.

(23.3.1, DE575233, 33488559)

- **Symptom:** Kafka event publication is configured in the Event Manager service and the Event Manager process shows as started, but the Event Manager service is not publishing events to a Kafka topic and the Event Manager data source is not available. The `EMService.log` file shows a message like the following during startup:

```
Caused by: org.apache.kafka.common.config.ConfigException: No resolvable bootstrap urls given in
bootstrap.servers
```

**Resolution:** This problem is due to manual misconfiguration of the `em.properties` event manager file or a DNS failure. With this fix, the exception is caught and logged so that the Event Manager service can continue to run. The Event Manager Publication Service now shows as Failed with a status description stating loss of contact with Kafka. The administrator must ensure that the bootstrap servers are properly defined in the `em.properties` file and that the bootstrap server hosts are accessible from the NetOps Portal server by way of the provided hostnames, and then restart the Event Manager service.

(23.3.1, DE575379)

- **Symptom:** The MUI Datagrid utils does not return the expected results after package version updates.  
**Resolution:** With this fix, the MUI Datagrid library is updated to version 6.11.2 which is the version compatible with other library updates.

(23.3.1, DE575413)

- **Symptom:** Customizable views do not work when configured with aggregated components at the item level.  
**Resolution:** With this fix, customizable views now work when configured with aggregated components at the item level.

(23.3.1, DE575695)

- **Symptom:** Client metrics data appears inconsistent with clients table data.  
**Resolution:** With this fix, client metric calculation is now based on number of relations (`IS_CONNECTED_TO`) for an access point.

(23.3.1, DE575888)

- **Symptom:** When configuring the data aggregator to use HTTPS using a self-signed certificate and you run the SSL Configuration tool for the data aggregator (the `sslConfig.sh` script) on the data aggregator, the script complains about not finding the `truststore.jks` file, and then fails to configure HTTPS fully.

**Resolution:** With this fix, the `sslConfig.sh` script now builds the `truststore.jks` file, and then completes the HTTPS setup.

(23.3.1, DE577060)

## Known Issues

Review the list of known issues in this release.

The following known issues have been identified in this release of DX NetOps Performance Management:

- [Alarms Tab on Customized Context Pages](#)
- [Apostrophes in Custom Attribute Descriptions](#)
- [Broken Links in Integrations with Spectrum](#)
- [Disabled Trend Charts with Events](#)
- [Extended DX NetOps Network Flow Analysis Views](#)
- [Flow Administration Issues with IPv6](#)
- [Group Scorecard Table Metric Fields](#)
- [Invalid Certificates](#)
- [NFA Drill-Down Page Searches](#)
- [Revert View Settings Limitations](#)
- [Screen Reader Limitations](#)
- [SDN/NFV Dashboard Limitations](#)
- [SD-WAN Monitoring Limitations](#)
- [Viptela Inventory Includes Only Devices with Valid Certificates](#)
- [Virtual and SNMP Interface Reconciliation](#)
- [Wi-Fi Device Visibility with VNA and Spectrum Integrations](#)
- [Consul Port for Fault Tolerant Data Aggregators with Consul as HTTPS for Upgrades to 23.3.4 and Higher](#)

### **Alarms Tab on Customized Context Pages**

The **Alarms** tab is automatically added to most device-level context pages. For customized device-level context pages, you must manually add the **Alarms** tab. If you have many tenants with customized device-level context pages, see the ReadMe file in the following location to streamline the addition of the **Alarms** tab:

```
<installation_directory>/PerformanceCenter/SQL/plugins/custom_context_spectrum
```

- ***installation\_directory***  
The default installation directory for NetOps Portal.  
**Default:** /opt/CA

### **Apostrophes in Custom Attribute Descriptions**

When a custom attribute description contains an apostrophe, the tooltip for the attribute does not render correctly in some situations.

### **Broken Links in Integrations with Spectrum**

#### **Issue:**

If you have integrated with DX NetOps Spectrum (Spectrum), you might find broken links from Spectrum OneClick to NetOps Portal.

#### **Workaround:**

Complete the following steps:

1. Specify the fully-qualified domain name of the NetOps Portal host for all access to the NetOps Portal website.  
For more information about how to specify this host, see [Configure the Basic Security Settings Using the SSO Configuration Tool](#).
2. Synchronize the Spectrum data source.  
For more information, see [Synchronize Data Sources](#).

### **Disabled Trend Charts with Events**

Trend charts with events are disabled when the time range is greater than three months.

#### **Device Context Page:**

- Availability Trend with Events
- Average CPU Utilization Trend with Events
- Average Memory Utilization Trend with Events

#### **Interface Context Page:**

- Interface Utilization/Discard Out Trend with Events
- Interface Utilization/Discard In Trend with Events
- Interface Utilization Out Trend/Baseline Detail with Events
- Interface Utilization In Trend/Baseline Detail with Events

### **Extended DX NetOps Network Flow Analysis Views**

Extended DX NetOps Network Flow Analysis views have the following known limitations:

- The search box on related table views does not filter the content in the data results.
- You cannot export the extended flow views on the interface context page in a scheduled report.
- You cannot export the extended flow views on the dashboard page in a scheduled report containing all pages.
- The maximum rows are limited for performance. By default, the **Top Selection** table views are 1000 rows and the **Trend Table** views are 25 rows.
- For upgrades only, you must manually add extended flow views to customized interface context pages. If you have many tenants with customized interface context pages, do the following post-installation steps:
  - a. On the NetOps Portal server, change to the following directory by issuing the following command:
 

```
cd <installation_directory>/PerformanceCenter/SQL/plugins/reporter
```

    - **installation\_directory**  
The default installation directory for NetOps Portal.  
**Default:** /opt/CA
  - b. Follow the steps described in the `ReadMe` file.
- The links in the charts and tables in the **IP Performance** page open the corresponding network flow pages within NetOps Portal.

For more information about the DX NetOps Network Flow Analysis interface context views, see [Network Flow Analysis Views in NetOps Portal](#).

### **DX NetOps Network Flow Analysis Administration Issues with IPv6**

If the pages for flow administration are empty, check the `PCService.log` file.

If the following message appears, the DX NetOps Network Flow Analysis console server might be configured with an IPv6 address:

```
Failed in sending REST: http://[0000:0000:0000:0000:0000:0000:0000:0001]:8981/odata/api/AuthToken
java.net.ConnectException
```

#### **Follow these steps:**

1. [Disable the IPv6 address](#).
2. Ensure that only one Network Interface Card (NIC) is present on the DX NetOps Network Flow Analysis Console server.
3. Issue the following command and verify that there are no IPv6 addresses:

ipconfig

For more information about flow administration, see [Flow Administration](#).

### **Group Scorecard Table Metric Fields**

The first time that you edit the view settings for a group scorecard table view, you might encounter an issue with the selected metrics. If you change between Hierarchy Calculate Levels and Metric Calculate Levels, your selected metric fields might clear. For example, if you select several metric fields, and then change the selection from **Device Hierarchy** to **by Device**, your selected metric fields clear.

### **Invalid Certificates**

For the first discovery, invalid certificates might cause devices, interfaces, and tunnels to be missing. After discovery, invalid certificates remove all items that belong to any invalidated devices (routers, interfaces, tunnels, and possibly sites).

### **NFA Drill-Down Page Searches**

Search on the name or description fields. When searching the hosts, provide the host name without braces. For example, if the host name is 10.10.1.2, search the host name using 10.10.1.2. For conversation searches, use the Server IP or the Client IP. For example, when searching for the conversation 10.19.19.179 - 10.20.20.179, use the search value 10.19.19.179 or 10.20.20.179. Search does not work using 10.19.19.179 - 10.20.20.179. For ToS searches, search without parentheses () of ToS description. For example, if the ToS description is ToS 20 (ToS 20), use the search value ToS20. For a ToS format like AF12 (DSCP12) ECT=0;CE=1 (ToS 49), use the search value AF12 (DSCP12) ECT=0.

Search does not work on 'other' as it is group of items. Protocol search does not work on the group protocols (ip,ipv6,tcp,udp) as they are not single protocols.

### **Revert View Settings Limitations**

You cannot convert views to reports at a group level if a context item filter is applied when adding on-demand reports or dynamic trend views to a page within the Dashboard Builder.

If you are editing a view that is locked to a device or interface that no longer exists, avoid reverting the view settings (by clicking **Use Defaults**) at the All Tenant Users level.

If you revert the view settings, do the following tasks:

1. Restore the metric selection by editing the view twice.
2. Add a new device or interface to the view before rendering data.

### **Screen Reader Limitations**

When using screen readers to read the contents of charts, you cannot use the up and down arrow keys. Usually, you can move focus in the chart using the screen reader key in conjunction with the left and right arrow keys.

### **SDN/NFV Dashboard Limitations**

The **VNF Count by Type** stacked chart in the **SDN/NFV Virtual Inventory Overview** dashboard does not show data for the last period. The data for the same period appears when you export the data to a CSV file.

For more information about this dashboard, see [Monitor SDN/NFV Virtual Inventory](#).

### **SD-WAN Monitoring Limitations**

The following limitations apply to SD-WAN monitoring:

- NetOps Portal processes SD-WAN tunnel reporting at the device component level. Parent devices do not appear in the reported raw data. The edge device source, and destination, can be manually added to the tabular related reports. Pick the inventory columns on the rendered grid.
- Direct reconciliation of the SD-WAN edge devices to SNMP physical devices does not occur. Therefore, reporting for the CPU and Memory utilization of edge devices is done from the virtual host metric family.
- Direct reconciliation of virtual interfaces to SNMP physical interfaces does not occur.
- The parented edge device of the tunnel must report the virtual interface of the SD-WAN tunnels. Direct queries by the tunnel item are unsupported.
- The VNA Domain Sites are initially shown as numeric values. You can alter the site group using the Group Admin UI.
- When including metrics to display in on-demand report templates, for the SD-WAN metric families (Tunnels and Application Paths), percentile and projection metrics are available, but are unsupported and the views render incorrectly in the report template.  
For more information about how to include metrics to display in on-demand report templates, see [Manage On-Demand Report Templates](#).
- If you set a custom time range to 30 minutes on an SD-WAN dashboard, the time bar charts and trend views have data gaps.

For more information, see [Monitor SD-WAN Devices](#).

### **Viptela Inventory Includes Only Devices with Valid Certificates**

The Viptela plug-in can accurately interpret vEdge devices that have valid certificate states. If the certificate of a vEdge device is invalidated, the Viptela plug-in no longer discovers the vEdge device in the inventory. It no longer reports associated interfaces, however it continues to report tunnels and references the interface underlay.

vEdge devices must have valid security certificates to participate in Viptela networks. You can configure the certificate state from the vManage Certificates administration page.

For more information about how to administer a certificate for a vEdge device, see [the Cisco documentation](#).

### **Virtual and SNMP Interface Reconciliation**

DX NetOps Performance Management reconciles interfaces from VNA with existing SNMP interfaces. When the VNA Gateway restarts (for example, during an upgrade), DX NetOps Performance Management reconciles the virtual interfaces with existing SNMP interfaces, and the virtual interface is deactivated. DX NetOps Performance Management does not migrate performance data on the virtual interface over to the SNMP interface. New performance data goes to the SNMP interface. The performance data on the virtual interface is not easily accessible.

### **Wi-Fi Device Visibility with VNA and Spectrum Integrations**

#### **Issue:**

If you have integrated with VNA and DX NetOps Spectrum (Spectrum), you have configured a plug-in that monitors Wi-Fi devices, and Spectrum is integrated with the same VNA data source, VNA and Spectrum contribute the same inventory to NetOps Portal. The data aggregator reconciles this inventory, and consolidates these items on the **Wi-Fi Device** page (the **Inventory, Items, Wi-Fi Devices** menu).

Deleting the VNA Gateway (the connection to VNA) removes the **Wi-Fi Devices** menu. Spectrum-contributed Wi-Fi devices remain in the NetOps Portal inventory, but you cannot access them from NetOps Portal because the **Wi-Fi Device** page is not available.

#### **Workaround:**

Perform a full synchronization of the Spectrum data source. This re-evaluates the device sub-type in the device inventory tables, and adds back the **Wi-Fi Devices** menu.

For more information, see [Synchronize Data Sources](#).



## **Consul Port for Fault Tolerant Data Aggregators with Consul as HTTPS for Upgrades to 23.3.4 and Higher**

### **issue:**

If you are using fault-tolerant data aggregators and Consul is configured as HTTPS, after an upgrade to 23.3.4 and higher, you cannot start the fault-tolerant data aggregators.

### **Workaround:**

To allow you to start the fault-tolerant data aggregators, do *one* of the following based on your position in the upgrade process *for each data aggregator*:

- If you have *not yet upgraded* the fault-tolerant data aggregators, edit the `/etc/consul-ext.cfg` file, and set the `consul_port` parameter to 8443.
- If you have *already upgraded* the fault-tolerant data aggregators, edit the following files:
  - The `<installation_directory>/consul-ext/bin/start-consul-ext.sh` script file, and set the `CONSUL_PORT` parameter to 8443.
    - ***installation\_directory***  
Specifies the default installation for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
  - The `/etc/consul-ext.cfg` script file, set the `consul_port` parameter to 8443.

For more information, see [Configure Consul as HTTPS](#) and [Verify the Upgrade](#).

## **Data Source Compatibility**

To view a list of the known interoperable combinations of DX NetOps components and data sources for this release, see the [DX NetOps Interoperability](#).

## **Language Support**

DX NetOps Performance Management supports the following locales:

- English (US)
- French (France)
- Japanese

For more information about how to customize the language for your user account, see [Customize Your User Settings](#).

In this article:

### **Untranslated Items in DX NetOps Performance Management**

The following items are not localized.

#### **Component and Device Type Names**

Component types are not localized, and the following device type names are not localized when you view them in report dashboards, groups, and inventory views:

- **Device types**
  - Other Devices
  - Pingable Devices
  - Call Server

The **Device Component** label and description in the **Group Rule** dialog are not localized.



---

## **Custom Context Types from Data Sources**

Custom context types synchronized by data sources, such as the context types display in the **Dashboard Editor**, are not localized.

## **Data Collector Installer Download Page**

The data collector installer download page ([http://<da\\_host>:<port>/dcm/install.htm](http://<da_host>:<port>/dcm/install.htm)) is not localized. For more information about how to install the data collectors, see [Install the Data Collectors](#).

## **Data Collector Names**

The names of data collectors are not localized.

## **Data Source Localization Limitations**

The names, descriptions, and other aspects of the predefined monitoring profiles and threshold event profiles that you can view in DX NetOps Performance Management administration are not localized.

DX NetOps Network Flow Analysis is only translated into French, Chinese (Simplified), and Japanese.

Event information from CA Application Delivery Analysis, DX NetOps Network Flow Analysis, and CA Unified Communications Monitor data sources, such as event descriptions and event types, is not localized.

## **Custom Item Types from Data Sources**

Some data sources support unique managed item types. For example, the data aggregator data source supports a "Device Component" managed item type. When data sources synchronize managed items with unique item types, the items appear in the Inventory, but their names and types are not localized.

## **Data Source Role Rights**

Some of the supported data sources have their own sets of role rights. In some cases, the data source user interface is not localized into all of the supported languages. Role rights that are synchronized from such data sources that appear in the **Edit Role Rights** dialog are not localized.

Additional data sources, such as CA Application Delivery Analysis and CA Unified Communications Monitor also have limitations.

Role rights that are synchronized from CA Application Delivery Analysis, CA Unified Communications Monitor, and DX NetOps Network Flow Analysis are not localized.

## **Direction Terms**

The direction terms "In" and "Out" are not localized in the **Interfaces Over Threshold** view.

## **DNS Names**

DNS names are not localized.

## **English String in NetOps Portal Installer**

When you run the installer on a server or in the command line with a locale set to a language other than English, the string `DEFAULT` appears in English. This string is not localized.

## **How To Videos**

Videos that supplement the documentation are not translated.

### **Installation Scripts**

The `dr_validate.sh` and `dr_install.sh` data repository scripts are not localized.

### **Limitations on Custom Strings**

You cannot provide multiple translations for strings that you customize, such as the following strings:

- Group Name
- Tenant Name
- Domain Name
- Role Name
- View Title
- Monitoring Profile Name
- Threshold Event Profile Name
- Discovery Profile Name

### **Monitoring Profiles**

The names and descriptions of the product default monitoring profiles are not localized.

### **MIB Compiler Errors**

MIB compiler errors are not localized.

### **Overview Tab in NetOps Portal**

CA Unified Communications Monitor uses some tabs on the top-level dashboards. The Overview tab is not localized. The workaround is to manually create a new dashboard that contains the CA Unified Communications Monitor views.

### **Theme Names**

Theme names are not localized.

### **XML Tag Names**

The XML tag names for vendor certification and metric family files are not localized.

### **Third-Party Information**

The following third-party information is not localized:

### **License Agreements for Third-Party Products**

The license agreements for third-party products are not localized.

### **Third-Party Scripts**

Any third-party scripts included in the product, such as the capabilities of a script, are not localized.

### **Known Issues**

The date format is incorrect in Asian languages on the **Performance** tab on device context pages.

## Third-Party Software Acknowledgments

DX NetOps Performance Management uses third-party software in accordance with the terms and conditions for use, reproduction, and distribution as defined by the applicable license agreements.

To view the third-party software license agreements, download the file pertaining to the release:

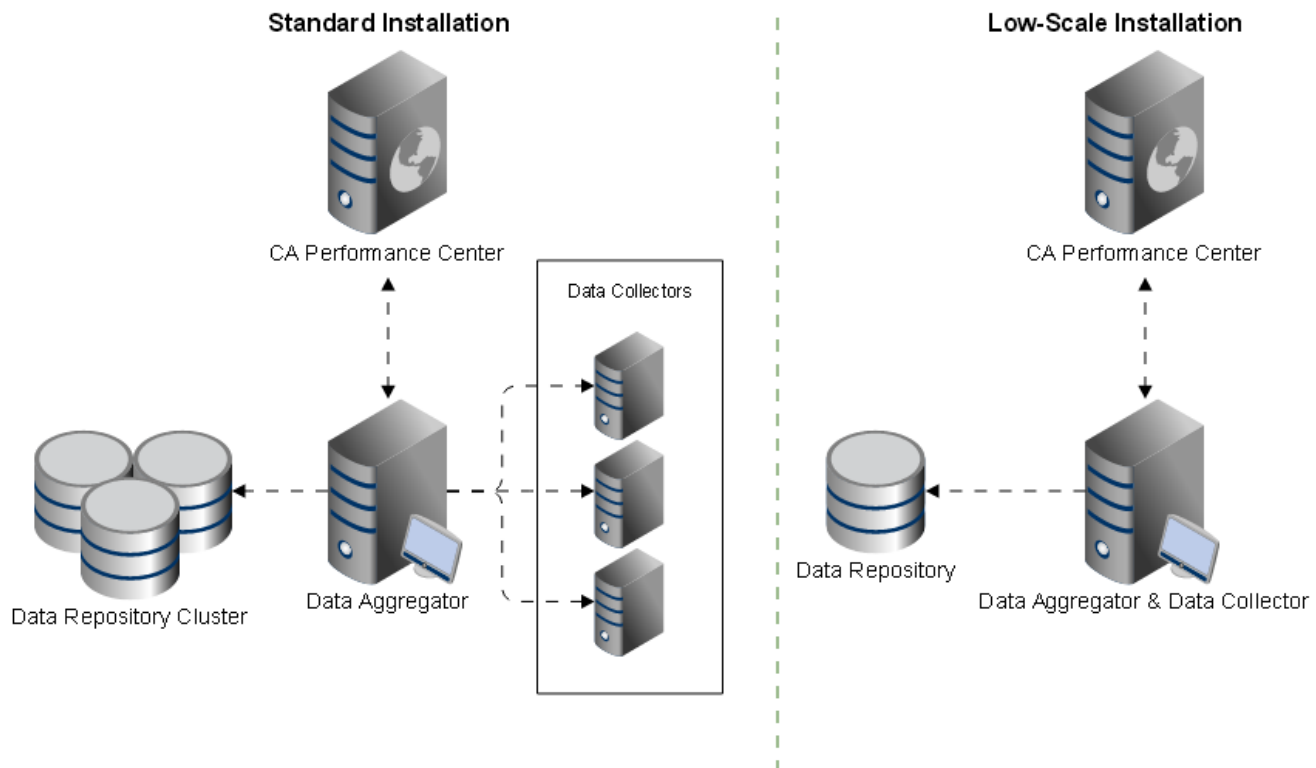
- **23.3.6:** [PM\\_TPSRs\\_2336.zip](#)
- **23.3.5:** [PM\\_TPSRs\\_2335.zip](#)
- **23.3.4:** [PM\\_TPSRs\\_2334.zip](#)
- **23.3.3:** [PM\\_TPSRs\\_2333.zip](#)
- **23.3.2:** [PM\\_TPSRs\\_2332.zip](#)
- **23.3.1:** [PM\\_TPSRs\\_2331.zip](#)

# Installing

DX NetOps Performance Management is a distributed application that includes multiple components across several servers. The deployment strategy you use depends on the number of devices, location of these devices, and which metrics you want to monitor.

Install DX NetOps Performance Management using one of the following installation options:

**Figure 2: Performance Management Component Installation Options**



A successful DX NetOps Performance Management deployment includes installing NetOps Portal, the data repository, the data aggregator, and the data collectors.

Use the following process for a *standard* installation:

**NOTE**

For a *low-scale* installation, see [Install a Low-Scale System](#).

1. [Review Installation Requirements and Considerations](#).
2. Install NetOps Portal:
  - a. [Prepare to install NetOps Portal](#).
  - b. [Install NetOps Portal](#).
3. Install the data repository:
  - a. [Prepare to install the data repository](#).
  - b. [Install the data repository](#).

4. Install the data aggregator:
  - a. [Prepare to install the data aggregator.](#)
  - b. [Install the data aggregator.](#)
5. Install the data collectors:
  - a. [Prepare to install the data collectors.](#)
  - b. [Install the data collectors.](#)

**IMPORTANT**

Deploy DX NetOps Performance Management with the data aggregator and the data collector on *separate* dedicated nodes.

6. Complete *one* of the following:
  - (If you are licensed to use DX NetOps under a Portfolio License Agreement (PLA) subscription) Configure it to send product-specific usage data.
  - (If you are licensed to use DX NetOps under a standard license) Optionally consent to having DX NetOps send this data to Broadcom.

For more information, see [the DX NetOps documentation](#).

7. [Complete the post-installation configuration.](#)
8. (Optional) Install the following:
  - If you require minimal data disruption/loss and you want to minimize software disruption, [configure the data collectors for fault tolerance](#).
  - If you plan to monitor software-defined networking (SDN) and network functions virtualization (NFV) using DX NetOps Virtual Network Assurance (VNA), [integrate with DX NetOps Virtual Network Assurance](#).

**IMPORTANT**

**Best Practice:** Install VNA on a *separate* host as the data collector.

- If you plan to monitor the performance of non-SNMP based devices using DX NetOps Mediation Manager, such as mobile wireless, fiber-optic switch, radio access, and 3G or 4G voice data, [install DX NetOps Mediation Manager](#).
- If you plan to use FIPS-compliant encryption and hashing algorithms (where applicable) for user passwords and Single Sign-On, [enable FIPS-compliant encryption](#).

## Installation Requirements and Considerations

Review the following information before installing DX NetOps Performance Management.

In this article:

- [Operating System Requirements](#)
- [Package Requirements](#)
- [NetOps Portal Access Requirements](#)
- [Virtual and SAN Environment Requirements](#)
- [Common Considerations](#)
- [Multi-tenant Deployment Considerations](#)
- [DX NetOps Performance Management Connectivity](#)

### **Operating System Requirements**

To ensure a successful installation of DX NetOps Performance Management, verify that your environment meets the operating system (OS) requirements. To verify the sizing requirements, use the [DX NetOps Sizing Tool](#).

DX NetOps Performance Management supports the following operating systems and versions unless otherwise specified:

Operating System	Supported
Red Hat Enterprise Linux (RHEL)	RHEL 8.x RHEL 7.3 and higher
Oracle Linux (OL)	OL 7.3 and higher
SUSE Linux Enterprise Server (SLES)	SLES 12 SP2/3/4/5 openSUSE 42.3
Community Enterprise Operating System (CentOS) Linux	CentOS Linux 8.x CentOS Linux 7.3 and higher
Rocky Linux	Rocky Linux 8.x

For the OS versions that previous DX NetOps Performance Management versions support, use the **Versions** drop-down.

### IMPORTANT

**Data Repository Operating System Requirements:** The data repository requires a Vertica-supported OS, but the `dr_install.sh` script (the data repository installer) is limited to the RHEL/CentOS x86\_64, SLES x86\_64 12 SP2, openSUSE x86\_64 42.3, and OL (Red Hat-compatible kernels only) x86\_64 operating systems.

For more information about the Vertica-supported operating systems, see [the 10.1 Vertica documentation](#).

### Package Requirements

The installer for each component requires the following packages:

Component (Operating System)	Packages
All (SLES/openSUSE)	<ul style="list-style-type: none"> <li>• dialog</li> <li>• mcelog</li> <li>• zip</li> <li>• unzip</li> <li>• libcap-progs</li> <li>• sysvinit-tools</li> <li>• python36-base</li> </ul>
All (RHEL/CentOS Linux 7.x and 8.x, Rocky Linux 8.x, and OL)	<ul style="list-style-type: none"> <li>• dialog</li> <li>• mcelog</li> <li>• zip</li> <li>• unzip</li> <li>• chrony</li> <li>• (RHEL/CentOS Linux 7.x and 8.x, Rocky Linux 8.x only) libcap</li> </ul>
All (RHEL/CentOS Linux and Rocky Linux 8.x)	<ul style="list-style-type: none"> <li>• Install one of the following: <ul style="list-style-type: none"> <li>– python36</li> <li>– python38</li> <li>– python39</li> <li>– python3.11</li> </ul> </li> </ul>
All (RHEL/CentOS Linux 7.x and OL)	<ul style="list-style-type: none"> <li>• python3</li> </ul>

Component (Operating System)	Packages
NetOps Portal (SLES/openSUSE)	<ul style="list-style-type: none"> <li>fontconfig</li> <li>libaiol</li> <li>libnuma1</li> <li>libuser</li> <li>wget</li> </ul>
NetOps Portal (RHEL/CentOS Linux 7.x and 8.x, Rocky Linux 8.x, and OL)	<ul style="list-style-type: none"> <li>fontconfig</li> <li>libaio</li> <li>libaio-devel</li> <li>numactl-libs</li> <li>wget</li> </ul>
NetOps Portal (OL)	<ul style="list-style-type: none"> <li>libncurses</li> </ul>
NetOps Portal (RHEL/CentOS Linux 7.x)	<ul style="list-style-type: none"> <li>ncurses-libs</li> </ul>
NetOps Portal (RHEL/CentOS Linux/Rocky Linux 8.x)	<ul style="list-style-type: none"> <li>ncurses-compat-libs</li> </ul>
Data repository (All)	<ul style="list-style-type: none"> <li>bc</li> <li>pstack For RHEL/CentOS Linux 7.x and 8.x, and Rocky Linux 8.x, this package is included in the gdb package.</li> <li>gstack For RHEL/CentOS Linux 7.x and 8.x, and Rocky Linux 8.x, this package is included in the gdb package.</li> <li>(23.3.3 and higher) sshpass</li> </ul>
Data repository (RHEL/CentOS Linux/ Rocky Linux 8.x)	<ul style="list-style-type: none"> <li>libnsl</li> </ul>
Data collectors (RHEL/CentOS Linux 7.x and 8.x, Rocky Linux 8.x, SLES/openSUSE, and OL)	<ul style="list-style-type: none"> <li>at</li> </ul>

### NetOps Portal Access Requirements

NetOps Portal has the following access requirements:

- [Supported Web Browsers.](#)
- [Other Requirements.](#)

### Supported Web Browsers

NetOps Portal supports the following web browsers:

- Google Chrome 92.x and higher
- Microsoft Edge version 92.x and higher
- Mozilla Firefox 92.x and higher

Use the latest production-level versions of these browsers whenever possible.

### Other Requirements

NetOps Portal supports a minimum screen resolution of 1280x1024.

### Virtual and SAN Environment Requirements

Review the policies for installing and operating Infrastructure Management products on virtualized servers or Storage Array Networks (SAN).

For more information, see [the CA Support Statement for Running CA Infrastructure Management Products in Virtualization and SAN Environments document](#).

**NOTE**

To view this document, log in as a registered user.

**Common Considerations**

- Install each DX NetOps Performance Management component on a separate system.
- Administrative privileges are required to install the software. Typically, the root user installs the software. In some environments, unrestricted root user access is not available.  
If root user access is not available, configure a non-root user using sudo with access to a limited set of commands.  
If you install the components as a non-root user using sudo, add the `sudo` prefix to commands that require the same user as the service owner, such as restart commands and SSL set up.
- Verify that all your servers meet the minimum requirements and sizing guidelines.

**TIP**

To provide high availability for your data, future scalability, and best end-user experience, deploy the data repository as a cluster.

For information about the sizing requirements, see the [DX NetOps Sizing Tool](#).

**NOTE**

If the sizing tool indicates a low-scale deployment, see [Install a Low-Scale System](#).

- If you plan to stand up DX NetOps Performance Management in the cloud, see [Review Cloud Sizing Guidelines](#).
- If you are configuring the data collectors for fault tolerance, ensure that you have [reviewed the hardware requirements for fault-tolerant data collectors](#).
- Time synchronization using the Network Time Protocol (NTP) networking protocol is required. If it is not running, start the NTP daemon. All machines must use the same NTP server.
  - [Verify on SUSE Linux Enterprise Server](#)
  - [Verify on Red Hat Enterprise Linux 7.x/8.x and Oracle Linux](#)

**Verify on SUSE Linux Enterprise Server****Follow these steps:**

- Open a console and issue the following command:  

```
$ systemctl status ntpd
```
- Verify that the NTP daemon is in an active (running) state.
- Start and enable the NTP daemon manually by issuing the following command:  

```
$ systemctl start ntpd  
$ systemctl enable ntpd
```

The daemon is started.

**Verify on Red Hat Enterprise Linux 7.x/8.x and Oracle Linux**

RHEL 7.x/8.x and OL 7.x run NTP with `chronyd`.

**Follow these steps:**

- Open a console and issue the following command:  

```
$ systemctl status chronyd
```
- Verify that the `chrony` daemon is in an active (running) state.
- Start and enable the `chrony` daemon by issuing the following command:  

```
$ systemctl start chronyd  
$ systemctl enable chronyd
```

The daemon is started.

- The DX NetOps Performance Management components require an IPv4 address (for example, 127.0.0.1/8) on the loopback interface. The IPv4 loopback address is the only IPv4 requirement in IPv6 deployments; other addresses (loopback or otherwise) can use IPv6.



**NOTE**

Except for anti-virus, system management, and time-synchronization software, do not install third-party software, especially third-party network monitoring software, on the same server as DX NetOps Performance Management components. Third-party software can interfere with the monitoring abilities of the Broadcom system, and could void the warranty.

If you install third-party software on a Broadcom system, Broadcom Support might ask you to uninstall this software before troubleshooting an issue on the server.

**Multi-tenant Deployment Considerations**

In a multi-tenant deployment, note the following information:

- DX NetOps Performance Management shares the data aggregator between tenants. The information for each tenant is secure and other tenants cannot view this information.
- In a standard tenant deployment, each tenant has a dedicated data collector. A tenant can have more than one data collector. For multiple tenants that reside in the same IP routing space, you can configure DX NetOps Performance Management to use fewer data collectors.  
For more information, see [Tenant-Agnostic Data Collectors](#).
- Where a managed service provider (MSP) is monitoring devices for multiple tenants, you can install the data collector at the MSP site.

**NOTE**

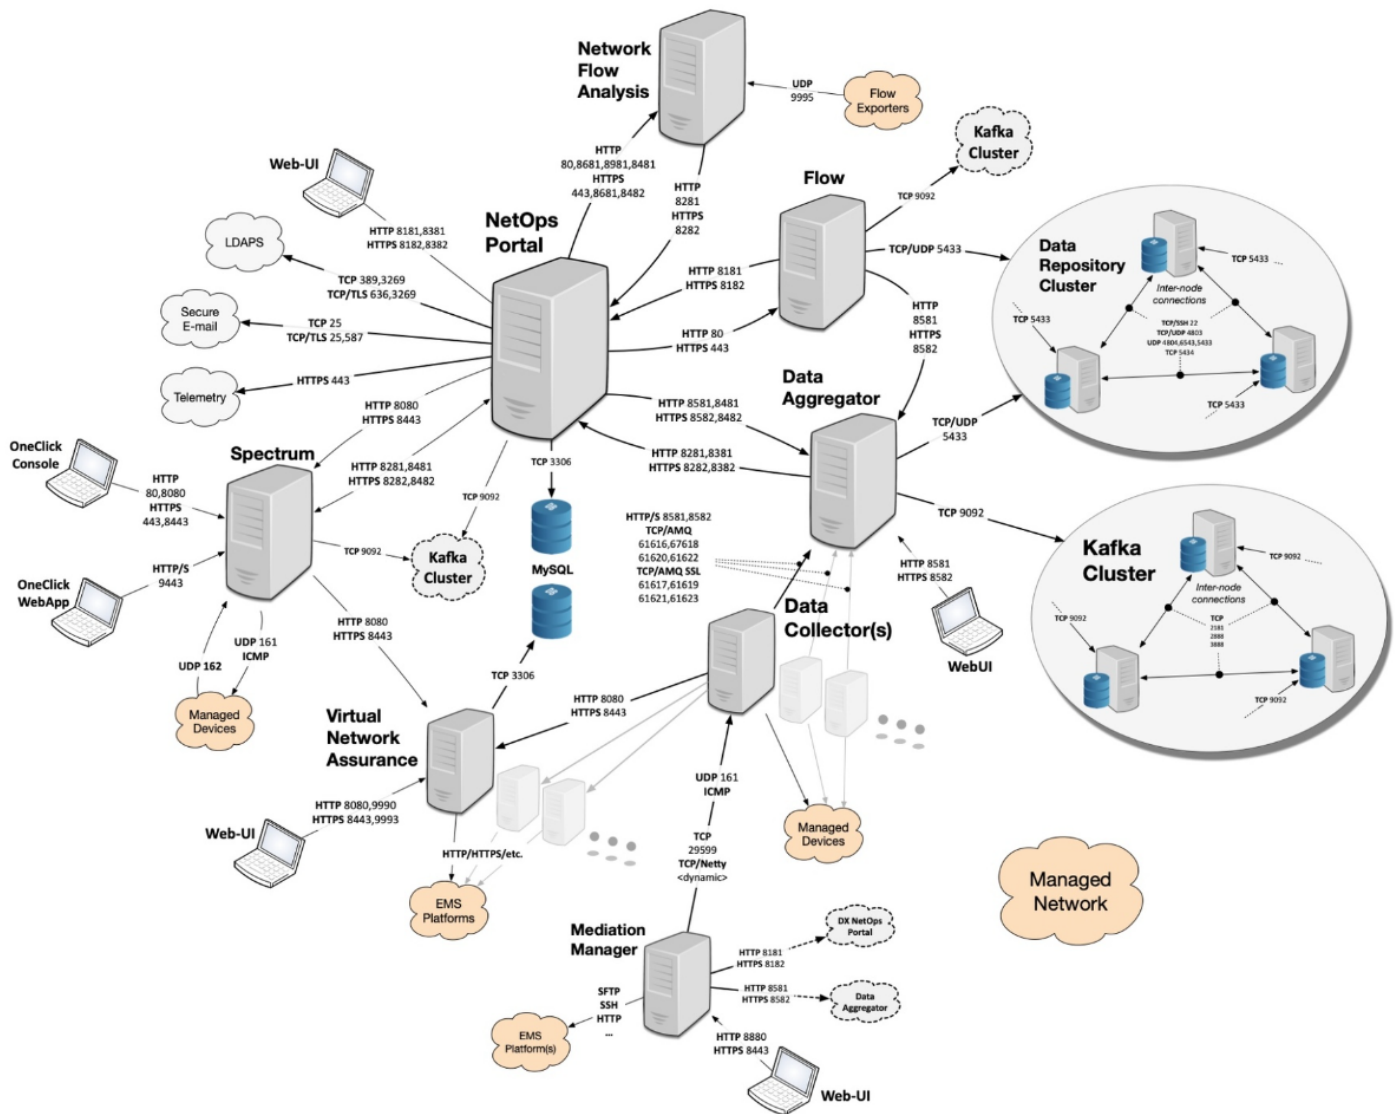
This setup requires the data collector to gain access through a tenant firewall to poll the devices that are being managed.

**DX NetOps Performance Management Connectivity**

In a firewall-protected environment, certain ports must be open.

The following diagrams show the required ports for a hardened environment with a single data aggregator:

Figure 3: Diagram of required ports for 23.3



Allow DX NetOps Performance Management communications to function properly by opening the following ports:

- Ports for the NetOps Portal Server
- Ports for the Data Aggregator Server
- Ports for the Fault-Tolerant Data Aggregator Proxy Server
- Ports for the Data Repository Server
- Ports for the DX NetOps Virtual Network Assurance (VNA) Server
- Ports for the DX NetOps Spectrum Server
- Ports for the DX NetOps Mediation Manager (DX NetOps MM) Server
- Ports for the DX NetOps Network Flow Analysis (NFA) Server
- Ports for the CA Application Delivery Analysis (ADA) Server

**NOTE**

Throughout the documentation, ports 8182, 8382, 8582, 61617, 61619, 61621, and 61623 appear as suggested port numbers for secured communications. In the instances where these ports appear, you can use any value that you want as long as no other processes are using it.

**NetOps Portal Server**

Function	Protocol and Port	Clients
Web Client and data aggregator authorization	HTTP 8181/HTTPS 8182 The Web Client default port for NetOps Portal. Also used by the SSO Service, Event Manager, and Device Manager. See Web Clients.	NetOps Portal services
MySQL	TCP 3306 Enables communication to the MySQL database (inbound) from the NetOps Portal services for the Single Sign-on (SSO) Service, Event Manager, and Device Manager.	NetOps Portal services
Single Sign-On Service	HTTP 8381/HTTPS 8382 Enables communication between NetOps Portal and the SSO Service.	NetOps Portal services
Device Manager	HTTP 8481/HTTPS 8482 Enables communication between NetOps Portal and the Device Manager.	NetOps Portal services
Event Manager	HTTP 8281/HTTPS 8282 Enables communication between NetOps Portal and the Event Manager.	NetOps Portal services
Web Client and data aggregator authentication	HTTP 8381/HTTPS 8382 Enables communication between client computers and the NetOps Portal server. Also enables log in using the SSO authentication component SSO Service.	Web Clients
DX NetOps Performance Management	HTTP 8281/( HTTPS 8282 Enables communication between the Event Manager and the data aggregator.	Data aggregator
Spectrum	HTTP 8281/HTTPS 8282 Enables communication between Spectrum OneClick server and NetOps Portal.	Spectrum
Spectrum	HTTP 8481/HTTPS 8482 Enables communication between Spectrum OneClick server and the Device Manager.	Spectrum
Lightweight Directory Application Protocol (LDAP)	TCP 389 Enables clear-text communication between NetOps Portal and the LDAP server.	NetOps Portal
LDAP	TCP 3268 Enables clear-text communication between NetOps Portal and the LDAP server.	NetOps Portal

Function	Protocol and Port	Clients
LDAP	TCP 636 Enables clear-text communication between NetOps Portal and the Secure LDAP server.	NetOps Portal
LDAP	TCP 3269 Enables secure communication between NetOps Portal and the LDAP server.	NetOps Portal
Secure E-mail Server	TCP 25, TCP/TLS 25, 587 Enables communication with secure email server.	NetOps Portal

### **Data Aggregator Server**

Function	Protocol and Port	Clients
Data aggregator API and fault-tolerant data aggregator proxy	HTTP 8581/HTTPS 8582 Enables communication for the data aggregator.	Web clients, data collector
Fault-tolerant data aggregator proxy	TCP 8300, TCP/UDP 8301, TCP 8500 In fault-tolerant environments, enables communication between the fault-tolerant proxy server and the data aggregators. The port must be open on the fault-tolerant proxy and on the data aggregators.	Fault-tolerant data aggregator proxy
Apache ActiveMQ (AMQ)	TCP/AMQ 61616/AMQ SSL 61617 Enables ActiveMQ traffic between the data collector and the data aggregator.	Data collectors
AMQ	TCP/AMQ 61618/AMQ SSL 61619 Enables poll response delivery traffic between the data collector and the data aggregator.	Data collectors
AMQ	TCP/AMQ 61620/AMQ SSL 61621 Enables distributed IREP traffic between the data collector and the data aggregator.	Data collectors
AMQ	TCP/AMQ 61622/AMQ SSL 61623 Enables large data transfers between the data collector and the data aggregator.	Data collectors

### **Fault-Tolerant Data Aggregator Proxy Server**

Function	Protocol and Port	Clients
Fault-tolerant DX NetOps Performance Management	HTTP 8581/HTTPS 8582 In fault-tolerant environments, enables communication between NetOps Portal and the fault-tolerant proxy.	NetOps Portal

**Data Repository Server**

Function	Protocol and Port	Clients
DX NetOps Performance Management	TCP/UDP 5433 Enables communication between the data aggregator and the data repository for Java database connectivity.	Data aggregator
DX NetOps Performance Management	TCP/SSH 22 Enables Vertica administration and backup utilities to run and enabled Vertica communication between nodes.	Data repository
DX NetOps Performance Management	TCP/UDP 4803 Enables Vertica communication between nodes.	Data repository
DX NetOps Performance Management	UDP 4804, TCP 5434, UDP 6543 Enables communication between the data repository and the Vertica database.	Vertica
DX NetOps Performance Management	TCP 50000 Enables the data repository host to access the custom <code>rsynch/ssh</code> on the backup and disaster recovery hosts.	Data repository

**DX NetOps Virtual Network Assurance Server**

Function	Protocol and Port	Clients
Spectrum	HTTP 8080/HTTPS 8443 Enables communication between Spectrum and VNA.	Spectrum
DX NetOps Performance Management	HTTP 8080/HTTPS 8443 Enables communication between the data collector and VNA.	Data collector
Web clients	HTTP 8080,9990/HTTPS 8443 Enables communication between web clients and VNA.	Web clients

**DX NetOps Spectrum Server**

Function	Protocol and Port	Clients
DX NetOps Performance Management	HTTP 8080/HTTPS 8443 Enables communication between NetOps Portal and Spectrum.	NetOps Portal
Spectrum	HTTP 80, 8080/HTTPS 443,8443 Enables the OneClick web client to access Spectrum.	Web client
Spectrum	HTTP/S 9443	Web app client

Function	Protocol and Port	Clients
Spectrum	UDP 162 Enables managed devices to traps to Spectrum.	Managed devices

### **DX NetOps Mediation Manager Server**

Function	Protocol and Port	Clients
DX NetOps MM	HTTP 8880/HTTPS 8443 Enables communication from web clients to DX NetOps MM.	Web clients
DX NetOps Performance Management	TCP 29599 / Netty Dynamic	Data collectors

### **DX NetOps Network Flow Analysis Server**

Function	Protocol and Port	Clients
DX NetOps Performance Management	HTTP 80 Enables communication between NetOps Portal and NFA.	NetOps Portal
DX NetOps Performance Management	HTTP 8681 Enables communication between NetOps Portal and NFA.	NetOps Portal
DX NetOps Performance Management	HTTP 8981 Enables communication between NetOps Portal and NFA.	NetOps Portal
DX NetOps Performance Management	HTTPS 443/8681 Enables communication between NetOps Portal and NFA.	NetOps Portal
DX NetOps Performance Management	HTTP 8281/HTTPS 8282 Enables communication between NFA and NetOps Portal.	NetOps Portal
Device management	UDP 9995 Enables communication from flow exporters.	Network devices

### **CA Application Delivery Analysis Server**

Function	Protocol and Port	Clients
DX NetOps Performance Management	HTTP 80 Enables communication between NetOps Portal and ADA.	NetOps Portal
DX NetOps Performance Management	HTTP 8681 Enables communication between NetOps Portal and ADA.	NetOps Portal

Function	Protocol and Port	Clients
Network device monitoring	UDP 161/ICMP Enables Simple Network Management Protocol (SNMP) and Internet Control Message Protocol (ICMP) between data collectors and devices.	Data collectors

## Review Cloud Sizing Guidelines

If you plan to stand up DX NetOps Performance Management in the cloud, review the following cloud sizing guidelines. You can choose from a variety of cloud platforms, including Amazon Web Services (AWS) and Google Cloud Platform (GCP). The current guidelines focus on AWS.

### Follow these steps:

1. Determine the estimated requirements of your environment using the DX NetOps Sizing Tool.  
For more information, see the [DX NetOps Sizing Tool](#).
2. Review the estimated disk requirements from the sizing tool and consider your data retention rates. Your data retention rates impact disk requirements greatly.  
For more information, see [Configure Data Retention Rates](#).
3. Review the other entries in the sizing tool, which impact the requirements of your environment.
4. Adjust the entries in the sizing tool until they most accurately reflect the needs of your environment.
5. Review the guidelines for your cloud platform.

**AWS**

Components	Instance Types	Descriptions	Added Storage	Notes	Examples
Data Repository / Vertica	r5.4, r5.8, r5.12	<p>r5 instance types are memory-optimized instances. These instances deliver fast performance for workloads that process large data sets in memory. Vertica supports r5 instance types for the data repository. For more information, see the <a href="#">Vertica documentation</a>.</p> <p>The core to memory ratio best matches what the DX NetOps Sizing Tool recommends. Storage does not factor into this recommendation. Use one of the recommended added storage options from EBS to cover the recommended disk size from the DX NetOps Sizing Tool.</p>	IOPS SSD, Throughput HDD	<p>To select the instance type best suited for your data repository, review the estimated CPU requirements in the sizing tool.</p> <p>The Data Repository requires storage beyond the boot disk. Amazon EBS allows you to create storage volumes and attach them to Amazon EC2 instances. Use EBS for the catalog directory and the data directory.</p>	<p>Using the sizing tool with the default values and 50,000 devices, the expected data repository sizing is 14 cores and 224 GiB RAM. The r5.8 has 32 hyperthreaded cores and 256 GiB RAM. You can deselect the hyperthreaded option to run at 16 cores and 256 GiB RAM, which gives the expected ratio by the sizing tool. You are not required to deselect the hyperthreaded option, as the benefit can vary from deployment to deployment.</p>
Data Aggregator	m5.1, m5.2, m5.4	<p>m5 instance types provide a balance of compute, memory, and networking resources. The 1:4 CPU to memory ratio best matches the recommendations from the DX NetOps Sizing Tool.</p>	No	<p>To select the instance type best suited for the data aggregator, review the estimated memory requirements in the sizing tool.</p>	<p>Using the DX NetOps Sizing Tool with the default values and 50,000 devices, the expected data aggregator sizing is 18 cores and 48 GiB RAM. The m5.4 is the closest fitting instance type.</p>



Data Collector	c5d.1, c5d.2, c5d4	c5 instance types are compute-optimized instances. These instances are for compute-bound applications that benefit from high performance processors. The 1:2 CPU to memory ratio best matches the recommendations from the DX NetOps Sizing Tool.	No	<p>To select the instance type best suited for the data collectors, consider the following guidance:</p> <ul style="list-style-type: none"> <li>• Select c5d.1 in a small environment monitoring around 100K items.</li> <li>• Select c5d.2 in a medium environment monitoring around 500K items.</li> <li>• Select c5d.4 in a large environment monitoring around 1000K items or an environment with intensive metrics or polling rates.</li> </ul>	Using the DX NetOps Sizing Tool with the default values and 50,000 devices, the expected data collector sizing is 8 cores and 16 GiB RAM. The c5d.2 fits the DX NetOps Sizing Tool recommendations exactly.
NetOps Portal MySQL Database Node	m5.1, m5.2, c5	<p>m5 instance types provide a balance of compute, memory, and networking resources.</p> <p>c5 instance types are compute-optimized instances. These instances are for compute-bound applications that benefit from high-performance processors.</p>	IOPS SSD	<p>To select the instance type best suited for the NetOps Portal MySQL database node, review the estimated memory requirements in the sizing tool.</p> <p>The MySQL Database Node requires storage beyond the boot disk. Amazon EBS allows you to create storage volumes and attach them to Amazon EC2 instances. Use EBS for the MySQL data directory.</p>	Using the DX NetOps Sizing Tool with the default values and 50,000 devices, the expected NetOps Portal MySQL node sizing is 12 cores and 32 GiB RAM. This example uses an m5.2. If you have a high user load, consider using a c5.4 to help with concurrent requests made by users.

NetOps Portal Core Services Node	c5.4	c5 instance types are compute-optimized instances. These instances are for compute-bound applications that benefit from high performance processors.	No	To select the instance type best suited for the NetOps Portal core services node, review the estimated CPU requirements in the sizing tool.	Using the DX NetOps Sizing Tool with the default values and 50,000 devices, the expected NetOps Portal MySQL node sizing is 12 cores and 32 GiB RAM. This example uses an m5.2. If you have a high user load, consider using a c5.4 to help with concurrent requests made by users.
CA Virtual Network Assurance	r5.2	r5 instance types are memory-optimized instances. These instances deliver fast performance for workloads that process large data sets in memory. DX NetOps Virtual Network Assurance, varying based on the plugins deployed, is a memory intensive application. The r5 instance type scales best for both smaller and larger VNA deployments.	No	Use an r5.2 instance type for DX NetOps Virtual Network Assurance.	
DX NetOps Spectrum SS Node	c5.2, c.5.4	c5 instance types are compute-optimized instances. These instances are for compute-bound applications that benefit from high performance processors. The 1:2 CPU to memory ratio best matches the recommendations from the DX NetOps Sizing Tool.	No	To select the instance type best suited for your DX NetOps Spectrum SS Node, review the estimated memory requirements in the <a href="#">DX NetOps Spectrum Sizing Tool</a> .	

DX NetOps Spectrum OC Node	c5.1, c5.2, c5.4	c5 instance types are compute-optimized instances. These instances are for compute-bound applications that benefit from high performance processors. The 1:2 CPU to memory ratio best matches the recommendations from the DX NetOps Sizing Tool.	No	To select the instance type best suited for your DX NetOps Spectrum OC Node, review the estimated memory requirements in the <a href="#">Spectrum Sizer</a> .	
----------------------------	------------------	---	----	---	--

## Fault Tolerance

Fault tolerance enables DX NetOps Performance Management to continue operating properly when a hardware failure or network issue occurs.

In fault-tolerant environments, the components can switch over to standbys in the event of a failure, thus minimizing disruption and potential data loss.

In this article:

- [Fault-Tolerant Data Aggregators](#)
- [Fault-Tolerant Data Collectors](#)

### Fault-Tolerant Data Aggregators

In fault-tolerant data aggregator environments, when the active data aggregator ("Active" status) goes offline, a secondary inactive data aggregator that is ready to take over ("Ready" status) automatically becomes active ("Active" status), and takes over to organize and feed data to NetOps Portal and to the data repository. This newly-active data aggregator retains state information from the previously-active data aggregator. The previously-active data aggregator is available for failover when it becomes available ("Ready" status).

You install fault tolerant data aggregators during the DX NetOps Performance Management installation. For more information about this step in the process, see [Installing](#).

Review the following information about fault-tolerant data aggregator environments:

- [System Architecture](#)
- [Hardware Requirements](#)
- [Data Loss Comparison](#)

### The System Architecture of Fault-Tolerant Data Aggregator Environments

The following diagram illustrates the system architecture of a fault-tolerant data aggregator environment:

**Figure 4: Fault-tolerant data aggregator environment**

#### **NOTE**

Træfik is a modern HTTP reverse proxy and load balancer that you can use to deploy microservices with ease. You use the Consul tool to manage services in DX NetOps Performance Management.

For more information:

- About Træfik, see [the Træfik site](#).
- About Consul, see [the Consul site](#).

### Hardware Requirements for Fault-Tolerant Data Aggregator Environments

The following extra hardware is required for a fault-tolerant data aggregator environment:

- An additional data aggregator server.
- A proxy server. The proxy server works as the third node of the service management cluster for fault tolerance. The service management cluster includes the proxy server, the active data aggregator, and the inactive data aggregator.

For more information, see [Install Fault-Tolerant Data Aggregators](#).

### Data Loss Comparison for Fault-Tolerant Data Aggregator Environments

Some data loss might occur in a fault-tolerant environment when a hardware failure or network issue occurs. However, the amount of data loss is less than in a non-fault-tolerant data aggregator environment. The following table compares the data loss from a hardware failure or network outage:

	Hardware Failure		Network Outage	
Is fault tolerance configured?	No	Yes	No	Yes
What happens to rollups?	Pending rollups are lost and never recovered.	The other available data aggregator consumes the pending rollups when it becomes active.	Pending backups are consumed when the network is restored.	The other available data aggregator consumes the pending rollups when it becomes active.
What is lost in memory?	For 10K polls in memory at scale, loss should not exceed 1 poll cycle. Max loss would be 10K items per metric family.	For 10K polls in memory at scale, loss should not exceed 1 poll cycle. Max loss would be 10K items per metric family.	For 10K polls in memory at scale, loss should not exceed 1 poll cycle. Max loss would be 10K items per metric family.	For 10K polls in memory at scale, loss should not exceed 1 poll cycle. Max loss would be 10K items per metric family.
What happens to data transfer object (DTO) files?	If the hardware failure is the disk, all files are lost. Otherwise, whole DTO files are consumed when the hardware is restarted after repair. Incomplete files are discarded.	Whole DTO files are processed and partially written DTO files are discarded. A DTO file is 1 metric family over 1 poll cycle.	Whole DTO files are processed and partially written DTO files are discarded. The data aggregator attempts to shut down gracefully and close any DTO file in flight.	Whole DTO files are processed and partially written DTO files are discarded. A DTO file is 1 metric family over 1 poll cycle.
What happens with the ActiveMQ Broker?	For 600-MB cache in memory and an average message size of 1.3K, approximately 470K messages could be lost.	For 600-MB cache in memory and an average message size of 1.3K, approximately 470K messages could be lost.	For 600-MB cache in memory and an average message size of 1.3K, approximately 470K messages could be lost.	For 600-MB cache in memory and an average message size of 1.3K, approximately 470K messages could be lost.
What happens with thresholding?	Data loss does not exceed 1 poll cycle.	Data loss does not exceed 1 poll cycle.	Data loss does not exceed 1 poll cycle.	Data loss does not exceed 1 poll cycle.

### Fault-Tolerant Data Collectors

Data collector fault-tolerant environments use an N+1 redundancy model. In fault-tolerant data collector environments, when the active data collector in the standby group (“Active” status) goes down, a standby data collector (“Standby” status) in the standby group becomes active (“Active” status). This newly-active data collector takes over the polling for

the previously-active data collector, and the previously-active data collector becomes a standby data collector (“Standby” status) in the standby group. You configure a data collector for fault tolerance *after* the DX NetOps Performance Management installation. For more information about this step in the process, see [Installing](#).

Review the following information about fault-tolerant data collector environments:

- [System Architecture](#)
- [Hardware Requirements](#)

### **The System Architecture of Fault-Tolerant Data Collector Environments**

The following diagram illustrates the system architecture of a fault-tolerant data collector environment:

**Figure 5: Fault-tolerant data collector environment**

### **Hardware Requirements for Fault-Tolerant Data Collector Environments**

The following extra hardware is required for a fault-tolerant data collector environment:

- A separate host for each data collector instance that you are making a standby.
- A data collector server for each fault-tolerant data collector that you want to use as a standby data collector, running on a separate host, with equivalent specifications. Size each host the same for each data collector that you assign as a standby for an active data collector.  
For example, you can configure a standby data collector for two or more active data collectors. Or, less commonly, you can configure two standby data collectors for an active data collector.

For more information:

- About how to install a data collector, see [Install the Data Collectors](#).
- About how to configure data collectors for fault tolerance, see [Configure the Data Collectors for Fault Tolerance](#).

## **Prepare to Install NetOps Portal**

Ensure that you can install NetOps Portal by preparing for the installation.

Complete the following procedures before installing NetOps Portal:

1. [Verify the Prerequisites](#)
2. [Set the Limit on the Number of Open Files](#)
3. [Verify the Communication Ports](#)
4. [Verify Time Synchronization](#)
5. (If you do not have root access to install and run NetOps Portal) [Configure the Sudo User Account for NetOps Portal](#)
6. [Configure UTF-8 Support](#)
7. [Install the Required and Non-English Fonts](#)
8. (If you plan to configure the Event Manager service (`caperfcenter_eventmanager`) to publish events to a Kafka topic (enable the Event Manager-Kafka integration)) [Install NetOps Kafka](#)

### **Verify the Prerequisites**

Complete the following prerequisites before you install NetOps Portal:

- [You have reviewed the installation requirements and considerations](#). This includes verifying that you have the packages that the installer for NetOps Portal requires.
- You can install the software in any file system to which the root user has write access. `/opt/CA` is the default installation directory for NetOps Portal. You can select another location while running the setup program.
- You have verified that Security-Enhanced Linux (SELinux) is disabled or permissive on the computer where you are going to install NetOps Portal. By default, some Linux distributions have this feature enabled, which does not allow NetOps Portal to function properly. Disable SELinux, set to permissive mode, or create a policy to exclude NetOps Portal processes from SELinux restrictions.

For information more about how to configure SELinux security policies, see [the Red Hat documentation](#).

- You have verified that the selected file system has enough allocated disk space to support a database.
- To prevent scanning by a local instance of an antivirus client and scanning by a remote antivirus instance which can result in database corruption, you have excluded the installation directory for NetOps Portal, the default installation directory for the MySQL database, and the following MySQL subdirectories, from antivirus scans:

– `<installation_directory>/PerformanceCenter/MySQL/`

The default installation directory for the MySQL database.

#### NOTE

You can select another directory during installation.

– `<installation_directory>/PerformanceCenter/MySQL/bin`

– `<installation_directory>/PerformanceCenter/MySQL/data`

– `<installation_directory>/PerformanceCenter/MySQL/tmp`

#### NOTE

`/opt/CA` is the default installation directory for NetOps Portal.

- NetOps Portal requires DNS resolution. If you do not have DNS configured, you have added system entries to the `/etc/hosts` file on your server.
- You have verified that the `/tmp` location has at least 4 GB of space available.
- You have verified that Perl is installed. Some of the scripts require this utility.

### **Set the Limit on the Number of Open Files**

Verify that the user account that is installing NetOps Portal has a value of at least 65536 on the number of open files. Set this value permanently.

#### NOTE

For systems where a sudo user installs NetOps Portal, the user installing might not have the required permissions to complete this procedure. Work with the system administrator to configure the limit.

#### **Follow these steps:**

1. Log in to the NetOps Portal host.
2. Edit the `/etc/security/limits.conf` file.
3. Add the following lines to the file:
 

```
# Added by Performance Center
* soft nofile 65536
# Added by Performance Center
* hard nofile 65536
```
4. Restart the login session.
5. Verify that the number of open files is set properly by issuing the following command:

```
ulimit -n
```

The command returns the limit that you have specified.

## **Verify the Communication Ports**

DX NetOps Performance Management uses multiple ports to communicate with various components, particularly data sources. In addition, some of the products and components that integrate with DX NetOps Performance Management have specific port requirements. For DX NetOps Performance Management to work correctly in a firewall-protected environment, certain ports must be open.

Verify the ports that must be open:

- The ports that allow NetOps Portal services to communicate with NetOps Portal.
- The ports that allow users to contact NetOps Portal.
- The ports that allow other data sources to contact NetOps Portal for eventing and OpenAPI single sign-on.
- For any firewall that protects the DX NetOps Performance Management server, the ports and protocols for the data sources that you are deploying.
- The ports that the other data sources use.

For more information about the list of required ports and protocols, see the documentation for each data source.

For a detailed list of these ports, see [Review Installation Requirements and Considerations](#).

## **Verify Time Synchronization**

NetOps Portal requires time synchronization using the Network Time Protocol (NTP) daemon. All data source consoles must use the NTP daemon. On Linux servers, the NTP daemon ensures that the clocks on the hosts are synchronized for timing purposes. Verify that the daemon is running on the NetOps Portal host server based on your version:

- [Verify on SUSE Linux Enterprise Server \(SLES\)](#)
- [Verify on Red Hat Enterprise Linux \(RHEL\) 7.x/8.x and Oracle Linux \(OL\)](#)

### **Verify Time Synchronization on SUSE Linux Enterprise Server**

**Follow these steps:**

1. Open a console and issue the following command:  

```
$ systemctl status ntpd
```
2. Verify that the NTP daemon is in an active (running) state.
3. Start and enable the NTP daemon manually by issuing the following command:  

```
$ systemctl start ntpd
$ systemctl enable ntpd
```

The daemon is started.

### **Verify Time Synchronization on Red Hat Enterprise Linux 7.x/8.x and Oracle Linux**

RHEL 7.x/8.x and OL 7.x run NTP with `chronyd`.

**Follow these steps:**

1. Open a console and issue the following command:  

```
$ systemctl status chronyd
```
2. Verify that the `chrony` daemon is in an active (running) state.
3. Start and enable the `chrony` daemon by issuing the following command:  

```
$ systemctl start chronyd
$ systemctl enable chronyd
```

The daemon is started.

**(Optional) Configure the Sudo User Account for NetOps Portal**

If you do not have root access to install and run NetOps Portal, configure the sudo user account.

With the sudo user configured, you can add the sudo prefix to all commands to install NetOps Portal.

**Example:**

```
sudo ./CAPerfCenterSetup.bin
```

**Follow these steps:**

1. Locate and open the `/etc/sudoers` file on the NetOps Portal host.
2. Add a command alias with the following permissions to the file:
  - `/tmp/CAPerfCenterSetup.bin`
  - `/etc/init.d/mysql`
  - **(RHEL 7.x/8.x, SLES, OL)** `<installation_directory>/PerformanceCenter/PC/bin/caperfcenter_console`
  - **(RHEL 7.x/8.x, SLES, OL)** `<installation_directory>/PerformanceCenter/DM/bin/caperfcenter_devicemanager`
  - **(RHEL 7.x/8.x, SLES, OL)** `<installation_directory>/PerformanceCenter/EM/bin/caperfcenter_eventmanager`
  - **(RHEL 7.x/8.x, SLES, OL)** `<installation_directory>/PerformanceCenter/sso/bin/caperfcenter_sso`
  - `<installation_directory>/PerformanceCenter/Tools/bin/npcshell.sh`
  - `<installation_directory>/PerformanceCenter/SsoConfig`
  - `<installation_directory>/PerformanceCenter/Uninstall_MySql`
  - `<installation_directory>/PerformanceCenter/Uninstall_PerformanceCenter`
  - `<installation_directory>/PerformanceCenter/Uninstall_SSO`
  - `/sbin/service`
  - `<installation_directory>/MySQL/bin/mysql`
  - `<installation_directory>/MySQL/bin/mysqldump`
  - `<installation_directory>/PerformanceCenter/sso`
  - `<installation_directory>/PerformanceCenter/PC`
  - `<installation_directory>/PerformanceCenter/PC/webapps/pc/apps`
  - `<installation_directory>/PerformanceCenter/PC/webapps/pc/css/CA-Blue/images`
  - `<installation_directory>/PerformanceCenter/PC/webapps/pc/css/CA-White/images`
  - `/usr/bin/vim`
  - `<installation_directory>/jre/bin/keytool`
  - `<installation_directory>/PerformanceCenter/RemoteEngineer/re.sh`
  - `<installation_directory>/PerformanceCenter/SslConfig`

**NOTE**

`/opt/CA` is the default installation directory for NetOps Portal.

Separate the permissions with commas and place them all on a single line.

**Example:**

```
Cmnd_Alias CAPERFCENTER = /tmp/CAPerfCenterSetup.bin,/opt/CA/PerformanceCenter/PC/bin/caperfcenter_console,/opt/CA/PerformanceCenter/DM/bin/caperfcenter_devicemanager,/opt/CA/PerformanceCenter/EM/bin/caperfcenter_eventmanager,/opt/CA/PerformanceCenter/sso/bin/caperfcenter_sso,/etc/init.d/mysql,/opt/CA/PerformanceCenter/Tools/bin/npcshell.sh,/opt/CA/PerformanceCenter/SsoConfig,/opt/CA/PerformanceCenter/Uninstall_MySql,/opt/
```



```
CA/PerformanceCenter/Uninstall_PerformanceCenter,/opt/CA/PerformanceCenter/
Uninstall_SSO,/sbin/service,/opt/CA/MySQL/bin/mysql,/opt/CA/MySQL/bin/mysqldump,/
opt/CA/PerformanceCenter/sso,/opt/CA/PerformanceCenter/PC,/opt/CA/PerformanceCenter/
PC/webapps/pc/apps,/opt/CA/PerformanceCenter/PC/webapps/pc/css/CA-Blue/images,/opt/
CA/PerformanceCenter/PC/webapps/pc/css/CA-White/images,/usr/bin/vim,/opt/CA/jre/bin/
keytool,/opt/CA/PerformanceCenter/RemoteEngineer/re.sh,/opt/CA/PerformanceCenter/
SslConfig
sudouser ALL = CAPERFCENTER
```

– **sudouser**

Specify the user who can run the sudo commands.

3. Save your changes.

The sudo user account is configured.

### **Configure UTF-8 Support**

Configure the NetOps Portal host to support UTF-8 encoding (set the language variable). Enable UTF-8 encoding to ensure that characters display properly during the installation.

The appropriate language packs are also required to support localized deployments.

**NOTE**

The installation scripts of selected components are not localized and run in English.

For more information about the languages that DX NetOps Performance Management supports, see [Language Support](#).

Issue the following command:

```
LANG=<LANG_value> ; export LANG ;
LC_ALL=<LANG_value> ; export LC_ALL
```

**Example:**

```
LANG=fr_FR.utf8 ; export LANG ;
LC_ALL=fr_FR.utf8 ; export LC_ALL
```

- **LANG\_value**

Specifies the language that you want DX NetOps Performance Management to support.

**Values:**

- **English:** en\_US.utf8
- **French:** fr\_FR.utf8
- **Japanese:** ja\_JP.utf8

### **Install the Required and Non-English Fonts**

The installer and the NetOps Portal PDF file-generation process require the necessary fonts. Follow the standard instructions for installing the necessary fonts on your operating system.

NetOps Portal and its data sources support multiple languages. As the administrator, you can select a preferred language for each unique user. Language packs take advantage of operating system support for localized environments.

**NOTE**

By default, for some language preferences, you might not be able to view dashboard data in reports. To view dashboard data in reports for those language preferences, install the fonts on the NetOps Portal host.

**Follow these steps:**

1. [Log in to the NetOps Portal host.](#)
2. Install the fonts by issuing the following command based on your operating system (OS):
  - **(RHEL, OL, Community Enterprise Operating System (CentOS) Linux, and Rocky Linux)**

```
yum groupinstall fonts
```
  - **(SLES)**

```
zypper install dejavu-sans-fonts dejavu-fonts-common arphic-ukai-fonts arphic-uming-fonts ipa-ex-mincho-fonts ipa-mincho-fonts ipa-pmincho-fonts xano-mincho-fonts baekmuk-bitmap-fonts baekmuk-ttf-fonts liberation-fonts
```
3. Rebuild the font caches by issuing the following command based on your OS:
  - **(RHEL, OL, CentOS Linux, and Rocky Linux)**

```
fc-cache -v
```
  - **(SLES)**

```
fc-cache -fv
```
4. [Restart NetOps Portal.](#)

The required and non-English fonts are installed on the NetOps Portal host.

**(Optional) Install NetOps Kafka**

If you plan to configure the Event Manager service (`caperfcenter_eventmanager`) to publish events to a Kafka topic (enable the Event Manager-Kafka integration), install the NetOps Kafka that DX NetOps Performance Management includes prior to installing NetOps Portal. The NetOps Portal installer (`CAPerfCenterSetup.bin`) prompts whether to configure the Event Manager service to publish events to a Kafka topic. If you choose this option, the NetOps Portal installer prompts for information about one or more installed Kafka brokers, and tests the connectivity.

For more information about how to install NetOps Kafka, see [Install NetOps Kafka](#).

**Next Steps**

- [Install NetOps Portal.](#)

## Install NetOps Portal

Install NetOps Portal after you have met the requirements.

Use the following process to install NetOps Portal and its services:

1. [Verify the prerequisites.](#)
2. [Install NetOps Portal.](#)
3. [Verify the installation.](#)
4. (If DX NetOps Network Flow Analysis (NFA) is registered as a data source) [Restart the NFA OData service.](#)
5. (Optional) [Review the log files.](#)

The following video shows the installation process:

The event manager is part of NetOps Portal, and is installed with NetOps Portal.

During the installation from the command line, you can configure the Event Manager service (`caperfcenter_eventmanager`) to publish events to a Kafka topic (enable the Event Manager-Kafka integration) at the time of the install, or later by way of the `em.properties` event manager properties file for the Event Manager-Kafka integration.

**TIP**

If you do not already have a Kafka deployment, you can use NetOps Kafka, which is provided as a separate download.

For more information about how to install NetOps Kafka, see [Install NetOps Kafka](#).

For more information about how to update the Event Manager properties file, see [Change the Event Manager Properties](#).

**Verify the Prerequisites**

Prior to running the installation, ensure that you have completed the [prerequisites](#).

**Install NetOps Portal**

Install NetOps Portal using one of the following options:

- [Install from the command line](#).
- [Install in silent mode](#).

**Install from the Command Line**

From the command line, you install and configure the database and website using the NetOps Portal Setup program. You can use this program for the following types of installations:

- **Complete:** Installs NetOps Portal services and the MySQL database on a single node.
- **Advanced:** For better performance, installs NetOps Portal services and the MySQL database on separate nodes. First, install the database by running the Setup program. Then, install the services by running the Setup program on a different node.

**Follow these steps:**

1. Log in to the NetOps Portal host as the root or the `sudo` user.
2. Copy the `CAPerfCenterSetup.bin` file to the `/tmp` directory.

**NOTE**

Verify that your `/tmp` location has at least 4 GB of space available.

3. Change to the `/tmp` directory by issuing the following command:  
`cd /tmp`
4. Change the permissions for the installation file by issuing the following command:  
`chmod +x CAPerfCenterSetup.bin`
5. Run the installation file by issuing the following command:  
`./CAPerfCenterSetup.bin -i console`
6. Follow the instructions in the console.

You are prompted to specify an install owner. You can specify a non-root user.

The installation checks to see whether the partition with the MySQL data directory has enough disk space to handle storage engine upgrades. If there does not appear to be enough disk space to complete the installation successfully, exit the installer, allocate more space for the data partition, and reinstall NetOps Portal.

To enhance security, you are prompted set a custom MySQL password. Ensure that this password meets the following requirements:

- Excludes the user names "root" or "netqos".
- Is not the same as the username.
- Minimum length of 8 characters.
- Maximum length of 30 characters.
- Contains at least three of the following types of characters:

- Special Characters (You can use only the !#&? characters)
- Uppercase
- Lowercase
- Numbers (0-9)

The installation runs.

You have installed NetOps Portal from the command line. A message states that the program has been installed successfully.

### **Install in Silent Mode**

In silent mode, you install NetOps Portal avoiding the prompts and using the values that you set in the `silent.properties` file.

#### **Follow these steps:**

1. Log in to the NetOps Portal host as the root or the sudo user.
2. Copy the `CAPerfCenterSetup.bin` file to the `/tmp` directory.

#### **NOTE**

Verify that your `/tmp` location has at least 4 GB of space available.

3. Change to the `/tmp` directory by issuing the following command:

```
cd /tmp
```

4. Change the permissions for the installation file by issuing the following command:

```
chmod +x CAPerfCenterSetup.bin
```

5. Create the `silent.properties` file in the `/tmp` directory by completing the following steps:

- a. Issue the following command on all servers where you want to install NetOps Portal:

```
./CAPerfCenterSetup.bin -r /tmp/silent.properties
```

- b. Follow the prompts until you get to the summary, type `quit`, and then press the Return key on your keyboard.

6. Open the `/tmp/silent.properties` file, and if present, confirm the values for the following variables:

#### – **USER\_INPUT\_INSTALL\_OWNER**

Designates a user as the install owner. You can specify a non-root user.

**Default:** `root`

#### – **USER\_INSTALL\_DIR**

Designates the directory where the application is installed.

**Default:** `/opt/CA`

#### – **MYSQL\_DATA\_FOLDER**

Designates the location for the MySQL data directory.

**Default:** `/opt/CA/MySQL/data`

#### – **MYSQL\_TEMP\_FOLDER**

Designates the location for the directory to store MySQL temporary files.

**Default:** `/opt/CA/MySQL/tmp`

#### – **DB\_PASSWORD\_VARIABLE**

Designates the MySQL password.

#### – **DB\_PASSWORD\_CONFIRM**

Confirms the MySQL password.

7. Issue the following command on all servers where you want to install NetOps Portal:

```
./CAPerfCenterSetup.bin -i silent -f /tmp/silent.properties
```

The installation begins.

An empty prompt indicates that NetOps Portal is installed.

You have installed NetOps Portal in silent mode.

## Verify the Installation

The following Linux daemons (the services) are created and started during the installation:

- **caperfcenter\_console**  
The NetOps Portal Console daemon. This service uses port 8181.
- **caperfcenter\_devicemanager**  
The Device Manager daemon. This service uses port 8481.
- **caperfcenter\_eventmanager**  
The Event Manager daemon. This service uses port 8281.
- **caperfcenter\_sso**  
The Single Sign-On (SSO) daemon. This service uses port 8381.
- **mysql**  
The MySQL database daemon. This service uses port 3306.

For more information about the ports DX NetOps Performance Management requires to work properly, see [Review Installation Requirements and Considerations](#).

### Follow these steps:

1. Check the status of a daemon by issuing the following command:

```
systemctl status <service_name>
```

#### Example:

```
systemctl status mysql
```

#### – **service\_name**

The Linux daemon name.

#### Options:

- **caperfcenter\_console** : The NetOps Portal Console service
- **caperfcenter\_devicemanager** : The Device Manager service
- **caperfcenter\_eventmanager** : The Event Manager service
- **caperfcenter\_sso** : The SSO service
- **mysql** : The MySQL Database service

2. Access NetOps Portal at the following URL:

```
http://<PC_host>:8181/pc/desktop/page
```

#### – **PC\_host**

The NetOps Portal hostname alias, the FQHN, or the IP.

**Example:** my\_hostname.domain.com

**Default:** The FQDN of the server.

If the NetOps Portal login screen appears, NetOps Portal has installed successfully.

### IMPORTANT

At the initial login, you must change the **admin** and **user** passwords. Change these passwords immediately after a fresh install.

## (Optional) Restart the NFA OData Service

If NFA is registered as a data source, restart the NFA OData service.

### Follow these steps:

1. Click **Administration Services**, and then **Services**.
2. Right-click the **CA NFA OData Service**, and then click **Restart**.

The NFA OData Service is restarted.

## **(Optional) Review the Log Files**

You can track events that occur during the installation using the following log files:

- **Installation Errors and Configuration Events**

`<installation_directory>/PerformanceCenter/InstallLogs`

During the installation, a history file that indicates the installed version is generated in this directory.

- **Device Manager Daemon**

`<installation_directory>/PerformanceCenter/DM/logs`

- **Website and Console Errors**

`<installation_directory>/PerformanceCenter/PC/logs`

- **MySQL Database Errors**

`<installation_directory>/MySQL/data/<hostname>.err`

- **MySQL Initialization Results**

`<installation_directory>/MySQL/mysql_initialize_results.txt`

During the installation, this initialization results log file is generated in this directory.

- **Event Manager (Other Events)**

`<installation_directory>/PerformanceCenter/EM/logs`

- **User Authentication (SSO)**

`<installation_directory>/PerformanceCenter/sso/logs`

**NOTE**

`/opt/CA` is the default installation directory.

## **Prepare to Install the Data Repository**

Ensure that you can install the data repository successfully by preparing for the installation.

**IMPORTANT**

The content in this article is for preparing to install the data repository *if you are installing 23.3.3 and higher*. If you are installing 23.3.2 and lower, see [Prepare to Install the Data Repository for 23.3.2 and Lower](#).

### **(23.3.3 and higher)**

Complete the following procedures before installing the data repository:

1. [Verify the Prerequisites](#)
2. [Verify the Prerequisites to a Data Repository Installation on VMware Virtual Machines](#)
3. [Verify the Prerequisites to a Data Repository Installation on Shared Storage](#)
4. [Set a Unique Hostname for the Data Repository Host](#)
5. [Verify Time Synchronization](#)
6. [Enable the rc-local Service](#)
7. (If you cannot use the root user account to install and upgrade the data repository) [Configure the Sudo User Account Privileges to Install and Upgrade the Data Repository](#)
8. (If you cannot use the root user account to manage the data repository) [Configure the Sudo User Account Privileges for Vertica Management](#)
9. (If you do not want to provide an SSH password for the install user) [Set up the install user for passwordless SSH](#).
10. [Next Steps](#).

### **Verify the Prerequisites**

Before you install the data repository, ensure that you have met the following prerequisite steps:

- (Cluster installations only) The `dr_install.sh` installation script and the `dr_validate.sh` validation script require SSH to run commands remotely. SSH requires authentication. The scripts accept SSH authentication using one of the following methods:
  - Prompt or provide the SSH password in the `drinstall.properties` file.
  - Prompt or provide the private key for passwordless SSH (public key) authentication.
  - [Set up full passwordless SSH](#).
 You have determined which SSH authentication method you will use.
- [You have reviewed the installation requirements and considerations](#). This includes verifying the following:
  - Your environment is using a Vertica-supported operating system (OS) that the `dr_install.sh` script (the data repository installer) supports.
  - You have the packages that the data repository installer requires.
  - You have verified that the ports required to allow the data repository to communicate and function properly are open to remote access.
- You have reviewed [the Vertica 10.1 documentation](#). The data aggregator stores inventory, configuration, and polled data in a Vertica database using the data repository. NetOps Portal includes the following Vertica versions:

NetOps Portal Version	Included Vertica version
23.3.x	10.1.1-20

- If you require use of an elevated privileges tool for the data repository installation, Vertica requires that you use the `sudo` command and not any other elevated privileges tool.  
For more information about this requirement, see [the Vertica 10.1 documentation](#).
- The data repository installation and upgrade requires passwordless SSH. The installation process sets up passwordless SSH for you, or [you can set this up](#). The Vertica administrative tools, such as `admintools`, `backup`, `migration`, and the `copycluster` operations, also require passwordless SSH.

#### TIP

- For maximum security and maximum control of the allowable commands, install and upgrade the data repository by configuring and using a non-root user using `sudo` as the passwordless SSH user.
- You can enrich the security that is already afforded by passwordless SSH and ensure stronger protection by firewalling the Vertica cluster from other systems.
- You have verified that you have at least 2 GB of swap space on the data repository host.
- You have verified that the data repository hosts use the `ext3`, `ext4`, or `XFS` file systems.

#### IMPORTANT

The default file system for Red Hat Enterprise Linux (RHEL) 7.x/8.x and Oracle Linux (OL) 7.x is `XFS`. The default file system for SUSE Linux Enterprise Server (SLES) is `btrfs`. The disks with the Vertica files for the `data` and `catalog` directories can use the `ext3`, `ext4`, or `XFS` file systems.

**Best practice:** For best database performance, use the `ext4` file system.

- You have verified that the data repository processes have been excluded from Security-Enhanced Linux (SELinux) restrictions by disabling SELinux or by setting it to permissive mode on the computer where you are going to install the data repository. By default, some Linux distributions have SELinux enabled, which does not allow the data repository to function properly.

For more information:

- About SELinux configuration in Vertica, see [the Vertica 10.1 documentation](#).
- About SELinux modes and RHEL, including how to write SELinux security policies, see [the Red Hat documentation](#).
- You have enabled TCP forwarding by setting `AllowTcpForwarding = Yes` in the `/etc/ssh/sshd_config` file on the Vertica hosts, on both source and destination systems. Backing up the data repository and the `copycluster` command within Vertica using the `vbr` utility requires that TCP forwarding be enabled. This allows the utility to forward connections from database hosts to backup hosts.

For more information, see [the Vertica 10.1 documentation](#).

- To avoid database corruption and to prevent scanning by a local instance of an antivirus client and scanning by a remote antivirus instance, you have excluded the installation directory, and the following subdirectories:

- /opt/vertica/\*
- /opt/vconsole/\*
- The specified data directory.

**Default:** /data

#### IMPORTANT

Ensure that the `data` directory is on a separate mount from the `catalog` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

- The specified `catalog` directory.

**Default:** /catalog

#### IMPORTANT

Ensure that the `catalog` directory is on a separate mount from the `data` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

- The Vertica `/tmp/vbr/*` temporary file.
- The directory where you back up the data repository.

For more information about disk locations, see the [Vertica 10.1 documentation](#).

- To avoid the data repository installation from failing, you have ensured that a file named `release` is not in the `/etc` directory. Remove the file if it exists.
- You have verified the access according to your installation type:
  - **Single node:** Root access is required to install the data repository. Determine whether you have this access level.
  - **Cluster:** Verify that the root user or sudo user can create database administrator user accounts, or can have an administrator create these accounts.
- You have verified that central processing unit (CPU) frequency scaling is disabled. Disable CPU frequency scaling through the host system basic input/output system (BIOS) and operating system (OS) settings.

#### NOTE

If CPU frequency scaling is enabled, you might experience inconsistent performance for similar queries in Vertica. CPU frequency scaling can cause observable slowness and variation in dashboard loading.

- You have verified that you are not using Logical Volume Manager (LVM) for the `data` and `catalog` directories.
- (Cluster installations only) You have verified that all the hosts in the cluster are in the same subnet.
- (Cluster installations only installing using root user) You have verified that the root user can use Secure Shell (SSH) to log in (`ssh`) to all the hosts in the cluster.

#### NOTE

Set up SSH for the root user.

- (Cluster installations only installing using non-root user using sudo) You have verified that the non-root user can use SSH to log in (`ssh`) to all the hosts in the cluster.

#### NOTE

Set up SSH for the non-root user using `sudo`.

- You have verified that the default shell environment is `bash`.
- (Cluster installations only) You have selected the hosts where you plan to install the data repository nodes.

#### NOTE

You deploy database software on each participating host in a cluster. The software represents a node in the cluster. A three-node cluster represents the simplest configuration that can tolerate the loss of a single node. You can, however, include more than three hosts in the cluster.



**IMPORTANT**

If more than one node fails or shuts down, the data repository is no longer available for use and the data aggregator shuts down automatically.

**Verify the Prerequisites to a Data Repository Installation on VMware Virtual Machines**

For best performance, install the data repository in a bare-metal environment. However, if you plan to install the data repository on virtual machines (VMs), verify the following requirements:

- You are using VMware version 5.5 or later.
- The number of VMs per host does not exceed the number of physical processors.
- Pre-allocate and reserve 4 GB of memory for each of the VMs.
- Each VM has a dedicated 10 GB NIC.
- You have disabled CPU frequency scaling at the host level and for each VM.
- You have disabled VMotion. VMotion can disrupt communication, and can cause the data repository to shut down.
- You have set the VMware parameters for hugepages to the VMware 5.5 default values.
- You have verified the hardware and network performance.

**TIP**

You can verify performance using the `vioperf` Vertica utility.

For more information about this utility, see [the Vertica 10.1 documentation](#).

For more information about how to run Vertica on VMs, see [the Vertica documentation](#).

**Verify the Prerequisites to a Data Repository Installation on Shared Storage**

Before installing the data repository on storage area network (SAN), verify the following requirements:

- The hosts have no contention for disk space or bandwidth.
- Each host has a unique catalog and data location. The hosts cannot share the location for these directories.
- The storage has enough input/output (I/O) bandwidth for each node to access the storage independently.

**TIP**

You can verify the I/O bandwidth by simultaneously running the Vertica `vioperf` utility from all hosts in the data repository cluster.

For more information, see the following procedures.

**Set a Unique Hostname for the Data Repository Host**

Set a unique hostname for each data repository host in the cluster.

**Follow these steps:**

1. As the root user, log in to each data repository host, and verify the unique hostname.  
The hostname must be associated with the IP address and *not* with the loopback address of 127.0.0.1.
2. Verify that the following lines appear in the `/etc/hosts` file on each computer:  
Do not remove the following line, or various programs  
# that require network functionality will fail.  
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4  
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6  
<IP address of your host> <YourHostName YourHostName.domain>
3. If you change the file, issue the following command:  
`systemctl restart network`  
The `/etc/hosts` file is configured correctly.  
The unique hostname is set.

#### 4. (Cluster installations only) Complete the following:

- a. The hostnames of all hosts in the cluster must resolve correctly. If the hostname resolution is incorrect, the data repository cluster does not install or work properly. All participating hosts in the cluster must use static IP or permanently leased DHCP addresses. Set up the `/etc/hosts` file on each of the hosts you selected for the cluster. The hosts file must contain entries for all hosts in the cluster.

**Example:** This example `/etc/hosts` file is for a cluster where the hosts are named `host01`, `host02`, and `host03`:

```
127.0.0.1 localhost.localdomain localhost
192.168.13.128 host01.domain host01
192.168.13.129 host02.domain host02
192.168.13.130 host03.domain host03
```

#### IMPORTANT

- Do not remove the loopback address (`127.0.0.1`) line.
- The local data repository hostname cannot be on the `127.0.0.1` line.
- Do not use the loopback address or `localhost` name when you are defining hosts in the cluster.

- b. Verify that hostname resolution works for each host in the cluster.

For example, on `host01`, the following syntax is correct:

```
$ /bin/hostname -f
host01
```

Hostname resolution is configured.

A unique hostname is set for each data repository host in the cluster.

### Verify Time Synchronization

The data repository requires time synchronization using the Network Time Protocol (NTP) daemon. All data source consoles must use the NTP daemon. On Linux servers, the NTP daemon ensures that the clocks on the hosts are synchronized for timing purposes. Verify that the daemon is running on all DX NetOps Performance Management host servers based on your version:

- [Verify on SUSE Linux Enterprise Server \(SLES\)](#)
- [Verify on Red Hat Enterprise Linux \(RHEL\) 7.x/8.x and Oracle Linux \(OL\)](#)

### Verify Time Synchronization on SUSE Linux Enterprise Server

#### Follow these steps:

1. Open a console and issue the following command:
 

```
$ systemctl status ntpd
```
2. Verify that the NTP daemon is in an active (running) state.
3. Start and enable the NTP daemon by issuing the following command:
 

```
$ systemctl start ntpd
$ systemctl enable ntpd
```

The daemon is started.

### Verify Time Synchronization on Red Hat Enterprise Linux 7.x/8.x and Oracle Linux

RHEL 7.x/8.x and OL 7.x run NTP with `chronyd`.

#### Follow these steps:

1. Open a console and issue the following command:
 

```
$ systemctl status chronyd
```

2. Verify that the `chrony` daemon is in an active (running) state.
3. Start and enable the `chrony` daemon by issuing the following command:

```
$ systemctl start chronyd
$ systemctl enable chronyd
```

The daemon is started.

## **Enable the rc-local Service**

Follow these steps:

1. Edit the `/usr/lib/systemd/system/rc-local.service` file.
2. Change to the following based on your installation:

– **(RHEL 7.x/8.x, OL 7.x, CentOS Linux 7/8, Rocky Linux 8.x)**

```
[Unit]
Description=/etc/rc.local
ConditionPathExists=/etc/rc.local
```

```
[Service]
Type=forking
ExecStart=/etc/rc.local start
TimeoutSec=0
StandardOutput=tty
RemainAfterExit=yes
```

```
[Install]
WantedBy=multi-user.target
```

– **(SLES)**

```
[Unit]
Description=/etc/init.d/boot.local Compatibility
ConditionFileIsExecutable=/etc/init.d/boot.local
After=basic.target
```

```
[Service]
Type=forking
ExecStart=/etc/init.d/boot.local start
TimeoutSec=0
RemainAfterExit=yes
```

```
[Install]
WantedBy=multi-user.target
```

3. Enable the service, and then check if the service is active by issuing the following commands:

```
systemctl enable rc-local
systemctl start rc-local
systemctl is-active rc-local
```

The `is-active` command should return "active".

The `rc-local` service is enabled.

## **Configure the Sudo User Account Privileges to Install and Upgrade the Data Repository**

Due to security policies, in some environments, you cannot use the root user to install and upgrade the data repository. In this case, configure a non-root user to use `sudo` to install and upgrade the data repository.

For cluster environments, complete this procedure on *each* host in the cluster.

### Follow these steps:

1. Locate the `/etc/sudoers` file.
2. Add command aliases with the following permissions to the file based on your installation, *replacing* the user who will install and manage the Vertica node (`sudouser`) and the database administrator system account that the Vertica install creates, and who will own and run Vertica (`dbadmin`) with your values:

#### – (RHEL 7.x/8.x)

```
Cmnd_Alias CA_DATAREP=/opt/vertica/sbin/install_vertica,/tmp/installDR.bin,/opt/
CA/IMDataRepository_vertical0/dr_validate.sh,/opt/CA/IMDataRepository_vertical0/
dr_install.sh,/usr/bin/vim,/usr/bin/reboot,/opt/CA/IMDataRepository_vertical0/
RemoteEngineer/re.sh,/bin/mkdir*,/usr/bin/whoami,/bin/echo,/sbin/service,/bin/
grep,/usr/bin/test,/sbin/iptables,/opt/vertica/oss/python/bin/python,/usr/bin/
tee,/usr/sbin/ntpd,/etc/init.d/ntpd,/sbin/blockdev,/etc/init.d/sshd,/etc/sysconfig/
sshd,/etc/ssh/sshd_config,/bin/su,/usr/sbin/sshd restart,/usr/bin/ssh,/bin/df,/bin/
mv,/bin/rm,/usr/bin/install,/usr/bin/env
Cmnd_Alias VERTICA = /opt/vertica/bin/,/opt/vertica/sbin/,/opt/vertica/oss/python/
bin/
Cmnd_Alias VERTICA_INSTALL = /bin/echo,/bin/ps -A,/bin/cp /opt/vertica/config/
admintools.conf /opt/vertica/config/admintools.conf.bak.*,/bin/rm -rf /tmp/
dbRPM.rpm,/bin/df --portability /tmp,/usr/bin/install --owner * --mode 700 -d *,/
bin/mv -f /tmp/vstage-*/file /tmp/*,/bin/rm -rf /tmp/vstage-*/,/usr/bin/id *,/bin/
cp -T /opt/vertica/* /tmp/vstage-*/,/bin/su --login <dbadmin> *,/bin/mkdir -p /
opt/vertica/*,/bin/touch /opt/vertica/*,/bin/rm -rf /opt/vertica/*,/bin/mv -f /
tmp/vstage-* /opt/vertica/*,/bin/mkdir -p /opt/vertica/*,/bin/touch /opt/vertica/
config/users/<dbadmin>/agent.conf,/bin/su <dbadmin> *,/bin/sh -c *,/usr/bin,/opt/
vertica/share/binlib/test/*,/usr/bin/su <dbadmin>,/bin/test [ -e /* ],/usr/bin/[ -
e /* ],/sbin/usermod,/usr/sbin/sshd,/sbin/sysctl,/bin/ls,/bin/free,/bin/cat,/bin/
getent,/bin/stat,/sbin/lvdisplay
Cmnd_Alias USEFUL = /usr/bin/lshw,/usr/bin/yum,/bin/rpm,/sbin/reboot,/sbin/
shutdown,/usr/bin/cpan,/bin/chgrp,/bin/chmod,/bin/chown,/bin/mnt,/usr/bin/test,/
bin/[,/sbin/service,/bin/systemctl
## Allows the Data Repository user to manage the Data Repository
<sudouser> ALL = CA_DATAREP, VERTICA , VERTICA_INSTALL , USEFUL
Defaults env_keep += "VERT_DBA_USR VERT_DBA_HOME VERT_DBA_GRP VERT_DBA_DATA_DIR
_ENV_VPWD_VAR"
```

#### • **dbadmin**

The database administrator system account that the Vertica install creates, and who will own and run Vertica.

#### – (SLES 12)

```
Cmnd_Alias CA_DATAREP =/opt/vertica/sbin/install_vertica,/tmp/installDR.bin,/opt/
CA/IMDataRepository_vertical0/dr_validate.sh,/opt/CA/IMDataRepository_vertical0/
dr_install.sh,/usr/bin/vim,/usr/bin/reboot,/opt/CA/IMDataRepository_vertical0/
RemoteEngineer/re.sh,/usr/bin/mkdir, /sbin/SuSEfirewall2 off *,/usr/bin/whoami,/
usr/bin/echo,/usr/bin/id,/usr/bin/env,/usr/sbin/service,/usr/bin/grep,/usr/bin/
test,/sbin/iptables,/opt/vertica/oss/python/bin/python,/usr/bin/tee,/usr/sbin/
```

```

ntpd,/etc/init.d/ntpd,/sbin/blockdev,/etc/init.d/sshd,/etc/sysconfig/sshd,/etc/ssh/
sshd_config,/usr/bin/su,/usr/sbin/sshd restart,/usr/bin/ssh,/usr/bin/sh,/usr/bin/
install,/usr/bin/env
Cmnd_Alias VERTICA = /opt/vertica/bin/,/opt/vertica/sbin/,/opt/vertica/oss/python/
bin/
Cmnd_Alias VERTICA_INSTALL = /usr/bin/echo,/usr/bin/ps -A,/usr/bin/cp /opt/vertica/
config/admintools.conf /opt/vertica/config/admintools.conf.bak.*,/usr/bin/rm -rf /
tmp/dbRPM.rpm,/usr/bin/df --portability /tmp,/usr/bin/install --owner * --mode 700
-d *,/usr/bin/mv -f /tmp/vstage-*/file /tmp/*,/usr/bin/rm -rf /tmp/vstage-*,/usr/
bin/id *,/usr/bin/cp -T /opt/vertica/* /tmp/vstage-*,/usr/bin/su --login <dbadmin>
*,/usr/bin/mkdir -p /opt/vertica/*,/usr/bin/touch /opt/vertica/*,/usr/bin/rm -
rf /opt/vertica/*,/usr/bin/mv -f /tmp/vstage-* /opt/vertica/*,/usr/bin/mkdir -p /
opt/vertica/*,/usr/bin/touch /opt/vertica/config/users/<dbadmin>/agent.conf,/usr/
bin/su <dbadmin> *,/usr/bin/sh -c *,/opt/vertica/share/binlib/test/*,/usr/bin/su
<dbadmin>,/usr/bin/test [ -e /* ],/usr/bin/[ -e /* ],/usr/sbin/usermod,/usr/sbin/
sshd,/sbin/sysctl,/usr/bin/ls,/usr/bin/free,/usr/bin/cat,/usr/bin/getent,/usr/bin/
stat,/sbin/SuSEfirewall2,/usr/sbin/lvdisplay,/usr/bin/df -Th *,/usr/bin/[ -d /*
Cmnd_Alias USEFUL = /usr/bin/lshw,/usr/bin/yum,/bin/rpm,/sbin/reboot,/sbin/
shutdown,/usr/bin/cpan,/bin/chgrp,/bin/chmod,/bin/chown,/bin/mnt,/usr/bin/test,/
bin/[,/sbin/service,/bin/systemctl
## Allows the Data Repository user to manage the Data Repository
<sudouser> ALL = CA_DATAREP, VERTICA , VERTICA_INSTALL , USEFUL
Defaults env_keep += "VERT_DBA_USR VERT_DBA_HOME VERT_DBA_GRP VERT_DBA_DATA_DIR
_ENV_VPWD_VAR"

```

- **sudouser**  
The user who will install and manage the Vertica node.
- **dbadmin**  
The database administrator system account that the Vertica install creates, and who will own and run Vertica.

The sudo user account privileges are configured for the install and upgrade of the data repository.

### Configure the Sudo User Account Privileges for Vertica Management

To manage the Vertica installation as a non-root user using sudo, configure the sudo user account privileges for the non-root user.

For cluster environments, complete this procedure on *each* host in the cluster.

#### **Follow these steps:**

1. Locate the `/etc/sudoers` file.
2. Insert the following line, adding a command alias that details the commands that the sudo user can issue:
 

```

Cmnd_Alias CA_MGMT = /opt/CA/IMDataRespository_vertical0/RemoteEngineer/re.sh
## Allows the Data Repository user to manage the Data Repository
      
```
3. Below the `Cmnd_Alias` line, do *one* of the following:
  - If you configured the sudo user account privileges to install and upgrade the data repository, insert the following `sudouser` line, replacing `sudouser` with the user who will install and manage the Vertica node, and appending `CA_MGMT` to the line:
 

```

<sudouser> ALL = CA_DATAREP, VERTICA , VERTICA_INSTALL , USEFUL, CA_MGMT
          
```

- **sudouser**

The user who will install and manage the Vertica node. This command alias details the commands that the sudo user can issue.

- Insert the following sudouser line and replace `sudouser` with the user who will install and manage the Vertica node:

```
<sudouser> ALL = CA_MGMT
```

- **sudouser**

The user who will install and manage the Vertica node. This command alias details the commands that the sudo user can issue.

4. Save your changes.

The sudo user account privileges are configured for Vertica management.

### **Set Up the Install User for Passwordless SSH**

If you do not want to have to provide the SSH password or the private key for public authentication, during the data repository installation or upgrade, set up either the root or sudo user for passwordless SSH.

Repeat this procedure for each pair of hosts.

#### **NOTE**

Passwordless SSH is automatically set up for the data repository admin user when you install the data repository.

#### **Follow these steps:**

1. Open a console and log in to the data repository host as the root or sudo user.
2. Issue the following commands:

```
ssh-keygen -N "" -t rsa -f ~/.ssh/id_rsa
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys2
chmod 644 ~/.ssh/authorized_keys2
```

A public and private key are set up for the user, and the public key is appended to the authorization file.

3. Copy the user's public key into the list of authorized keys on the remote hosts by issuing the following command:

```
ssh-copy-id -i <user>@<remotehost>
```

- **remotehost**

A host in the cluster where you are copying the SSH ID.

4. At the prompt, enter the user's password.  
The user is set up for passwordless SSH.
5. To verify that you have set up the user for passwordless SSH, log in to the remote host from the local host by issuing the following command:

```
ssh <user>@<remotehost> ls
```

- **remotehost**

A host in the cluster where you are copying the SSH ID.

You have set up the user for passwordless SSH successfully if you are not prompted for a password. You also see a directory listing.

### **Next Steps**

- [Install the data repository.](#)

## **Prepare to Install the Data Repository for 23.3.2 and Lower**

Ensure that you can install the data repository successfully by preparing for the installation.

**IMPORTANT**

The content in this article is for preparing to install the data repository *if you are installing 23.3.2 and lower*. If you are installing 23.3.3 and higher, see [Prepare to Install the Data Repository](#).

**(23.3.2 and lower)**

Complete the following procedures before installing the data repository:

1. [Verify the Prerequisites](#)
2. [Verify the Prerequisites to a Data Repository Installation on VMware Virtual Machines](#)
3. [Verify the Prerequisites to a Data Repository Installation on Shared Storage](#)
4. [Set a Unique Hostname for the Data Repository Host](#)
5. [Verify Time Synchronization](#)
6. [Enable the rc-local Service](#)
7. If you are installing or upgrading the data repository using a non-root user using sudo, complete one of the following based on your implementation:
  - (If the root passwordless SSH user is set up, and you need to use sudo to install and run the data repository) [Configure the Sudo User Account Privileges for a Root Passwordless SSH SetUp](#)
  - (If the root passwordless SSH user *is not* set up, and you need to use sudo to install and run the data repository) [Configure the Sudo User Account Privileges when Root Passwordless SSH is not Set Up](#)
8. [Next Steps](#)

**Verify the Prerequisites**

Before you install the data repository, ensure that you have met the following prerequisite steps:

- [You have reviewed the installation requirements and considerations](#). This includes verifying the following:
  - Your environment is using a Vertica-supported operating system (OS) that the `dr_install.sh` script (the data repository installer) supports.
  - You have the packages that the data repository installer requires.
  - You have verified that the ports required to allow the data repository to communicate and function properly are open to remote access.
- You have reviewed [the Vertica 10.1 documentation](#). The data aggregator stores inventory, configuration, and polled data in a Vertica database using the data repository. NetOps Portal includes the following Vertica versions:

NetOps Portal Version	Included Vertica version
23.3.x	10.1.1-20

- If you require use of an elevated privileges tool for the data repository installation, Vertica requires that you use the `sudo` command and not any other elevated privileges tool.  
For more information about this requirement, see [the Vertica 10.1 documentation](#).
- The data repository installation and upgrade requires passwordless SSH. The installation process sets up passwordless SSH for you, or [you can set this up manually](#). The Vertica administrative tools, such as `admintools`, `backup`, `migration`, and the `copycluster` operations, also require passwordless SSH.

**TIP**

- For maximum security and maximum control of the allowable commands, install and upgrade the data repository by configuring and using a non-root user using `sudo` as the passwordless SSH user.
  - You can enrich the security that is already afforded by passwordless SSH and ensure stronger protection by firewalling the Vertica cluster from other systems.
- You have verified that you have at least 2 GB of swap space on the data repository host.
- You have verified that the data repository hosts use the `ext3`, `ext4`, or `XFS` file systems.

**IMPORTANT**

The default file system for Red Hat Enterprise Linux (RHEL) 7.x/8.x and Oracle Linux (OL) 7.x is XFS. The default file system for SUSE Linux Enterprise Server (SLES) is btrfs. The disks with the Vertica files for the `data` and `catalog` directories can use the ext3, ext4, or XFS file systems.

**Best practice:** For best database performance, use the ext4 file system.

- You have verified that the data repository processes have been excluded from Security-Enhanced Linux (SELinux) restrictions by disabling SELinux or by setting it to permissive mode on the computer where you are going to install the data repository. By default, some Linux distributions have SELinux enabled, which does not allow the data repository to function properly.

For more information:

- About SELinux configuration in Vertica, see [the Vertica 10.1 documentation](#).
  - About SELinux modes and RHEL, including how to write SELinux security policies, see [the Red Hat documentation](#).
  - You have enabled TCP forwarding by setting `AllowTcpForwarding = Yes` in the `/etc/ssh/sshd_config` file on the Vertica hosts, on both source and destination systems. Backing up the data repository and the `copycluster` command within Vertica using the `vbr` utility requires that TCP forwarding be enabled. This allows the utility to forward connections from database hosts to backup hosts.
- For more information, see [the Vertica 10.1 documentation](#).

- To avoid database corruption and to prevent scanning by a local instance of an antivirus client and scanning by a remote antivirus instance, you have excluded the installation directory, and the following subdirectories:

- `/opt/vertica/*`
- `/opt/vconsole/*`
- The specified `data` directory.

**Default:** `/data`

**IMPORTANT**

Ensure that the `data` directory is on a separate mount from the `catalog` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

- The specified `catalog` directory.

**Default:** `/catalog`

**IMPORTANT**

Ensure that the `catalog` directory is on a separate mount from the `data` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

- The Vertica `/tmp/vbr/*` temporary file.
- The directory where you back up the data repository.

For more information about disk locations, see the [Vertica 10.1 documentation](#).

- To avoid the data repository installation from failing, you have ensured that a file named `release` is not in the `/etc` directory. Remove the file if it exists.
- You have verified the access according to your installation type:
  - **Single node:** Root access is required to install the data repository. Determine whether you have this access level.
  - **Cluster:** Verify that the root user or sudo user can create database administrator user accounts, or can have an administrator create these accounts.
- You have verified that central processing unit (CPU) frequency scaling is disabled. Disable CPU frequency scaling through the host system basic input/output system (BIOS) and operating system (OS) settings.



**NOTE**

If CPU frequency scaling is enabled, you might experience inconsistent performance for similar queries in Vertica. CPU frequency scaling can cause observable slowness and variation in dashboard loading.

- You have verified that you are not using Logical Volume Manager (LVM) for the `data` and `catalog` directories.
- (Cluster installations only) You have verified that all the hosts in the cluster are in the same subnet.
- (Cluster installations only) You have verified that the root user can use Secure Shell (SSH) to log in (`ssh`) to all the hosts in the cluster.

**NOTE**

Set up SSH for the root user.

- You have verified that the default shell environment is `bash`.
- (Cluster installations only) You have selected the hosts where you plan to install the data repository nodes.

**NOTE**

You deploy database software on each participating host in a cluster. The software represents a 'node' in the cluster. A three-node cluster represents the simplest configuration that can tolerate the loss of a single node. You can, however, include more than three hosts in the cluster.

**IMPORTANT**

If more than one node fails or shuts down, the data repository is no longer available for use and the data aggregator shuts down automatically.

### **Verify the Prerequisites to a Data Repository Installation on VMware Virtual Machines**

For best performance, install the data repository in a bare-metal environment. However, if you plan to install the data repository on virtual machines (VMs), verify the following requirements:

- You are using VMware version 5.5 or later.
- The number of VMs per host does not exceed the number of physical processors.
- Pre-allocate and reserve 4 GB of memory for each of the VMs.
- Each VM has a dedicated 10 GB NIC.
- You have disabled CPU frequency scaling at the host level and for each VM.
- You have disabled VMotion. VMotion can disrupt communication, and can cause the data repository to shut down.
- You have set the VMware parameters for hugepages to the VMware 5.5 default values.
- You have verified the hardware and network performance.

**TIP**

You can verify performance using the `vioperf` Vertica utility.

For more information about this utility, see [the Vertica 10.1 documentation](#).

For more information about how to run Vertica on VMs, see [the Vertica documentation](#).

### **Verify the Prerequisites to a Data Repository Installation on Shared Storage**

Before installing the data repository on storage area network (SAN), verify the following requirements:

- The hosts have no contention for disk space or bandwidth.
- Each host has a unique catalog and data location. The hosts cannot share the location for these directories.
- The storage has enough input/output (I/O) bandwidth for each node to access the storage independently.

**TIP**

You can verify the I/O bandwidth by simultaneously running the Vertica `vioperf` utility from all hosts in the data repository cluster.

For more information, see the following procedures.

## Set a Unique Hostname for the Data Repository Host

Set a unique hostname for each data repository host in the cluster.

### Follow these steps:

1. As the root user, log in to each data repository host, and verify the unique hostname.  
The hostname must be associated with the IP address and *not* with the loopback address of 127.0.0.1.
2. Verify that the following lines appear in the `/etc/hosts` file on each computer:  
Do not remove the following line, or various programs  
# that require network functionality will fail.  
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4  
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6  
<IP address of your host> <YourHostName YourHostName.domain>
3. If you change the file, issue the following command:  
`systemctl restart network`  
The `/etc/hosts` file is configured correctly.  
The unique hostname is set.
4. (Cluster installations only) Complete the following:
  - a. The hostnames of all hosts in the cluster must resolve correctly. If the hostname resolution is incorrect, the data repository cluster does not install or work properly. All participating hosts in the cluster must use static IP or permanently leased DHCP addresses. Set up the `/etc/hosts` file on each of the hosts you selected for the cluster. The hosts file must contain entries for all hosts in the cluster.  
**Example:** This example `/etc/hosts` file is for a cluster where the hosts are named host01, host02, and host03:  
127.0.0.1 localhost.localdomain localhost  
192.168.13.128 host01.domain host01  
192.168.13.129 host02.domain host02  
192.168.13.130 host03.domain host03

### IMPORTANT

- Do not remove the loopback address (127.0.0.1) line.
- The local data repository hostname cannot be on the 127.0.0.1 line.
- Do not use the loopback address or localhost name when you are defining hosts in the cluster.

- b. Verify that hostname resolution works for each host in the cluster.  
For example, on host01, the following syntax is correct:

```
$ /bin/hostname -f
host01
```

Hostname resolution is configured.

A unique hostname is set for each data repository host in the cluster.

## Verify Time Synchronization

The data repository requires time synchronization using the Network Time Protocol (NTP) daemon. All data source consoles must use the NTP daemon. On Linux servers, the NTP daemon ensures that the clocks on the hosts are synchronized for timing purposes. Verify that the daemon is running on all DX NetOps Performance Management host servers based on your version:

- [Verify on SUSE Linux Enterprise Server \(SLES\)](#)
- [Verify on Red Hat Enterprise Linux \(RHEL\) 7.x/8.x and Oracle Linux \(OL\)](#)

## **Verify Time Synchronization on SUSE Linux Enterprise Server**

### **Follow these steps:**

1. Open a console and issue the following command:  

```
$ systemctl status ntpd
```
2. Verify that the NTP daemon is in an active (running) state.
3. Start and enable the NTP daemon manually by issuing the following command:  

```
$ systemctl start ntpd
```

```
$ systemctl enable ntpd
```

The daemon is started.

## **Verify Time Synchronization on Red Hat Enterprise Linux 7.x/8.x and Oracle Linux**

RHEL 7.x/8.x and OL 7.x run NTP with `chronyd`.

### **Follow these steps:**

1. Open a console and issue the following command:  

```
$ systemctl status chronyd
```
2. Verify that the `chrony` daemon is in an active (running) state.
3. Start and enable the `chrony` daemon by issuing the following command:  

```
$ systemctl start chronyd
```

```
$ systemctl enable chronyd
```

The daemon is started.

## **Enable the rc-local Service**

### **Follow these steps:**

1. Edit the `/usr/lib/systemd/system/rc-local.service` file.
2. Change to the following based on your installation:

#### **– (RHEL 7.x/8.x, OL 7.x, CentOS Linux 7/8, Rocky Linux 8.x)**

```
[Unit]
Description=/etc/rc.local
ConditionPathExists=/etc/rc.local
```

```
[Service]
Type=forking
ExecStart=/etc/rc.local start
TimeoutSec=0
StandardOutput=tty
RemainAfterExit=yes
```

```
[Install]
WantedBy=multi-user.target
```

#### **– (SLES)**

```
[Unit]
Description=/etc/init.d/boot.local Compatibility
ConditionFileIsExecutable=/etc/init.d/boot.local
After=basic.target
```

```
[Service]
```

```
Type=forking
ExecStart=/etc/init.d/boot.local start
TimeoutSec=0
RemainAfterExit=yes

[Install]
WantedBy=multi-user.target
```

3. Enable the service, and then check if the service is active by issuing the following commands:

```
systemctl enable rc-local
systemctl start rc-local
systemctl is-active rc-local
```

The `is-active` command should return "active".

The `rc-local` service is enabled.

### **Configure the Sudo User Account Privileges for a Root Passwordless SSH SetUp**

If you have set up the root passwordless SSH user, and you will install the data repository as the non-root user using `sudo`, configure the `sudo` user account privileges.

For cluster environments, complete this procedure on each host in the cluster.

#### **Follow these steps:**

1. Locate the `/etc/sudoers` file.
2. Add a command alias that details the commands that the `sudo` user can issue by issuing the following command, *replacing* `sudouser` with the user who will install and manage the Vertica node:

```
Cmnd_Alias CA_DATAREP = /tmp/installDR.bin,/opt/CA/IMDataRepository_vertica10/
dr_validate.sh,/opt/CA/IMDataRepository_vertica10/dr_install.sh,/usr/bin/vim,/usr/
bin/reboot,/usr/bin/yum,/opt/CA/IMDataRepository_vertica10/RemoteEngineer/re.sh
## Allows the Data Repository user to manage the Data Repository
<sudouser> ALL = CA_DATAREP
```

#### **– *sudouser***

The user who will install and manage the Vertica node. This command alias details the commands that the `sudo` user can issue.

The `sudo` user account privileges are configured.

### **Configure the Sudo User Account Privileges when Root Passwordless SSH is not Set Up**

Due to security policies, in some environments, you cannot use the root user as the passwordless SSH user on the host servers. On RHEL 7.x or SLES 12, you can install the data repository without requiring that level of access by using the `sudo` user account.

#### **Follow these steps:**

1. Locate the `/etc/sudoers` file.
2. Add command aliases with the following permissions to the file based on your installation, *replacing* the user who will install and manage the Vertica node (`sudouser`) and the database administrator system account that the Vertica install creates, and who will own and run Vertica (`dbadmin`) with your values:

#### **– (RHEL 7.x/8.x)**

```
Cmnd_Alias CA_DATAREP=/opt/vertica/sbin/install_vertica,/tmp/installDR.bin,/opt/
CA/IMDataRepository_vertica10/dr_validate.sh,/opt/CA/IMDataRepository_vertica10/
dr_install.sh,/usr/bin/vim,/usr/bin/reboot,/opt/CA/IMDataRepository_vertica10/
```

```

RemoteEngineer/re.sh,/bin/mkdir*,/usr/bin/whoami,/bin/echo,/sbin/service,/bin/
grep,/usr/bin/test,/sbin/iptables,/opt/vertica/oss/python/bin/python,/usr/bin/
tee,/usr/sbin/ntpd,/etc/init.d/ntpd,/sbin/blockdev,/etc/init.d/sshd,/etc/sysconfig/
sshd,/etc/ssh/sshd_config,/bin/su,/usr/sbin/sshd restart,/usr/bin/ssh,/bin/df,/bin/
mv,/bin/rm,/usr/bin/install,/usr/bin/env
Cmnd_Alias VERTICA = /opt/vertica/bin/,/opt/vertica/sbin/,/opt/vertica/oss/python/
bin/
Cmnd_Alias VERTICA_INSTALL = /bin/echo,/bin/ps -A,/bin/cp /opt/vertica/config/
admintools.conf /opt/vertica/config/admintools.conf.bak.*,/bin/rm -rf /tmp/
dbRPM.rpm,/bin/df --portability /tmp,/usr/bin/install --owner * --mode 700 -d *,/
bin/mv -f /tmp/vstage-*/file /tmp/*,/bin/rm -rf /tmp/vstage-*,/usr/bin/id *,/bin/
cp -T /opt/vertica/* /tmp/vstage-*,/bin/su --login <dbadmin> *,/bin/mkdir -p /
opt/vertica/*,/bin/touch /opt/vertica/*,/bin/rm -rf /opt/vertica/*,/bin/mv -f /
tmp/vstage-*/opt/vertica/*,/bin/mkdir -p /opt/vertica/*,/bin/touch /opt/vertica/
config/users/<dbadmin>/agent.conf,/bin/su <dbadmin> *,/bin/sh -c *,/usr/bin,/opt/
vertica/share/binlib/test/*,/usr/bin/su <dbadmin>,/bin/test [ -e /* ],/usr/bin/[ -
e /* ]
Cmnd_Alias USEFUL = /usr/bin/lshw,/usr/bin/yum,/bin/rpm,/sbin/reboot,/sbin/
shutdown,/usr/bin/cpan,/bin/chgrp,/bin/chmod,/bin/chown,/bin/mnt,/usr/bin/test,/
bin/[,/sbin/service
## Allows the Data Repository user to manage the Data Repository
<sudouser> ALL = CA_DATAREP, VERTICA , VERTICA_INSTALL , USEFUL
Defaults env_keep += "VERT_DBA_USR VERT_DBA_HOME VERT_DBA_GRP VERT_DBA_DATA_DIR
_ENV_VPWD_VAR"

```

- **dbadmin**

The database administrator system account that the Vertica install creates, and who will own and run Vertica.

- (SLES 12)

```

Cmnd_Alias CA_DATAREP =/opt/vertica/sbin/install_vertica,/tmp/installDR.bin,/opt/
CA/IMDataRepository_vertica10/dr_validate.sh,/opt/CA/IMDataRepository_vertica10/
dr_install.sh,/usr/bin/vim,/usr/bin/reboot,/opt/CA/IMDataRepository_vertica10/
RemoteEngineer/re.sh,/usr/bin/mkdir, /sbin/SuSEfirewall2 off *,/usr/bin/whoami,/
usr/bin/echo,/usr/bin/id,/usr/bin/env,/usr/sbin/service,/usr/bin/grep,/usr/bin/
test,/sbin/iptables,/opt/vertica/oss/python/bin/python,/usr/bin/tee,/usr/sbin/
ntpd,/etc/init.d/ntpd,/sbin/blockdev,/etc/init.d/sshd,/etc/sysconfig/sshd,/etc/ssh/
sshd_config,/usr/bin/su,/usr/sbin/sshd restart,/usr/bin/ssh,/usr/bin/sh,/usr/bin/
install,/usr/bin/env
Cmnd_Alias VERTICA = /opt/vertica/bin/,/opt/vertica/sbin/,/opt/vertica/oss/python/
bin/
Cmnd_Alias VERTICA_INSTALL = /usr/bin/echo,/usr/bin/ps -A,/usr/bin/cp /opt/vertica/
config/admintools.conf /opt/vertica/config/admintools.conf.bak.*,/usr/bin/rm -rf /
tmp/dbRPM.rpm,/usr/bin/df --portability /tmp,/usr/bin/install --owner * --mode 700
-d *,/usr/bin/mv -f /tmp/vstage-*/file /tmp/*,/usr/bin/rm -rf /tmp/vstage-*,/usr/
bin/id *,/usr/bin/cp -T /opt/vertica/* /tmp/vstage-*,/usr/bin/su --login <dbadmin>
*,/usr/bin/mkdir -p /opt/vertica/*,/usr/bin/touch /opt/vertica/*,/usr/bin/rm -
rf /opt/vertica/*,/usr/bin/mv -f /tmp/vstage-*/opt/vertica/*,/usr/bin/mkdir -p /

```

```

opt/vertica/*,/usr/bin/touch /opt/vertica/config/users/<dbadmin>/agent.conf,/usr/
bin/su <dbadmin> *,/usr/bin/sh -c *,/opt/vertica/share/binlib/test/*,/usr/bin/su
<dbadmin>,/usr/bin/test [ -e /* ],/usr/bin/[ -e /* ]
Cmnd_Alias USEFUL = /usr/bin/lshw,/usr/bin/yum,/bin/rpm,/sbin/reboot,/sbin/
shutdown,/usr/bin/cpan,/bin/chgrp,/bin/chmod,/bin/chown,/bin/mnt,/usr/bin/test,/
bin/[,/sbin/service
## Allows the Data Repository user to manage the Data Repository
<sudouser> ALL = CA_DATAREP, VERTICA , VERTICA_INSTALL , USEFUL
Defaults env_keep += "VERT_DBA_USR VERT_DBA_HOME VERT_DBA_GRP VERT_DBA_DATA_DIR
_ENV_VPWD_VAR"

```

- **sudouser**

The user who will install and manage the Vertica node.

- **dbadmin**

The database administrator system account that the Vertica install creates, and who will own and run Vertica.

The sudo user account privileges are configured.

### **Set Up the User for Passwordless SSH Manually**

During the data repository installation or upgrade, the hosts in a data repository cluster require passwordless SSH for either the root or sudo user. The `dr_validate.sh` validation script sets up passwordless SSH, but it requests the password for the root or sudo user many times. You can avoid the validation script repeatedly asking for the password by setting up passwordless SSH for one of these users before running the script.

Repeat this procedure for each pair of hosts.

#### **NOTE**

Passwordless SSH is automatically set up for the data repository admin user when you install the data repository.

#### **Follow these steps:**

1. Open a console and log in to the data repository host as the root or sudo user.
2. Issue the following commands:

```

ssh-keygen -N "" -t rsa -f ~/.ssh/id_rsa
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys2
chmod 644 ~/.ssh/authorized_keys2

```

A public and private key are set up for the user, and the public key is appended to the authorization file.

3. Copy the user's public key into the list of authorized keys on the remote hosts by issuing the following command:

```
ssh-copy-id -i <user>@<remotehost>
```

- **remotehost**

A host in the cluster where you are copying the SSH ID.

4. At the prompt, enter the user's password.  
The user is set up for passwordless SSH.
5. To verify that you have set up the user for passwordless SSH, log in to the remote host from the local host by issuing the following command:

```
ssh <user>@<remotehost> ls
```

- **remotehost**

A host in the cluster where you are copying the SSH ID.

You have set up the user for passwordless SSH successfully if you are not prompted for a password. You also see a directory listing.

## Next Steps

- [Install the data repository.](#)

# Install the Data Repository

Install the data repository *after* preparing for its installation.

### IMPORTANT

The content in this article is for installing the data repository *if you are installing 23.3.3 and higher*. If you are installing 23.3.2 and lower, see [Install the Data Repository for 23.3.2 and Lower](#).

### (23.3.3 and higher)

After you have met the prerequisites that are described in [Prepare to Install the Data Repository](#), use the following process to install the data repository:

1. [Verify the Prerequisites](#)
2. [Install the Database](#)
3. [Verify the Database Installation](#)

After you have installed the data repository, complete the [Next Steps](#).

## Installation Files and Users

The `installDR.bin` installation file includes the following files:

- The data repository RPM Package Manager (RPM) installation package.
- The Vertica license file.
- `dr_validate.sh`  
This validation script verifies the OS settings and modifies the settings if necessary.
- `dr_install.sh`  
This installation script (the data repository installer) installs the data repository, creates the database, and disables unnecessary Vertica processes on the hosts in the cluster. If the Vertica database administrator system account (`dradmin`) does not already exist, the data repository installer creates this user.
- `migrate_sdn_device_metrics.sh`  
This migration script migrates virtual disk usage data on existing SDN devices from the Virtual Disk metric family to the SDN Devices Metrics metric family, and unifies the data into a single metric family for VNA to send in the data aggregator.

The data repository installation creates the following users:

- **dradmin**  
This is the Linux user that serves as the Vertica database administrator system account (the database administrator user). This user can run data repository processes and the Vertica Administration Tools utility, `adminTools`. This user owns the data repository files in the `catalog` and `data` directories.  
The data repository installer creates this user first. It also creates the `verticadba` group for tighter control over filesystem access in the `/opt/vertica/` directories, and adds this user to this group.  
This user has permissions set to `775`. This setting grants full privileges to the `verticadba` group and read/execute privileges to the other users. The modified permissions are located in the `/opt/vertica/log` and `/opt/vertica/config` directories.  
**Example password:** `drpass`  
**User Account:** Operating System and Vertica Database
- **dauser**  
The data aggregator connects and interacts with the database using this user. The data repository installer creates this user during the data aggregator installation.

**Example password:** dbpassword

**User Account:** Vertica Database

### Verify the Prerequisites

Prior to running the installation, ensure that you have completed the [prerequisites](#).

### Install the Database

#### Follow these steps:

1. Do one of the following:
  - For a data repository cluster, log in to any host in the cluster as the root user or as the non-root user using sudo user.

#### **NOTE**

For a cluster installation, initiate the data repository installation from any of the hosts that participates in the cluster. The installation pushes the required software components to the additional nodes.

- For a single-node installation, log in to the host as the root user or as the non-root user using sudo.
2. Copy the `installDR.bin` installation file locally.

#### **NOTE**

This file is included with the data aggregator package.

For more information, see [Installation Files and Users](#).

3. Change permissions for the installation file by issuing the following command:

```
chmod u+x installDR.bin
```

4. Extract the files by issuing the following command based on your user:

#### – **Root user**

```
./installDR.bin
```

#### – **Non-root user using sudo**

```
sudo ./installDR.bin
```

5. When prompted, specify the data repository installation directory to which to extract the files, and then press the **Return/Enter** key on your keyboard twice.

**Default:** `/opt/CA/IMDataRepository_verticaVersion/`

6. Change directories to the location where you extracted the installation scripts by issuing the following command:

```
cd <data_repository_directory>
```

#### – **data\_repository\_directory**

The default installation directory for the data repository.

**Default:** `/opt/CA/IMDataRepository_verticaVersion`

7. Adjust the following parameters in the `<data_repository_directory>/drinstall.properties` file to reflect installation-specific values:

#### – **DbAdminLinuxUser**

The Linux user that the data repository installer creates to serve as the Vertica database administrator.

**Default:** `dradmin`

#### – **DbAdminLinuxUserHome**

The Vertica database administrator user home directory. If the data repository installer creates the `dradmin` user, it will also create this directory.

#### **IMPORTANT**

Ensure that the directory leading up to the home account (for example, the `/export` directory) already exists on the system.

**Default:** `/export/dradmin`

#### – **DbDataDir**



The location of the `data` directory.

**IMPORTANT**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `catalog` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

**Default:** `/data`

– **DbCatalogDir**

The location of the `catalog` directory.

**IMPORTANT**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `data` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

**Default:** `/catalog`

– **DbHostNames**

The comma-delimited list of hostnames for the data repository.

**Default:** `yourhostname1,yourhostname2,yourhostname3`

– **DbName**

The database name.

**Default:** `drdata`

**Case sensitive:** Yes

– **DbPwd**

The database password. The data repository installer uses this password during the installation of the data aggregator. You can use special characters (except for single quotation marks) in passwords. If the installer does not find the `DbPwd` property or if it is blank, it prompts for this information at runtime.

**Default:** `dbpass`

– **SudoPassword**

The sudo password to use if installing as a non-root user using `sudo`. Leave value blank if not needed.

If you provide **SshPassword**, and you do not provide **SudoPassword**, then the scripts use **SshPassword** as the sudo password.

If you provide **PublicKeyFile**, provide **SudoPassword**.

If you do not provide **SudoPassword** and **SshPassword**, the scripts prompt for the sudo password.

– Provide *one* of the following for SSH calls:

• **SshPassword**

The SSH password to use for SSH calls. Leave value blank if not needed.

SSH authentication requires either **SshPassword** or **PublicKeyFile**.

If you do not provide **SshPassword**, the scripts prompt for the SSH password.

• **PublicKeyFile**

The directory and private key file name for the SSH public key authentication to use for SSH calls, for example

`<directory>/<filename>`.

SSH authentication requires either **PublicKeyFile** or **SshPassword**.

If you do not provide **PublicKeyFile**, the scripts prompt for the private key file for the SSH public key authentication.

If you provide **PublicKeyFile**, and you are installing as a non-root user using `sudo`, the installation script prompts for the non-root user's sudo password.

8. From the installation directory, run the validation script with the `-p` option by issuing the following command:

```
./dr_validate.sh -p drinstall.properties
```

**TIP**

You can use the `-n` option with the command, which skips database connectivity checks.

9. Review on-screen output for failures or warnings. You can run the validation script multiple times after you fix any failures or warnings. The script automatically corrects many failures or warnings. Proceed to the next step only if the final status is "PASSED". If the final status is not "PASSED", contact Broadcom Support.

The validation script might ask you to reboot.

#### NOTE

The validation and installation scripts generate a log file in the `<data_repository_directory/logs` directory on the data repository host from which you run the scripts. These log files include the step-by-step output of the scripts. Validate successful/failed script runs by reviewing the script output.

The following example shows the script output and lists what settings the script verifies and changes:

Log File: `logs/install_log_validate_10-29-2023_11-14-11.log`

```
=====
Checking Passwordless SSH to all hosts: verticahost-dr
=====
```

```
Passwordless SSH from verticahost-dr to root@verticahost-dr .....[ OK ]
=====
```

```
Beginning Data Repository Prerequisite Compliance Enforcement on host verticahost-dr
=====
```

```
Red Hat Enterprise Linux Major Release: 7 .....[ OK ]
Processor Type: AMD .....[ OK ]
Checking to see if required 'dialog' package is available. ....[ OK ]
CPU frequency scaling not available on this system .....[ OK ]
Maximum number of open files allowed for dradmin is 65536 .....[ OK ]
User home dir /home/dradmin for group is not writable .....[ OK ]
User home dir /home/dradmin for world is not writable .....[ OK ]
Max number of threads / processes for dradmin is 255472 .....[ OK ]
Maximum number of file handles >= 65536 .....[ OK ]
Maximum number of memory maps is Total Mem(KB)/16 .....[ OK ]
Page reclaim threshold value is 22900 .....[ OK ]
Disabling necessary firewall settings. ....[ OK ]
Starting chronyd. ....[ OK ]
Ensuring tuned is off for RHEL7+... .....[ OK ]
Readahead parameter for /dev/sda is 2048 .....[ OK ]
Block Size for /dev/sda is 4096 .....[ OK ]
Readahead parameter for /dev/sda1 is 2048 .....[ OK ]
Block Size for /dev/sda1 is 512. Expected value >= 4096 .....[WARN]
Readahead parameter for /dev/sda2 is 2048 .....[ OK ]
Block Size for /dev/sda2 is 4096 .....[ OK ]
Swappiness is already set to 1 .....[ OK ]
Huge Page Compaction is enabled .....[ OK ]
Huge Page Compaction Defrag is disabled .....[ OK ]
Disk Scheduler for sda is not mq-deadline .....[WARN]
Set Disk Scheduler for sda to mq-deadline .....[ OK ]
SELinux is disabled .....[ OK ]
Verifying Swap Space. ....[ OK ]
No Logical Volumes exist. ....[ OK ]
Root entry exists in /etc/sudoers file. ....[ OK ]
Verifying ext4 or xfs filesystem used for data directory. ....[ OK ]
Verifying ext4 or xfs filesystem used for catalog directory. ....[ OK ]
Verifying that sshd AllowTcpForwarding is enabled. ....[ OK ]
Ensuring /etc/rc.local is executable. ....[ OK ]
```

```
Existing Vertica package discovered - need to perform database connectivity testing.
Database is running and available for connectivity verification. ....[ OK ]
Validating connection to database host: verticahost-dr
Verifying database connection properties. ....[ OK ]
Checking existing DB for ETL health.
Checking ETL health. ....[ OK ]

Data Repository Prerequisite Compliance Status on host verticahost-dr -- PASSED
=====
Script finished - ./dr_validate.sh
=====
```

10. Run the data repository installer with the `-p` option by issuing the following command:

```
./dr_install.sh -p drinstall.properties
```

11. If you are installing as a non-root user using `sudo`, and you provided **PublicKeyFile**, Vertica prompts during the install for the sudo password (the **Password:** prompt). At the prompt, enter the sudo password, and then press the **Return/Enter** key on your keyboard.

The following image shows an example of this prompt:

The data repository is installed, the database is created, and the unnecessary Vertica processes on the hosts in the cluster are disabled. If the `dradmin` user does not already exist, the data repository installer creates this user, and you are prompted to assign a new password.

### Verify the Database Installation

Verify that the installation script has installed the data repository successfully. Use `adminTools`.

#### **Follow these steps:**

1. Log in to the database server as the database administrator (`dradmin`) user by issuing the following command:  
`su - dradmin`
2. Open `adminTools` from the `/opt/vertica/bin` directory.
3. Select option **1 (View Database Cluster State)** from the main menu of the **Administration Tools** dialog.
4. Select **OK** or press the **Return/Enter** key on your keyboard.  
The database name appears and the **State** is reported as "UP".
5. Acknowledge that the database is "UP" by selecting **OK**.
6. Select option **E (Exit)**, and then press the **Return/Enter** key on your keyboard.

#### **NOTE**

If the database does not start automatically, to avoid data aggregation installation failure, start the database manually by selecting **Start DB**.

### Next Steps

After you have installed the data repository, you can do the following:

- (If you want to limit the users who can log in to the database to only the data repository administrative account (`dradmin`) and the root user) [Secure the Data Repository](#)
- (If you want to prevent the underlying `vertica.log` data repository log file from becoming too large) [Configure Log Rotation for the Data Repository](#)
- (If you want to preserve your data against failures) [Set Up Automatic Backups of the Data Repository](#)

## Secure the Data Repository

If you want to limit the users who can log in to the database to only the data repository administrative account (`dradmin`) and the root user, lock down the database.

### Follow these steps:

1. Modify the `/etc/pam.d/sshd` file by adding the following entry, for the PAM access module, after the "account required pam\_nologin.so" entry:

```
account required pam_access.so accessfile=/etc/security/sshd.conf
```

#### NOTE

If this file is missing, create it.

For more information, see [the SSHD documentation](#).

2. If the following line from the `/etc/security/access.conf` file exists, remove it:

```
 -:ALL EXCEPT <database_admin_user> root:LOCAL
```

#### Example:

```
 -:ALL EXCEPT dradmin root:LOCAL
```

The data repository is secured.

## Configure Log Rotation for the Data Repository

If you want to prevent the underlying `vertica.log` data repository log file from becoming too large, configure log rotation for the data repository. Configure a daily log rotation with logs retained for 21 days.

### Follow these steps:

1. Log in to the database server for the data repository as the database administrator user (`dradmin`) by issuing the following command:

```
su - dradmin
```

2. Issue the following command, with the following options:

```
/opt/vertica/bin/admintools -t logrotate -d database_name -r frequency -k number
```

#### – -d

Indicates the database name.

**Case sensitive:** Yes

#### – -r

Specifies how often to rotate the daily logs.

**Values:** daily, weekly, monthly

#### – -k

Specifies how many logs to keep according to the frequency. For example, if the frequency is weekly, a value of 3 keeps three weeks of daily log files.

#### Example:

```
/opt/vertica/bin/admintools -t logrotate -d drdata -r daily -k 14
```

3. (Optional) To verify that you have configured the data repository log file rotation correctly, look at the new `vertica.log` gzipped files in the Vertica catalog directory for previous days. The log files use the following filename format:

```
vertica.log.YYYYMMDD.gz
```

The data repository log rotation is configured.

## Set Up Automatic Backups of the Data Repository

If you want to preserve your data against failures, [set up automatic backups of the data repository](#).

## Install the Data Repository for 23.3.2 and Lower

Install the data repository *after* preparing for its installation.

### IMPORTANT

The content in this article is for installing the data repository *if you are installing 23.3.2 and lower*. If you are installing 23.3.3 and higher, see [Install the Data Repository](#).

### (23.3.2 and lower)

After you have met the prerequisites that are described in [Prepare to Install the Data Repository](#), use the following process to install the data repository:

1. [Verify the Prerequisites](#)
2. [Install the Database](#)
3. [Verify the Database Installation](#)

After you have installed the data repository, complete the [Next Steps](#).

### Installation Files and Users

The `installDR.bin` installation file includes the following files:

- The data repository RPM Package Manager (RPM) installation package.
- The Vertica license file.
- `dr_validate.sh`  
This validation script verifies the OS settings and modifies the settings if necessary.
- `dr_install.sh`  
This installation script (the data repository installer) installs the data repository, creates the database, and disables unnecessary Vertica processes on the hosts in the cluster. If the Vertica database administrator system account (`dradmin`) does not already exist, the data repository installer creates this user.
- `migrate_sdn_device_metrics.sh`  
This migration script migrates virtual disk usage data on existing SDN devices from the Virtual Disk metric family to the SDN Devices Metrics metric family, and unifies the data into a single metric family for VNA to send in the data aggregator.

The data repository installation creates the following users:

- **dradmin**  
This is the Linux user that serves as the Vertica database administrator system account (the database administrator user). This user can run data repository processes and the Vertica Administration Tools utility, `adminTools`. This user owns the data repository files in the `catalog` and `data` directories.  
The data repository installer creates this user first. It also creates the `verticadba` group for tighter control over filesystem access in the `/opt/vertica/` directories, and adds this user to this group.  
This user has permissions set to 775. This setting grants full privileges to the `verticadba` group and read/execute privileges to the other users. The modified permissions are located in the `/opt/vertica/log` and `/opt/vertica/config` directories.  
**Example password:** `drpass`  
**User Account:** Operating System and Vertica Database
- **dauser**  
The data aggregator connects and interacts with the database using this user. The data repository installer creates this user during the data aggregator installation.  
**Example password:** `dbpassword`  
**User Account:** Vertica Database

## Verify the Prerequisites

Prior to running the installation, ensure that you have completed the [prerequisites](#).

## Install the Database

Use the following process to install the data repository:

1. Install the data repository using *one* of the following options:
  - [Install as the root user in the following cases](#):
    - You have the root user's password.
    - The root passwordless SSH user is already set up.
  - Install as a non-root user using sudo based on your configuration:
    - [You do not have the root password to set up the root user for passwordless SSH.](#)
    - [The root passwordless SSH user is already set up.](#)

The following video shows the data repository installation process:

2. [Verify the database installation.](#)
3. [Complete the next steps.](#)

## Install as Root User with Passwordless SSH Configured

In a cluster installation, initiate the data repository installation from any of the hosts that participates in the cluster. The installation pushes the required software components to the additional nodes.

### Follow these steps:

1. Log in to *any* host in the data repository cluster as the root user.
2. Copy the `installDR.bin` installation file locally.
3. Change permissions for the installation file by issuing the following command:
4. Extract the files by issuing the following command:
5. Follow the instructions in the console.
6. When prompted, specify the installation directory to which to extract the files, and then press the **Return/Enter** key on your keyboard twice

**Default:** `/opt/CA/IMDataRepository_verticaVersion/`

7. Adjust the following parameters in the `<installation_directory>/drinstall.properties` file to reflect installation-specific values:

- **DbAdminLinuxUser**

The Linux user that the data repository installer creates to serve as the Vertica database administrator.

**Default:** `dradmin`

- **DbAdminLinuxUserHome**

The Vertica database administrator user home directory. If the data repository installer creates the `dradmin` user, it will also create this directory.

**IMPORTANT**

Ensure that the directory leading up to the home account (for example, the `/export` directory) already exists on the system.

**Default:** `/export/dradmin`

- **DbDataDir**

The location of the `data` directory.

**IMPORTANT**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `catalog` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

**Default:** `/data`

– **DbCatalogDir**

The location of the `catalog` directory.

**IMPORTANT**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `data` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

**Default:** `/catalog`

– **DbHostNames**

The comma-delimited list of hostnames for the data repository.

**Default:** `yourhostname1,yourhostname2,yourhostname3`

– **DbName**

The database name.

**Default:** `drdata`

**Case sensitive:** Yes

– **DbPwd**

The database password. The data repository installer uses this password during the installation of the data aggregator. You can use special characters (except for single quotation marks) in passwords. If the installer does not find the `DbPwd` property or if it is blank, it prompts for this information at runtime.

**Default:** `dbpass`

8. Run the validation script with the `-p` option using the same location by issuing the following command:

```
./dr_validate.sh -p drinstall.properties
```

**TIP**

You can use the following options with the command:

- `-l` : Allows `localhost` as the value for the `DbHostNames` property.
- `-n` : Skips database connectivity checks.

The root user is established with passwordless Secure Shell (SSH) across the hosts in the cluster. If this user account is not configured with passwordless SSH, you are prompted, sometimes multiple times, for a password.

9. Review on-screen output for failures or warnings. You can run the validation script multiple times after you fix any failures or warnings. The script automatically corrects many failures or warnings. Proceed to the next step only if the final status is "PASSED". If the final status is not "PASSED", contact Broadcom Support.

The validation script might ask you to reboot.

**NOTE**

The validation and installation scripts generate a log file in the `installation_directory/logs` directory on the data repository host from which you run the scripts. If the data repository installation fails early enough in the process, the log file might be available in the home directory of the root or sudo user. These log files include the step-by-step output of the scripts. Validate successful/failed script runs by reviewing the script output.

The following example shows the script output and lists what settings the script verifies and changes:

```
Log File: logs/install_log_validate_10-29-2015_11-14-11.log
```

```
=====
```

```
Checking Passwordless SSH to all hosts: verticahost-dr
```

```
=====
```

```
Passwordless SSH from verticahost-dr to root@verticahost-dr .....[ OK ]
```

```
=====
Beginning Data Repository Prerequisite Compliance Enforcement on host verticahost-dr
=====
```

```
Red Hat Enterprise Linux Major Release: 6 .....[ OK ]
Processor Type: Intel .....[ OK ]
CPU frequency scaling not available on this system .....[ OK ]
DR Administrative User dradmin does not exist. It will be created during vertica
  installation. [ OK ]
Maximum number of file handles >= 65536 .....[ OK ]
Detected incorrect maximum number of memory maps .....[WARN]
Set maximum number of memory maps to Total Mem(KB)/16 .....[ OK ]
Detected incorrect page reclaim threshold value .....[WARN]
Set page reclaim threshold value to 7924 .....[ OK ]
Disabling necessary firewall settings. ....[ OK ]
Enabling NTP daemon. ....[ OK ]
Starting the NTP daemon. ....[ OK ]
Detected incorrect readahead parameter for /dev/sda .....[WARN]
Set readahead parameter for /dev/sda to 2048 .....[ OK ]
Block Size for /dev/sda is 4096 .....[ OK ]
Readahead parameter for /dev/sda1 is 2048 .....[ OK ]
Block Size for /dev/sda1 is 1024. Expected value >= 4096 .....[WARN]
Readahead parameter for /dev/sda2 is 2048 .....[ OK ]
Block Size for /dev/sda2 is 4096 .....[ OK ]
Readahead parameter for /dev/sda3 is 2048 .....[ OK ]
Block Size for /dev/sda3 is 4096 .....[ OK ]
Detected incorrect swappiness setting .....[WARN]
Set swappiness to 0 .....[ OK ]
Transparent hugepages in /sys/kernel/mm/redhat_transparent_hugepage/enabled are
  enabled [WARN]
Disabled Huge Page Compaction .....[ OK ]
Huge Page Compaction Defrag in /sys/kernel/mm/redhat_transparent_hugepage/defrag is
  enabled [WARN]
Disabled Huge Page Compaction Defrag .....[ OK ]
Disk Scheduler for sda is not deadline .....[WARN]
Set Disk Scheduler for sda to deadline .....[ OK ]
Reloading sysctl.conf .....[WARN]
SELinux is disabled .....[ OK ]
Verifying Swap Space. ....[ OK ]
No Logical Volumes exist. ....[ OK ]
Root entry exists in /etc/sudoers file. ....[ OK ]
Verifying ext3 or ext4 filesystem used for data directory. ....[ OK ]
Verifying ext3 or ext4 filesystem used for catalog directory. ....[ OK ]
Fresh install of Vertica is being performed - skipping database connectivity testing.
Data Repository Prerequisite Compliance Status on host verticahost-dr -- PASSED
=====
Script finished - /user/home/verticahost/dr_validate.sh
```



10. Run the data repository installer with the `-p` option by issuing the following command:

```
./dr_install.sh -p drinstall.properties
```

The data repository is installed, the database is created, and the unnecessary Vertica processes on the hosts in the cluster are disabled. If the `dradmin` user does not already exist, the data repository installer creates this user, and you are prompted to assign a new password.

### **Install as Non-Root User Using Sudo**

You set up the data repository by installing and configuring the Vertica database as the `sudo` user.

#### **Follow these steps:**

1. Log in to *each* node in the data repository cluster as the `sudo` user.
2. Copy the `installDR.bin` file locally.

#### **NOTE**

This file is included with the data aggregator package.

3. Change permissions for the installation file by issuing the following command:

```
chmod u+x installDR.bin
```

4. Extract the installation file as the `sudo` user by issuing the following command:

```
sudo ./installDR.bin
```

5. Follow the instructions in the console, and then press the **Return/Enter** key on your keyboard twice.
6. Adjust the following parameters in the `<installation directory>/drinstall.properties` file to reflect installation-specific values. This file applies to the validation and installation scripts.

#### – **DbAdminLinuxUser**

The Linux user that the data repository installer created to serve as the Vertica database administrator.

**Default:** `dradmin`

#### – **DbAdminLinuxUserHome**

The Vertica database administrator user home directory.

**Default:** `/export/dradmin`

#### **NOTE**

This directory is created if the data repository installer creates the `dradmin` user. Ensure that the directory leading up to the home account already exists on the system. For example, if you are using the `/export/dradmin` directory, ensure that the `/export` directory exists.

#### – **DbDataDir**

The location of the data directory.

**Default:** `/data`

#### **NOTE**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `catalog` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

#### – **DbCatalogDir**

The location of the `catalog` directory.

**Default:** `/catalog`

#### **NOTE**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `data` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

#### – **DbHostNames**

The list of hostnames for the data repository.

**Default:** `yourhostname1,yourhostname2,yourhostname3`

**NOTE**

List only the local hostname. You add all other nodes in a later step.

– **DbName**

The database name.

**Default:** drdata

**Case sensitive:** Yes

– **DbPwd**

The database password. You can use special characters (except for single quotation marks) in passwords.

**Default:** dbpass

**NOTE**

The data repository installer uses this password during the installation of the data aggregator. If the installer does not find the `DbPwd` property or if it is blank, it prompts for this information at runtime.

7. Run the validation script, with the `-sp` option, on *each* node in the data repository cluster using the same location by issuing the following command:

```
sudo ./dr_validate.sh -sp drinstall.properties
```

**TIP**

You can use the following options with the command:

- To allow `localhost` as the value for the `DbHostNames` property, use the `-l` option.
- To skip database connectivity checks, use the `-n` option.

The sudo user is established with passwordless Secure Shell (SSH) across the hosts in the cluster. If SSH without a password does not exist for this account, you are prompted, sometime multiple times, for a password.

8. Review the on-screen output for failures or warnings. You can run the validation script multiple times after you fix failures or warnings. The script automatically corrects many failures or warnings. Proceed only if the final status is "PASSED". If the final status is not "PASSED", contact Broadcom Support.

The validation script might ask you to reboot.

**NOTE**

The validation and installation scripts generate a log file in `installation_directory/logs` on the data repository host from which you run the scripts. If the installation fails early enough in the process, the log file might be available in the home directory of the root or sudo user. These log files include the step-by-step output of the scripts. To validate successful/failed script runs, review the script output.

The following example shows the script output and lists what settings the script verifies and changes:

```
Log File: logs/install_log_validate_10-29-2015_11-14-11.log
```

```
=====
Checking Passwordless SSH to all hosts: verticahost-dr
```

```
=====
Passwordless SSH from verticahost-dr to root@verticahost-dr .....[ OK ]
=====
```

```
Beginning Data Repository Prerequisite Compliance Enforcement on host verticahost-dr
=====
```

```
Red Hat Enterprise Linux Major Release: 6 .....[ OK ]
```

```
Processor Type: Intel .....[ OK ]
```

```
CPU frequency scaling not available on this system .....[ OK ]
```

```
DR Administrative User dradmin does not exist. It will be created during vertica
installation. [ OK ]
```

```
Maximum number of file handles >= 65536 .....[ OK ]
```

```
Detected incorrect maximum number of memory maps .....[WARN]
```

```
Set maximum number of memory maps to Total Mem(KB)/16 .....[ OK ]
```

```
Detected incorrect page reclaim threshold value .....[WARN]
```

```

Set page reclaim threshold value to 7924 .....[ OK ]
Disabling necessary firewall settings. ....[ OK ]
Enabling NTP daemon. ....[ OK ]
Starting the NTP daemon. ....[ OK ]
Detected incorrect readahead parameter for /dev/sda .....[WARN]
Set readahead parameter for /dev/sda to 2048 .....[ OK ]
Block Size for /dev/sda is 4096 .....[ OK ]
Readahead parameter for /dev/sda1 is 2048 .....[ OK ]
Block Size for /dev/sda1 is 1024. Expected value >= 4096 .....[WARN]
Readahead parameter for /dev/sda2 is 2048 .....[ OK ]
Block Size for /dev/sda2 is 4096 .....[ OK ]
Readahead parameter for /dev/sda3 is 2048 .....[ OK ]
Block Size for /dev/sda3 is 4096 .....[ OK ]
Detected incorrect swappiness setting .....[WARN]
Set swappiness to 0 .....[ OK ]
Transparent hugepages in /sys/kernel/mm/redhat_transparent_hugepage/enabled are
enabled [WARN]
Disabled Huge Page Compaction .....[ OK ]
Huge Page Compaction Defrag in /sys/kernel/mm/redhat_transparent_hugepage/defrag is
enabled [WARN]
Disabled Huge Page Compaction Defrag .....[ OK ]
Disk Scheduler for sda is not deadline .....[WARN]
Set Disk Scheduler for sda to deadline .....[ OK ]
Reloading sysctl.conf .....[WARN]
SELinux is disabled .....[ OK ]
Verifying Swap Space. ....[ OK ]
No Logical Volumes exist. ....[ OK ]
Root entry exists in /etc/sudoers file. ....[ OK ]
Verifying ext3 or ext4 filesystem used for data directory. ....[ OK ]
Verifying ext3 or ext4 filesystem used for catalog directory. ....[ OK ]
Fresh install of Vertica is being performed - skipping database connectivity testing.
Data Repository Prerequisite Compliance Status on host verticahost-dr -- PASSED
=====
Script finished - /user/home/verticahost/dr_validate.sh
=====

```

9. Go to the first node and edit the DbHostnames parameter in the drinstall.properties file to include *all* the nodes in the cluster.

10. Run the installation script with the `-sp` option by issuing the following command:

```
sudo ./dr_install.sh -sp drinstall.properties
```

#### TIP

You can run the script as sudo by setting up passwordless SSH (the public key) for the sudo account between the data repository hosts.

For more information, see [Prepare to Install the Data Repository](#).

The data repository is installed, the database is created, and the unnecessary Vertica processes on all the hosts in the cluster are disabled. You are prompted for the sudo user password for *each* node during this process. If the database

administrator user (dradmin) does not already exist, the user is created, and you are prompted to assign a new password for *each* node during this process.

### **Install as a Non-Root User using Sudo for a Root Passwordless SSH Setup**

In a cluster installation, initiate the data repository installation from any of the hosts that participates in the cluster. The installation pushes the required software components to the additional nodes. If the database administrator user (dradmin) does not already exist, the user is created, and you are prompted to assign a new password. The installation configures passwordless SSH for the database administrator user (dradmin). Vertica cluster communication requires the usage of passwordless SSH for the database administrator user (dradmin).

#### **Follow these steps:**

1. Log in to *any* host in the data repository cluster as the non-root user.
2. Copy the `installDR.bin` file locally.
3. Change permissions for the installation file by issuing the following command:  

```
chmod u+x installDR.bin
```
4. Extract the files by issuing the following command:  

```
sudo ./installDR.bin
```
5. Follow the instructions in the console.
6. When prompted, specify the installation directory to which to extract the files, and then press the **Return/Enter** key on your keyboard twice

**Default:** `/opt/CA/IMDataRepository_verticaVersion/`

7. Adjust the following parameters in the `<installation_directory>/drinstall.properties` file to reflect installation-specific values:

- **DbAdminLinuxUser**

The Linux user that is created to serve as the Vertica database administrator.

**Default:** `dradmin`

- **DbAdminLinuxUserHome**

The Vertica database administrator user home directory. If the installation script creates the Vertica database administrator user, it will also create this directory.

**IMPORTANT**

Ensure that the directory leading up to the home account (for example, the `/export` directory) already exists on the system.

**Default:** `/export/dradmin`

- **DbDataDir**

The location of the `data` directory.

**IMPORTANT**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `catalog` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

**Default:** `/data`

- **DbCatalogDir**

The location of the `catalog` directory.

**IMPORTANT**

Do not use the LVM for this directory. Ensure that this directory is on a separate mount from the `data` directory. This isolates those file systems from performance and space interference so that they are unencumbered from any other disk usage or performance considerations, including each other.

**Default:** `/catalog`

- **DbHostNames**

The comma-delimited list of hostnames for the data repository.

**Default:** yourhostname1,yourhostname2,yourhostname3

– **DbName**

The database name.

**Default:** drdata

**Case sensitive:** Yes

– **DbPwd**

The database password. The installation script uses this password during the installation of the data aggregator. You can use special characters (except for single quotation marks) in passwords. If the script does not find the DbPwd property or if it is blank, the script prompts for this information at runtime.

**Default:** dbpass

8. Run the validation script with the `-p` option using the same location by issuing the following command:

```
sudo ./dr_validate.sh -p drinstall.properties
```

**TIP**

You can use the following options with the command:

- `-l`: Allows `localhost` as the value for the `DbHostNames` property.
- `-n`: Skips database connectivity checks.

The script validates that the settings that Vertica requires are in conformance.

9. Review on-screen output for failures or warnings. You can run the validation script multiple times after you fix any failures or warnings. The script automatically corrects many failures or warnings. Proceed to the next step only if the final status is "PASSED". If the final status is not "PASSED", contact Broadcom Support.

The validation script might ask you to reboot.

**NOTE**

The validation and installation script generate a log file in the `installation_directory/logs` directory on the data repository host from which you run the scripts. If the installation script fails early enough in the process, the log file might be available in the home directory of the root or sudo user. These log files include the step-by-step output of the scripts. Validate successful/failed script runs by reviewing the script output.

The following example shows the script output and lists what settings the script verifies and changes:

```
Log File: logs/install_log_validate_10-29-2015_11-14-11.log
```

```
=====
Checking Passwordless SSH to all hosts: verticahost-dr
```

```
=====
Passwordless SSH from verticahost-dr to root@verticahost-dr .....[ OK ]
=====
```

```
Beginning Data Repository Prerequisite Compliance Enforcement on host verticahost-dr
=====
```

```
Red Hat Enterprise Linux Major Release: 6 .....[ OK ]
```

```
Processor Type: Intel .....[ OK ]
```

```
CPU frequency scaling not available on this system .....[ OK ]
```

```
DR Administrative User dradmin does not exist. It will be created during vertica
installation. [ OK ]
```

```
Maximum number of file handles >= 65536 .....[ OK ]
```

```
Detected incorrect maximum number of memory maps .....[WARN]
```

```
Set maximum number of memory maps to Total Mem(KB)/16 .....[ OK ]
```

```
Detected incorrect page reclaim threshold value .....[WARN]
```

```
Set page reclaim threshold value to 7924 .....[ OK ]
```

```
Disabling necessary firewall settings. ....[ OK ]
```

```
Enabling NTP daemon. ....[ OK ]
```

```
Starting the NTP daemon. ....[ OK ]
```

```

Detected incorrect readahead parameter for /dev/sda .....[WARN]
Set readahead parameter for /dev/sda to 2048 .....[ OK ]
Block Size for /dev/sda is 4096 .....[ OK ]
Readahead parameter for /dev/sda1 is 2048 .....[ OK ]
Block Size for /dev/sda1 is 1024. Expected value >= 4096 .....[WARN]
Readahead parameter for /dev/sda2 is 2048 .....[ OK ]
Block Size for /dev/sda2 is 4096 .....[ OK ]
Readahead parameter for /dev/sda3 is 2048 .....[ OK ]
Block Size for /dev/sda3 is 4096 .....[ OK ]
Detected incorrect swappiness setting .....[WARN]
Set swappiness to 0 .....[ OK ]
Transparent hugepages in /sys/kernel/mm/redhat_transparent_hugepage/enabled are
enabled [WARN]
Disabled Huge Page Compaction .....[ OK ]
Huge Page Compaction Defrag in /sys/kernel/mm/redhat_transparent_hugepage/defrag is
enabled [WARN]
Disabled Huge Page Compaction Defrag .....[ OK ]
Disk Scheduler for sda is not deadline .....[WARN]
Set Disk Scheduler for sda to deadline .....[ OK ]
Reloading sysctl.conf .....[WARN]
SELinux is disabled .....[ OK ]
Verifying Swap Space. ....[ OK ]
No Logical Volumes exist. ....[ OK ]
Root entry exists in /etc/sudoers file. ....[ OK ]
Verifying ext3 or ext4 filesystem used for data directory. ....[ OK ]
Verifying ext3 or ext4 filesystem used for catalog directory. ....[ OK ]
Fresh install of Vertica is being performed - skipping database connectivity testing.
Data Repository Prerequisite Compliance Status on host verticahost-dr -- PASSED
=====
Script finished - /user/home/verticahost/dr_validate.sh
=====

```

10. Run the installation script with the `-p` option by issuing the following command:

```
sudo ./dr_install.sh -p drinstall.properties
```

The data repository is installed, the database is created, and the unnecessary Vertica processes on the hosts in the cluster are disabled.

### **Verify the Database Installation**

Verify that the installation script has installed the data repository successfully. Use `adminTools`.

#### **Follow these steps:**

1. Log in to the database server as the database administrator (`dradmin`) user by issuing the following command:  

```
su - dradmin
```
2. Open `adminTools` from the `/opt/vertica/bin` directory.
3. Select option **1 (View Database Cluster State)** from the main menu of the **Administration Tools** dialog.
4. Select **OK** or press the **Return/Enter** key on your keyboard.  
The database name appears and the **State** is reported as "UP".

5. Acknowledge that the database is "UP" by selecting **OK**.
6. Select option **E (Exit)**, and then press the **Return/Enter** key on your keyboard.

**NOTE**

If the database does not start automatically, to avoid data aggregation installation failure, start the database manually by selecting **Start DB**.

**Next Steps**

After you have installed the data repository, you can do the following:

- (If you want to limit the users who can log in to the database to only the data repository administrative account (dradmin ) and the root user) [Secure the Data Repository](#)
- (If you want to prevent the underlying `vertica.log` data repository log file from becoming too large) [Configure Log Rotation for the Data Repository](#)
- (If you want to preserve your data against failures) [Set Up Automatic Backups of the Data Repository](#)

**Secure the Data Repository**

If you want to limit the users who can log in to the database to only the data repository administrative account (dradmin ) and the root user, lock down the database.

**Follow these steps:**

1. Modify the `/etc/pam.d/sshd` file by adding the following entry, for the PAM access module, after the "account required pam\_nologin.so" entry:

```
account required pam_access.so accessfile=/etc/security/sshd.conf
```

**NOTE**

If this file is missing, create it.

For more information, see [the SSHD documentation](#).

2. If the following line from the `/etc/security/access.conf` file exists, remove it:

```
-:ALL EXCEPT <database_admin_user> root:LOCAL
```

**Example:**

```
-:ALL EXCEPT dradmin root:LOCAL
```

The data repository is secured.

**Configure Log Rotation for the Data Repository**

If you want to prevent the underlying `vertica.log` data repository log file from becoming too large, configure log rotation for the data repository. Configure a daily log rotation with logs retained for 21 days.

**Follow these steps:**

1. Log in to the database server for the data repository as the database administrator user (dradmin) by issuing the following command:

```
su - dradmin
```

2. Issue the following command, with the following options:

```
/opt/vertica/bin/admintools -t logrotate -d database_name -r frequency -k number
```

– **-d**

Indicates the database name.

**Case sensitive:** Yes

– **-r**

Specifies how often to rotate the daily logs.

**Values:** daily, weekly, monthly

– **-k**

Specifies how many logs to keep according to the frequency. For example, if the frequency is weekly, a value of 3 keeps three weeks of daily log files.

**Example:**

```
/opt/vertica/bin/admintools -t logrotate -d drdata -r daily -k 14
```

3. (Optional) To verify that you have configured the data repository log file rotation correctly, look at the new `vertica.log` gzipped files in the Vertica catalog directory for previous days. The log files use the following filename format:

```
vertica.log.YYYYMMDD.gz
```

The data repository log rotation is configured.

### **Set Up Automatic Backups of the Data Repository**

If you want to preserve your data against failures, [set up automatic backups of the data repository](#).

## **Prepare to Install the Data Aggregator**

Ensure that you can install the data aggregator successfully by preparing for the installation.

Complete the following procedures before installing the data aggregator:

1. [Verify the Prerequisites](#)
2. [Verify Time Synchronization](#)
3. (If you do not have root access to install and run the data aggregator) [Configure the Sudo User Account for the Data Aggregator](#)
4. [Configure the Limit on the Number of Open Files on the Data Aggregator](#)
5. [Configure UTF-8 Support](#)

### **Verify the Prerequisites**

Ensure that you have met following prerequisites before installing a fault-tolerant data aggregator or non-fault-tolerant data aggregators:

- [You have reviewed the installation requirements and considerations](#). This includes verifying that you have the packages that the installer for the data aggregator requires, and verifying that the ports required to allow the data collector and the data aggregator to communicate and function properly are open to remote access.
- The installer requires that the data repository be installed and that the service is running.

#### **TIP**

You can verify that the data repository is running by completing the following steps:

1. Open a console and log in to the data repository host as the root or sudo user.
2. Issue the following command:

```
su - dr_admin_user -c "/opt/vertica/bin/admintools -t show_active_db"
```

The following response is expected:

```
version
```

```
-----
```

```
Vertica Analytic Database vx.x.x-x
```

- You have verified that Security-Enhanced Linux (SELinux) is disabled or permissive on the computer where you are going to install the data aggregator. By default, some Linux distributions have this feature enabled, which does not allow the data aggregator to function properly. Disable SELinux, set to permissive mode, or create a policy to exclude data aggregator processes from SELinux restrictions.



For information more about how to configure SELinux security policies, see [the Red Hat documentation](#).

- To avoid potential corruption of data, you have completed the following:
  - Excluded the installation directory, the backup directory, and all subdirectories from antivirus scans.
  - Prevented scanning by a local instance of an antivirus client and scanning by a remote antivirus instance.
- You have verified that the directory where you are going to install has write privileges for your data aggregator user.
- (Fault-tolerant environments) You have verified the following:
  - [The hardware requirements](#).
  - You have a shared data directory (for example, /DASharedRepo ) and the same user ID is shared between data aggregator hosts. DX NetOps Performance Management stores the data from the active data aggregator in this directory.

For more information about the sizing requirements, see the [DX NetOps Sizing Tool](#).

#### **NOTE**

If you are using Network File System (NFS), DX NetOps Performance Management supports only NFSv4 and higher because of the ActiveMQ Kaha locking requirements.

#### **IMPORTANT**

To allow data to be loaded and to prevent data loss, this shared data directory must be up and accessible at all times.

### **Verify Time Synchronization**

The data aggregator requires time synchronization using the Network Time Protocol (NTP) daemon. All data source consoles must use the NTP daemon. On Linux servers, the NTP daemon ensures that the clocks on the hosts are synchronized for timing purposes. Verify that the daemon is running on all DX NetOps Performance Management servers based on your version:

- [Verify on SUSE Linux Enterprise Server \(SLES\)](#)
- [Verify on Red Hat Enterprise Linux \(RHEL\) 7.x/8.x and Oracle Linux \(OL\)](#)

### **Verify Time Synchronization on SUSE Linux Enterprise Server**

#### **Follow these steps:**

1. Open a console and issue the following command:
 

```
$ systemctl status ntpd
```
2. Verify that the NTP daemon is in an active (running) state.
3. Start and enable the NTP daemon manually by issuing the following command:
 

```
$ systemctl start ntpd
$ systemctl enable ntpd
```

The daemon is started.

### **Verify Time Synchronization on Red Hat Enterprise Linux 7.x/8.x and Oracle Linux**

RHEL 7.x/8.x and OL 7.x run NTP with `chronyd`.

#### **Follow these steps:**

1. Open a console and issue the following command:
 

```
$ systemctl status chronyd
```
2. Verify that the `chrony` daemon is in an active (running) state.
3. Start and enable the `chrony` daemon by issuing the following command:
 

```
$ systemctl start chronyd
$ systemctl enable chronyd
```

The daemon is started.

### **(Optional) Configure the Sudo User Account for the Data Aggregator**

If you do not have root access to install and run the data aggregator, configure the sudo user account.

#### **Follow these steps:**

1. Locate the `/etc/sudoers` file on the data aggregator host.

2. Add the following command aliases with the permissions to the file:

```
Cmd_Alias CA_DATAAGG = /tmp/installDA.bin,/usr/sbin/service dadaemon
*,/usr/sbin/service activemq *,<installation_directory>/Uninstall/
Uninstall,<installation_directory>/RemoteEngineer/re.sh,<installation_directory>/
scripts/sslConfig.sh
## Allows the Data Aggregator user to manage the Data Aggregator
dasudouser_name ALL = CA_DATAAGG
```

#### **Example:**

```
Cmd_Alias CA_DATAAGG = /tmp/installDA.bin,/usr/sbin/service dadaemon *,/usr/sbin/
service activemq *,/opt/IMDataAggregator/Uninstall/Uninstall,/opt/IMDataAggregator/
RemoteEngineer/re.sh,/opt/IMDataAggregator/scripts/sslConfig.sh
## Allows the Data Aggregator user to manage the Data Aggregator
dasudouser_name ALL = CA_DATAAGG
```

#### **– *installation\_directory***

Specify the installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

#### **– *dasudouser***

Specify the user who can run the sudo commands.

3. Add the sudo prefix to the data aggregator install commands by issuing the following command:

#### **Example:**

```
sudo ./installDA.bin
```

### **Configure the Limit on the Number of Open Files on the Data Aggregator**

Configure the limit on the number of open files to *at least* 65536 for the user that is installing the data aggregator. This configuration is permanent.

#### **Follow these steps:**

1. As the root user or a sudo user, log in to the data aggregator host.

2. Issue the following command:

```
ulimit -n ulimit_number
```

#### **Example:**

```
ulimit -n 65536
```

3. Open the `/etc/security/limits.conf` file, and add the following lines:

```
# Added by Data Aggregator
* soft nfile 65536
# Added by Data Aggregator
* hard nfile 65536
```

4. Restart the data aggregator.

**NOTE**

For upgrades, the upgrade process automatically restarts the data aggregator.

5. Verify that the number of open files is set properly by issuing the following command:

```
ulimit -n
```

The command returns the limit that you specified.

**Configure UTF-8 Support**

Configure the data aggregator host to support UTF-8 encoding. If UTF-8 encoding is not enabled, characters might not display properly during the installation.

The appropriate language packs are also required to support localized deployments.

**NOTE**

Some scripts that are used in the installation of selected components are not localized and run in English.

For more information, see [Language Support](#).

Issue the following commands based on your installation:

- **(Korn or bash shell)**

```
export LANG=<LANG_value> ; export LC_ALL=$LANG
```

**Example:**

```
export LANG=ja_JP.utf8 ; export LC_ALL=$LANG
```

- **LANG\_value**

Specifies the value of the language you want the product to support. The following variables are supported:

**English:** en\_US.utf8

**French:** fr\_FR.utf8

**Japanese:** ja\_JP.utf8

- **(Bourne shell)**

```
LANG=<LANG_value> ; export LANG
```

```
LC_ALL=<LANG_value> ; export LC_ALL
```

**Example:**

```
LANG=ja_JP ; export LANG
```

```
LC_ALL=ja_JP ; export LC_ALL
```

The data aggregator host is configured to support UTF-8 encoding.

**Next Steps**

- [Install the data aggregator.](#)

## Install the Data Aggregator

After you have met the requirements to install the data aggregator, complete the installation.

The data aggregator aggregates network item data from the data collectors into the data repository.

You can install the data aggregator in one of the following environments:

- [Install in a non-fault-tolerant environment.](#)
- [Install in a fault-tolerant environment.](#)

For more information about fault tolerance, see [Fault Tolerance](#).

## Install Fault-Tolerant Data Aggregators

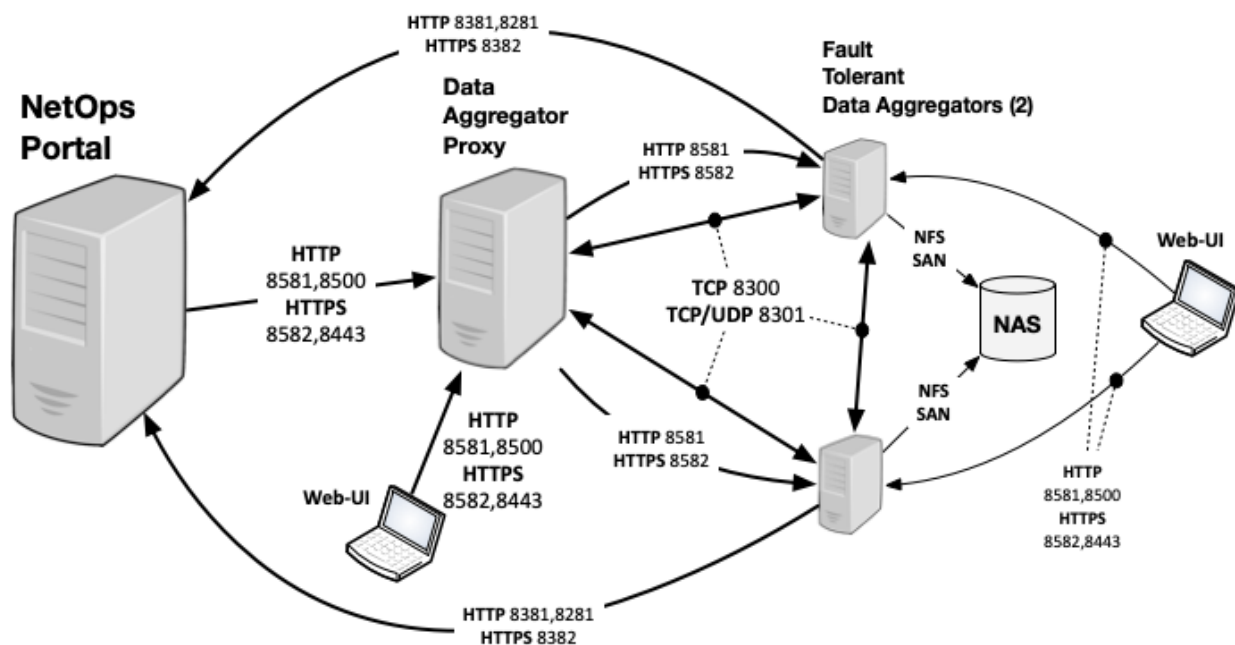
You can install a fault-tolerant environment.

When you first install or upgrade the data aggregator, you are prompted if you want to configure a fault-tolerant environment. Fault-tolerant environments require a shared data directory (for example, /DASharedRepo ) to help limit data loss. This shared data directory stores customized metric families, DTO files, and the ActiveMQ Kaha database. When a hardware failure or network issue occurs, the second data aggregator accesses the shared data directory. The data aggregator picks up where the first data aggregator left off. The user ID that the shared drive is created with must be synced to both of the data aggregators. Then both data aggregators have read and write permissions to that directory.

The following diagram illustrates the configuration of a fault-tolerant environment:

**Figure 6: High Availability**

### Fault-Tolerant Data Aggregator Configuration



Use the following process to install the environment:

1. [Meet the requirements to install a data aggregator](#)
2. [Install the proxy server](#)
3. (If you want to use Nginx instead of Traefik as your reverse proxy) [Install Nginx](#), and then [configure the connection from NetOps Portal to the proxy server using Nginx as the reverse proxy](#)
4. [Install the fault-tolerant data aggregator pair](#)
5. [Verify the data aggregator installation](#)
6. [Register the data aggregator \(the proxy server\) as a data source so that you can connect NetOps Portal to the data aggregator](#)

For more information about fault tolerance, see [Fault Tolerance](#).

## **Install a Fault-Tolerant Data Aggregator**

### **Follow these steps:**

1. Log in to the data aggregator host as the root user or `sudo` user.
2. Copy the `installDA.bin` file to the `/tmp` directory.
3. Change to the `/tmp` directory by issuing the following command:  

```
cd /tmp
```
4. Change permissions for the installation file by issuing the following command:  

```
chmod a+x installDA.bin
```
5. Do *one* of the following steps:
  - Install as the root user from the command line:  

```
./installDA.bin
```
  - Install as the root user in silent mode (silent installation):  

```
./installDA.bin -i console
```
  - Install as the `sudo` user from the command line:  

```
sudo ./installDA.bin
```
  - Install as the `sudo` user in silent mode (silent installation):  

```
sudo ./installDA.bin -i console
```

### **NOTE**

You can generate a response file for silent installation by completing the following steps:

1. Issue the install command with the following argument:  

```
-r <response_file>
```

  - ***response\_file***  
 Specifies the directory path and file name for the response file.  
**Example:** `/temp/installer.properties`
2. Follow the prompts until you get to the summary.
3. Type `quit`.
4. Press the Return/Enter key on your keyboard.

**Tip:** To install as the root user in silent mode (silent installation), complete the following:

1. Generate a response file for silent installation.
2. Issue the command with the following arguments:  

```
./installDA.bin -i silent -f <response_file>
```

  - ***response\_file***  
 Specifies the directory path and file name of the previously-generated response file. You can use this response file to run future installations in silent mode.

6. When prompted, specify the following parameters for the data repository:

### **IMPORTANT**

When you are prompted for the data directory, use the default directory. Do not use the `<installation_directory>/apache-karaf/data` directory.

- ***installation\_directory***  
 The default installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`

- **Data Repository server hostname/IP**

Defines either a name or an IP address for the data repository server host.

For a data repository cluster, specify the name or the IP address of any one host in the cluster. The installer automatically determines the name and IP address of the remaining nodes.

– **Data Repository server port**

Defines the port number for the data repository server.

**Default:** 5433

– **Database name**

Defines the database name of the data repository.

– **Data Repository username**

Specifies the username that the data aggregator uses to connect to the database. The username and the password cannot match. This username and password combination is added to the database during the installation.

**Example:** dauser

– **Data Repository admin username**

Specify the Linux user account that was used to install the data repository. This username is needed for administration, such as backing up and restoring the data repository, or updating the database schema.

**Example:** dradmin

– **Data Repository admin password**

Defines the password for the data repository admin username.

**NOTE**

This database user account password was specified when you created the database after the data repository installation.

**Example:** dbpassword

**TIP**

You can change the data repository admin password in the `dbconnection.cfg` file using the `doEncryption.sh` script.

7. When prompted, specify the following parameters for fault tolerance.

– **Configure Data Aggregator For Fault Tolerance**

Specify the number for the **Yes** option to configure this data aggregator for fault tolerance.

**Options:**

- **1- No:** Does not configure the data aggregator for fault tolerance (configures a non-fault tolerant environment).
- **2- Yes:** Configures the data aggregator for fault tolerance.

**Default:** 1- No

– **Data Aggregator Proxy Host**

Specify the hostname/IP address of the proxy server.

– **Consul HTTP port**

Specify the port for communication with Consul.

For more information about this tool, including how you manage services in DX NetOps Performance Management using this tool, see [Fault Tolerance](#).

**Default:** 8500

– **Choose host IP address for Consul**

**NOTE**

This prompt appears only when multiple public IP addresses are configured.

Specify the bind address that the Consul agents use to communicate with each other. The Consul agents include the proxy host and both data aggregators in the cluster. If prompted for an address, specify an address that the other two hosts in the Consul cluster can reach.

– **Data Aggregator Data Home**

Specify the path of the shared data aggregator data directory for fault tolerance.

The fault-tolerant data aggregator is installed and started.

## Verify the Data Aggregator Installation

### NOTE

The data aggregator installation is located in the static `/etc/DA.cfg` directory and file. This location also includes the chosen installation options for upgrade.

### Follow these steps:

1. Verify that the Data Aggregator service is running by issuing the following command:

```
systemctl status dadaemon
```

2. Review the `<installation_directory>/Logs/CA_Performance_Management_Data_Aggregator_Install_<timestamp>.log` file.

#### – **installation\_directory**

The installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

If the installation is successful, the log shows 0 Warnings, 0 NonFatalErrors, and 0 FatalErrors.

### NOTE

If the installation fails early enough in the process, the log file might be available in the home directory of the root or sudo user.

3. Verify that the ActiveMQ broker is running by issuing the following command:

```
systemctl status activemq
```

4. Open a browser on a computer where you have HTTP access to the data aggregator. Navigate to the following address:

```
http://<da_host>:<port>/rest
```

#### – **da\_host**

Specifies the data aggregator host name.

#### – **port**

Specifies the data aggregator required port number.

**Default:** (HTTPS) 8582 (HTTP) 8581

The return is a list of hyperlinks for available web services. When you click a link, the XML content describing the selection displays.

The data aggregator installation is verified.

## Install a Non-Fault-Tolerant Data Aggregator

You can install a non-fault-tolerant data aggregator.

Use the following process to install a non-fault-tolerant data aggregator:

1. [Meet the Requirements to Install a Data Aggregator](#)
2. [Install the Data Aggregator](#)
3. [Verify the Data Aggregator Installation](#)
4. [Register the Data Aggregator as a Data Source](#)

The following video shows the installation process:

### **Meet the Requirements to Install a Data Aggregator**

For more information, see [Prepare to Install the Data Aggregator](#).

## Install the Data Aggregator

### Follow these steps:

1. Log in to the data aggregator host as the root user or `sudo` user.
2. Copy the `installDA.bin` file to the `/tmp` directory.
3. Change to the `/tmp` directory by issuing the following command:  

```
cd /tmp
```
4. Change permissions for the installation file by issuing the following command:  

```
chmod a+x installDA.bin
```

5. Do *one* of the following steps:
  - Install as the root user from the command line:  

```
./installDA.bin
```
  - Install as the root user in silent mode (silent installation):  

```
./installDA.bin -i console
```
  - Install as the `sudo` user from the command line:  

```
sudo ./installDA.bin
```
  - Install as the `sudo` user in silent mode (silent installation):  

```
sudo ./installDA.bin -i console
```

### NOTE

You can generate a response file for silent installation by completing the following steps:

1. Issue the install command with the following argument:  

```
-r <response_file>
```

  - **response\_file**  
 Specifies the directory path and file name for the response file.  
**Example:** `/temp/installer.properties`
2. Follow the prompts until you get to the summary.
3. Type `quit`.
4. Press the Return/Enter key on your keyboard.

**Tip:** To install as the root user in silent mode (silent installation), complete the following:

1. Generate a response file for silent installation.
2. Issue the command with the following arguments:  

```
./installDA.bin -i silent -f <response_file>
```

  - **response\_file**  
 Specifies the directory path and file name of the previously-generated response file. You can use this response file to run future installations in silent mode.

6. When prompted, specify the following parameters for the data repository:

### IMPORTANT

When you are prompted for the data directory, use the default directory. Do not use the `<installation_directory>/apache-karaf/data` directory.

- **installation\_directory**  
 The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`

- **Data Repository server hostname/IP**  
 Defines either a name or an IP address for the data repository server host.



For a data repository cluster, specify the name or the IP address of any one host in the cluster. The installer automatically determines the name and IP address of the remaining nodes.

- **Data Repository server port**

Defines the port number for the data repository server.

**Default:** 5433

- **Database name**

Defines the database name of the data repository.

- **Data Repository username**

Specifies the username that the data aggregator uses to connect to the database. The username and the password cannot match. This username and password combination is added to the database during the installation.

**Example:** dauser

- **Data Repository admin username**

Specify the Linux user account that was used to install the data repository. This username is needed for administration, such as backing up and restoring the data repository, or updating the database schema.

**Example:** dradmin

- **Data Repository admin password**

Defines the password for the data repository admin username.

**Example:** dbpassword

**NOTE**

This database user account password was specified when you created the database after the data repository installation.

**Example** dbpassword

**TIP**

You can change this password in the `dbconnection.cfg` file using the `doEncryption.sh` script.

The non-fault-tolerant data aggregator is installed and started.

## Verify the Data Aggregator Installation

**NOTE**

The data aggregator installation is located in the static `/etc/DA.cfg` directory and file. This location also includes the chosen installation options for upgrade.

### Follow these steps:

1. Verify that the Data Aggregator service is running by issuing the following command:

```
systemctl status dadaemon
```

2. Review the `<installation_directory>/Logs/`

`CA_Performance_Management_Data_Aggregator_Install_<timestamp>.log` log file on the data aggregator host:

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

If the installation is successful, the log shows 0 Warnings, 0 NonFatalErrors, and 0 FatalErrors.

**NOTE**

If the installation fails early enough in the process, the log file might be available in the home directory of the root or sudo user.

3. Verify that the ActiveMQ broker is running by issuing the following command:

```
systemctl status activemq
```

4. Open a browser on a computer where you have HTTP access to the data aggregator. Navigate to the following address:

```
http://<da_host>:<port>/rest
```

- **da\_host**  
The data aggregator host name or IP address.
- **port**  
Specifies the data aggregator required port number.

**Default:** (HTTPS) 8582 (HTTP) 8581

The return is a list of hyperlinks for available web services. When you click a link, the XML content describing the selection displays.

The data aggregator installation is verified.

### **Register the Data Aggregator as a Data Source**

Registering, or configuring, a data source binds it to NetOps Portal.

The following video shows how to bind the data aggregator to NetOps Portal:

For more information, see [Configure a Data Source](#).

## **Prepare to Install the Data Collectors**

Ensure that you can install the data collectors successfully by preparing for the installation.

Complete the following procedures before installing the data collectors:

1. [Review the Data Collector Considerations](#)
2. [Verify the Prerequisites](#)
3. [Verify Time Synchronization](#)
4. (If you do not have root access to install and run the data collector) [Configure the sudo user account for the data collector](#)
5. [Configure UTF-8 Support](#)
6. [Set a unique hostname for the data collector host](#)
7. (If you are configuring the data collectors for fault tolerance) [Review the fault-tolerant deployment considerations](#)
8. (If you plan to monitor software-defined networking (SDN) and network functions virtualization (NFV) using DX NetOps Virtual Network Assurance (VNA)) [Plan for a VNA installation on the same host as the data collector](#)

### **Review Data Collector Considerations**

In a standard tenant deployment, each tenant has a dedicated data collector, and the data collector supports only one data aggregator. In a multi-tenant environment, a managed service provider monitors devices for multiple tenants. For multi-tenant deployments, where tenants reside in the same Internet Protocol (IP) routing space, you can configure DX NetOps Performance Management to use fewer data collectors.

For more information, see [Configure Tenant-Agnostic Data Collectors](#).

Review the following considerations:

- [Standard tenant deployment considerations](#).
- [Multi-tenant deployment considerations](#).
- [Data collector disk usage considerations](#).

### **Standard Tenant Deployment Considerations**

The following considerations apply to a standard tenant deployment:

- [Review the installation requirements and considerations.](#)
- You can install one or more data collectors. Install each data collector on a separate host.
- If the data aggregator is only IPv6, the data collector must support the IPv6 protocol.

**TIP**

You can verify that the data collector supports IPv6 by taking the the following steps on the data aggregator host:

1. Find the IPv6 address by issuing the following command:

```
ifconfig
```

2. Ensure that the data collector can contact the data aggregator using its IPv6 address by issuing the following command:

```
ping6 ipv6_address_of_data_aggregator
```

### **Multi-tenant Deployment Considerations**

The following consideration applies to a multi-tenant deployment. You can do *one* of the following:

- Install the data collector at the managed service provider (MSP) site.

**NOTE**

This setup requires that the data collector gain access through a tenant firewall to poll managed devices.

- Install the data collector at each tenant site.

### **Data Collector Disk Usage Considerations**

The data collector uses a disk cache. The default disk cache size is equal to 50% of the maximum data collector memory. By default, this value is equal to 45 minutes, or 500K, of data when you use a 5-minute poll rate.

For more information about the default disk cache size, see [Modify the External ActiveMQ Memory Limit](#).

When the disk cache reaches 90% saturation of the disk cache size, the data collector rolls data off the back of the queue. The data collector uses disk space to cache poll messages (temporary messages). When the disk cache is full, the data collector drops the oldest polled messages, keeping the latest, and posts a log message to the `<DC_installation_directory>/apache-karaf/data/log/karaf.log` file.

- ***DC\_installation\_directory***

The installation directory of the data collector.

**Default:** `/opt/IMDataCollector`

The following is an example of the log message:

```
DATE TIME | WARN | pool-15-thread-1 | PRQCleanupService |
e.jms.health.PRQCleanupService$2 135 | 175 - com.ca.im.common.core.jms -
X.X.X.RELEASE-XXX | | JMS Health: dropped 178895/178895 messages from PRQ
(dropRate=10%, maxDiskUsage=2566M)
```

When the data aggregator and data collector connect, the data collector sends the oldest messages first. Cached poll responses can take some time before they display in NetOps Portal.

**NOTE**

If you restart ActiveMQ, the cached data is lost.

**TIP**

To monitor the cache burndown, locate the `<DC_installation_directory>/apache-karaf/etc/org.ops4j.pax.logging.cfg` file, and then create the following entries in the file:

```
log4j.logger.com.ca.im.core.jms.health.JmsBrokerHealthAnalyser=DEBUG,sift
log4j.additivity.com.ca.im.core.jms.health.JmsBrokerHealthAnalyser=false
```

- ***DC\_installation\_directory***  
The installation directory for the data collector.  
**Default:** `/opt/IMDataCollector`

The `<DC_installation_directory>/apache-karaf/data/log/com.ca.im.common.core.jms.log` file is created.

### **Verify the Prerequisites**

Meet the following prerequisites before you install data collector:

- You have verified that the ports required to allow the data collector and the data aggregator to communicate and function properly are open. This includes verifying that you have the packages that the installer for the data collector requires.

#### **TIP**

You can change port 616xx to another port after you install the data aggregator.  
For more information, see [Complete the Post-Installation Configuration](#).

#### **IMPORTANT**

Ports 1099 and 11099 must be open only locally for internal communication, and blocked from external access.

- You have verified that Security-Enhanced Linux (SELinux) is disabled or permissive on the computer where you are going to install the data collector. By default, some Linux distributions have this feature enabled, which does not allow the data collector to function properly. Disable SELinux, set to permissive mode, or create a policy to exclude the data collector processes from SELinux restrictions.  
For information more about how to configure SELinux security policies, see [the Red Hat documentation](#).
- To avoid database corruption of ActiveMQ broker files, you have excluded the installation directory, and all its subdirectories, from antivirus scans. Prevent scanning by a local instance of an antivirus client and scanning by a remote antivirus instance.
- You have installed and provisioned the tenants and corresponding IP domain in NetOps Portal to which you will be assigning the data collectors.

### **Verify Time Synchronization**

The data collectors require time synchronization using the Network Time Protocol (NTP) daemon. All data source consoles must use the NTP daemon. On Linux servers, the NTP daemon ensures that the clocks on the hosts are synchronized for timing purposes. Verify that the daemon is running on all DX NetOps Performance Management servers based on your version:

- [Verify on SUSE Linux Enterprise Server \(SLES\)](#)
- [Verify on Red Hat Enterprise Linux \(RHEL\) 7.x/8.x and Oracle Linux \(OL\)](#)

### **Verify Time Synchronization on SUSE Linux Enterprise Server**

#### **Follow these steps:**

1. Open a console and issue the following command:  

```
$ systemctl status ntpd
```
2. Verify that the NTP daemon is in an active (running) state.
3. Start and enable the NTP daemon manually by issuing the following command:  

```
$ systemctl start ntpd  
$ systemctl enable ntpd
```

The daemon is started.

## **Verify Time Synchronization on Red Hat Enterprise Linux 7.x/8.x and Oracle Linux**

RHEL 7.x/8.x and OL 7.x run NTP with `chronyd`.

### **Follow these steps:**

1. Open a console and issue the following command:  

```
$ systemctl status chronyd
```
2. Verify that the `chrony` daemon is in an active (running) state.
3. Start and enable the `chrony` daemon by issuing the following command:  

```
$ systemctl start chronyd
```

```
$ systemctl enable chronyd
```

The daemon is started.

### **(Optional) Configure the Sudo User Account for the Data Collector**

If you do not have root access to install and run the data collector, configure the sudo user account for the data collector. As a sudo user, you can add the sudo prefix to all commands to install the data collector, for example `sudo ./install.bin`.

### **Follow these steps:**

1. Locate the `/etc/sudoers` file on the data collector host.
2. Add the following command aliases with the following permissions to the file:  

```
## Allows the user to install the data collector
Cmdnd_Alias CA_DATACOLL_INSTALL = /tmp/install.bin

## Allows the product to be restarted when assigning a data collector
Cmdnd_Alias CA_DATACOLL_RESTART = /usr/sbin/service dcmd *, /usr/bin/systemctl *
dcmd, <DC_installation_directory>/scripts/dcmd *, <DC_installation_directory>/apache-
karaf/bin/restart 5, /usr/bin/at now

## Allows the user to issue commands to restart the dcmd and activemq processes, as
well as running remote engineer and the uninstaller
Cmdnd_Alias CA_DATACOLL = /usr/sbin/service activemq *, /usr/bin/systemctl *
activemq, <DC_installation_directory>/scripts/activemq *, /usr/sbin/service
icmcpd *, /usr/bin/systemctl * icmcpd, <DC_installation_directory>/ICMPD/
icmcpd *, /usr/sbin/service atd *, <DC_installation_directory>/Uninstall/
Uninstall, <DC_installation_directory>/re.sh, /usr/bin/crontab

## Assigns the commands to the sudo user
## Assigns the general service commands to the user
<sudo_user> ALL = CA_DATACOLL

## Assigns the install commands to the user and enables them to set the environment
with the command (e.g. setting the DCM_ID when replacing a data collector)
<sudo_user> ALL = (ALL) SETENV: CA_DATACOLL_INSTALL

## Assigns the restart commands to the user and allows them to be run without a
password - needed for when the process issues a restart during DC assignment
<sudo_user> ALL = (ALL) NOPASSWD: CA_DATACOLL_RESTART
```

### **Example:**

```
Cmdnd_Alias CA_DATACOLL_INSTALL = /tmp/install.bin
```

```

Cmnd_Alias CA_DATACOLL_RESTART = /usr/sbin/service dcmd *, /usr/bin/systemctl *
dcmd, /opt/IMDataCollector/scripts/dcmd *, /opt/IMDataCollector/apache-karaf/bin/
restart 5, /usr/bin/at now
Cmnd_Alias CA_DATACOLL = /usr/sbin/service activemq *, /usr/bin/systemctl *
activemq, /opt/IMDataCollector/scripts/activemq *, /usr/sbin/service icmpd *, /usr/
bin/systemctl * icmpd, /opt/IMDataCollector/ICMPD/icmpd *, /usr/sbin/service atd *, /
opt/IMDataCollector/Uninstall/Uninstall, /opt/IMDataCollector/re.sh, /usr/bin/crontab
<sudo_user> ALL = CA_DATACOLL
<sudo_user> ALL = (ALL) SETENV: CA_DATACOLL_INSTALL
<sudo_user> ALL = (ALL) NOPASSWD: CA_DATACOLL_RESTART

```

#### – **DC\_installation\_directory**

The installation directory of the data collector.

**Default:** /opt/IMDataCollector

#### – **sudouser**

Specify the user who can run the sudo commands.

This command alias details the commands that the sudo user must be able to run.

### **Configure UTF-8 Support**

Configure the data collector host to support UTF-8 encoding. If UTF-8 encoding is not enabled, characters might not display properly during the installation.

The appropriate language packs are also required to support localized deployments.

#### **NOTE**

Some of the scripts used in the installation of selected components are not localized and run in English.

For more information, see the *Localization Status Readme* file.

### **Follow these steps:**

1. Issue the following command based on your installation:

#### – **(Korn or bash shell)**

```
export LANG=LANG_value ; export LC_ALL=$LANG
```

- **LANG\_value** specifies the value of the language you want the product to support. The following variables are supported:

**English:** en\_US.utf8

**French:** fr\_FR.utf8

**Japanese:** ja\_JP.utf8

#### **Example:**

```
export LANG=fr_FR.utf8 ; export LC_ALL=$LANG
```

#### – **(Bourne shell)**

```
LANG=LANG_value ; export LANG
```

```
LC_ALL=LANG_value ; export LC_ALL
```

#### **Example:**

```
LANG=fr_FR ; export LANG
```

```
LC_ALL=fr_FR ; export LC_ALL
```

The language variable is set.

### **Set a Unique Hostname for the Data Collector Host**

Set a unique hostname for the computer where you plan to install data collector.

**Follow these steps:**

1. As the root user, log in to the data collector host.
2. Verify the unique hostname on the computer.  
The hostname for the computer must be associated with the IP address and *not* the loopback address of 127.0.0.1.
3. Verify that the following lines appear in the `/etc/hosts` file on the computer:

```
Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
address of your host YourHostName YourHostName.ca.com
```

4. If you changed the hostname, restart the network by issuing the following command:

```
systemctl restart network
```

A unique hostname for the data collector host is set.

**(Optional) Review the Fault-Tolerant Deployment Considerations**

If you are configuring the data collectors for fault tolerance, ensure that you have [reviewed the hardware requirements for fault-tolerant data collectors](#).

**(Optional) Plan for a DX NetOps Virtual Network Assurance Installation on the Same Host as the Data Collector**

If you plan to monitor software-defined networking (SDN) and network functions virtualization (NFV) using VNA, and you plan to install VNA on the *same* host as the data collector, ensure that you allocate enough memory to the host.

**NOTE**

**Best Practice:** Install VNA on a *separate* host as the data collector.

Determine how much memory the data collector will require using [the DX NetOps sizing tool](#).

**Next Steps**

- [Install the data collectors](#).

## Install the Data Collectors

Install the data collectors *after* you install the data aggregator and after you have prepared to install the data collectors.

The data collector collects network item inventory and performance data.

If you are installing multiple data collectors, install each data collector instance on a separate host. If you are reinstalling a data collector, follow the steps for [Update the Data Collector](#).

Use this process to install the data collectors:

1. [Prepare to Install the Data Collectors](#)
2. [Install the Data Collector](#)
3. [Verify the Installation](#)
4. (If you are installing the data collector for fault tolerance) [Configure the Data Collector for Fault Tolerance](#)

The following video examines the data collector installation:

## Install the Data Collector

### Follow these steps:

1. Log in to the data collector host either as the root user or the sudo user.
2. Access the data collector installation package based on your configuration:  
**(If you have HTTP access to the data aggregator host and you are running an X Window System)** Complete the following steps:
  - a. Open a web browser on the data collector host.
  - b. Navigate to the following address, and then download the installation package:  
`http://<da_host>:<port>/dcm/install.htm`
    - **da\_host**  
Specifies the data aggregator host name.
    - **port**  
Specifies the data aggregator required port number.  
**Default:** 8581  
 For more information about the data aggregator server ports that should be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).
  - c. Save the installation package to the /tmp directory.**(If you have HTTP access to the data aggregator host and you are not running an X Window System)** Complete the following steps:
  - a. Open a command prompt on the data collector host.
  - b. Download the installation package to the /tmp directory based on your configuration:
    - (If wget is available) Issue the following command
      - (HTTP data aggregator)  
`wget -P /tmp -nv http://<da_host>:<port>/dcm/InstData/Linux/VM/install.bin`
        - **da\_host**  
Specifies the data aggregator host name.
        - **port**  
Specifies the data aggregator required port number.  
**Default:** 8581
      - (HTTPS-enabled data aggregator) Issue the command with the `--no-check-certificate` option:  
`wget -P /tmp --no-check-certificate -nv https://<da_host>:<port>/dcm/InstData/Linux/VM/install.bin`
        - **da\_host**  
Specifies the data aggregator host name.
        - **port**  
Specifies the data aggregator required port number.  
**Default:** 8582
    - (If wget is unavailable) Issue the following command:
      - (HTTP data aggregator)  
`curl -o /tmp/install.bin http://<da_host>:<port>/dcm/InstData/Linux/VM/install.bin`
        - **da\_host**  
Specifies the data aggregator host name.
        - **port**  
Specifies the data aggregator required port number.



**Default:** 8581

- (HTTPS-enabled data aggregator) Issue the command with the -k option:

```
curl -k -o /tmp/install.bin https://<da_host>:<port>/dcm/InstData/Linux/VM/
install.bin
```

- **da\_host**  
Specifies the data aggregator host name.
- **port**  
Specifies the data aggregator required port number.  
**Default:** 8582

**(If you do not have HTTP access to the data aggregator host)** Open a command prompt on a computer that *does* have HTTP access, and then complete the following steps:

- Download the installation package to your desktop directory using one of the following commands:

- (If `wget` is available) Issue the following command:

- (HTTP data aggregator)

```
wget -P ~/Desktop -nv http://<da_host>:<port>/dcm/InstData/Linux/VM/
install.bin
```

- **da\_host**  
Specifies the data aggregator host name.
- **port**  
Specifies the data aggregator required port number.  
**Default:** 8581

- (HTTPS-enabled data aggregator) Issue the command with the --no-check-certificate option:

```
wget -P ~/Desktop --no-check-certificate -nv https://<da_host>:<port>/dcm/
InstData/Linux/VM/install.bin
```

- **da\_host**  
Specifies the data aggregator host name.
- **port**  
Specifies the data aggregator required port number.  
**Default:** 8582

- (If `wget` is unavailable) Issue the following command:

- (HTTP data aggregator)

```
curl -o /tmp/install.bin http://<da_host>:<port>/dcm/InstData/Linux/VM/
install.bin
```

- **da\_host**  
Specifies the data aggregator host name.
- **port**  
Specifies the data aggregator required port number.  
**Default:** 8581

- (HTTPS-enabled data aggregator) Issue the command with the -k option:

```
curl -k -o /tmp/install.bin https://<da_host>:<port>/dcm/InstData/Linux/VM/
install.bin
```

- **da\_host**  
Specifies the data aggregator host name.
- **port**  
Specifies the data aggregator required port number.  
**Default:** 8582

- Transfer the `install.bin` installation file to the `/tmp` directory on the data collector host.

- Change to the `/tmp` directory by issuing the following command:

```
cd /tmp
```

4. Change the permissions for the installation file by issuing the following command:

```
chmod a+x install.bin
```

5. Run the console installation by issuing the following command based on your installation:

- (As the root user):

```
./install.bin -i console
```

- (As the sudo user):

```
sudo ./install.bin -i console
```

- (In silent mode):

#### NOTE

In silent mode, you generate a response file, and then use this response file to run future installations in silent mode. Add the response file argument.

```
./install.bin -i silent -f <response_file>
```

- **response\_file**

The directory path and file name of the response file.

#### TIP

Generate a response file for a silent installation by adding the following argument, following the prompts until you get to the summary, typing **quit**, and then pressing the **Return/Enter** key on your keyboard:

```
-r <response_file>
```

- **response\_file**

Specifies the directory the directory path and file name for the response file.

**Example:** /temp/installer.properties

- (If you are migrating the data collector as a sudo user):

```
sudo DCM_ID=<data_collector_id> ./install.bin -i console
```

- **data\_collector\_id**

The DCM\_ID that you noted for the data collector.

6. Follow the instructions in the console.

7. When the installer prompts for the data aggregator host information, specify the IP address or the host name for the data aggregator.

#### IMPORTANT

To prevent the data collector from shutting down after the installation, logging an error to the `<DC_installation_directory>/apache-karaf/shutdown.log` file, and requiring an uninstall and reinstall of the data collector, specify the data aggregator host information correctly.

- **DC\_installation\_directory**

The installation directory of the data collector.

**Default:** /opt/IMDataCollector

8. When you are prompted whether to associate this data collector with the Default Tenant, enter **y** or **n**:

#### Options:

- **No:** You are planning to deploy multi-tenancy. You can then associate each data collector installation with a tenant.
- **Yes:** You are *not* planning on deploying multi-tenancy.

9. (Optional) When prompted, specify the following parameters for fault tolerance:

- **Is the data aggregator configured with fault tolerance?**

If you have configured a fault-tolerant environment, specify **2** for Yes.

**Default:** 1 (the data aggregator is not configured for fault tolerance).

- **Second Data Aggregator Host/IP Address**

Specify the host name or IP address of the second data aggregator responsible for this data collector.

**TIP**

You can view this information on the **System Status** page.  
For more information, see [View System Status](#).

For more information, see [Fault Tolerance](#).

The data collector is installed, started, and connects to the data aggregator.

**Verify the Installation****NOTE**

The data collector installation location and the chosen installation options for upgrade are stored in the `/opt/DCM.cfg` file.

**Follow these steps:**

1. Verify that the data collector is running on the data collector host by issuing the following command:  
`systemctl status dcmd`
2. Review the `<DC_installation_directory>/Logs/Performance_Management_Data_Collector_install_<timestamp>.log` file on the data collector host. If the installation is successful, the log shows:  
`0 Warnings, 0 NonFatalErrors, and 0 FatalErrors`

**NOTE**

If the installation fails early enough in the process, this log file might be available in the home directory of the root or sudo user.

– ***DC\_installation\_directory***

The installation directory of the data collector.

**Default:** `/opt/IMDataCollector`

3. Verify that the data collector connection is successful and appears in the list on the **System Status** page.

**NOTE**

The list can take several minutes to refresh and show the data collector that you just installed.

If the data collector is not assigned to a tenant or IP domain (the **Tenant** and **IP Domain** fields are blank), [assign the data collector to a tenant and IP domain](#).

**NOTE**

While you can assign more than one data collector to a single IP domain, you can assign a data collector instance that does discovery requests to only *one* IP domain.

If you are not deploying multi-tenancy, assign the data collector to the Default Tenant and to the Default Domain.

The installation is verified.

**Configure the Data Collector for Fault Tolerance**

If you are installing the data collector for fault tolerance, [put it in standby mode by assigning it to a standby group](#).

**Complete the Post-Installation Configuration**

Perform the following procedures after you install DX NetOps Performance Management:

- [Set Up Autostart on the Data Repository](#)
- [Configure the Automatic Recovery for the Data Aggregator Process](#)
- (Optional) [Modify the External ActiveMQ Memory Limit](#)
- (Optional) [Change the Opened Port Number on the Data Aggregator Host](#)
- (Optional) [Disable the ActiveMQ Admin Console for the Data Aggregator or the Data Collector](#)
- (Optional) [Update Access to the ActiveMQ Admin Console for the Data Aggregator or the Data Collectors](#)
- [Authenticate and Encrypt ActiveMQ Communication](#)

### **Set Up Autostart on the Data Repository**

You can set up autostart on the data repository so that the data repository starts automatically when you reboot the computer where it is installed.

#### **IMPORTANT**

Autostart might not work if the data repository does not shut down properly. If it does not shut down properly, the database might require you to restore the last good epoch manually during startup. If the Vertica database does not start automatically after an improper shutdown, start it using the Vertica Administration Tools (the `adminTools` utility).

The data aggregator stops automatically when the data repository becomes inaccessible. Restart the data aggregator manually after the data repository is online again.

#### **Follow these steps:**

1. Do one of the following steps:
  - Start the Data Aggregator service by issuing the following command:
 

```
systemctl start dadaemon
```
  - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:
 

```
<installation_directory>/scripts/dadaemon activate
```

    - ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
2. Become the Linux user account for the database administrator user by issuing the following command:
 

```
su <dradmin>
```
3. Verify that the Linux user account for the database administrator user is set up with a passwordless ssh key by completing the following steps:
  - a. Verify that the passwordless ssh key is already set up by issuing the following command:
 

```
ssh <dr_host> ls
```
  - b. If the passwordless ssh key is set up, you are *not* prompted for a password, and you can skip setting up the Linux user account for the database administrator (step 4) and enabling ssh (step 5), and continue with setting the restart options (step 6). If the passwordless ssh key is *not* set up, and you *are* prompted for a password, ignore the prompt, press the **Ctrl+C** keys on your keyboard, and then skip setting up the Linux user account for the database administrator (step 4), and continue with enabling ssh (step 5).
4. (Optional) Set up the Linux user account for the database administrator user with a passwordless ssh key by completing the following steps:

#### **IMPORTANT**

The ability to configure autostart on the data repository requires that you set up the passwordless ssh key.

- a. Generate a public key by issuing the following command. In a cluster installation, issue this command on each host that is participating in the cluster:

```
ssh-keygen -N "" -t rsa -f ~/.ssh/id_rsa
```

- b. Copy the contents of the public key to the `authorized_keys2` file on the same computer by issuing the following command. In a cluster installation, copy the contents of the public key to the `authorized_keys2` file on each host in the cluster:
 

```
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys2
```
  - c. (Cluster installation only) Copy the contents of the node's from *each* host to *each of the other* hosts by completing one of the following steps:
    - a. As the database administrator user on the *first* host, complete the following steps:
      - a. Issue the following commands:
 

```
vi ~/.ssh/id_rsa.pub
vi ~/.ssh/authorized_keys2
vi ~/.ssh/authorized_keys2
```
      - b. Copy the content of the `id_rsa.pub` file, and then paste them to the end of the `authorized_keys2` file on the *second* host.
    - b. As the database administrator user on the *third* host, complete the following steps:
      - a. Issue the following command:
 

```
vi ~/.ssh/authorized_keys2
```
      - b. Paste the contents from the `id_rsa.pub` file from the *first* host to the end of the `authorized_keys2` file on the *third* host.
5. Enable ssh from one host to another without being prompted for a password by repeating the following steps for all hosts in the cluster:
  - a. Set permissions for the `authorized_keys2` file by issuing the following command. In a cluster environment, issue these commands on each host in the cluster:
 

```
chmod 644 ~/.ssh/authorized_keys2
```
  - b. As the root user, restart the ssh daemon by issuing the following command. In a cluster environment, issue this command on each host in the cluster:
 

```
su root
systemctl restart sshd
```
  - c. (Single-node installations only) Confirm that you are not prompted for a password by issuing the following commands:
 

```
su <dradmin>
ssh dradmin@<hostname> ls /tmp
```

    - **hostname**  
Specifies the hostname.
  - d. (Cluster installations only) Confirm that you are not prompted for a password by issuing the following commands on each host in the cluster:
 

```
su <dradmin>
ssh dradmin@<host1> ls /tmp
ssh dradmin@<host2> ls /tmp
ssh dradmin@<host3> ls /tmp
```
6. Set the restart options by completing the following steps:
  - a. From an open command prompt, open the `adminTools` utility by issuing the following command:
 

```
/opt/vertica/bin/adminTools
```

The **Administration Tools** dialog opens.
  - b. Select **(6) Configuration Menu**, and then click **OK**.
  - c. Select **(4) Set Restart Policy**, and then click **OK**.  
The **Select Database** dialog opens.
  - d. Select the database name, and then click **OK**.

The **Select policy** dialog opens.

- e. For the **Restart Policy** setting, choose one of the following options, and then click **OK**:
  - **always**: In a single-node data repository installation, the data repository automatically restarts when the system restarts.
  - **ksafe**: In a cluster installation, upon the system restarting, the data repository node automatically restarts if the database still has a status of 'UP'.

The **Restart Policy** setting is saved.

- f. Click **OK** again.
  - g. Select **(M) Main Menu**.
  - h. Select **(E) Exit**.
7. (Optional) Test that the data repository starts when you reboot the computer where the data repository is installed by completing the following steps:
- a. Reboot the computer where the data repository is installed as the root user or sudo user.
  - b. Become the Linux user account for the database administrator user by issuing the following command:
 

```
su <dradmin>
```
  - c. From an open command prompt, open the `adminTools` utility by issuing the following command:
 

```
/opt/vertica/bin/adminTools
```

The **Administration Tools** dialog opens.
  - d. Select **(1) View Database Cluster State**, and then click **OK**.  
The state is "UP."
  - e. Click **OK** again.

#### NOTE

The data repository can take several minutes to start up after you reboot.

Autostart is set up on the data repository.

### Configure the Automatic Recovery for the Data Aggregator Process

If the database server runs out of memory, or if the data repository is unavailable for a time, the data aggregator shuts down automatically to ensure that data consistency is maintained. When the data aggregator shuts down, an audit message is logged in the `<installation_directory>/apache-karaf/shutdown.log` file.

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

When the data aggregator is unavailable, the data collectors continue polling. The data collector caches the poll responses in memory, up to a configurable limit. When the data aggregator host becomes available, the cached polled data is sent to the data aggregator.

**Best Practice:** Disable this cron job before you upgrade the data aggregator. If you stop the Data Aggregator service manually (by issuing the `systemctl stop dadaemon` command), the cron job does not restart the data aggregator automatically. DX NetOps Performance Management can perform maintenance without having the cron job disrupt the system when it is expected to be down.

#### NOTE

In a fault-tolerant environment, Consul manages the start and stop state of the data aggregator. You can skip this procedure.

For more information, see [Fault Tolerance](#).

### Follow these steps:

1. Log in to the computer where the data aggregator is installed as the root user.
2. From an open console, issue the following command:

```
crontab -e
```

A vi session opens. The file opens with the existing cron job definitions for the database administrator user. If there are no cron jobs for this user, an empty file opens.

3. Add the following lines to the file for the cron job:

```
EXECUTED_BY_CRON=1
* * * * * systemctl start dadaemon > /dev/null
```

The cron job issues a start command to the data aggregator every minute. If the data aggregator is running, the `start` command is ignored.

The automatic recovery for the data aggregator process is configured.

### **(Optional) Modify the External ActiveMQ Memory Limit**

The data aggregator installer calculates the memory that is needed on your system to accommodate the Apache ActiveMQ process. However, you can manually modify the memory limit settings to fine tune ActiveMQ on the data aggregator system. For example, you can modify the settings under the following circumstances:

- When the system memory has changed.
- When the number of data collector systems have changed.
- To optimize the memory settings.
- When you have determined that the performance of ActiveMQ is degraded. Monitor the performance through the JConsole or the DX NetOps Performance Management custom chart with ActiveMQ metrics.

#### **Follow these steps:**

1. Calculate the amount of memory for ActiveMQ based on the following settings:
  - **Maximum java heap size**  
Default: 20%  
Minimum: 512M
  - **Initial minimum java heap size**  
50% of maximum java heap size
  - **Young generation java heap size**  
25% of the maximum java heap size
  - **Memory limit for all messages**  
50% of the maximum java heap size
  - **Memory limit per queue**  
Calculate based on how many data collector installations you have.  
**Example:** The memory per queue  
(system memory for all messages)/5/(Data collector count)
2. Log in to the computer where the data aggregator is installed as the root user or a sudo user with access to a limited set of commands.
3. Stop the ActiveMQ broker by issuing the following command:
 

```
systemctl stop activemq
```
4. Modify the java heap size for ActiveMQ by completing the following steps:
  - a. Access the `<installation_directory>/broker/apache-activemq/bin/activemq` file.
    - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
  - b. Locate the line that defines `ACTIVEMQ_OPTS_MEMORY`.
  - c. Change – `Xms` to be the Initial minimum java heap size.
  - d. Change – `Xmx` to be the Maximum java heap size.

- e. Change `-Xmn` to be the Young generation java heap size.
- f. Save the file.
5. Modify the ActiveMQ memory limit for the producer flow control by completing the following steps:
  - a. Access the `<installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml` file.
    - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
    - **apache-activemq-\***  
The installation directory of ActiveMQ.  
**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`
  - b. Locate the following line, and then change the value to memory limit for all messages:  
`<memoryUsage limit="value"/>`
  - c. Locate the following line, and then change the value to memory limit per queue:  
`<policyEntry queue="" producerFlowControl="true" memoryLimit="value"/>`
  - d. Save the file.
6. Start the ActiveMQ broker on the data aggregator system by issuing the following command:  
`systemctl start activemq`

The external ActiveMQ memory limit is modified.

### **(Optional) Change the Opened Port Number on the Data Aggregator Host**

#### **NOTE**

You opened port 61616 before you installed the data aggregator and the data collector.

#### **Follow these steps:**

1. Log in to the computer where the data aggregator is installed as the root user or a sudo user with access to a limited set of commands.
2. Do one of the following steps:
  - Stop the Data Aggregator service by issuing the following command:  
`systemctl stop dadaemon`
  - (Fault-tolerant environment) If the local data aggregator is running, shut it down and prevent it from restarting until maintenance is complete by issuing the following command:  
`<installation_directory>/scripts/dadaemon maintenance`
    - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
3. Remove the data directory and the `local-jms-broker.xml` file from the `deploy` directory by issuing the following commands:  
`rm -rf <installation_directory>/apache-karaf/data`  
`rm -rf <installation_directory>/apache-karaf/deploy/local-jms-broker.xml`
  - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
4. Change the opened port number on the data aggregator host by completing the following steps:
  - a. Access the `<installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml` file.
    - **installation\_directory**  
The installation directory of the data aggregator.



**Default:** /opt/IMDataAggregator

- **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

- Locate the following lines, and then change the 61616, 61618, 61620, 61622 values with the ports that you want to use for incoming connections on the data aggregator:

```
<transportConnectors>
<transportConnector name="openwire" uri="tcp://0.0.0.0:61616"/>
<transportConnector name="PRQ" uri="tcp://0.0.0.0:61618"/>
<transportConnector name="IREP" uri="tcp://0.0.0.0:61620"/>
<transportConnector name="blob" uri="tcp://0.0.0.0:61622"/>
</transportConnectors>
```

**NOTE**

In a fault-tolerant environment, ensure that both data aggregators use the same ActiveMQ ports.

- Save the file.

- Do one of the following steps:

- Start the Data Aggregator service by issuing the following command:  
systemctl start dadaemon
- (Fault- tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing one of the following command:

```
<installation_directory>/scripts/dadaemon activate
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

- Wait a few minutes, and then verify that the port change is successful by issuing the following command:

```
netstat -a | grep <port>
```

- **port**

The port number that you specified previously for incoming connections on the data aggregator.

If the port change is successful, the data aggregator waits for incoming connections on that port.

**NOTE**

If the data aggregator is not waiting for incoming connections, complete the following steps:

- Review the karaf.log file for errors by issuing the following command:

```
grep ERROR karaf.log
```

- Resolve the errors.

- Log in to the computer where the data collector is installed as the root user or a sudo user with access to a limited set of commands.

- From a command prompt, stop the Data Collector service by issuing the following command:

```
systemctl stop dcmd
```

- Remove the data directory and the deploy/local-jms-broker.xml file from the deploy directory by issuing the following commands:

```
rm -rf <installation_directory>/apache-karaf/data
rm -rf <installation_directory>/apache-karaf/deploy/local-jms-broker.xml
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

10. Change the opened port number on the data collector host by completing the following steps:

a. Access the `<DC_installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml` file.

- **DC\_installation\_directory**

The installation directory for the data collector.

**Default:** `/opt/IMDataCollector`

- **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

b. Locate the following lines, and then change the 61616, 61618, 61620, and 61622 values with the ports that you specified previously in the `activemq.xml` file on the data aggregator host:

```
<networkConnector name="da_manager" uri="static:(tcp://dahostname:61616)"
  duplex="true" suppressDuplicateTopicSubscriptions="false">
<networkConnector name="da_manager-PRQ" uri="static:(tcp://dahostname:61618)"
  duplex="true" suppressDuplicateTopicSubscriptions="false">
<networkConnector name="da_manager-IREP" uri="static:(tcp://dahostname:61620)"
  duplex="true" suppressDuplicateTopicSubscriptions="false">
<networkConnector name="da_manager-blob" uri="static:(tcp://dahostname:61622)"
  duplex="true" suppressDuplicateTopicSubscriptions="false">
```

c. Save the file.

11. From a command prompt, start the Data Collector service by issuing the following command:

```
systemctl start dcmd
```

12. Wait a few minutes, and then verify that each port change is successful by issuing the following command:

```
netstat -a | grep <port>
```

– **port**

The port number that you specified in a previous step for incoming connections on the data aggregator.

If the port change is successful, the console shows a connection between the data aggregator and the data collector.

**NOTE**

If you do not see a connection, complete the following steps:

1. Review the `karaf.log` file for errors by issuing the following command:

```
grep ERROR karaf.log
```

2. Resolve the errors.

The opened port numbers on the data aggregator host are changed.

### **(Optional) Disable the ActiveMQ Admin Console for the Data Aggregator or the Data Collector**

Generally, the ActiveMQ admin console should not be available on the network. Therefore, disable it for the data aggregator or the data collector.

#### **Follow these steps:**

1. Disable the ActiveMQ admin console by completing the following steps:

a. Open the following file based on the component for which you want to disable the ActiveMQ admin console:

- **Data Aggregator**

`<installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml`

- **installation\_directory**

The installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

- **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- **Data Collector**

`<DC_installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml`

- **DC\_installation\_directory**

The installation directory for the data collector.

**Default:** `/opt/IMDataCollector`

- **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- Comment out the `<import resource="jetty.xml"/>` line in the file.
  - Save the file.
- Shut down the ActiveMQ broker by issuing the following command based on the component for which you want to disable the ActiveMQ admin console:
    - **Data Aggregator**  
Issue the following command:  
`systemctl stop activemq`
    - **Data Collector**  
Issue the following command on each data collector:  
`systemctl stop activemq`
  - Start the ActiveMQ broker by issuing the following command based on the component for which you want to disable the ActiveMQ admin console:
    - **Data Aggregator**  
Issue the following command:  
`systemctl start activemq`

**TIP**  
If you do not, the data aggregator starts the broker automatically.  
The data collectors automatically restart the ActiveMQ brokers.

  - **Data Collector**  
Issue the following command on each data collector:  
`systemctl start activemq`

The ActiveMQ admin console is disabled for the data aggregator or the data collector.

### **(Optional) Update Access to the ActiveMQ Admin Console for the Data Aggregator or the Data Collectors**

Generally, the ActiveMQ admin console should not be available on the network. However, if certain users require access to the console, you can grant them access.

#### **Follow these steps:**

- Encrypt the user passwords by issuing the following command based on the component for which you want to update access to the ActiveMQ admin console:
  - **Data aggregator**  
`java -cp <installation_directory>/broker/<apache-activemq-*>/lib/web/<jetty-all-*>.jar org.eclipse.jetty.util.security.Password <password>`
  - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
  - **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- **jetty-all-\***

The installation directory of Jetty.

**Example:** `9.4.39.v20210325-uber`

- **password**

The password for ActiveMQ admin console that you want to use for the admin account.

- **Data collector**

```
java -cp <DC_installation_directory>/broker/<apache-activemq-*>/lib/web/<jetty-all-
*>.jar org.eclipse.jetty.util.security.Password <password>
```

- **DC\_installation\_directory**

The installation directory for the data collector.

**Default:** `/opt/IMDataCollector`

- **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- **jetty-all-\***

The installation directory of Jetty.

**Example:** `9.4.39.v20210325-uber`

- **password**

The password for ActiveMQ admin console that you want to use for the admin account.

2. Update user access to the ActiveMQ admin console by completing the following steps:

- a. Open the file based on the component for which you want to update access to the ActiveMQ admin console:

- **Data Aggregator**

```
<installation_directory>/broker/<apache-activemq-*>/conf/jetty-realm.properties
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

- **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- **Data Collector**

```
<DC_installation_directory>/broker/<apache-activemq-*>/conf/jetty-realm.properties
```

- **DC\_installation\_directory**

The installation directory for the data collector.

**Default:** `/opt/IMDataCollector`

- **apache-activemq-\***

The installation directory of ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- b. Replace the MD5: `<encryptedpassword>` entry for the admin account with the new encrypted user password that you specified in the previous step.

- c. Save the file.

3. Shut down the ActiveMQ broker by issuing the following command based on the component for which you want to update access to the ActiveMQ admin console:

- **Data Aggregator**

Issue the following command:

```
systemctl stop activemq
```

- **Data Collector**

Issue the following command on each data collector:

```
systemctl stop activemq
```

4. Start the ActiveMQ broker by issuing the following command based on the component for which you want to update access to the ActiveMQ admin console:

- **Data Aggregator**

Issue the following command:

```
systemctl start activemq
```

**TIP**

If you do not, the data aggregator starts the broker automatically.

The data collectors automatically restart the ActiveMQ brokers.

- **Data Collector**

Issue the following command on each data collector:

```
systemctl start activemq
```

User access to the ActiveMQ admin console on the network has been updated.

### **Authenticate and Encrypt ActiveMQ Communication**

By default, the communication between the data aggregator and the data collector is unencrypted and unauthenticated. To secure communications, [secure the communication between the ActiveMQ brokers on these servers](#).

## **Determine Monitoring Requirements**

Before you configure monitoring in a new environment, use the following guidelines to determine the monitoring requirements of the environment.

To determine the monitoring requirements of the environment, complete the following steps:

1. [Identify data sources](#).
2. [Plan discovery](#).
3. [Plan collections and monitoring profiles](#).
4. [Plan events](#).
5. [Plan dashboards](#).
6. [Plan reports](#).

After you determine the monitoring requirements, [configure monitoring](#).

### **Identify Data Sources**

During the installation, you register the data aggregator as a data source. Many implementations of DX NetOps Performance Management include other data sources, such as DX NetOps Spectrum and CA Application Delivery Analysis. To import inventory from other data sources, register those data sources.

### **Plan Discovery**

Discovery is the process that DX NetOps Performance Management uses to build an inventory of devices in the network. You can define SNMP profiles with the authentication credentials to communicate with devices in the network and use discovery profiles to limit discovery.

### **Plan Collections and Monitoring Profiles**

Collections are system groups that group devices for monitoring. Monitoring profiles control how often to poll devices and which metrics to collect. Associating a collection with a monitoring profile causes DX NetOps Performance Management

to monitor the devices according to the parameters in that profile. You can identify a basic set of monitoring profiles and collections up front to reduce maintenance overhead.

This includes:

- [Plan Metrics](#).
- [Plan Poll Rates](#).
- [Plan Monitoring Profiles](#).
- [Plan Collections](#).
- [Plan Monitoring Profile Filters](#).
- [Plan Group Rules for Collections](#).

### **Plan Metrics**

You can assign metric families to each monitoring profile to specify the metrics that the system should collect. Consider the metrics that are most useful for monitoring the environment.

**Example:** You might plan to collect the following metric families:

- Interface
- CPU
- Memory
- Disk
- IPSLA
- QoS

### **Plan Poll Rates**

In each monitoring profile, you can specify how frequently to poll for specific metrics.

**Example:** You might plan to collect metrics at the following poll rates:

- Interface (1 minute)
- Interface (5 minutes)
- CPU (5 minutes)
- Memory (5 minutes)
- Disk (5 minutes)
- IPSLA (5 minutes)
- QoS (5 minutes)

### **Plan Monitoring Profiles**

You can create a matrix like the following example to determine the necessary monitoring profiles:

Polling Rate	Interface	CPU	Memory	Disk	IPSLA	QoS
1 minute	<b>yes</b>	no	no	no	no	no
5 minutes	<b>yes</b>	<b>yes</b>	<b>yes</b>	<b>yes</b>	yes	yes
15 minutes	no	no	no	no	no	no

You can then use the planned metrics and poll rates to plan the monitoring profiles.

In this example, the administrator plans the following monitoring profiles:

- Fast Interfaces (includes Interface metrics that are polled at 1 minute)
- Standard (includes Interface, CPU, Memory, and Disk metrics that are polled at 5 minutes)
- IPSLA (includes IPSLA metrics that are polled at 5 minutes)
- QoS (includes QoS metrics that are polled at 5 minutes)

You can use monitoring profile filters to refine which managed items the monitoring profile applies.

### **Plan Collections**

Consider each of the managed items as you plan the collections.

In this example, the administrator plans the following collections:

- Fast Interfaces (including WAN links) for association with the Fast Interfaces monitoring profile
- Active Devices (including all active devices with some exceptions) for association with the Standard monitoring profile
- Core Routers (including routers in core) for association with the IPSLA monitoring profile
- Distributed Routers (including routers in the distributed network) for association with the QoS monitoring profile

You can use group rules to refine which managed items the collection applies.

### **Plan Monitoring Profile Filters**

Monitoring profile filters specify criteria that governs which components are monitored. Only the component items that match the filter criteria are polled for the associated metric family. Filtering limits SNMP traffic and ensures that the system monitors only relevant components. The filters of each monitoring profile are assessed independently.

### **Plan Group Rules for Collections**

Use rules to keep the collections up-to-date when systems and networks change. Newly discovered items that meet rule specifications are added to collections. If existing items no longer meet rule requirements, they are removed from collections. After you create a rule, you can modify it by deleting filters or adding subrules.

### **Plan Events**

An event is a message that provides information about what is happening in DX NetOps Performance Management. Events provide information for monitoring the health and status of your system and your environment. All events include basic information, such as related devices and the time of the occurrence that triggered the event.

For more information, see [Events](#).

This includes:

- [Plan Threshold Profiles](#).
- [Plan Monitoring Profile Event Rules](#).

### **Plan Threshold Profiles**

Threshold profiles trigger events when specified conditions occur in associated groups. Event rules define the conditions that trigger events. Each event rule is set to a single metric family, and determines the conditions that cause or clear a violation. Each threshold profile requires at least one event rule.

### **Plan Monitoring Profile Event Rules**

For thresholds that apply broadly to devices in the network, you can add event rules at the monitoring profile level. For example, a rule that creates an event whenever CPU utilization exceeds 95 percent could apply to any device.

## **Plan Dashboards**

Dashboards contain sets of views that show you polled data as meaningful information. DX NetOps Performance Management includes several out-of-the-box dashboards that provide basic information about your infrastructure. To set up dashboards that match your specific monitoring requirements, create a custom dashboard or edit an out-of-the-box dashboard.

## **Plan Reports**

You can access and share On-Demand reports, which dynamically retrieve the most recent data sets from specific sets of items or groups. You can also access and share dashboards and views.

## **Configure Monitoring in a New Environment**

Determine the monitoring requirements of your environment, and configure the system to monitor your network.

After you install DX NetOps Performance Management, complete the following steps:

1. [Determine the monitoring requirements of your environment.](#)
2. Configure the system to monitor your network.

This article discusses the basic procedures and best practices for monitoring. These procedures represent the simplest methods to begin monitoring. Each procedure includes links to topics that provide more information and details about complex configurations. Many of these procedures require the Administrator role.

## **Register Data Sources**

During the installation, you registered the data aggregator as a data source. If your implementation of DX NetOps Performance Management includes other data sources, such as DX NetOps Spectrum and DX NetOps Network Flow Analysis, you can import inventory from those data sources. If your installation includes only the data aggregator data source, skip this procedure.

For more information about how to register a data source, see [Configure a Data Source](#).

## **Discover Devices**

If you do not use quick discovery to discover Simple Network Management Protocol (SNMP) devices, you can configure SNMP profiles and discovery profiles manually.

The following video examines the detailed steps that are required to discover devices:

For more information, see [Discovery](#).

## **Configure SNMP Profiles**

SNMP profiles provide authentication credentials to communicate with devices in your network.

For more information, see [SNMP Profiles](#).

## **Create Discovery Profiles**

Discovery profiles specify which devices NetOps Portal discovers in your network. You can create granular discovery profiles that reduce unnecessary SNMP requests. Create them for devices with different SNMP credentials or different rediscovery schedule.

For more information, see [Manage Discovery Profiles](#).



## **Run Discovery**

To build your inventory, run discovery using the discovery profiles. NetOps Portal discovers the devices within the specified IP addresses and hostnames, and adds the devices to the system inventory.

For more information about how to run discovery, see [Run Device Discovery](#).

## **Configure Monitoring Profiles**

Monitoring profiles control how often to poll devices and which information to collect. Metric families are related sets of metrics that are collected across several technologies. The metric definitions determine how to report the values for the metrics. Metric families normalize performance data from different devices and device types. Assigned metric families determine which metrics the system collects.

For more information, see [Manage Monitoring Profiles](#).

## **Organize Devices into a Device Collection**

Device collections are system groups that group devices for monitoring.

Consider the following best practices for organizing devices into device collections for monitoring:

- Create custom device collections that match the monitoring requirements in the environment:
  - Consider the different layers of the network, access, distribution, and core. Devices in different layers might require different levels of monitoring.
  - Consider which technologies and metric families are required. Metric families that would be applied to all devices, such as CPU and memory, apply to broad device collections. Targeted monitoring, such as QoS and IPSLA, apply to limited device collections.
- Create device collections that enable the flexibility to break out monitoring:
  - Some devices are included in multiple device collections so that specific metric families are polled at different rates.
  - Devices in different device collections have different filtering criteria.
  - Different monitoring requirements depending on importance of device

For more information, see [Device Collections](#).

## **Next Steps**

1. [Assign the device collection to the monitoring profile](#).

Assigning a device collection to a monitoring profile causes NetOps Portal to monitor the devices according to the parameters in that profile.

2. Configure reporting.

For more information, see [Configure Reporting in a New Environment](#).

## **Configure Reporting in a New Environment**

Reporting in DX NetOps Performance Management includes dashboards and threshold alerts. Thresholds and dashboards scope reporting to items in groups.

**Prerequisite:** You have configured monitoring.

In this article:

- [Group Polled Items](#)
- [Build Dashboards](#)
- [Configure Threshold Profiles](#)

## **Group Polled Items**

Groups organize items logically for reporting and thresholding. When you view a dashboard, the data is scoped to your selected group. Individual user profiles are also scoped to particular groups. Each profile is mapped to an initial group that is the selected context when that user logs in.

### **NOTE**

Do not assign All Groups as the default for any users. Because this group includes all items in the system, this context causes dashboards to load slowly.

Organize your group structure according to business and reporting needs. To create a regional structure that represents regions, countries, and locations, use site groups. Use custom groups for other types of organizations, such as customers, services, or technologies.

Threshold profiles apply threshold rules to all items in a group. The group hierarchy requirements for thresholding are probably different from the requirements for reporting. Create separate groups that address both sets of requirements. Consider the different layers of the network and how to create thresholds for components in those layers. For example, you might threshold on CPU, memory, and interface metrics on the core network differently to the distribution layer. Create multiple groups to apply threshold rules appropriately.

For more information, see [Manage Groups](#), [Manage Group Rules](#), and [Configure Threshold Profiles](#).

## **Build Dashboards**

Dashboards contain sets of views that show you polled data as meaningful information. DX NetOps Performance Management includes several out-of-the-box dashboards that provide basic information about your infrastructure. To set up dashboards that match your specific monitoring requirements, create a custom dashboard or edit an out-of-the-box dashboard.

For more information, see [Manage Dashboards](#).

## **Configure Threshold Profiles**

Threshold profiles trigger events when specified conditions occur in associated groups. Event rules define the conditions that trigger events. Each event rule is set to a single metric family, and determines the conditions that cause or clear a violation. Each threshold profile requires at least one event rule.

### **NOTE**

For thresholds that apply broadly to devices in your network, add event rules at the monitoring-profile level. For example, a rule that creates an event whenever CPU utilization exceeds 95 percent could apply to any device.

For more information, see [Configure Monitoring Profiles](#).

For more information, see [Configure Threshold Profiles](#).

## **Install a Disaster Recovery System**

If a large-scale disaster occurs, you can switch over to a recovery (target) system using a disaster recovery plan for DX NetOps Performance Management.

This plan involves provisioning a secondary system as a recovery environment and regularly transferring data from the primary (source) system.

The following video introduces how DX NetOps Performance Management establishes a detailed disaster recovery plan meant to re-establish normal operations in the event of a major disruption, such as a hurricane or fire:

Use the following process to provision a recovery (target) system:

1. [Install the Components for the Recovery System](#)

2. [Configure Incremental Data Transfer for Each Component](#)
3. [Prepare the Disaster Recovery Scripts](#)
4. (Optional) [Test the Recovery System](#)

Then, when required to switch over to the recovery (target) system, [activate it](#).

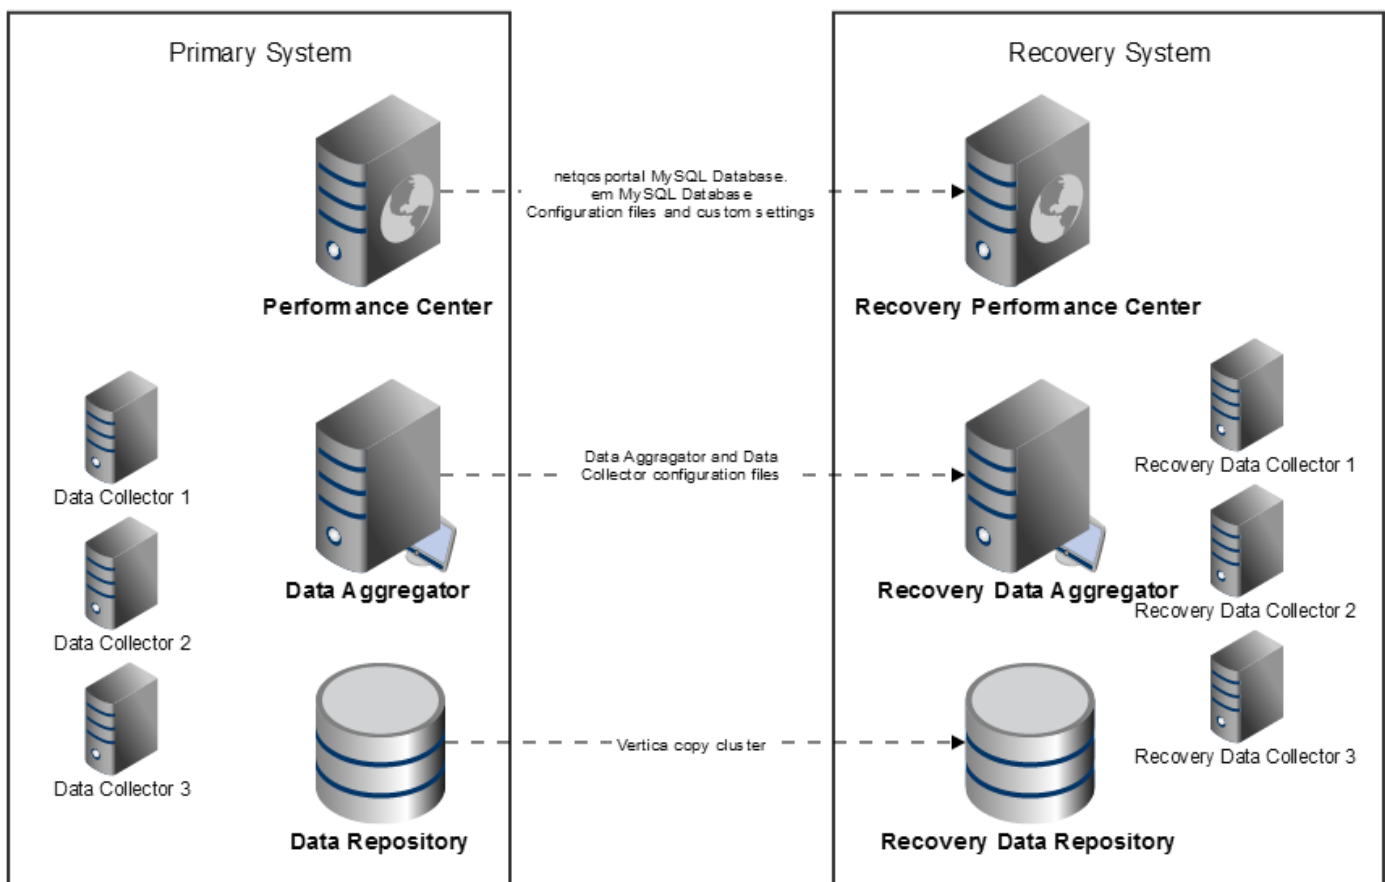
### IMPORTANT

Switching to a recovery (target) system is not a temporary measure. The recovery (target) system completely replaces the primary (source) system. You can reverse this process and return monitoring to the original site by going through the same process with the recovery (target) and primary (source) hosts switched.

If you deploy new hardware due to the disaster, start by reinstalling the DX NetOps Performance Management components on the new hardware. If the original hardware is available and no upgrades have occurred on the active recovery (target) system, start by [configuring incremental data transfer](#).

The following diagram shows the primary (source) and recovery (target) systems, and the files that are copied regularly:

**Figure 7: Disaster Recover Architecture**



### Install the Components for the Recovery System

The recovery (target) system is a secondary system that contains all the components for DX NetOps Performance Management. Under normal operations, the recovery (target) system is offline. It has the same requirements as the primary (source) system.

For more information, see [Review Installation Requirements and Considerations](#).

**IMPORTANT**

If you have upgraded your primary (source) system, upgrade the components in the recovery (target) system. Each recovery component must be running the same DX NetOps Performance Management version as the primary (source) system.

Use the following process to install the components for the recovery (target) system:

1. [Install the Data Repository](#)
2. [Install the Data Aggregator](#)
3. [Install the Data Collectors](#)
4. [Install NetOps Portal](#)

**Install the Data Repository**

Before you install Vertica on the recovery (target) cluster, prepare the environment for the installation.

For more information about how to prepare the recovery (target) cluster, see [Prepare to Install the Data Repository](#).

Ensure that the recovery (target) system has the same settings as the primary (source) cluster for the following settings:

- Database version
- Node names

**TIP**

To get the node name, issue the following command on each node:

```
/opt/vertica/bin/admintools -t list_allnodes
```

This command also returns the installed Vertica version and the database name.

- Database name
- Database administrator
- Database user
- Catalog directory

**TIP**

To get the catalog directory configuration, issue the following Vertica `admintools` command on each node:

```
/opt/vertica/bin/admintools -t list_db -d <database name>
```

- Data directory

The recovery (target) cluster has the following requirements:

- It is accessible from the primary (source) cluster.
- It has the same number of nodes as the primary (source) cluster.
- For each node in the primary (source) cluster, you have verified that the Database Administrator user account (default: `dradmin`) has been set up with passwordless SSH. The `copycluster` command requires that it be set up.

**TIP**

If you have not yet, you can verify that passwordless SSH is working by issuing the following Vertica `admintools` command from any node in the primary (source) cluster:

```
ssh dradmin@<paired-hostname> '/opt/vertica/bin/admintools -t list_allnodes
```

This command executes on the paired-hostname of a node on the target cluster.

- ***paired-hostname***

The hostname of a node in the target cluster.

If the Database Administrator user is prompted for a password, passwordless SSH is not set up.

For more information, see the "[Configure Passwordless SSH for the Database Administrator User](#)" section.

- Port 50000 is open between all the data repository nodes and disaster recovery (target) hosts.

The installation process is the same as a normal data repository installation.

For more information, see [Install the Data Repository](#).

### IMPORTANT

Use the same configuration for the new cluster as for the primary (source) cluster. For example, the Vertica version, node count, database name, administrator, user, catalog directory, and data directory must be the same as the original data repository.

## Install the Data Aggregator

Prepare the host and install the data aggregator for the recovery (target) system.

During installation, use the details for the data repository for the recovery (target) system.

For more information, see [Prepare to Install the Data Aggregator](#) and [Install the Data Aggregator](#).

## Install the Data Collectors

For each data collector in the primary (source) system, install a data collector in the recovery (target) system. The DCM ID identifies the data collector to the system.

### TIP

This scenario assumes that the data collectors are centrally located with the other components. Some deployments use remote data collectors, which are deployed close to the monitored infrastructure in other data centers or geographical locations. To continue using a remote data collector if a disaster occurs, update it to communicate to the recovery data aggregator.

For more information, see [Configure Data Collector When the Data Aggregator IP Address Changes](#).

Use the following process to install the data collectors:

1. Prepare the hosts for the recovery (target) system data collectors.  
For more information, see [Prepare to Install the Data Collectors](#).
2. Record the DCM ID values for each data collector in the primary (source) system:
  - a. Go to the following URL:  
`https://primary_da_host:8581/rest/dcms`
  - b. For each data collector, note the value in the `<DcmID>` tag.  
The following XML shows an example:

```
<DataCollectionMgrInfoList>
<DataCollectionMgrInfo version="1.0.0">
<ID>4077</ID><DcmID>primary-dc: 69658898-a48c-44a6-9cba-963bb9c09684</DcmID>
<Enabled>true</Enabled>
<IPAddress>10.237.1.67</IPAddress>
<RelatedDeviceItem>4078</RelatedDeviceItem>
...
```

3. For *each* recovery (target) data collector, export the DCM ID for the corresponding primary (source) data collector, and install the component in the recovery (target) system:
  - a. On the data collector host for the recovery (target) system, export the DCM ID for the primary (source) system data collector by issuing the following command:

```
export DCM_ID=<DATA_COLLECTOR_DCM_ID>
```

**Example:**

```
export DCM_ID=primary-dc:69658898-a48c-44a6-9cba-963bb9c09684
```

- b. From the same session, install the data collector.  
During the installation, specify the details for the data aggregator for the recovery (target) system.  
For more information, see [Install the Data Collectors](#).
4. To verify the installation, look at the DCM ID on the recovery data collector:
  - a. Open the `<DC_installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml` file.

**Example:**

```
/opt/IMDataCollector/broker/apache-activemq-5.18.3/conf/activemq.xml
```

- **DC\_installation\_directory**

The installation directory for the data collector.

**Default:** `/opt/IMDataCollector`

- **apache-activemq-\***

The installation directory of Apache ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- b. Find the broker name property, and verify the DCM\_ID.  
The following example shows the section of the `activemq.xml` file that includes the broker name:

```
...
<broker
  xmlns="http://activemq.apache.org/schema/core"
  brokerName="dc_broker_69658898-a48c-44a6-9cba-963bb9c09684"
  dataDirectory="${activemq.data}"
  useShutdownHook="false"
  useJmx="true">
...

```

If the broker name does not match the expected DCM\_ID UUID from the originating system, update the file with the correct DCM\_ID UUID from the DCM\_ID of the originating data collector.

5. Stop the data collector services by issuing the following commands:

```
systemctl stop dcmd
systemctl stop activemq
```

or

```
sudo systemctl stop dcmd
sudo systemctl stop activemq
```

6. Stop the data aggregator services by issuing the following command:
  - a. Log in to the data aggregator host for the recovery (target) system.
  - b. Do one of the following steps:
    - Stop the data aggregator and ActiveMQ services by issuing the following commands:
 

```
systemctl stop dadaemon
systemctl stop activemq
```
    - (Fault-tolerant environments) If the local data aggregator is running, shut it down and prevent it from restarting until maintenance is complete by issuing the following command:
 

```
<installation_directory>/scripts/dadaemon maintenance
```

- ***installation\_directory***  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`

## **Install NetOps Portal**

Prepare the host and install NetOps Portal for the recovery (target) system.

### **IMPORTANT**

Do not configure LDAP integration or HTTPS on the recovery (target) system. The settings are inherited from the primary (source) system.

For more information, see [Prepare to Install NetOps Portal](#) and [Install NetOps Portal](#).

## **Configure Incremental Data Transfer for Each Component**

Copy data provides the recovery (target) system everything that it requires to continue operation when the primary (source) system is down. Devise a plan with a regular interval that occurs often enough to duplicate the required data. Use the same frequency for all components and start the transfer from each component simultaneously. Transfer all components daily.

### **TIP**

The data aggregator and NetOps Portal require regular file copies between the primary (source) system and the recovery (target) system. You can use any file copy and schedule method as required in your system. In our lab environment, we configured SSH between the primary (source) and recovery (target) system, and used crontab to invoke the SCP command from the recovery (target) system.

### **Example:**

The following crontab example is configured on the secondary data aggregator host to copy a backup directory from the primary (source) data aggregator. The copy occurs daily at 12:30 AM:

```
30 0 * * * scp -r root@primary_DA:/tmp/backup /tmp/backup
```

Use the following process to configure data transfer:

1. [Configure data transfer for the data repository.](#)
2. [Configure data transfer for the data aggregator.](#)
3. [Configure data transfer for the data collectors.](#)
4. [Configure data transfer for NetOps Portal.](#)

## **Configure Data Transfer for the Data Repository**

For the data repository, duplicate the primary (source) database to the recovery (target) database using the `vbr` script with the `--copycluster` option. `copycluster` is an incremental backup that copies all updates to the database. Because this data transfer is the longest transfer, the backup frequency to the recovery system is limited by the runtime of `copycluster`. Issue the command multiple times before you schedule a regular transfer to verify the runtime. Ensure that the backup frequency is at least twice the runtime of `copycluster`.

### **NOTE**

For existing large databases, `copycluster` takes as long as a full backup to complete. To minimize the performance impact to the system:

1. Restore a backup of the primary (source) system to the recovery system.
2. Configure and issue the `vbr` script with the `--copycluster` option.

**TIP**

For a large database, an incremental `copycluster` for one day takes about one hour. Run an incremental `copycluster` at least daily.

Use the following process to configure data transfer for the data repository:

1. [Configure passwordless SSH for the Database Administrator user.](#)
2. [Create the configuration files for copy cluster.](#)
3. [Stop the database on the target cluster.](#)
4. [Copy the historical data.](#)
5. [Verify the copy of the historical data.](#)
6. [Create a cron Job.](#)

**Configure Passwordless SSH for the Database Administrator User**

The `copycluster` command requires that passwordless SSH is enabled for the Database Administration user account between the nodes on the primary (source) cluster and the target cluster.

For more information, see [Configure Passwordless SSH](#).

**Create the Configuration Files for Copy Cluster**

Create the configuration files for the `copycluster` command.

**Follow these steps:**

1. Log in to the data repository host as the Database Administrator user.
2. Create a password file, for example `/opt/vertica/config/password.txt`.

**NOTE**

You can choose a different location for the password file.

```
[Passwords]
; Specified password for db admin account
dbPassword = DBpassword
; Specifies password for rsync user account - if different than DB admin
; serviceAccessPass = rsyncpwd
; Specifies password for the dest_dbuser Vertica account. Used only for restoring to
; alternate cluster.
; dest_dbPassword = DestinationPwd
```

3. Create a `copycluster.ini` configuration file, specifying the *full* names of the nodes in the source system (source node names) and the host names of the nodes in the target cluster (target host names) in the `[Mapping]` section, as separate lines, as shown in the following example. Create this file in any location on the source system, and give it any name with a `.ini` extension.

**TIP**

You can find the full node names in the `/opt/vertica/config/admintools.conf` file.

**Example:**

The following example configuration file copies a database on a three node cluster (`v_drdata_node0001`, `v_drdata_node0002`, and `v_drdata_node0003`) to another cluster consisting of three nodes (`recovery-host01`, `recovery-host02`, and `recovery-host03`):

**NOTE**

The `dbName` parameter is case-sensitive.

```
[Misc]
snapshotName = Copydrdata
restorePointLimit = 5
```



```
tempDir = /tmp/vbr
retryCount = 5
retryDelay = 1

[Database]
dbName = <drdata>
dbUser = <dradmin>
dbPromptForPassword = False

[Transmission]
encrypt = False
checksum = False
port_rsync = 50000

[Mapping]
; backupDir is not used for cluster copy
v_drdata_node0001= recovery-host01:/data
v_drdata_node0002= recovery-host02:/data
v_drdata_node0003= recovery-host03:/data
```

- **tempDir**  
The directory for the `vbr` utility. Ensure that this directory has read and write permissions.  
**Example:** `/tmp/vbr`
- **backupdir**  
The location of the `password.txt` file.
- **drdata**  
The name of the database to copy.  
**Case-sensitive:** yes
- **dradmin**  
The database administrator account name.

The configuration files are created.

### **Stop the Database on the Target Cluster**

Before you start the migration, shut down the database on the target cluster.

#### **Follow these steps:**

1. Log in to the target cluster as the Database Administrator user.
2. Start the Vertica Administration Tools utility (`adminTools`):  
`/opt/vertica/bin/adminTools`
3. Select **(4) Stop Database**.

Wait for the shutdown to complete.

### **Copy the Historical Data**

Copy the data from the primary (source) database using the `copycluster` command. This command simultaneously backs up the existing source database and restores the data to the target cluster. It configures, starts, and runs the copy cluster on the primary (source) system to point to the destination (target) system.

Issuing the `copycluster` command copies the data in the data repository from *before* you run the command. Because DX NetOps Performance Management continues to collect data while the command is running, the process requires that you issue the command multiple times.

#### NOTE

The `copycluster` command requires that passwordless SSH is enabled for the Database Administrator user account. You enabled this when you set up the data repository.

#### Follow these steps:

1. Log in to the primary (source) cluster as the Database Administrator user account.
2. Issue the following commands:
 

```
chown dradmin /opt/vertica/config/password.txt
chmod 600 /opt/vertica/config/password.txt
```
3. From any node on the primary (source) system, copy the historical data for the source database by issuing the following command:
 

```
vbr.py --task copycluster --config-file CopyClusterConfigurationFile.ini
```

The following message is displayed:

```
> vbr.py --config-file CopyClusterConfigurationFile.ini --task copycluster
Preparing...
Copying...
1871652633 out of 1871652633, 100%
All child processes terminated successfully.
copycluster done!
```

The historical data is copied over to all nodes on the recovery (target) system.

#### Verify the Copy of the Historical Data

After the historical data is copied, ensure the integrity of the data.

#### Follow these steps:

1. >Start the database on the target cluster:
  - a. While logged in to the target cluster as the Database Administrator user, start the Vertica Administration Tools utility (`adminTools`):
 

```
/opt/vertica/bin/adminTools
```
  - b. Select **(3) Start Database**.
2. From any node in the cluster, open the Vertica SQL prompt by issuing the following command:
 

```
/opt/vertica/bin/vsql -U dauser
```
3. Verify the timestamp of these key database tables by issuing the following queries:
 

```
SELECT to_timestamp(max(tstamp)) from dauser.reach_rate;
SELECT to_timestamp(max(tstamp)) from dauser.ifstats_rate;
```

The date and time must correspond to the time when you started the copy.

The copy of the historical data is verified.

#### Create a cron Job

Create a cron job to schedule `copycluster` from the primary (source) system on a regular interval. The following command initiates the transfer:

```
vbr.py --task copycluster --config-file home/dradmin/CopyClusterConfigurationFile.ini
```

### Example:

```
vbr.py --task copycluster --config-file home/dradmin/copycluster.ini
```

## Configure Data Transfer for the Data Aggregator

Schedule a regular copy of the following files from the primary (source) system to the recovery (target) system:

- `<installation_directory>/apache-karaf/deploy/*.xml`

### NOTE

Do not copy this file from the `local-jms-broker.xml` directory. This directory might not initially contain other files.

#### – **installation\_directory**

The installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

- `<installation_directory>/apache-karaf/etc/org.ops4j.pax.web.cfg`

#### – **installation\_directory**

The installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

- `<installation_directory>/data/custom/devicetypes/DeviceTypes.xml`

#### – **installation\_directory**

The installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

In a fault-tolerant environment, a shared directory (for example, `/DASharedRepo`) is defined to help limit data loss.

Therefore, in fault-tolerant environments, the file is located in the `DASharedRepo/custom/devicetypes` directory.

For more information, see [Fault Tolerance](#).

The data transfer for the data aggregator is configured.

## Configure Data Transfer for the Data Collectors

The data collectors do not require a regular backup. All relevant information is stored on the data aggregator and the data repository.

### NOTE

If the primary (source) data collectors include custom memory settings, configure the recovery data collectors as required.

## Configure Data Transfer for NetOps Portal

Create a database dump of the netqosportal and em databases, and back up custom settings.

For more information, see [Back Up NetOps Portal](#).

## Prepare the Disaster Recovery Scripts

The disaster recovery scripts replace hostname and IP address references to match the components in the recovery system.

For each of script, do the following:

1. Create a copy.
2. Provide the relevant information for your system:
  - a. [Update the Data Repository Disaster Recovery Script](#)

## b. Update the NetOps Portal Disaster Recovery

### Update the Data Repository Disaster Recovery Script

On the data repository host in the recovery system, update the bold sections of the `<data_repository_directory>/update_da_dc_database_references.sh` data repository disaster recovery script to match your system:

- **data\_repository\_directory**

The installation directory for the data repository.

**Default:** `/opt/CA/IMDataRepository_verticaVersion`

```
#####
# UPDATE DAUSER/DAPASS BELOW TO REFLECT THE NON-ADMIN
# VERTICA USERNAME/PASSWORD FOR THIS SYSTEM
#####
DAUSER=dauser
DAPASS=dapass
#####
# UPDATE TO REFLECT THE NEW/RECOVERY DATA AGGREGATOR'S IP ADDRESS BELOW
#####
RECOVERY_DA_IP_ADDRESS="<Recovery/New IP Address for the Data Aggregator>"
#####
# UPDATE TO REFLECT THE NEW/RECOVERY DATA AGGREGATOR'S HOSTNAME BELOW
#####
SOURCE_DA_HOSTNAME="<Source/Original Hostname for the Data Aggregator>"
RECOVERY_DA_HOSTNAME="<Recovery/New Hostname for the Data Aggregator>"
#####
# UPDATE THE FOLLOWING ARRAYS TO REFLECT THE SOURCE DATA
# COLLECTOR HOSTNAMES, NEW RECOVERY HOSTNAMES, AND NEW RECOVERY
# IP ADDRESSES RESPECTIVELY.
# IMPORTANT: THE ORDER OF THE ENTRIES BELOW IS CRITICAL FOR
# MAPPING PURPOSES. IN ADDITION, IF MULTIPLE VALUES
# ARE REQUIRED, SEPARATE VALUES WITH A SINGLE SPACE.
#####
declare -a SOURCE_DC_HOSTNAMES=(<Source/Original DC Hostname 1> <Source/Original DC
Hostname 2>)
declare -a RECOVERY_DC_HOSTNAMES=(<New/Recovery DC Hostname 1> <New/Recovery DC Hostname
2>)
declare -a RECOVERY_DC_IP_ADDRESSES=("<New/Recovery DC Hostname 1 IP Address>" "<New/
Recovery DC Hostname 2 IP Address>")
```

#### IMPORTANT

Ensure that the order of the data collectors for the source system and the recovery system is the same. The script uses the order of the list to map the primary (source) system components to the recovery (target) system.

### Update the NetOps Portal Disaster Recovery Script

On the NetOps Portal host in the recovery system, update the bold sections of the `<installation_directory>/Tools/bin/update_pc_da_database_references.sh` NetOps Portal disaster recovery script to match your system:

- **installation\_directory**

The installation directory for NetOps Portal.

**Default:** /opt/CA/PerformanceCenter

```
...
#####
# UPDATE THE FOLLOWING PC/DA VARIABLES TO REFLECT NEW ENVIRONMENT
#####
NEW_PC_IP_ADDRESS="<Recovery/New PC IP Address>"
NEW_PC_HOSTNAME="<Recovery/New PC Hostname>"
NEW_PC_EVENT_PRODUCER_PORT=8181
NEW_PC_EVENT_PRODUCER_PROTOCOL="http" # change to "https" if using SSL
NEW_DA_IP_ADDRESS="<Recovery/New DA IP Address>"
NEW_DA_HOSTNAME="<Recovery/New DA Hostname>"
NEW_DA_PORT_NUMBER=8581
...
```

### **(Optional) Test the Recovery System**

You can optionally test the recovery system.

#### **Follow these steps:**

1. Pause the incremental data transfer.
2. Start the data repository:
  - a. Log in to the recovery database cluster as the database admin user.
  - b. Open Vertica Administration Tools utility (adminTools):  
/opt/vertica/bin/adminTools
  - c. Select option **3 (Start Database)**.
  - d. Press the **Space** bar on your keyboard next to the database name, select **OK**, and then press the **Enter** key on your keyboard.  
You are prompted for the database password.
  - e. Enter the database password, and then press the **Enter** key on your keyboard.  
The data repository starts.
  - f. Select **Exit**, and then press the **Enter** key on your keyboard.
  - g. Run the `<installation_directory>/update_da_dc_database_references.sh` data repository disaster recovery script.
    - **installation\_directory**  
The installation directory for the data repository.  
**Default:** /opt/CA/IMDataRepository\_vertica
3. Start the data aggregator based on your installation:
  - Start the Start the ActiveMQ and Data Aggregator services by issuing the following commands:  
systemctl start activemq  
systemctl start dadaemon
  - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:  
`<installation_directory>/scripts/dadaemon activate`
  - **installation\_directory**  
The installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator starts. If the data repository is unavailable, the data aggregator shuts down.

4. Start NetOps Portal:
  - a. [Restore the NetOps Portal database.](#)
  - b. Run the `<installation_directory>/Tools/bin/update_pc_da_database_references.sh` NetOps Portal disaster recovery script.
    - **installation\_directory**  
The installation directory for NetOps Portal.  
**Default:** /opt/CA/PerformanceCenter
  - c. Start the SSO service by issuing the following command:  
`systemctl start caperfcenter_sso`
  - d. Wait one minute, then start the event manager and device manager by issuing the following commands:

```
systemctl start caperfcenter_eventmanager
systemctl start caperfcenter_devicemanager
```

- e. Wait one minute, then start the NetOps Portal Console service by issuing the following command:

```
systemctl start caperfcenter_console
```

5. Log in to the recovery NetOps Portal component and run reports against the recovery data repository and the data aggregator.
6. Verify that the data is available.
7. (Optional) If you have a set of recovery data collectors that you can double-poll for testing, start one or more data collectors by issuing the following command, then verify polling occurs and the data is stored in the database:

**NOTE**

To prevent the recovery system from issuing duplicate notifications and reports during testing, disable them beforehand.

```
systemctl start dcmd
```

The Data Collector service restarts. If the data aggregator is unavailable, the data collector shuts down.

**TIP**

For remote data collectors, update them to connect to the recovery data aggregator.

For more information, see [Configure Data Collector When the Data Aggregator IP Address Changes](#).

8. Shut down each component.

After the next incremental data transfer, the database is in sync with the primary (source) systems again. To fail over or test again, repeat these steps.

## Install DX NetOps Mediation Manager

DX NetOps Mediation Manager (DX NetOps MM) monitors the performance for non-SNMP based devices, such as mobile wireless, fiber-optic switch, radio access, and 3G or 4G voice data. It supports a wide range of protocols to access data, such as, SOAP, SSH, XML, SQL, JMS, SFTP, and HTTP. DX NetOps Mediation Manager is portable across all platforms.

**IMPORTANT**

- Install DX NetOps Mediation Manager after you install DX NetOps Performance Management. For more information about this step in the installation process, see [Installing](#).
- When DX NetOps Mediation Manager collects the same metric family data through SNMP, some of the data might not be correctly included in rollups. To avoid this problem, collect data from DX NetOps Mediation Manager through a dedicated data collector with a separate IP domain. The data from each IP domain is rolled up independently.

Use the following process to install DX NetOps Mediation Manager for infrastructure management:

1. Install the MultiController on the NetOps Portal host.
2. Install the LocalController on the data collector host.
3. Install the required device packs.  
DX NetOps MM monitors data from non-SNMP devices or obtains data from the Element Management System (EMS) using vendor-specific API plug-ins called device packs. To support a new device type or a new vendor, create and deploy device packs in your environment.

For more information, see the [DX NetOps Mediation Manager documentation](#).

## Install NetOps Kafka

Deploy NetOps Kafka as a cluster or as a single node for use with your installed DX NetOps solutions, such as the OI Connector.

NetOps Kafka is a package that provides an installer for Apache Kafka as well as an out-of-the-box configuration that you can use with DX NetOps.

The following diagram explains the data flow from the OI Connector and the data aggregator to NetOps Kafka:

**Figure 8: NetOps Kafka data flow**

Use the following process to install NetOps Kafka:

1. [Verify the prerequisites](#).
2. Install NetOps Kafka using *one* of the following methods:
  - [Install from the command line](#). With this method, the NetOps Kafka installation prompts you to answer questions for the installation.
  - [Install in silent mode \(silent installation\)](#). With this method, you install NetOps Kafka avoiding the prompts using the values that you define and provide in the `answer.properties` response file.
3. (If you plan to secure the connection) [Configure SSL for NetOps Kafka Using the Self-Signed Certificates](#).
4. (If you are installing NetOps Kafka on a multi-node cluster) After you have installed NetOps Kafka on the primary node, repeat steps 1, 2, and (if you have configured SSL for the primary node) 3 for each node.

The following diagram explains the installation:

**Figure 9: NetOps Kafka installation**

## Verify the Prerequisites for a NetOps Kafka Installation

Before installing NetOps Kafka on a single-node or multi-node cluster for use with your installed DX NetOps solutions, verify the prerequisites.

Complete the following prerequisites before installing NetOps Kafka:

- You have copied the `<netops-kafka-*>-RELEASE.txe` installer file (the installer) to a local directory on the host machine, for example:

```
/root/installer/netops-kafka-3.5.1.1-RELEASE.txe
```

### TIP

You can verify that you have the executable bit set for the installer by issuing the following command:

```
ls -als
```

The following output is expected:

```
86264 -rwxr-xr-x. 1 root root 88331241 May 23 21:02 <file-name>.txe000
```

- You have set up adequate space on the machine on which you are installing NetOps Kafka.
- You are using another system for Kafka data storage other than Network File System (NFS).
- Depending on the number of nodes you want to have in the cluster, you have a corresponding number of independent machines with dedicated reserved memory (no shared memory).
- You have a minimum of three nodes in the cluster for high-availability in production. For best performance, disable swap for these nodes.
- You know which NetOps Kafka version is compatible with the DX NetOps Performance Management version that you have installed. This article includes information about how to install NetOps Kafka 3.5.1.1.
- You have installed one of the following operating systems:
  - Community enterprise Operating System (CentOS) Linux 7 and 8
  - Red Hat Enterprise Linux (RHEL) 7 and 8
  - Oracle Linux (OL) 8
  - Rocky Linux 8
- You have installed Java Runtime Environment 11 (JRE) 11 on the node/device on which you are installing NetOps Kafka. Ensure that you have the latest patched 11 version on the host machine that you will use to install NetOps Kafka. Specify `JAVA_HOME` as Java 11.
- You have ensured that the user account for the user installing NetOps Kafka exists.
- You have ensured that the user installing NetOps Kafka has sudo privileges.
- You have ensured that the following utilities are installed in and accessible to the individual host:
  - openssl
  - sed/awk
  - bash v4 or higher
  - Hostname
  - tar
- You have ensured that the following ports are accessible:

#### NOTE

You can configure these ports.

Function	
ZooKeeper/Kafka	TCP 2181 The port on which the ZooKeeper server will listen for client requests (from the Kafka prompt. In a silent installation (using the <code>answers.properties</code> file), this is the
ZooKeeper	TCP 2888 The leader port for the ZooKeeper cluster.
ZooKeeper	TCP 3888 The election port for the ZooKeeper cluster.
ZooKeeper/Kafka	TCP 2182 The secure client port for the ZooKeeper cluster. In an installation from the command prompt (using the <code>answers.properties</code> file), this is the <code>zookeeperSecureClientPort</code>
Kafka	TCP 9092 The server port for the Kafka brokers to use.
Kafka	TCP 10167 The JMX port.



Function	
Kafka	TCP 10168 The JMX RMI port.

- (If you plan to install NetOps Kafka on a multi-node cluster, and you plan to configure SSL for Kafka using the CA-signed certificates) The certificates are available on each host.
- (If you plan to install NetOps Kafka on a multi-node cluster) You have completed the following prerequisites:
  - You have identified an odd number of participating nodes for the cluster.
  - You know the fully-qualified domain name (FQDN) or IP addresses for the deployment.

## Install NetOps Kafka from the Command Line

Deploy NetOps Kafka answering questions for the installation at the prompts.

In an installation from the command line (interactive mode), the NetOps Kafka installer prompts you for answers.

### Follow these steps:

1. SSH/log into the machine.
2. Launch the installation by issuing the following command from console:

```
./<netops-kafka-*>-RELEASE.txe
```

#### Example:

```
./netops-kafka-3.6.1.4-RELEASE.txe
```

#### – **netops-kafka-\***

The version of NetOps Kafka.

**Example:** netops-kafka-3.6.1.4-RELEASE.txe

The installation is initiated.

#### Output:

```
Do you accept the terms of this License Agreement? [Y/N]
```

3. At the **Do you accept the terms of this License Agreement** prompt, to proceed further, enter **y** for Yes. The NetOps Kafka installation starts. Complete the prompts in the console. The following notification of installing from the command line (interactive mode) appears:

#### Output:

```
Do you accept the terms of this License Agreement? [Y/N]
```

```
Y
```

```
No answers file provided.
```

```
Stepping into interactive mode.
```

```
Installation directory (defaults to /opt/CA/netops-kafka)>
```

4. At the **Installation directory (defaults to /opt/CA/netops-kafka)** prompt, enter the top-level directory for the NetOps Kafka installation.

**Default:** /opt/CA/netops-kafka

#### Output:

```
The name of the user that will run netops-kafka (default install user group root)
```

5. At the **The name of the user that will run netops-kafka (default install user group root)** prompt, enter the user account for the user installing NetOps Kafka and who will own the installation directory and services.

**Default:** root

One of the following lines appear:

```
User <username> exists
```

```
Using default install group (root)
```

#### Output:

```
Java version 11 is installed
```

Do you want to configure netops-kafka with a multi-host cluster [y/n]?

6. At the **Do you want to configure netops-kafka with a multi-host cluster [y/n]?** prompt, enter **y** for Yes to configure a ZooKeeper cluster, or **n** for No to configure as a single node.

**Options:**

- **y:** Configure a multi-node ZooKeeper cluster.
- **n:** Configure a single-node ZooKeeper cluster.

**Default:** n

7. (If you have chosen to configure a multi-node ZooKeeper cluster) Complete the following prompts:
- At the **How many nodes in your cluster (default 3)** prompt, enter the number of nodes for the ZooKeeper cluster.  
**Default:** 3
  - At the **Host 1 of the cluster** prompt, enter the hostname for the node. Repeat this step for the number of nodes for the ZooKeeper cluster.
  - At the **Zookeeper leader port(default 2888)** prompt, enter the leader port for the ZooKeeper cluster.  
**Default:** 2888
  - At the **Zookeeper election port(default 3888)** prompt, enter the election port for the ZooKeeper cluster.  
**Default:** 3888
  - At the **Server id of this server** prompt, enter the numeric value (incremental for subsequent host installation).
  - At the **Zookeeper Memory (in #[m|M|g|G] format) (default 512M)** prompt, enter the maximum memory allowed for the ZooKeeper JVM. This value and maxMemory must not combine to more than 80% of total memory assuming a dedicated server.  
**Default:** 512M
  - At the **Zookeeper client port(default 2181)** prompt, enter the port on which the ZooKeeper server will listen for client requests (from the Kafka broker).  
**Default:** 2181
  - At the **Zookeeper logs directory (defaults to /opt/CA/my-new-dir/zookeeper-logs)** prompt, enter the location where ZooKeeper stores transaction logs. Dedicate a fast disk for this location. Do not use NFS for ZooKeeper data storage.  
**Default:** <installation\_directory>/zookeeper-logs
    - **installation\_directory**  
The installation directory for NetOps Kafka.  
**Default:** /opt/CA/netops-kafka
8. Complete the following prompts:
- At the **Zookeeper data directory (defaults to /opt/CA/my-new-dir/zookeeper-snapshots)** prompt, enter the location where ZooKeeper stores snapshots. This port must be a different port than the ZooKeeper port. Do not use NFS for ZooKeeper data storage.  
**Default:** <installation\_directory>/ZooKeeper-snapshots
    - **installation\_directory**  
The installation directory for NetOps Kafka.  
**Default:** /opt/CA/netops-kafka
  - At the **Max Broker memory (in #[m|M|g|G] format) (default 1G)** prompt, enter the maximum memory allowed for the Kafka broker JVM. This value and Zookeeper Memory should not combine to more than 80% of total memory assuming a dedicated server.  
**Default:** 1G
  - At the **Kafka server port(default 9092)** prompt, enter the server port the Kafka brokers to use.  
**Default:** 9092
  - At the **Data directory (defaults to /opt/CA/netops-kafka/kafka-data-logs)** prompt, enter the location of the Kafka persistence logging data. Dedicate a fast disk/RAID for this location. Do not use NFS for Kafka data storage.  
**Default:** <installationDirectory>/kafka-data-logs
    - **installation\_directory**  
The installation directory for NetOps Kafka.

**Default:** `/opt/CA/netops-kafka`

- At the **Auto-created topic partition count (default 10)** prompt, enter the total number of partitions for the topic that the broker creates automatically.  
**Default:** `10`
- At the **Auto-created topic retention in hours (default 1)** prompt, enter the retention period, in hours, to retain the data for the topic that the broker creates automatically.  
**Default:** `168`
- At the **Broker bind address** prompt, enter the host name/IP to which the broker will bind/listen on the Kafka server port.  
**Default:** The hostname, for example, `localhost`.
- At the **Broker client address** prompt, enter the host name/IP that is advertised to clients for accessing the broker. This address should be accessible to all expected Kafka clients.  
**Default:** The hostname, for example, `localhost`.
- At the **Do you want to configure netops-kafka to use SSL and TLS [y/n]?** prompt, enter **y** for Yes to configure SSL and TLS for Kafka and ZooKeeper, or **n** for No to skip the configuration.

**NOTE**

To generate a keystore and truststore, enter **y** for Yes, and then enter **n** for No at the **Do you have an existing keystore and truststore [y/n]?** prompt.

**Options:**

- **y:** You are configuring SSL and TLS for Kafka and ZooKeeper.
- **n:** You are *not* configuring SSL and TLS for Kafka and ZooKeeper, and you want to skip the configuration.

**Default:** `n`

9. (If you have chosen to configure SSL and TLS for Kafka and ZooKeeper) At the **Do you have an existing certificate** prompt, enter **y** for Yes if you have an existing certificate that a trusted certificate authority (CA) has signed (a CA-signed certificate), or **n** for No if you do not and you want to configure a certificate.

**Options:** `y` or `n`

**Default:** `n`

10. (If you have chosen to configure a certificate) Complete the following prompts:
  - At the **The number of days the certification should be considered valid** prompt, enter the number of days that you want the certificate to be valid.
  - At the **The organization for the certificate** prompt, enter the organization for the certificate.
  - At the **The organization unit for the certificate** prompt, enter the organization unit for the certificate.
  - At the **Email address** prompt, enter the email address for the certificate.
  - At the **The city for the certificate** prompt, enter the city for the certificate.
  - At the **The state for the certificate** prompt, enter the state for the certificate.
  - At the **The country for the certificate** prompt, enter the country code for the certificate.
  - At the **The two letter country code for the certificate** prompt, enter the country code for the certificate.
  - At the **IP Address for SAN** prompt, enter the IP address for Subject Alternative Name (SAN).
  - At the **Hostname for SAN** prompt, enter the host name for SAN.
  - At the **Certificate Authority Password** prompt, enter the password that the CA used to sign the certificate.
  - At the **Keystore password** prompt, enter the password that the CA uses for the keystore.
  - At the **Truststore password** prompt, enter the password that the CA uses for the truststore.
11. At the **Do you have an existing keystore and truststore [y/n]?** prompt, enter **y** for Yes if you have an existing keystore and truststore, or **n** for No if you do not. To generate a keystore and truststore, you must have entered **y** for Yes At the **Do you want to configure netops-kafka to use SSL and TLS [y/n]?** prompt, and then enter **n** for No at this prompt.

**Options:** `y` or `n`

12. (If you have chosen to generate a keystore and truststore) Complete the following prompts:

- At the **Keystore file** prompt, enter the full file path of the existing keystore.
- At the **Keystore password** prompt, enter the password for the keystore.
- (23.3.5 and higher) At the **Keystore type** prompt, enter the type of the existing keystore. (If you are installing NetOps Kafka on a multi-node cluster) Defines the type of the existing keystore for the node.  
**Default:** PKCS12
- At the **Truststore file** prompt, enter the full file path of the existing truststore. (If you are installing NetOps Kafka on a multi-node cluster) Defines the full file path of the existing truststore for the node.
- At the **Truststore password** prompt, enter the password for the truststore.
- (23.3.5 and higher) At the **Truststore type** prompt, enter the type of the existing truststore. (If you are installing NetOps Kafka on a multi-node cluster) Defines the type of the existing truststore for the node.  
**Default:** PKCS12

13. Complete the following prompts:

- At the **Zookeeper secure port** prompt, enter the secure client port for the ZooKeeper cluster. This port must be a different port than the port on which the ZooKeeper server will listen for client requests (from the Kafka broker), **Zookeeper client port(default 2181)**.  
**Default:** 2182
- At the **Do you want to configure JMX monitoring [y/n]?** prompt, specify whether to configure JMX monitoring.  
**Options:**
  - **y:** Configure JMX monitoring.
  - **n:** Do not configure JMX monitoring.

**Default:** n

(If you have chosen to skip the configuration of SSL and TLS for Kafka and ZooKeeper and to not configure JMX monitoring) Lines similar to the following appear:

```
Kafka 3.6.1.4 installation started at 08_12_2023_18_37_11...
Configuration Variables
#####
JDK_HOME = /usr/lib/jvm/java-11-openjdk-arm64
User = jjones
DESTINATION = /opt/CA/netops-kafka
Configure Cluster = Y
Data directory = /opt/CA/netops-kafka/kafka-data-logs
KAFKA Port = 9092
ZOOKEEPER Client Port = 2181
Maximum memory = 1G
Partitions = 10
Retention = 1
Zookeeper Max Memory = 512M
Zookeeper Data = /opt/CA/netops-kafka/zookeeper-snapshots
Zookeeper Logs = /opt/CA/netops-kafka/zookeeper-logs
Zookeeper host = localhost
Client Address = 192.168.105.60
Broker Address = 192.168.105.60
Configure SSL/TLS = n
JMX monitoring status = n
#####
```

14. (If you have chosen to configure JMX monitoring) Complete the following prompts:

- At the **JMX Port (default 10167)** prompt, enter the JMX port.

**Default:** 10167

- At the **JMX RMI Port (default 10168)** prompt, enter the JMX RMI port.

**Default:** 10168

- At the **Do you want to configure JMX authentication [y/n]?** prompt, specify whether to configure user/password-based authentication for JMX.

**Options:**

- **y:** Configure user/password-based authentication for JMX.
- **n:** Do not configure user/password-based authentication for JMX.

**Default:** n

**Optional:** yes

Lines similar to the following appear:

```
Kafka 3.6.1.4 installation started at 08_12_2023_18_41_07...
Configuration Variables
#####
JDK_HOME = /usr/lib/jvm/java-11-openjdk-arm64
User = jjones
DESTINATION = /opt/CA/netops-kafka
Configure Cluster = y
Data directory = /opt/CA/netops-kafka/kafka-data-logs
KAFKA Port = 9092
ZOOKEEPER Client Port = 2181
Maximum memory = 1G
Partitions = 10
Retention = 1
Zookeeper Max Memory = 512M
Zookeeper Data = /opt/CA/netops-kafka/zookeeper-snapshots
Zookeeper Logs = /opt/CA/netops-kafka/zookeeper-logs
Zookeeper host = localhost
Client Address = 192.168.105.60
Broker Address = 192.168.105.60
Configure SSL/TLS = y
JMX monitoring status = y
JMXPort = 10167
JMX RMI Port = 10168
JMX Authentication = n
#####
```

- At the **JMX Password** prompt, enter the password that the JMX authentication uses.

15. At the **Perform install with the above settings? (y/n)** prompt, enter **y** for Yes.

NetOps Kafka is installed.

The following image shows an example of selected options. The installation that follows is based on the options that you select:

```

Do you accept the terms of this License Agreement? [Y/N]
Y
No answers file provided.
Stepping into interactive mode.
Installation directory (defaults to /opt/CA/netops-kafka)> /opt/CA/my-new-dir
The name of the user that will run netops-kafka (default install user root)>
Using default install user (root)
The name of the user group that will run netops-kafka (default install user group root)>
Using default install group (root)
Java version 11 is installed
Do you want to configure netops-kafka with a multi-host cluster [y/n]?> y
How many nodes in your cluster (default 3)> 4
Host 1 of the cluster> host1.test
Host 2 of the cluster> host2.test
Host 3 of the cluster> host3.test
Host 4 of the cluster> host4.test
Zookeeper leader port(default 2888)>
Zookeeper election port(default 3888)>
Server id of this server>
Zookeeper Memory (in #[m|M|g|G] format) (default 512M)>
Zookeeper client port(default 2181)>
Zookeeper logs directory (defaults to /opt/CA/my-new-dir/zookeeper-logs)>
Zookeeper data directory (defaults to /opt/CA/my-new-dir/zookeeper-snapshots)>
Max Broker memory (in #[m|M|g|G] format) (default 1G)>
Kafka server port(default 9092)>
Data directory (defaults to /opt/CA/my-new-dir/kafka-data-logs)>
Auto-created topic partition count (default 10)>
Auto-created topic retention in hours (default 1)>
Broker bind address ([REDACTED]) >
Broker client address ([REDACTED]) >
Do you want to configure netops-kafka to use SSL and TLS [y/n]?> y
Do you have an existing keystore and truststore [y/n]?> n
Do you have an existing certificate [y/n]?> n
The number of days the certificate should be considered valid> 3065
The organization for the certificate> Broadcom
The organization unit for the certificate> AOD
Email address> test@test.com
The city for the certificate> [REDACTED]
The state for the certificate> [REDACTED]
The two letter country code for the certificate> US
IP Address for SAN> [REDACTED]
Hostname for SAN> [REDACTED]
Certificate Authority Password>
Keystore password>
Truststore password>
Zookeeper secure port(default 2182)>
Do you want to configure JMX monitoring [y/n]?> n
Kafka 3.6.1.3 installation started at 22_01_2024_14_23_10...
Configuration Variables
#####
JDK_HOME = /usr/lib/jvm/java-11-openjdk-11.0.21.0.9-2.el8_8.x86_64
User = root
Group = root
DESTINATION = /opt/CA/my-new-dir
Configure Cluster = y
Data directory = /opt/CA/my-new-dir/kafka-data-logs
KAFKA Port = 9092
ZOOKEEPER Client Port = 2181
Maximum memory = 1G
Partitions = 10
Retention = 1
Zookeeper Max Memory = 512M
Zookeeper Data = /opt/CA/my-new-dir/zookeeper-snapshots
Zookeeper Logs = /opt/CA/my-new-dir/zookeeper-logs

```

## Install NetOps Kafka in Silent Mode

Deploy NetOps Kafka avoiding the prompts using the values that you define and provide in the response file.

In a silent installation, the NetOps Kafka installer uses the values in the response file. You define the property values once, and then use the response file on multiple machines as required.

Use the following process to install NetOps Kafka in silent mode:

1. [Define the Response File](#)
2. [Install in Silent Mode](#)

### Define the Response File

Define and provide the following property values in the `answer.properties` response file. The following table lists the properties and property values in the response file. You can include or omit the default value for a property, or you can specify an alternate property value to override the default:

Property	
installationDirectory	Defines the top-level directory for the NetOps Kafka installation. <b>Default:</b> <code>/opt/CA/netops-kafka</code>
javaHome	Identifies the pre-installed Java 11 JRE to run NetOps Kafka. <b>Default:</b> <code>\$JAVA_HOME</code>
user	Defines the user account for the user installing NetOps Kafka and wh <b>Default:</b> <code>root</code>
maxMemory	Defines the maximum memory allowed for the Kafka broker JVM. Ass <b>Default:</b> <code>1G</code>
kafkaPort	Defines the server port the Kafka brokers to use. <b>Default:</b> <code>9092</code>
brokerBindHostAddress	Defines the host name/IP to which the broker will bind/listen on the s <b>Default:</b> <code>&lt;blank&gt;</code>
brokerClientHostAddress	Defines the host name/IP that will be advertised to clients for accessi <b>Default:</b> <code>&lt;blank&gt;</code>
dataDir	Defines the location of the Kafka persistence logging data. Dedicate <b>Default:</b> <code>&lt;installation_directory&gt;/kafka-data-logs</code>
logRetention	(When the broker creates a topic automatically) Defines the retention <b>Default:</b> <code>1 (hour)</code>
partitions	(When the broker creates a topic automatically) Defines the number o <b>Default:</b> <code>10</code>
zookeeperMaxMemory	Defines the max memory allowed to the ZooKeeper JVM. Assuming <b>Default:</b> <code>512M</code>
configCluster	Defines whether to install NetOps Kafka on a multi-node or single-no <b>Options:</b> <ul style="list-style-type: none"> <li>• <b>y</b>: You are installing NetOps Kafka on a multi-node cluster.</li> <li>• <b>n</b>: You are installing NetOps Kafka on a single-node cluster.</li> </ul> <b>Default:</b> <code>n</code>
clusterNodeCount	(If you have chosen to configure a cluster) Defines the number of no <b>Default:</b> <code>3</code>

Property	
zookeeperLeaderPort	(If you have chosen to configure a cluster) Defines the leader port for the ZooKeeper cluster. <b>Default:</b> 2888
zookeeperElectionPort	(If you have chosen to configure a cluster) The election port for the ZooKeeper cluster. <b>Default:</b> 3888
zookeeperHost(x)	Defines the host address the Kafka broker will use to communicate with the ZooKeeper cluster. <b>Default:</b> localhost
zookeeperClientPort	Defines the port on which the ZooKeeper server will listen for client requests. <b>Default:</b> 2181
zookeeperData	Defines the location where ZooKeeper is to store snapshots. Do not use a shared file system. <b>Default:</b> <installation_directory>/zookeeper-snapshots
zookeeperLogs	Defines the location where ZooKeeper will store transaction logs. Do not use a shared file system. <b>Default:</b> <installation_directory>/zookeeper-logs
configureSSL	Specifies whether to configure Secure Sockets Layer (SSL) and Transport Layer Security (TLS) for NetOps Kafka. <b>Options:</b> <ul style="list-style-type: none"> <li><b>y:</b> You are configuring SSL and TLS for NetOps Kafka.</li> <li><b>n:</b> You are <i>not</i> configuring SSL and TLS for NetOps Kafka.</li> </ul> <b>Default:</b> n
existingKSTS	Specifies whether there is an existing keystore and truststore. To generate a keystore and truststore, set this property to n for No, and y for Yes. <b>Options:</b> <ul style="list-style-type: none"> <li><b>y:</b> There is an existing keystore and truststore.</li> <li><b>n:</b> There is <i>not</i> an existing keystore and truststore.</li> </ul> <b>Default:</b> n
keystoreLocation	(If you have chosen to generate a keystore and truststore) Defines the location of the keystore. (If you are installing NetOps Kafka on a multi-node cluster) Defines the location of the keystore on each node.
keystorePassword	(If you have chosen to generate a keystore and truststore) Defines the password for the keystore.
(23.3.5 and higher) keystoreType	(If you have chosen to generate a keystore and truststore) Defines the type of the keystore. <b>Default:</b> PKCS12
truststoreLocation	(If you have chosen to generate a keystore and truststore) Defines the location of the truststore. (If you are installing NetOps Kafka on a multi-node cluster) Defines the location of the truststore on each node.
truststorePassword	(If you have chosen to generate a keystore and truststore) Defines the password for the truststore.
(23.3.5 and higher) truststoreType	(If you have chosen to generate a keystore and truststore) Defines the type of the truststore. <b>Default:</b> PKCS12
zookeeperSecureClientPort	Defines the secure client port for the ZooKeeper cluster. This port must be unique for each node. <b>Default:</b> 2182
validityInDays	(If you have chosen to configure a certificate) Defines the number of days the certificate is valid.
organization	(If you have chosen to configure a certificate) Defines the organization name.
organizationUnit	(If you have chosen to configure a certificate) Defines the organization unit name.
emailAddress	(If you have chosen to configure a certificate) Defines the email address.
location	(If you have chosen to configure a certificate) Defines the city for the certificate.
state	(If you have chosen to configure a certificate) The state for the certificate.
country	(If you have chosen to configure a certificate) Defines the country code.



Property	
ip	(If you have chosen to configure a certificate) Defines the IP address.
dns	(If you have chosen to configure a certificate) Defines the fully-qualified domain name.
caPassword	(If you have chosen to configure a certificate) Defines the password to the certificate authority.
jmxMonitoringEnabled	Specifies whether to configure JMX monitoring. <b>Options:</b> <ul style="list-style-type: none"> <li><b>y:</b> Configure JMX monitoring.</li> <li><b>n:</b> Do not configure JMX monitoring.</li> </ul> <b>Default:</b> n
jmxAuthenticationEnabled	(If you have chosen to configure JMX monitoring) Specifies whether to configure JMX authentication. <b>Options:</b> <ul style="list-style-type: none"> <li><b>y:</b> Configure user/password-based authentication for JMX.</li> <li><b>n:</b> Do not configure user/password-based authentication for JMX.</li> </ul> <b>Default:</b> y
jmxUser	(If you have chosen to allow configuration of JMX Monitoring) Defines the JMX user name. <b>Example:</b> kafkauser
jmxPassword	(If you have chosen to allow configuration of JMX Monitoring) Defines the JMX password. <b>Example:</b> changeit
jmxPort	(If you have chosen to allow configuration of JMX Monitoring) Defines the JMX port. <b>Example:</b> 10167
jmxRMIPort	(If you have chosen to allow configuration of JMX Monitoring) Defines the JMX RMI port. <b>Example:</b> 10168

**Example:**

The following example shows the `answers.properties` file for a multi-node cluster installation of four nodes (`configCluster=y`):

```

user=vagrant
javaHome=/usr/lib/jvm/java-11-openjdk-11.0.13.0.8-4.el8_5.x86_64
partitions=10
logRetention=168
installationDirectory=/opt/CA/netops-kafka
zookeeperData=/opt/CA/netops-kafka/zookeeper-snapshots
zookeeperLogs=/opt/CA/netops-kafka/zookeeper-logs
dataDirectory=/opt/CA/netops-kafka/kafka-data-logs
maxMemory=1024M
zookeeperMaxMemory=1024M
configCluster=y
clusterNodeCount=4
# Cluster Nodes
clusterHost1=oraclebrokers1.test
clusterHost2=oraclebrokers2.test
clusterHost3=oraclebrokers3.test
clusterHost4=oraclebrokers4.test
zookeeperClientPort=2181
zookeeperLeaderPort=2888

```

```

zookeeperElectionPort=3888
zookeeperSecureClientPort=2182
kafkaPort=9092
# Server ID will be automatically generated
serverId=1
brokerBindHostAddress=oraclebrokers1.test
brokerClientHostAddress=oraclebrokers1.test
configureSSL=y
existingKSTS=n
keystorePassword=changeit
truststorePassword=changeit
caPassword=changeit
validityInDays=365
organization=Broadcom
organizationUnit=AOD
emailAddress=test@test.com
location=HYD
state=TS
country=IN
ip=10.100.100.1
dns=oraclebrokers.test
# JMX properties
jmxMonitoringEnabled=y
jmxAuthenticationEnabled=y
jmxUser=kafka
jmxPassword=changeit
jmxPort=10167
jmxRMIPort=10168

```

### **Install in Silent Mode**

**Prerequisite:** You have defined and provided property values in the response file.

Issue the following command, using the response file as an argument:

```
./<netops-kafka-*>-Release.txe answers.properties
```

### **Example:**

```
./netops-kafka-3.6.1.4-Release.txe answers.properties
```

- ***netops-kafka-\****  
Defines the version of NetOps Kafka.

**Example:** netops-kafka-3.6.1.4-RELEASE.txe

The installer automatically deploys the instance, using the property values from the response file.

## **Install and Configure Log Analytics for Insights**

As an Administrator, establish the connection between DX Operational Intelligence - SaaS and DX NetOps Performance Management.

Establish the connection by installing and configuring the OI Connector and the log collector. This enables the Log Analytics in DX NetOps Insights capabilities with NetOps Portal.

For more information about log analytics, see Log Analytics in [the DX Operational Intelligence documentation](#).

Use the following process to install and configure log analytics for Insights:

1. [Verify the Prerequisites](#)
2. [Install the OI Connector](#)
3. [Download and Install the Log Collector](#)
4. [Configure the Log Collector](#)

After you have installed and configured log analytics for Insights, you can do the [next step](#).

### **Verify the Prerequisites**

Before installing and configuring log analytics for Insights, verify that you have completed the following the following steps:

- You have an DX Operational Intelligence - SaaS tenant with log analytics enabled for the selected tenant.
- You have installed DX NetOps Performance Management 23.3.x.

### **Install the OI Connector**

Install the OI Connector for the tenant. For new tenant provisioning based on your entitlement, log analytics is enabled by default. For existing tenants, you provision log analytics on-demand.

Complete the following steps as an Administrator.

#### **Follow these steps:**

1. Download and install the OI Connector. Complete the following steps:
  - a. Log in as a user with the Manage OI Connector Status role right.
  - b. Hover over **Administration, Configuration Settings**, and then click **OI Connector Status**.  
The **Manage Connector to DX Operational Intelligence** page appears. The **OI Connector Status** shows as Normal when the OI Connector has been installed, and as Unknown when it has not. The following image shows an example of this page when the OI Connector has been installed:



## Manage Connector to DX Operational Intelligence

OI Connector Status ✓ Normal

### Download and Install OI Connector

Connecting DX Operational Intelligence and DX NetOps together unleashes the power of analytics in NetOps.

It's just three steps to get started!

- 1 Enter your tenant URL:  Save  
Your DXI Administrator can provide you with the URL. You need to have an OI Administrator user account to be able to access your connector.
- 2 Download your connector. Download Connector  
The connector will be listed in your downloads page.
- 3 Install your connector.  
? Installation [instructions](#) for your connector.

#### TIP

When the OI Connector has not been installed already, and you attempt to view the logs for a device (from the **Log Events** tab), you can also get to the **Manage Connector to DX Operational Intelligence** page by clicking the link in the message that appears:

No log information available for this device. [The OI Connector] must be configured [here](#).

c. Complete the steps.

For more information about how to install the OI Connector, see [Install and Upgrade the OI Connector](#).

The OI Connector is installed.

### Download and Install the Log Collector

#### Follow these steps:

1. Login to DX SaaS.
2. Click **Settings**.
3. Click **Downloads** under the **External connections/integrations** section.
4. Click **DX Gateway** to download the package.

The DX Gateway package is downloaded, and the log collector is installed.

### Configure the Log Collector

For more information about how to install and configure the log collector, see Log Analytics in [the DX NetOps Insights documentation](#).

Use the following process to configure the log collector:

1. Identify the network devices for which you must enable syslog monitoring, and then configure syslog for those devices.

2. Install and configure rsyslog server. Configure the individual device's syslog to forward to a centralized rsyslog server. Update the rsyslog server config file to forward syslog to the log collector.

**NOTE**

You can manage the rsyslog server using DX NetOps Spectrum. DX NetOps Spectrum includes a rsyslog config file.

3. (Optional) Enable syslog monitoring by way of DX NetOps Spectrum.  
For more information about syslog monitoring, see [the DX NetOps Spectrum documentation](#).

**TIP**

You can check detailed DX NetOps syslog monitoring. For details, see the following video:

**Next Step**

After you have installed and configured log analytics for Insights, you can discover network devices, such as routers, switches, and firewalls, in NetOps Portal.

For more information, see [Use Log Analytics for Insights](#).

## Deploy NetOps Flow

Network devices, such as routers and switches, send flow data to NetOps Flow. NetOps Flow processes the data it receives and stores it into the data repository for NetOps Portal to consume.

Use the following process to deploy NetOps Flow:

1. [Verify the Prerequisites](#)
2. [Set up to Deploy NetOps Flow](#)
3. [Deploy NetOps Flow](#)

After you have deployed NetOps Flow, you can perform the [next steps](#).

The following video examines how to deploy NetOps Flow:

For more information about NetOps Flow, see [NetOps Flow](#).

**Verify the Prerequisites**

Before deploying NetOps Flow, ensure that you have completed the following prerequisite steps:

- You have verified that a Kubernetes cluster is installed, and that the Kubernetes client is installed and configured to manage that cluster.

**TIP**

You can list the Kafka client Kubernetes pods in all namespaces by issuing the `kubectl get pods` command with the following option:

```
kubectl get pods --all-namespaces
```

The Kubernetes configuration is good if this command does not fail.

- Helm 3 is installed.

**TIP**

You can verify that Helm is installed by issuing the following command:

```
helm list
```

The installation is good if this command does not fail.

- (If you want NetOps Flow-related status when you view system status in NetOps Portal) Ingress is configured.  
For more information about how to view system status, see [View System Status](#).

## **Set up to Deploy NetOps Flow**

Use the following process to set up to deploy NetOps Flow:

1. [Download the NetOps Flow Distribution File](#)
2. [Extract the NetOps Flow Distribution File](#)
3. [Import the Container Images for NetOps Flow into the Docker Registry](#)
4. [Configure Kubernetes to use the Docker Registry Credentials](#)
5. [Create the Kafka Topics for NetOps Flow](#)

The following video examines how to prepare to deploy NetOps Flow:

## **Download the NetOps Flow Distribution File**

Download the NetOps Flow distribution file into the file management system:

```
NetOps-Flow-<VERSION>-RELEASE.tar.gz
```

### **Example:**

```
NetOps-Flow-23.3.1-RELEASE.tar.gz
```

- **VERSION**  
The version for NetOps Flow.  
**Example:** 23.3.1

For more information, contact Broadcom Support.

## **Extract the NetOps Flow Distribution File**

Issue the following command:

```
tar -zxvf NetOps-Flow-<VERSION>-RELEASE.tar.gz
```

### **Example:**

```
tar -zxvf NetOps-Flow-23.3.1-RELEASE.tar.gz
```

- **VERSION**  
The version for NetOps Flow.  
**Example:** 23.3.1

From the current path, the `netops-flow-k8s-distribution` directory contains the files that are required to deploy NetOps Flow.

## **Import the Container Images for NetOps Flow into the Docker Registry**

The `flow-images.tgz` file includes the NetOps Flow container images. The `flow-images.tgz` file is in the `NetOps-Flow-<VERSION>-RELEASE.tar.gz` file. Import the container images for NetOps Flow into the Docker registry.

### **Follow these steps:**

1. Change directory by issuing the following command:  

```
cd netops-flow-k8s-distribution/netops-k8s-installer/registry/ansible/assets/netops/dist/
```
2. Issue the following command:  

```
export REGISTRY=<REGISTRY_URL>;  
output="$(docker image load --input flow-images.tgz)"  
while IFS= read -r line; do
```

```

image=$(echo $line | sed 's/Loaded image: //g')
target_image="$REGISTRY/$(echo $image | sed 's/[^\/*\.\$]/\1/' )"
docker tag $image $target_image
docker image rm $image
docker push $target_image
docker image rm $target_image
done <<< "$output"

```

**Example:**

```

export REGISTRY=myregistry.broadcom.net:5000
output="$(docker image load --input flow-images.tgz)"
while IFS= read -r line; do
    image=$(echo $line | sed 's/Loaded image: //g')
    target_image="$REGISTRY/$(echo $image | sed 's/[^\/*\.\$]/\1/' )"
    docker tag $image $target_image
    docker image rm $image
    docker push $target_image
    docker image rm $target_image
done <<< "$output"

```

**– REGISTRY\_URL**

The URL for Docker registry.

**Example:** myregistry.broadcom.net

The container images for NetOps Flow are imported into the Docker registry.

**Configure Kubernetes to use the Docker Registry Credentials****Follow these steps:**

1. Create the `netops-docker-registry` Kubernetes secret for the Docker registry that includes the container images. Issue the following `kubectl create secret` command for every Kubernetes namespace to where NetOps Flow will be deployed, with the following options:

```

kubectl -n <NAMESPACE> create secret docker-registry
netops-docker-registry --docker-server=<REGISTRY_URL>
--docker-username=<REGISTRY_USER>
--docker-password=<REGISTRY_PASSWORD> --docker-email=<ANY_EMAIL>

```

**Example:**

```

kubectl -n ns1 create secret docker-registry
netops-docker-registry --docker-server=myregistry.broadcom.net
--docker-username=vhall
--docker-password=password --docker-email=vhall@mycompany.com

```

**– NAMESPACE**

The Kubernetes namespace to where NetOps Flow will be deployed. To use the default namespace, `default`, omit the `-n <NAMESPACE>` option.

**Example:** ns1

**– REGISTRY\_URL**

The URL for Docker registry.

**Example:** myregistry.broadcom.net

**– REGISTRY\_USER**

The username for Docker registry.

– **REGISTRY\_PASSWORD**

The password for the user for Docker registry.

– **ANY\_EMAIL**

The email for the user for Docker registry.

2. (Optional) Patch the default Kubernetes service account by issuing the following `kubectl patch serviceaccount` commands for every Kubernetes namespace to where NetOps Flow will be deployed, with the following options:

**NOTE**

If you choose not to patch the default account, then set the `global.imagePullSecrets` value when you deploy the NetOps Flow Helm chart (issue the `helm install` command).

For more information about this value, see [the "Deploy NetOps Flow" section](#).

```
kubectl -n <NAMESPACE> patch serviceaccount default -p '{"secrets":
[{"name": "netops-docker-registry"}]}'
kubectl patch serviceaccount default -p '{"imagePullSecrets":
[{"name": "netops-docker-registry"}]}'
```

**Example:**

```
kubectl -n ns1 patch serviceaccount default -p '{"secrets":
[{"name": "netops-docker-registry"}]}'
kubectl patch serviceaccount default -p '{"imagePullSecrets":
[{"name": "netops-docker-registry"}]}'
```

– **NAMESPACE**

The Kubernetes namespace to where the NetOps Flow components will be deployed. To use the default namespace, `default`, omit the `-n <NAMESPACE>` option.

**Default:** `ns1`

Kubernetes is configured to use the Docker registry credentials.

### Create the Kafka Topics for NetOps Flow

By default, NetOps Flow attempts to create the Kafka topics automatically. However, if the Kafka server in use *does not* allow Kafka clients to create Kafka topics, create these topics, and then when deploying the NetOps Flow Helm chart by issuing a `helm install` command, set the NetOps Flow Helm chart `global.flow.initFlowTopics` value to `false`.

For more information about this value, see [the "Deploy NetOps Flow" section](#).

If the Kafka server in use *allows* Kafka clients to create Kafka topics, you can skip this procedure.

**NOTE**

`kafka-topics` is a Kafka CLI tool which is available in most Kafka distributions. You can create the required Kafka topics using the Kafka management interface of your choice, or you can use a free and commercial user interface (UI) and command-line interface (CLI). In this procedure, `kafka-topics` is used.

Create the Kafka topics using the `kafka-topics` tool by issuing the following command:

```
kafka-topics --bootstrap-server <KAFKA_ADDRESS> --create --topic <TOPIC_NAME> --config
retention.ms=<RETENTION_TIME_MS> --partitions <NUMBER_OF_PARTITIONS>
```

**Example:**

The following command creates a `flow.data-record` Kafka topic with one partitions and a 600000 retention time:

```
kafka-topics --bootstrap-server kafka-server.mycompany.net:9092 --create --topic 2-flow.data-record --config
retention.ms=600000 --partitions 1
```



For more information about the `kafka-topics` CLI tool, see [the Kafka documentation](#).

Create the following Kafka topic for NetOps Flow:

```
flow-item-facet-service
```

Create the following Kafka topics for each IP Domain from which NetOps Flow will consume flow data:

- `<IP_DOMAIN_ID>-flow-aggregated-record`
- (23.3.3 and lower) `<IP_DOMAIN_ID>-flow-aggregator-windowed-data-store-changelog`
- (23.3.3 and lower) `<IP_DOMAIN_ID>-flow-aggregator-windowed-data-suppress-store-changelog`
- `<IP_DOMAIN_ID>-flow-application-mapping-record`
- `<IP_DOMAIN_ID>-flow-data-record`

#### NOTE

- **KAFKA\_ADDRESS**  
Defines the address that the Kafka topic will use to connect to Kafka. Set this value to a hostname/FQDN or IP address, followed by the port. If you are issuing this command from a system which hosts the Kafka service, set this value to `localhost:9092`.  
**Example:** `kafka-server.mycompany.net:9092`
- **TOPIC\_NAME**  
Defines the name of the Kafka topic that you are creating.  
**Example:** `flow.data-record`
- **RETENTION\_TIME\_MS**  
Defines the retention period, in milliseconds, to retain the Kafka topic messages. Set this value to `600000`.
- **NUMBER\_OF\_PARTITIONS**  
Defines the number of partitions for the topic. Set this value to `1`.
- **IP\_DOMAIN\_ID**  
Defines the ID for the IP domain that the deployment will use as the prefix for the Kafka topic name.  
**Example:** `10`

#### TIP

You can find the IDs for an IP domain by opening the following URL:

```
http://<DA_HOST>:<DA_PORT>/rest/ipdomains
```

Look for the ID XML tags.

## Deploy NetOps Flow

Deploying NetOps Flow involves deploying it into a Kubernetes cluster using the Helm chart that is included with the NetOps Flow distribution, in the `netops-flow-k8s-distribution/helm/netops-flow-helm.tgz` file.

#### IMPORTANT

If you are deploying NetOps Flow to more than one IP domains, deploy the NetOps Helm chart once per IP domain.

For more information, see [the "NetOps Flow Helm Chart Deployment Examples" section](#).

Deploy NetOps Flow into a Kubernetes cluster using the NetOps Flow Helm chart by issuing a `helm install` command, using the following Helm chart values (23.3.3 and lower) *in addition to* the values in the `netops-flow-k8s-distribution/helm/netops-flow-override.yaml` NetOps Flow Helm chart override file:

- **global.netops.ipDomainID**  
Defines the IP domain ID to which to deploy NetOps Flow, and from which NetOps Flow can receive flow data. To allow NetOps Flow to discover inventory across *all* IP domains, set this value to `0`.

#### NOTE

If you are deploying NetOps Flow configured to discover inventory across all IP domains, ensure that there are no overlapping IP addresses on devices in different IP domains.

For example, you are deploying to two IP domains, **IP domain A**, which has an IP domain ID of 10, and **IP domain B**, which has an IP domain ID of 20. To have NetOps Flow monitor both IP domains, do the following:

1. Create a set of Kafka topics with the prefix 10- and another set with the prefix 20- .
2. Issue the `helm install` command with this value set to 10 .
3. Issue the `helm install` command with this value set to 20 .

**Default:** 2

**TIP**

You can find existing NetOps Portal IP domain IDs using the `ipdomains` endpoint for the Data Aggregator REST server (by entering the `http://<DA_HOST>:<DA_PORT>/rest/ipdomains` URL in your browser), and then looking for the ID XML tags.

- **global.da.tls.enabled**

Specifies whether the communication between NetOps Flow and the data aggregator is secured with transport layer security (TLS) encryption for HTTPS.

**Options:**

- **true:** TLS is enabled. NetOps Flow will communicate with the data aggregator using HTTPS with TLS encryption.
- **false:** TLS is not enabled. NetOps Flow will communicate with the data aggregator using HTTP.

**Default:** `false`

- **global.da.tls.secure**

Defines whether NetOps Flow validates the data aggregator server certificate.

**Options:**

- **true:** NetOps Flow validates the data aggregator server certificate.

**NOTE**

If you set this value, you must set the `global.da.tls.trustStore.file` value (provide a truststore file).

- **false:** NetOps Flow will not validate the data aggregator server certificate.

**Default:** `true`

- **global.da.tls.trustStore.file**

Defines the path to the truststore (PKCS12) file containing the certificate for NetOps Flow to access the data aggregator.

For more information about this certificate/truststore file, see [Enable HTTPS for the Data Aggregator](#).

**NOTE**

You set this value using the `--set-file` Helm argument.

- **global.da.tls.trustStore.password**

Defines the password for the truststore for the data aggregator.

- **global.da.host**

Defines the data aggregator hostname/FQDN or IP address.

**Required:** Yes

- **global.da.port**

Defines the data aggregator port.

**Default:** 8581

- **global.pc.tls.enabled**

Specifies whether the communication between NetOps Flow and NetOps Portal will be secured with TLS encryption for HTTPS.

**Options:**

- **true:** TLS is enabled. NetOps Flow will communicate with NetOps Portal using HTTPS with TLS encryption.
- **false:** TLS is disabled. NetOps Flow will communicate with NetOps Portal using HTTP.

**Default:** `false`

- **global.pc.tls.secure**

Specifies whether NetOps Flow will validate the NetOps Portal server certificate.

**Options:**

- **true:** NetOps Flow will validate the NetOps Portal server certificate.

**NOTE**

If you set this value, you must set the `global.pc.tls.trustStore.file` value (provide a truststore file).

- **false:** NetOps Flow will not validate the NetOps Portal server certificate.

**Default:** `true`

- **global.pc.tls.trustStore.file**

Defines the path to the truststore (PKCS12) file containing the certificate for NetOps Flow to access NetOps Portal. For more information about this certificate/truststore file, see [Enable HTTPS for NetOps Portal](#).

**NOTE**

You set this value using the `--set-file` Helm argument.

- **global.pc.tls.trustStore.password**

Defines the password for the truststore for NetOps Portal.

- **global.pc.host**

Defines the NetOps Portal hostname/FQDN or IP address.

**Required:** Yes

- **global.pc.port**

Defines the NetOps Portal port.

**Default:** `8181`

- **global.pc.adminUser**

Defines the NetOps Portal administrator username.

**Default:** `admin`

- **global.pc.adminPassword**

Defines the NetOps Portal administrator password.

**Default:** `admin`

- **global.dr.hosts**

Defines the data repository address.

**NOTE**

- You can enter hostnames/FQDNs or IPs as the value.
- Enter a single host, or to provide multiple hosts for high-availability purposes, in the command line, use commas (,) and surround the value by double quotes (") and curly brackets ({}), for example, `--set global.dr.hosts="{host1,host2}"`.

**Required:** Yes

- **global.dr.port**

Defines the port for communication between the data aggregator and the data repository.

**Default:** `5433`

- **global.dr.dbUser**

Defines the username for the data repository database into which NetOps Flow stores the flow data it processes.

**NOTE**

Set a value that is the same as the credentials that the data aggregator uses to connect to and interact with the data repository database.

For more information about this user, see [Install the Data Repository](#).

**Default:** `dauser`

- **global.dr.dbPassword**

Defines the password for the data repository database into which NetOps Flow stores the flow data it processes.

**Default:** `dapass`

- **global.dr.dbName**

Defines the name of the data repository database into which NetOps Flow stores the flow data it processes.

**Default:** `drdata`

- **global.kafka.bootstrapServers**

Defines the addresses for communicating with Kafka.

**NOTE**

- The format of this value is `host:port`, where the host can be a hostname/FQDN or IP address.
- Enter a single server address, or to provide multiple addresses for high-availability purposes, use commas, for example, `kafka-0.mycompany.net:9092,kafka-1.mycompany.net:9092`.

**Required:** Yes

- **global.kafka.exposedBootstrapServers**

Defines the Kafka address to be shared with the NetOps data aggregator.

**NOTE**

- If the Kafka address is reachable from NetOps Flow and the data aggregator, set a value that is the same as the `global.kafka.bootstrapServers` value. Otherwise, enter a Kafka address to which the data aggregator can reach.
- The format of this value is `host:port`, where the host can be a hostname/FQDN or IP address.
- Enter a single server address, or to enter multiple addresses, use commas, for example, `kafka-0.mycompany.net:9092,kafka-1.mycompany.net:9092`.

**Required:** Yes

- **global.kafka.security.protocol**

Defines the security protocol for Kafka.

**Options:**

- **PLAINTEXT:** No data transmission encryption. Data is transmitted as plain text.
- **SSL:** Enables data transmission encryption.

**NOTE**

If you set this value, set the following values:

- `global.kafka.ssl.keyPassword`
- `global.kafka.ssl.trustStore.file`
- `global.kafka.ssl.trustStore.password`
- `global.kafka.ssl.keyStore.file`
- `global.kafka.ssl.keyStore.password`

**Default:** `PLAINTEXT`

- **global.kafka.ssl.keyPassword**

Defines the Kafka SSL keystore password.

**NOTE**

If you set the `global.kafka.security.protocol` value to `SSL`, set this value.

- **global.kafka.ssl.trustStore.file**

Defines the path to the trust store (PKCS12) file containing the certificate to access the Kafka service.

**NOTE**

- You set this value using the `--set-file` Helm argument.
- If you set the `global.kafka.security.protocol` value to `SSL`, set this value.

- **global.kafka.ssl.trustStore.password**

Defines the password for the Kafka SSL truststore.

**NOTE**

If you set the `global.kafka.security.protocol` value to `SSL`, set this value.

- **global.kafka.ssl.keyStore.file**

Defines the path to the key store (PKCS12) file containing the key to access the Kafka service.

**NOTE**

- You set this value using the `--set-file` Helm argument.
- If you set the `global.kafka.security.protocol` value to `SSL`, set this value.

- **global.kafka.ssl.keyStore.password**

Defines the password for the Kafka SSL keystore.

**NOTE**

If you set the `global.kafka.security.protocol` value to `SSL`, set this value.

- **flow-data-mgmt.enabled**

Defines TLS encryption as enabled for the Kubernetes Ingress, and whether the Kubernetes Ingress serves NetOps Flow status to NetOps Portal by way of HTTPS or HTTP.

**Options:**

- **true:** The deployment deploys the NetOps Flow Data Management service.

**IMPORTANT**

Ensure that there is only one NetOps Flow Data Management service across all deployments. If NetOps Flow will receive flow data from *more than one* IP domain, deploy only one NetOps Flow Data Management service across all IP domains. Do this by setting the value to `true` for the first deployment (the first IP domain), and then setting the value to `false` for consecutive deployments (the other IP domains). For example, if this is the first IP domain for which you are deploying the NetOps Flow Helm chart, do the following:

1. Set this value to `true`.
2. Give each NetOps Flow Helm chart deployment a unique release name.
3. Set the `global.netops.ipDomainID` value with the IP domain ID.

- **false:** The deployment does not deploy the NetOps Flow Data Management service.

**Default:** `true`

- **flow-data-mgmt.serviceAccessHostForPc**

Defines the address for NetOps Portal to retrieve data from NetOps Flow so that it can populate the NetOps Flow status in NetOps Portal (the **System Status** page).

For more information, see [View System Status](#).

**NOTE**

- Set the value to a hostname/FQDN or IP address. Use an address that is resolvable by DNS, but you can also set it to a Kubernetes node address, but this is not a high-availability alternative. A Kubernetes Ingress ensures routing to the appropriate component.
- If you set `flow-data-mgmt.enabled` to `true`, set this value.

- **flow-data-mgmt.serviceAccessPortForPc**

Defines the port that NetOps Portal uses to reach the NetOps Flow Data Management component.

**NOTE**

(23.3.4 and higher) If you set this value to 0, the port is set based on the value for `flow-data-mgmt.ingress.tls.enabled`. If `flow-data-mgmt.ingress.tls.enabled` is set to `true`, NetOps Portal uses port 443 to reach the NetOps Flow Data Management component, otherwise, it uses port 80.

**Default:** 80

- **flow-data-mgmt.dataRetentionPeriodDays**

Defines the number of days NetOps Flow stores the flow data it processes in the data repository.

**Default:** 45

- **flow-data-mgmt.ingress.tls.enabled** (23.3.3 and higher)

The `flow-data-mgmt` service serves NetOps Flow status to NetOps Portal using a Kubernetes Ingress. This Helm value defines whether TLS encryption is enabled for the `flow-data-mgmt` Kubernetes Ingress.

**Options:**

- **true:** TLS encryption is enabled for the `flow-data-mgmt` Kubernetes Ingress. The NetOps Flow status is available by way of HTTPS.

**NOTE**

(23.3.4 and higher) If you set to `true`, set the value for `global.flow.ingress.hostname` to be the same as the value for `flow-data-mgmt.serviceAccessHostForPc`.

- **false:** TLS encryption is not enabled for the `flow-data-mgmt` Kubernetes Ingress. The NetOps Flow status is available by way of HTTP.

**Default:** `false`

- **flow-data-mgmt.ingress.sslRedirect** (23.3.3 and higher)

Determines whether the `flow-data-mgmt` Kubernetes Ingress redirects HTTP clients to HTTPS.

**Options:**

- **true:** The `flow-data-mgmt` Kubernetes Ingress redirects HTTP clients to HTTPS.
- **false:** The `flow-data-mgmt` Kubernetes Ingress does not redirect HTTP clients to HTTPS.

**Default:** `false`

- **flow-data-mgmt.ingress.forceSslRedirect** (23.3.3 and higher)

Defines whether the `flow-data-mgmt` Kubernetes Ingress redirects HTTP clients to HTTPS, even when a TLS certificate is not available.

**Options:**

- **true:** The `flow-data-mgmt` Kubernetes Ingress redirects HTTP clients to HTTPS, even when a TLS certificate is not available.
- **false:** The `flow-data-mgmt` Kubernetes Ingress does not redirect HTTP clients to HTTPS when a TLS certificate is not available.

**Default:** `false`

- **flow.ingress.tls.certificate** (23.3.4 and higher)

Defines the path to the Kubernetes Ingress TLS certificate file.

**NOTE**

- You set this value using the `--set-file` Helm argument.
- If `flow.ingress.tls.certificate` and `flow.ingress.tls.privateKey` are set, a Kubernetes secret named `flow-tls-secret` containing both the certificate and the private key is created.

- **flow.ingress.tls.privateKey** (23.3.4 and higher)

Defines the path to the Kubernetes Ingress TLS private key file.

**NOTE**

- You set this value using the `--set-file` Helm argument.
- If `flow.ingress.tls.certificate` and `flow.ingress.tls.privateKey` are set, a Kubernetes secret named `flow-tls-secret` containing both the certificate and the private key is created.

- **flow-collector-nfa.enabled**

Defines whether the deployment uses sFlow and/or NetFlow v5/v7 processing.

**Options:**

- **true:** The deployment uses sFlow and/or NetFlow v5/v7 processing.

**NOTE**

If you set to `true`, then include `nfa-collector` flow data listener in the `flow-proxy.enabledListeners` list.

- **false:** The deployment does not use sFlow and/or NetFlow v5/v7 processing.

**Default:** `false`

- **flow-proxy.enabledListeners**

Defines that you want to use flow data listeners for NetOps Flow, and defines the protocols that you want sent to these listeners (to NetOps Flow). For example, if IPFIX, NetFlow v9, and sFlow are the protocols that you want sent to NetOps Flow, enter `ipfix,netflow-v9,nfa-collector` as the value. In the command line, enter the listeners using commas, and surround it by double quotes and curly brackets.

**Example:** `--set flow-proxy.enabledListeners="{ipfix,netflow-v9,nfa-collector}"`

**Values:**

- **ipfix:** Enables the listener for the IPFIX protocol.
- **netflow-v9:** Enables the listener for the NetFlow v9 protocol.
- **nfa-collector:** Enables the flow data listener (the listener for the sFlow, NetFlow v5, and NetFlow v7 protocols).

**Prerequisite:** You have set the `flow-collector-nfa.enabled` value to `true`.

**Default:** `{ipfix,netflow-v9}`

- **global.imagePullSecrets**

Defines the Kubernetes secrets that Kubernetes uses to authenticate with the Docker Registry so that it can pull the container images for NetOps Flow.

**NOTE**

- You can enter one or more secrets for this value.
- If you are using an external Docker container image registry, enter a value that is an appropriate Kubernetes secret for the Docker registry.
- If you have chosen to patch the default Kubernetes service account when you configured Kubernetes to use the Docker registry credentials, leave this value unset.

**Default:** not set

**Example:**

```
--set global.imagePullSecrets[0].name=my-registry-secret-a
--set global.imagePullSecrets[1].name=my-registry-secret-b
```

- **global.imagePullPolicy**

Defines the policy for Kubernetes to pull the container images for NetOps Flow from the Docker registry.

**Options:**

- **Always:** Kubernetes will always pull the container images from the Docker container image registry when a pod starts. Use this option when using a Docker registry or similar.
- **IfNotPresent:** Kubernetes will pull the container images from the Docker container image registry when a pod starts only if the container images are not already available in the local Kubernetes container runtime.
- **Never:** Kubernetes will not pull the container images from the Docker container image registry when a pod starts. Use this option in rare cases where the NetOps Flow container images were made available in the local container runtime of all Kubernetes nodes.

**Default:** (23.3.4 and higher) `Always` (23.3.3 and lower) `Never`

- **global.imageRegistry**

Defines the Docker container image registry address for Broadcom container images. The deployment uses this value as the prefix for the name of the Broadcom container images. Enter the Docker container image registry's FQDN or IP address.

**Default:** `esd-netops-docker-virtual.artifactory-lvn.broadcom.net`

- **global.dockerHubProxy**

Defines the Docker container image registry address for non-Broadcom container images. The deployment uses this value as a prefix for the name of the non-Broadcom container images. Enter the Docker container image registry's FQDN or IP address.

**NOTE**

This value and `global.imageRegistry` can have the same value, depending on the images that are available in the Docker container image registry in use. For example, if the registry to where you imported the NetOps Flow container images acts as a virtual repository, and the registry is also a proxy for the public Docker Hub, then both values can have the same value.

**Default:** `esd-netops-docker-virtual.artifactory-lvn.broadcom.net`

- **global.flow.ingress.hostname** (23.3.4 and higher)

Defines the host name that NetOps Flow uses for the flow-data-mgmt Ingress.

**NOTE**

If `flow-data-mgmt.ingress.tls.enabled` is set to `true`, set this to a value that is the same as the value for `flow-data-mgmt.serviceAccessHostForPc`.

**Default:** not set

- **global.flow.initFlowTopics**

Specifies whether NetOps Flow creates and configures the Kafka topics when you deploy the NetOps Flow Helm chart.

**Options:**

- **true:** NetOps Flow attempts to create and configure the required Kafka topics if they do not exist.

- **false:** NetOps Flow will not attempt to create and configure the Kafka topics.

**NOTE**

If the Kafka server in use does not allow Kafka clients to create Kafka topics, and you created these topics manually, then set this value to `false`, and then [manually create the topics](#).

**Default:** `true`

NetOps Flow is deployed into a Kubernetes cluster using the NetOps Flow Helm chart.

## **NetOps Flow Helm Chart Deployment Examples**

The following examples assume that the current path is `netops-flow-k8s-distribution/helm/`.

### **Example 1**

- Single IP domain (IP domain ID 2, the default).
- TLS is disabled.
- Default service account patched to use the appropriate Kubernetes secret for the Docker container image registry.
- IPFIX and NetFlow v9 protocols is enabled.
- Kafka topic initialization is disabled.

### **(23.3.4 and higher)**

```
helm install netops-flow netops-flow-helm.tgz \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.flow.initFlowTopics=false \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=10.220.22.14
```

### **(23.3.3 and lower)**

```
helm install netops-flow netops-flow-helm.tgz -f netops-flow-override.yaml \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
```



```
--set global.flow.initFlowTopics=false \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=10.220.22.14
```

## Example 2

- Single IP domain (IP domain ID 15).
- TLS is enabled (for NetOps Flow, the data aggregator, NetOps Portal, and Kafka).
- Default service account is not patched to use a Kubernetes secret for the Docker container image registry.
- IPFIX and NetFlow v9 protocols is enabled.
- Kafka topic initialization is disabled.

### (23.3.4 and higher)

```
helm install netops-flow netops-flow-helm.tgz \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=false \
--set global.netops.ipDomainID=15 \
--set global.da.host=da-host.mycompany.net \
--set global.da.tls.enabled=true \
--set global.da.port=8582 \
--set-file global.da.tls.trustStore.file=/tmp/da-server.pkcs12 \
--set global.da.tls.trustStore.password=changeit \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.pc.tls.enabled=true \
--set global.pc.port=8182 \
--set-file global.pc.tls.trustStore.file=/tmp/pc-server.pkcs12 \
--set global.pc.tls.trustStore.password=changeit \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=host0.mycompany.net \
--set flow-data-mgmt.serviceAccessPortForPc=0 \
--set flow-data-mgmt.ingress.tls.enabled=true \
--set flow-data-mgmt.ingress.sslRedirect=true \
--set flow-data-mgmt.ingress.forceSslRedirect=true \
--set global.flow.ingress.hostname=host0.mycompany.net \
--set-file flow.ingress.tls.certificate=/tmp/tls.crt \
--set-file flow.ingress.tls.privateKey=/tmp/tls.key
```

### (23.3.3 and lower)

```
helm install netops-flow netops-flow-helm.tgz -f netops-flow-override.yaml \
--set global.imageRegistry=image-registry.mycompany.net \
```

```

--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=false \
--set global.netops.ipDomainID=15 \
--set global.da.host=da-host.mycompany.net \
--set global.da.tls.enabled=true \
--set global.da.port=8582 \
--set-file global.da.tls.trustStore.file=/tmp/da-server.pkcs12 \
--set global.da.tls.trustStore.password=changeit \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.pc.tls.enabled=true \
--set global.pc.port=8182 \
--set-file global.pc.tls.trustStore.file=/tmp/pc-server.pkcs12 \
--set global.pc.tls.trustStore.password=changeit \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=10.220.22.14

```

### Example 3

- Two IP domains (IP domain IDs 15 and 27).
- TLS is disabled.
- Default service account is not patched to use a Kubernetes secret for the Docker container image registry.
- IPFIX and NetFlow v9 protocols is enabled.
- Kafka topic initialization is disabled.

### # Helm chart install command line for the IP domain ID 15

#### (23.3.4 and higher)

```

helm install netops-flow-15 netops-flow-helm.tgz \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=false \
--set flow-data-mgmt.enabled=true \
--set global.netops.ipDomainID=15 \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=10.220.22.14

```

#### (23.3.3 and lower)

```

helm install netops-flow-15 netops-flow-helm.tgz -f netops-flow-override.yaml \
--set global.imageRegistry=image-registry.mycompany.net \

```

```
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=false \
--set flow-data-mgmt.enabled=true \
--set global.netops.ipDomainID=15 \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=10.220.22.14
```

## # Helm chart install command line for the IP domain ID 27

### (23.3.4 and higher)

```
helm install netops-flow-27 netops-flow-helm.tgz \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=false \
--set flow-data-mgmt.enabled=false \
--set global.netops.ipDomainID=27 \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass
```

### (23.3.3 and lower)

```
helm install netops-flow-27 netops-flow-helm.tgz -f netops-flow-override.yaml \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=false \
--set flow-data-mgmt.enabled=false \
--set global.netops.ipDomainID=27 \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass
```

## Example 4

- Single IP domain (IP domain ID 2, the default).
- TLS is disabled.
- Default service account is not patched to use a Kubernetes secret for the Docker container image registry.
- IPFIX, sFlow, NetFlow v5, v7, and v9 protocols is enabled.
- Kafka topics initialization is enabled.

#### (23.3.4 and higher)

```
helm install netops-flow netops-flow-helm.tgz \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=true \
--set flow-collector-nfa.enabled=true \
--set flow-proxy.enabledListeners="{ipfix,netflow-v9,nfa-collector}" \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=10.220.22.14
```

#### (23.3.3 and lower)

```
helm install netops-flow netops-flow-helm.tgz -f netops-flow-override.yaml \
--set global.imageRegistry=image-registry.mycompany.net \
--set global.dockerHubProxy=image-registry.mycompany.net \
--set global.imagePullPolicy=Always \
--set global.imagePullSecrets[0].name=image-registry-secret \
--set global.flow.initFlowTopics=true \
--set flow-collector-nfa.enabled=true \
--set flow-proxy.enabledListeners="{ipfix,netflow-v9,nfa-collector}" \
--set global.da.host=da-host.mycompany.net \
--set global.pc.host=pc-host.mycompany.net \
--set global.pc.adminUser=admin \
--set global.pc.adminPassword=mypass \
--set global.dr.hosts="{dr-host-0.mycompany.net,dr-host-1.mycompany.net}" \
--set global.dr.dbUser=dauser \
--set global.dr.dbPassword=mypass \
--set flow-data-mgmt.serviceAccessHostForPc=10.220.22.14
```

### **Next Steps**

After you have deployed NetOps Flow, complete the following procedures:

1. [Check that the pods for NetOps Flow are ready.](#)
2. [Configure the devices.](#)
3. [Set up to view network flow.](#)

### **Check that the Pods for NetOps Flow are Ready**

**Prerequisite:** The NetOps Flow pods are in "Running" or "Completed" status.

Issue the following command:

```
kubectl -n <NAMESPACE> get pods
```

#### Example:

```
kubectl -n ns1 get pods
```

- **NAMESPACE**

Defines the Kubernetes namespace to where the NetOps Flow components are deployed.

**Example:** ns1

#### TIP

Get live output by including the `--watch` option with this command:

```
kubectl -n <NAMESPACE> get pods --watch
```

To stop watching the pods, press the **Ctrl+C** keys on the keyboard.

#### Example Output:

NAME	READY	STATUS	RESTARTS	AGE
flow-aggregator-bcbbcf595-wxbzp	2/2	Running	0	1m
flow-collector-ipfix-5f8949cd95-tkm7c	1/1	Running	0	1m
flow-collector-nfv9-967b49b5f-g6rqd	1/1	Running	0	1m
flow-data-mgmt-6dd5fc7479-rk4s9	1/1	Running	0	1m
flow-inventory-lookup-5c879c5cc8-2f7tx	1/1	Running	0	1m
flow-kafka-init-8nlrj	0/1	Completed	0	1m
flow-processor-644cc4f756-lp4v6	1/1	Running	0	1m
flow-proxy-264g2	1/1	Running	0	1m
flow-proxy-4x8qv	1/1	Running	0	1m
flow-proxy-gqsqq	1/1	Running	0	1m
flow-proxy-jx2dm	1/1	Running	0	1m
netops-hazelcast-0	1/1	Running	0	1m
netops-hazelcast-1	1/1	Running	0	1m
netops-hazelcast-2	1/1	Running	0	1m

#### NOTE

If you consistently see pods crashing and/or image pull related errors (for example, `ImagePullBackOff`), issue the following command, and then investigate the problem:

```
kubectl -n <NAMESPACE> describe pod <POD_NAME>
```

#### Example:

```
kubectl -n ns1 describe pod <POD_NAME>
```

- **NAMESPACE**

Defines the Kubernetes namespace to where the NetOps Flow components have been deployed. To check a pod for the default namespace, `default`, omit the `-n <NAMESPACE>` option.

**Example:** ns1

- **POD\_NAME**

Defines the name of the pod.

## Configure the Devices

Configure the devices to send flow data to NetOps Flow using the Internet Protocol Flow Information Export (IPFIX, also known as NetFlow v10), NetFlow v5/v7/v9, or sFlow protocols.

The NetOps Flow collector services are exposed as Kubernetes Node Ports at the following ports:

- **IPFIX (NetFlow v10):** 30739
- **NetFlow v9:** 30995
- **sFlow and NetFlow v5/v7:** 30996

### Follow these steps:

1. Find the IPs of a Kubernetes node by issuing the following command:

```
kubectl get nodes -o wide
```

#### Example output:

NAME	STATUS	ROLES	AGE	VERSION	INTERNAL-IP	EXTERNAL-IP	OS-IMAGE
		KERNEL-VERSION			CONTAINER-RUNTIME		
node1	Ready	master	159d	v1.19.10	10.110.10.101	<none>	CentOS
Linux 7 (Core)		3.10.0-1160.31.1.el7.x86_64			docker://19.3.14		
node2	Ready	master	159d	v1.19.10	10.110.10.102	<none>	CentOS
Linux 7 (Core)		3.10.0-1160.31.1.el7.x86_64			docker://19.3.14		
node3	Ready	<none>	159d	v1.19.10	10.110.10.103	<none>	CentOS
Linux 7 (Core)		3.10.0-1160.31.1.el7.x86_64			docker://19.3.14		
node4	Ready	<none>	159d	v1.19.10	10.110.10.104	<none>	CentOS
Linux 7 (Core)		3.10.0-1160.31.1.el7.x86_64			docker://19.3.14		

2. Look for the `INTERNAL-IP` column, for example, `10.110.10.101`.
3. Complete *one* of the following:
  - (Preferred) Use an external load balancer by completing the following steps:
    - a. Configure this load balancer to load balance the flow data to all Kubernetes nodes addresses at the ports exposed for NetOps Flow.
    - b. Configure each device to send flow data to the load balancer address.
  - Configure each device to send flow data to any Kubernetes node address. These nodes expose the same set of ports for NetOps Flow.

The devices are configured to send flow data to NetOps Flow.

### Set up to View Network Flow

Complete the following procedures after you have deployed NetOps Flow:

- [Ensure that a Discovery Profile is Discovering the Devices](#)
- [View Flow from the Flow Dashboards](#)
- [Manage Application Mappings](#)

### Ensure that a Discovery Profile is Discovering the Devices

Verify the following:

- That [a discovery profile is discovering the devices](#).
- That [the All Devices device collection is assigned to the Network Interface monitoring profile](#).  
For more information about device collections, see [Device Collections](#).

### View Flow from the Flow Dashboards

You can view flow from [the NetOps Flow \(flow\) dashboards](#).

## Manage Application Mappings

Application mappings group series of hosts by port (tie a port to multiple hosts) and define the hosts as an application. They define the application mapping of flow traffic in the Sankey diagram on the **Flow Dashboard**.

For more information about how to add application mappings, see [Manage Application Mappings](#).

## Enable the NetOps Business Reports

Enable the NetOps business reports so that you can run them in NetOps Portal.

(23.3.2 and higher)

### IMPORTANT

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

For more information, see [NetOps Business Reports](#).

Use the following process to enable the NetOps business reports:

1. [Verify the Prerequisites for Installing the NetOps Report Manager Service](#)
2. [Install the NetOps Report Manager Service](#)
3. [Verify the Installation](#)
4. [Set up to Run NetOps Business Reports](#)

After you have enabled the NetOps business reports, you can perform the [next step](#).

### Verify the Prerequisites for Installing the NetOps Report Manager Service

For more information, see [Verify the prerequisites for installing the NetOps Report Manager Service](#).

### Install the NetOps Report Manager Service

The NetOps Report Manager Service is a standalone application that is required to run NetOps business reports.

Install the NetOps Report Manager Service using *one* of the following methods:

- [Install from the Command Line](#). With this method, the NetOps Report Manager Service installation prompts you for answers (interactive mode).
- [Install in Silent Mode](#). With this method, you install avoiding the prompts using the values that you define and provide in the `silent.properties` response file.

### Verify the Installation

Follow these steps:

1. Check the status of the daemon by issuing the following command:

```
systemctl status srm-rib
```

The following codeblock shows an example of the expected output:

```
srm-rib.service - NetOps Report Manager Service
Loaded: loaded (/etc/systemd/system/srm-rib.service; disabled; vendor preset: disabled)
Active: active (running) since Fri 2023-06-30 16:34:37 UTC; 22s ago
Process: 13695 ExecStart=/opt/SRM/srm-rib/bin/srm-rib.sh start (code=exited, status=0/SUCCESS)
Main PID: 13758 (java)
```

```

CGroup: /system.slice/srm-rib.service
└─13758 /opt/SRM/jre/bin/java -Djetty.home=/opt/SRM/srm-rib/jetty -Djava.io.tmpdir=/tmp -
Djetty.home=/opt/SRM/srm-rib/jetty -Djetty.base=/opt/SRM/srm-rib -Xms256m -Xmx2646m -Djetty.http.port=8081
-DlogDir=/op...

Jun 30 16:34:31 vt046651-cabi srm-rib.sh[13695]: 2023-06-30 16:34:31.958:INFO :oejshC.srm_rib:main: 1
Spring WebApplicationInitializers detected on classpath
Jun 30 16:34:32 vt046651-cabi srm-rib.sh[13695]: 2023-06-30
16:34:32.429:INFO :oejss.DefaultSessionIdManager:main: Session workerName=node0
Jun 30 16:34:32 vt046651-cabi srm-rib.sh[13695]: 2023-06-30 16:34:32.440:INFO :oejshC.srm_rib:main: Set
web app root system property: 'webapp.root' = [/opt/SRM/srm-rib/webapps/srm-rib]
Jun 30 16:34:32 vt046651-cabi srm-rib.sh[13695]: 2023-06-30 16:34:32.440:INFO :oejshC.srm_rib:main:
Initializing Spring root WebApplicationContext
Jun 30 16:34:34 vt046651-cabi srm-rib.sh[13695]: . 2023-06-30
16:34:34.894:INFO :oejsh.ContextHandler:main: Started o.e.j.w.WebAppContext@3c9c0d96{srm-rib/,file:///
opt/SRM/srm-rib/webapps/srm-rib/,AVAILA...ebapps/srm-rib}
Jun 30 16:34:34 vt046651-cabi srm-rib.sh[13695]: 2023-06-30 16:34:34.947:INFO :oejs.RequestLogWriter:main:
Opened /opt/SRM/srm-rib/logs/2023_06_30.request.log
Jun 30 16:34:34 vt046651-cabi srm-rib.sh[13695]: 2023-06-30
16:34:34.956:INFO :oejs.AbstractConnector:main: Started ServerConnector@59e2d8e3{HTTP/1.1, (http/1.1)}
{0.0.0.0:8081}
Jun 30 16:34:34 vt046651-cabi srm-rib.sh[13695]: 2023-06-30 16:34:34.974:INFO :oejs.Server:main: Started
Server@1ebea008{STARTING}[10.0.13,sto=5000] @5323ms
Jun 30 16:34:37 vt046651-cabi srm-rib.sh[13695]: OK Fri Jun 30 16:34:37 UTC 2023
Jun 30 16:34:37 vt046651-cabi systemd[1]: Started NetOps Report Manager Service.
Hint: Some lines were ellipsized, use -l to show in full.

```

## 2. Access the NetOps Report Manager Service by entering the following URL in a web browser:

`http://<SRM_RIB_HOST>:<SRM_RIB_SERVER_PORT>/rib/service/status`

### Example:

`http://srm-rib.domain.com:8099/rib/service/status`

#### – Host Name

Defines the host name of the system where the NetOps Report Manager Service has been installed.

**Example:** srm-rib.domain.com

#### – Port

Defines the port that is configured for the NetOps Report Manager Service web server.

**Example:** 8099

The expected output is:

HTTP 401 Error (Unauthorized)

## Set up to Run NetOps Business Reports

For more information, see [Set up to Run NetOps Business Reports](#).

## Next Step

After you have enabled the NetOps business reports, you can:

- [Enable HTTPS for the NetOps Report Manager Service](#)
- [Manage and Run NetOps Business Reports](#)

## Verify the Prerequisites for Installing the NetOps Report Manager Service

Before installing the NetOps Report Manager Service, verify the prerequisites.



**(23.3.2 and higher)****IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

Before you install the NetOps Report Manager Service, ensure that you have completed the following prerequisite steps:

- [You have reviewed the system architecture for NetOps Business Reports.](#)
- You have ensured that the following ports are accessible:

Function	
DX NetOps Spectrum (Spectrum)	HTTP 8080/HTTPS 8443 Enables communication between the NetOps Report Manager Service and Spectrum.
The NetOps Report Manager Service	HTTP 8081/HTTPS Enables report details and data transfer between NetOps Portal and the NetOps Report Manager Service.
Spectrum's MySQL "reporting" database (SRMDB)	TCP 3306 Enables communication to SRMDB (inbound) from the NetOps Report Manager Service host.

- You have verified that the NetOps Report Manager Service host meets the following system requirements:
  - **Memory:** 8 GB
  - **Disk:** 500 MB
  - **Processing Cores:** Four or more
  - **Kernel:** 4.18.0.80+

**NOTE**

By default, the NetOps Report Manager Service installation file (the installer) sets one fourth (1/4) of total available system memory as the maximum Java heap size for the NetOps Report Manager Service process. For example, if the system where the NetOps Report Manager Service is installed has 8 GB of total system memory, then the NetOps Report Manager Service will use up to 2 GB.

- You have installed and integrated the following components:
  - DX NetOps Performance Management  
For more information, see [Installing](#).
  - DX NetOps Spectrum (Spectrum)  
To ensure that SRMDB is accessible for reporting purposes, the Spectrum OneClick instance that is integrated with NetOps Portal must have Spectrum Report Manager (during the OneClick installation process, you have installed OneClick with Spectrum Report Manager by specifying Spectrum Report Manager as a feature selection).  
For more information:
    - About how to install OneClick with Spectrum Report Manager, see [Install OneClick with Report Manager](#).
    - About how to integrate Spectrum with DX NetOps Performance Management, see [Integration with DX NetOps Performance Management](#) and [Integrate with DX NetOps Spectrum](#).
- (23.3.6 and higher) You have installed Java Runtime Environment 11 (JRE) 11 on the host where the Report Manager Service will be installed.
- You have a dedicated user for the NetOps Report Manager Service-to-Spectrum integration (SRMDB user), with permissions to accept connections from the NetOps Report Manager Service host. By default, the NetOps Report Manager Service uses `SRM_user` as the username.  
You can use one of the following users as the SRMDB user:

- The existing SRMDB user in MySQL. (If you plan to install the NetOps Report Manager Service on a system *separate* from SRMDB) [Grant this user the required permissions.](#)
- [An SRMDB user that you create.](#)
- [You have reviewed the common considerations for the NetOps Report Manager Service.](#)
- (If the user that will own the NetOps Report Manager Service installation and that can log in and manage the NetOps Report Manager Service and files (the installation owner user account) does not have root access to install the NetOps Report Manager Service) [You have configured the Sudo User Account.](#)
- [You have set up to install the NetOps Report Manager Service.](#)

## Create an SRMDB User

If you chose the option to create an SRMDB user, use the following procedure to create this user and grant this user permissions to SRMDB schema and to the Spectrum Report Manager Database API (SRMDBAPI ) database schemas.

### Follow these steps:

1. Connect to the SRMDB as the database administrator.
2. Create the SRMDB user and grant this user the required permissions by issuing the following commands, using `SRM_user` as the username:

```
CREATE USER 'SRM_user'@'<RMS_hostname>' IDENTIFIED BY '<password>';
GRANT SELECT,EXECUTE ON reporting.* TO 'SRM_user'@'<RMS_hostname>' ;
GRANT SELECT,EXECUTE ON srmdbapi.* TO 'SRM_user'@'<RMS_hostname>' ;
FLUSH PRIVILEGES;
```

– **password**

The password for this user.

– **RMS\_hostname**

The NetOps Report Manager Service host name.

**NOTE**

The NetOps Report Manager Service installation uses `SRM_user` as a default value during the installation process.

### Examples:

If you plan to install the NetOps Report Manager Service on a system *separate* from SRMDB, issue the following commands:

```
CREATE USER 'SRM_user'@'srm-rib-service.domain.com' IDENTIFIED BY '<password>';
GRANT SELECT,EXECUTE ON reporting.* TO 'SRM_user'@'srm-rib-service.domain.com';
GRANT SELECT,EXECUTE ON srmdbapi.* TO 'SRM_user'@'srm-rib-service.domain.com';
FLUSH PRIVILEGES;
```

If you plan to install the NetOps Report Manager Service on the *same* system as SRMDB, issue the following commands:

```
CREATE USER 'SRM_user'@'localhost' IDENTIFIED BY '<password>';
GRANT SELECT,EXECUTE ON reporting.* TO 'SRM_user'@'localhost';
GRANT SELECT,EXECUTE ON srmdbapi.* TO 'SRM_user'@'localhost';
FLUSH PRIVILEGES;
```

An SRMDB user is created and is granted the required permissions.

3. Validate this user by connecting to SRMDB using this user's credentials, for example:
 

```
mysql -uSRM_user -p<password>
mysql> SHOW GRANTS;
```
4. (Next step) (If you plan to install the NetOps Report Manager Service on a system *separate* from SRMDB) [Grant this user permissions to accept connections from the Report Manager Service host.](#)

**NOTE**

If you plan to install the NetOps Report Manager Service on the *same* system as SRMDB, the connections from this system do not require additional permissions, and therefore, this user already has the required permissions.

This user is now the SRMDB user, and has the required permissions.

**Grant the Existing SRMDB User in MySQL the Required Permissions**

If you chose the option to use the existing SRMDB user in MySQL, use the following procedure to grant this user the required permissions.

**Follow these steps:**

1. Ensure that this user has permissions to SRMDB schema and to the Spectrum Report Manager Database API (SRMDBAPI ) database schemas.  
For more information, contact the Spectrum Administrator.
2. (Next step) (If you plan to install the NetOps Report Manager Service on a system *separate* from SRMDB) [Grant this user permissions to accept connections from the Report Manager Service host.](#)

**NOTE**

If you plan to install the NetOps Report Manager Service on the *same* system as SRMDB, the connections from this system do not require additional permissions, and therefore, this user already has the required permissions.

The existing SRMDB user in MySQL is now the SRMDB user, and has the required permissions.

**Grant the SRMDB User Permissions to Accept Connections from the NetOps Report Manager Service Host**

(If you plan to install the NetOps Report Manager Service on a system *separate* from SRMDB) Grant this user permissions to accept connections from the NetOps Report Manager Service host.

For more information about SRMDB user management, contact a Spectrum Administrator.

**Common Considerations for the NetOps Report Manager Service**

Review the common considerations for the NetOps Report Manager Service:

- You can use the NetOps Report Manager Service in parallel with CA Business Intelligence (CABI).  
For more information about the options for running reports, see [Run Business Reports](#).

**NOTE**

The DX NetOps Performance Management integration with CA Business Intelligence (CABI) is End of Service as of May 18th, 2022.

For more information, see [the 22.2.1 release notes](#).

- Install the NetOps Report Manager Service on one of the following systems:
  - (Preferred) On a *separate* system from the other installed DX NetOps components.  
If you install using this option, [ensure that this system meets the operating system \(OS\) requirements](#).
  - On the *same* system as another installed DX NetOps component, such as the CABI instance.

**IMPORTANT**

Use this option with caution. Consider the performance impact of the already-installed DX NetOps component.

**Configure the Sudo User Account**

If the user that will own the NetOps Report Manager Service installation and that can log in and manage the NetOps Report Manager Service and files (the installation owner user account) does not have root access to install the NetOps

Report Manager Service, to upgrade the NetOps Report Manager Service, and to run the NetOps Report Manager Service, configure the sudo user account.

**Follow these steps:**

1. On the NetOps Report Manager Service host, locate and edit the `/etc/sudoers` file.
2. Add a line to configure an alias and list of folders and applications to be allowed to run using `sudo` command based on your version:

**(23.3.6 and higher)**

```
Cmnd_Alias <alias> = <installer_tmp_directory>/ReportManagerServiceSetup.txe,<installation_directory>/srm-rib,<installation_directory>/srm-rib/uninstall-netops-report-manager-service,<java_home_directory>/bin/keytool,/sbin/service,/bin/systemctl,/bin/vim,/bin/chmod,/bin/netstat
```

**(23.3.5 and lower)**

```
Cmnd_Alias <alias> = <installer_tmp_directory>/ReportManagerServiceSetup.bin,<installation_directory>/srm-rib,<installation_directory>/srm-rib/Uninstall_srm-rib,<installation_directory>/jre/bin/keytool,/sbin/service,/bin/systemctl,/usr/bin/vim,/bin/chmod
```

- **alias**  
An alias for the list of directories and applications that can run using the sudo command.
- **installer\_tmp\_directory**  
The directory where you placed the installer.
- **installation\_directory**  
The NetOps Report Manager Service installation directory.  
**Default:** `/opt/SRM`
- **java\_home\_directory**  
The Java home directory directory.  
**Default:** `/opt/SRM`

3. Add the following line with the sudo user name:

```
<sudo_username> ALL = <alias>
```

- **sudo\_username**  
The user who can run the sudo commands.  
**Example:** `sudouser`
- **alias**  
The alias for the list of directories and applications that can run using the sudo command. This must be the same alias name used in the `Cmnd_Alias` line.  
**Example:**

**(23.3.6 and higher)**

```
Cmnd_Alias SRMRIB = /tmp/ReportManagerServiceSetup.txe,/opt/SRM/srm-rib,/opt/SRM/srm-rib/uninstall-netops-report-manager-service,/usr/lib/jvm/java-11-openjdk-11.0.20.0.8-1.el7_9.x86_64/bin/keytool,/sbin/service,/bin/systemctl,/bin/vim,/bin/chmod,/bin/netstat
```

```
sudouser ALL = SRMRIB
```

**(23.3.5 and lower)**

```
Cmnd_Alias SRMRIB = /tmp/ReportManagerServiceSetup.bin,/opt/SRM/srm-rib,/opt/SRM/srm-rib/Uninstall_srm-rib,/opt/SRM/jre/bin/keytool,/sbin/service,/bin/systemctl,/usr/bin/vim,/bin/chmod
```

```
sudouser ALL = SRMRIB
```

4. Save your changes.

The sudo user account is configured.

## Set up to Install the NetOps Report Manager Service

### Follow these steps:

1. Download the `netops-srm-rib-service-<version>-Linux.tar.gz` installer package from the DX NetOps product page on [support.broadcom.com](https://support.broadcom.com), and then upload it onto the NetOps Report Manager Service host into a separate NetOps Report Manager Service installer location directory, for example, `/tmp`.

- **version**

The version of the NetOps Report Manager Service.

**Examples:** (for 23.3.5) 23.3.5.1 (for 23.3.2) 23.3.2.131

**NOTE**

If you do not have the NetOps Report Manager Service installer package, open a Support ticket.

2. Log into the NetOps Report Manager Service host as the root or non-root user.
3. Navigate to the directory where you placed the installer (for example, `/tmp`), and then unpack the installer package by issuing the following commands:

```
cd /tmp
tar -xzf netops-srm-rib-service-<version>-Linux.tar.gz
```

**Example:**

```
cd /tmp
tar -xzf netops-srm-rib-service-23.3.2.131-Linux.tar.gz
```

- **version**

The version of the NetOps Report Manager Service.

**Examples:** (for 23.3.5) 23.3.5.1 (for 23.3.2) 23.3.2.131

4. Change the permissions for the `ReportManagerServiceSetup.bin` installation file by issuing the following command based on your version:

- **(23.3.6 and higher)**

```
chmod +x ReportManagerServiceSetup.txe
```

- **(23.3.5 and lower)**

```
chmod +x ReportManagerServiceSetup.bin
```

The NetOps Report Manager Service is set up for installation.

**Next step:** [Install the NetOps Report Manager Service](#).

## Install the NetOps Report Manager Service from the Command Line

Install the NetOps Report Manager Service answering questions for the installation at the prompts.

**(23.3.2 and higher)**

**IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

Installing from the command line is also referred to as installing in interactive mode. With this method, you install the NetOps Report Manager Service using the NetOps Report Manager Service installation file (the installer).

**Prerequisite:** Before you install the NetOps Report Manager Service, [ensure that you have verified the prerequisites](#).

### Follow these steps:

1. Run the installer by issuing the following command (23.3.5 and lower) with the `-i` option, based on your version and your user account:

– (23.3.6 and higher)

- **For root user:**

```
./ReportManagerServiceSetup.txe
```

- **For non-root user:**

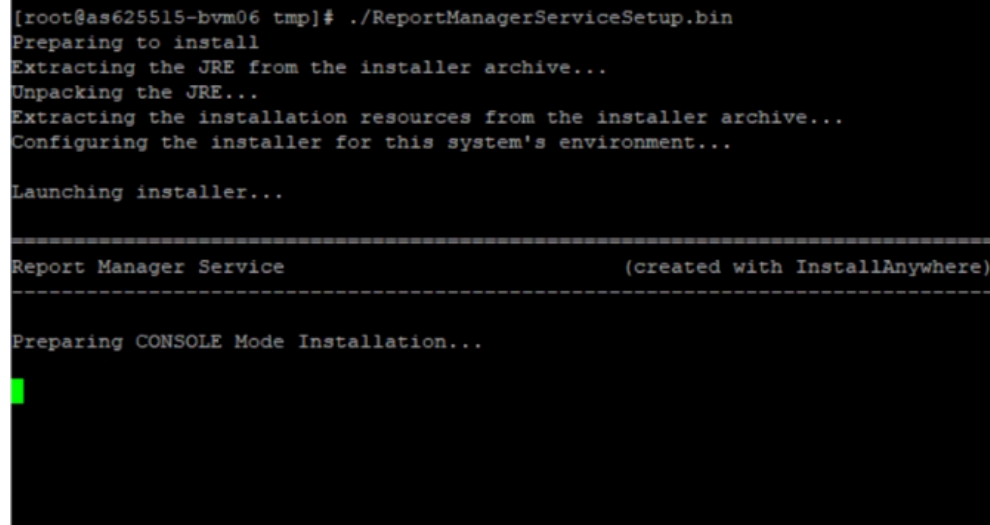
```
sudo ./ReportManagerServiceSetup.txe
```

– (23.3.5 and lower)

- **For root user:**

```
./ReportManagerServiceSetup.bin -i console
```

The following image show an example of issuing this command for root user:



```
[root@as625515-bvm06 tmp]# ./ReportManagerServiceSetup.bin
Preparing to install
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...

Launching installer...

=====
Report Manager Service                      (created with InstallAnywhere)
=====

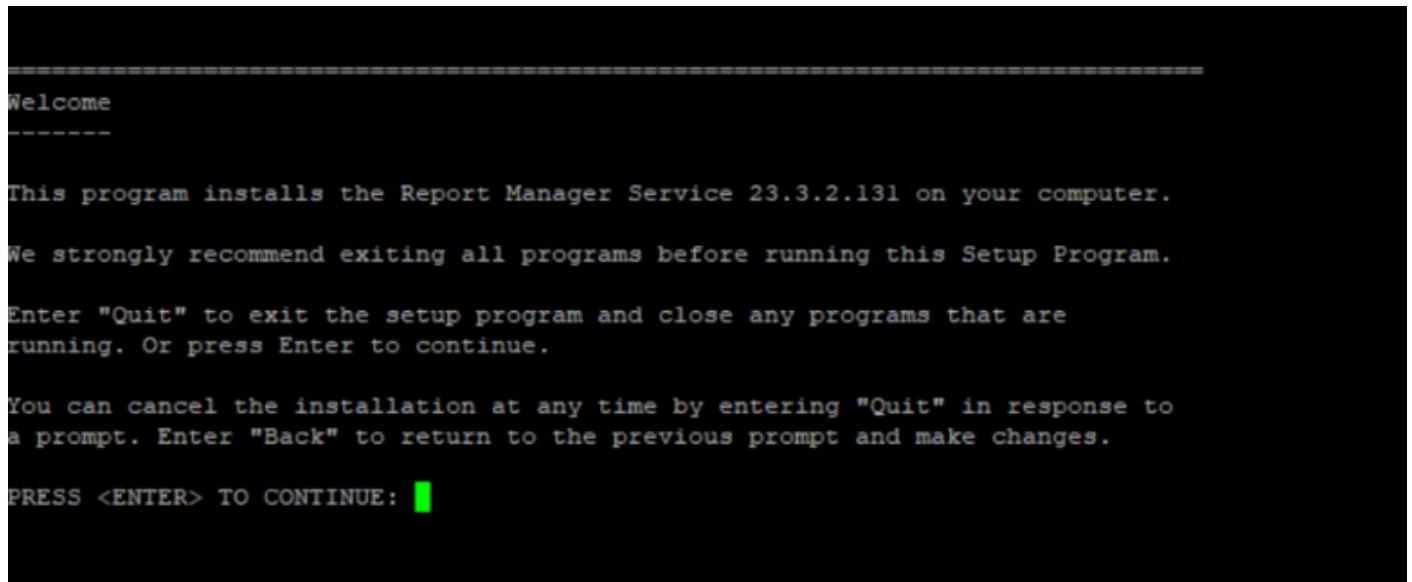
Preparing CONSOLE Mode Installation...
█
```

- **For non-root user:**

```
sudo ./ReportManagerServiceSetup.bin -i console
```

The **Welcome** message appears.

(23.3.5 and lower) The following image shows an example of this message:



```
=====
Welcome
-----

This program installs the Report Manager Service 23.3.2.131 on your computer.

We strongly recommend exiting all programs before running this Setup Program.

Enter "Quit" to exit the setup program and close any programs that are
running. Or press Enter to continue.

You can cancel the installation at any time by entering "Quit" in response to
a prompt. Enter "Back" to return to the previous prompt and make changes.

PRESS <ENTER> TO CONTINUE: █
```

2. Press the **Enter/Return** key on your keyboard to continue.
3. When prompted with the license agreement, press the **Enter/Return** key on your keyboard to continue.

4. When prompted to accept the terms of the license agreement, to proceed further, enter **Y** for Yes. The installer performs a disk check validation. If the disk check finds that the host where the Report Manager Service will be installed has met the following system requirements, the installation continues:
- 8 GB or more of total system memory is available
  - Four or more processing cores are available
  - 4.18.0.80+ Kernel is available

For more information, see [Verify the Prerequisites for Installing the NetOps Report Manager Service](#).

#### (23.3.6 and higher)

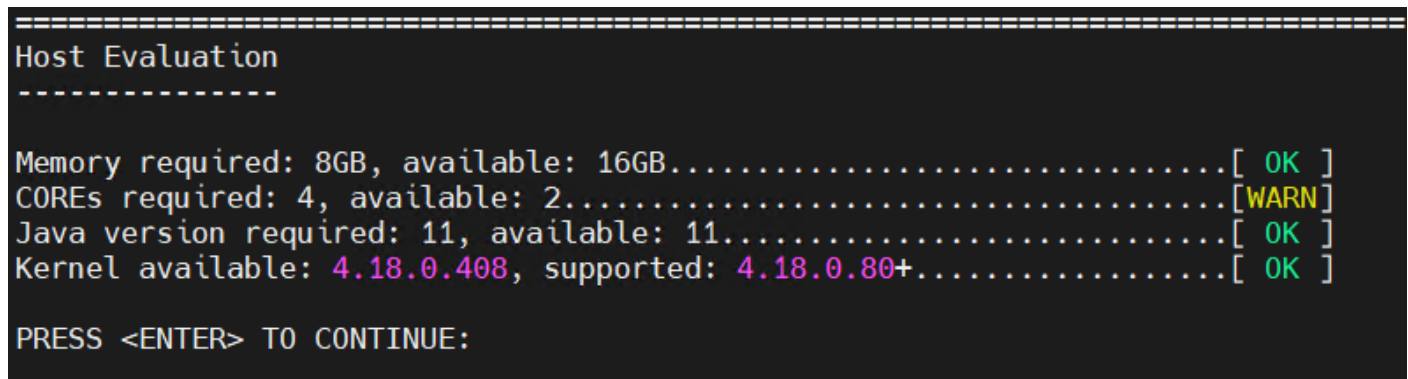
The evaluation displays below the **Host Evaluation** line.

Host Evaluation

-----

```
Memory required: 8 GB, available: 16GB.....[ OK ]
COREs required: 4, available: 2.....[ WARN ]
Java version required: 11, available: 11.....[ OK ]
Kernel available: 4.18.0.408, supported: 4.18.0.80+.....[ OK ]
```

The following image shows an example of the evaluation:



```
=====
Host Evaluation
-----

Memory required: 8GB, available: 16GB.....[ OK ]
COREs required: 4, available: 2.....[ WARN ]
Java version required: 11, available: 11.....[ OK ]
Kernel available: 4.18.0.408, supported: 4.18.0.80+.....[ OK ]

PRESS <ENTER> TO CONTINUE:
```

#### (23.3.5 and lower)

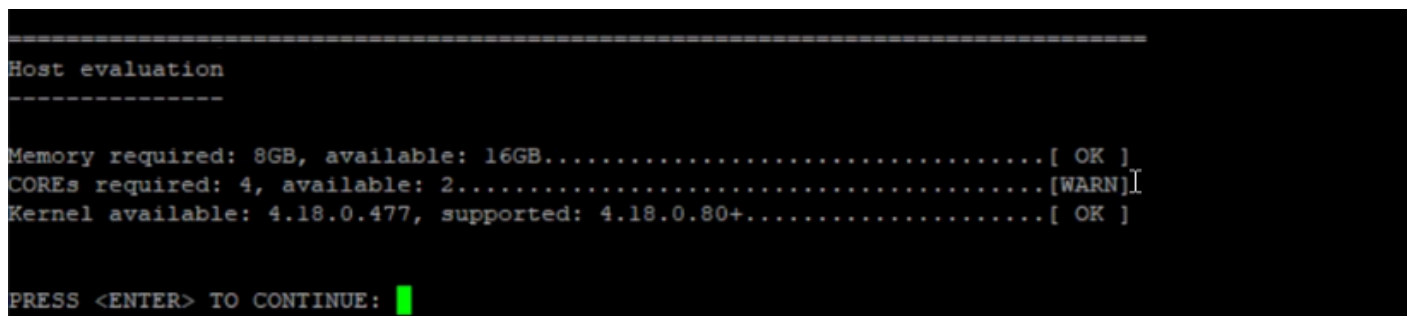
The evaluation displays below the **Host evaluation** line.

Host evaluation

-----

```
Memory required: 8 GB, available: 16GB.....[ OK ]
COREs required: 4, available: 2.....[ WARN ]
Kernel available: 4.18.0.477, supported: 4.18.0.80+.....[ OK ]
```

The following image shows an example of the evaluation:



```
=====
Host evaluation
-----

Memory required: 8GB, available: 16GB.....[ OK ]
COREs required: 4, available: 2.....[ WARN ]
Kernel available: 4.18.0.477, supported: 4.18.0.80+.....[ OK ]

PRESS <ENTER> TO CONTINUE: █
```

5. Press the **Enter/Return** key on your keyboard to continue.



(If you are installing as root user) The **Choose Installation Owner** line appears.  
 (23.3.6 and higher) The following image shows an example of this line, and the prompt that appears:

```
=====
Choose Installation Owner
-----

Please enter the name of the user you would like to run the Report Manager
Service as.

Username (Default: root):
```

6. At the **Username (Default: root):** prompt, accept the default (root) by pressing the **Enter/Return** key on your keyboard, or specify a non-root user that will own the NetOps Report Manager Service installation and that can log in and manage the NetOps Report Manager Service and files.

**IMPORTANT**

If you specify a non-root user, [update the sudo permissions for this user](#).

**Default:** root

The **Choose Installation Directory** line appears.

(23.3.6 and higher) The following image shows an example of this line, and the prompt that appears:

```
=====
Choose Installation Directory
-----

Installation Directory Path (Default: /opt/SRM):
```

7. At the **Installation Directory Path (Default: /opt/SRM):** prompt, accept the default (/opt/SRM ) by pressing the **Enter/Return** key on your keyboard, or specify another directory where the NetOps Report Manager Service is installed.

**Default:** /opt/SRM

The **Enter Report Manager Service Port** line appears.

(23.3.6 and higher) The following image shows an example of this line, and the prompt that appears:

```
=====
Enter Report Manager Service Port
-----

Port Number (Default: 8081): 9090
```

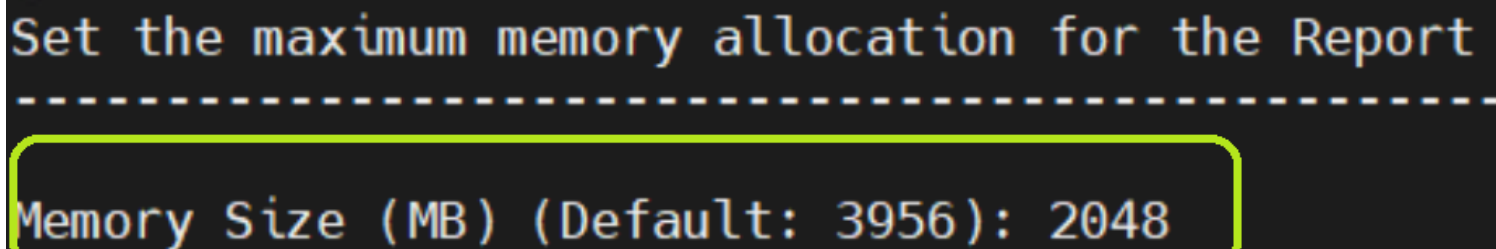
8. At the **Port Number (Default: 8081):** prompt, accept the default (8081 ) by pressing the **Enter/Return** key on your keyboard, or specify another HTTP connection port for the NetOps Report Manager Service.  
 The **Enter SRM Database Hostname** line appears.
9. At the **Hostname (Default: localhost):** prompt, accept the default (localhost ) by pressing the **Enter/Return** key on your keyboard, or specify another hostname for Spectrum's MySQL "reporting" database (SRMDB).

**Default:** localhost

The **Enter SRM Database Port** line appears.



10. At the **Port Number (Default: 3306)**: prompt, accept the default (3306 ) by pressing the **Enter/Return** key on your keyboard, or specify another port for SRMDB.  
**Default:** 3306  
The **Enter SRM Database User** line appears.
11. At the **User (Default: SRM\_user)**: prompt, accept the default (SRM\_user ) by pressing the **Enter/Return** key on your keyboard, or specify another username as the SRMDB user.  
**Default:** SRM\_user  
The **Enter SRM Database Password** line appears.
12. At the **Please Enter the Password**: prompt, enter the password for SRMDB, and then press the **Enter/Return** key on your keyboard.  
The **Set the maximum memory allocation for the NetOps Report Manager Service** line appears.  
The following image shows an example of this line, and the prompt that appears:



13. At the **Memory Size (MB) (Default: <value>)**: prompt, specify the maximum memory allocation for the NetOps Report Manager Service.  
**Default:** 1/4 of RAM or 2048 MB if the division result is less than 2048 MB

The installation begins. After the NetOps Report Manager Service has installed successfully, a message similar to the following appears:

```
Installation Complete
```

```
-----
```

```
Congratulations. The NetOps Report Manager Service has been successfully installed to:
```

```
/opt/SRM
```

```
PRESS <ENTER> TO EXIT THE INSTALLER:
```

**Next step:** [Verify the installation.](#)

## Install the NetOps Report Manager Service in Silent Mode

Install the NetOps Report Manager Service avoiding the prompts using the values that you define and provide in the response file.

(23.3.2 and higher)

### IMPORTANT

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make

Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

In a silent installation, the NetOps Report Manager Service installation file (the installer) uses the property values that you define in the `silent.properties` response file. You define these values once, and then use the response file on multiple machines as required.

Complete the following process to install the NetOps Report Manager Service in silent mode:

1. [Define the Response File](#)
2. [Check the Response File](#)
3. [Install in Silent Mode](#)
4. [Verify the Installation](#)

### **Define the Response File**

Define and provide the following property values in the `silent.properties` response file.

#### **Follow these steps:**

1. Run the installer by issuing the following command with the `-r` option to generate the response file with properties, based on your version and your user account:

- **(23.3.6 and higher)**

- **For root user:**

```
./ReportManagerServiceSetup.txe -r silent.properties
```

- **For non-root user:**

```
sudo ./ReportManagerServiceSetup.txe -r silent.properties
```

- **(23.3.5 and lower)**

- **For root user:**

```
./ReportManagerServiceSetup.bin -r silent.properties
```

- **For non-root user:**

```
sudo ./ReportManagerServiceSetup.bin -r silent.properties
```

The NetOps Report Manager Service installer runs in silent mode.

2. Follow the instructions in the console. Enter information for the following prompts:

- **Choose Installation Owner**

Specifies whether the user that will own the NetOps Report Manager Service installation and can log in and manage the NetOps Report Manager Service and files (the installation owner user account) has root access to install the NetOps Report Manager Service and to run the NetOps Report Manager Service or is a non-root user.

**Options:**

- **root:** The user has root access.
- **non-root:** The user is a non-root user.

**NOTE**

**Prerequisite:** [The sudo user account is configured.](#)

**Default:** `root`

- **Choose Installation Directory**

Specifies the directory path where the NetOps Report Manager Service is installed.

**Default:** `/opt/SRM`

- **Enter Report Manager Service Port**

Specifies the HTTP connection port for the NetOps Report Manager Service.

**Default:** `8081`

- **Enter SRM Database Hostname**

Specifies the hostname for Spectrum's MySQL "reporting" database (SRMDB).

**Default:** localhost

– **Enter SRM Database Port**

Defines the port for SRMDB.

**Default:** 3306

– **Enter SRM Database User**

Defines the username for the SRMDB user.

**Default:** SRM\_user

– **Enter SRM Database Password**

Specifies the password for SRMDB.

– **Skip Database Connection Check?** (23.3.6 and higher)

Defines whether the installer skips the database connection test.

- **N:** The installer performs the database connection test.
- **Y:** The installer skips the database connection test.

**Default:** N

– **Set the maximum memory allocation for the NetOps Report Manager Service**

Defines the maximum memory allocation for the NetOps Report Manager Service.

**Default:** 1/4 of RAM or 2048 MB if the division result is less than 2048 MB

The `silent.properties` response file is generated with properties.

### Check the Response File

Confirm the values for the following variables in the `silent.properties` response file:

• **USER\_INPUT\_INSTALL\_OWNER (Default: root)**

Specifies whether the user that will own the NetOps Report Manager Service installation and can log in and manage the NetOps Report Manager Service and files (the installation owner user account) has root access to install the NetOps Report Manager Service and to run the NetOps Report Manager Service or is a non-root user.

**Options:**

- **root:** The user has root access.
- **non-root:** The user is a non-root user.

**NOTE**

**Prerequisite:** [The sudo user account is configured.](#)

**Default:** root

• **USER\_INPUT\_INSTALL\_DIR (Default: /opt/SRM)**

Specifies the directory path where the NetOps Report Manager Service is installed.

**Default:** /opt/SRM

• **USER\_INPUT\_SERVER\_PORT (Default: 8081)**

Specifies the HTTP connection port for the NetOps Report Manager Service.

**Default:** 8081

• **USER\_INPUT\_DB\_HOST (Default: localhost)**

Specifies the hostname for Spectrum's MySQL "reporting" database (SRMDB).

**Default:** localhost

• **USER\_INPUT\_DB\_PORT (Default: 3306)**

Specifies the port for SRMDB.

**Default:** 3306

• **USER\_INPUT\_DB\_USER (Default: SRM\_user)**

Specifies the username for the SRMDB user.

**Default:** `SRM_user`

- **USER\_INPUT\_DB\_PASSWORD**

Specifies the password for SRMDB.

- **IGNORE\_DB\_CHECK** (23.3.6 and higher)

Defines whether the installer skips the database connection test.

- **N:** The installer performs the database connection test.
- **Y:** The installer skips the database connection test.

**Default:** `N`

- **MEM\_ALLOC\_CONSOLE\_INPUT**

Specifies the maximum memory allocation for the NetOps Report Manager Service.

**Default:** 1/4 of RAM or 2048 MB if the division result is less than 2048 MB

## **Install in Silent Mode**

**Prerequisite:** You have defined and provided property values in the response file.

Run the installer by issuing the following command (23.3.5 and lower) with the `-i` option and (all versions) the `-f` option with the response file as an argument, based on your version and your user account:

- **(23.3.6 and higher)**

- **For root user:**

```
./ReportManagerServiceSetup.txe -f silent.properties
```

- **For non-root user:**

```
sudo ./ReportManagerServiceSetup.txe -f silent.properties
```

- **(23.3.5 and lower)**

- **For root user:**

```
./ReportManagerServiceSetup.bin -i silent -f silent.properties
```

- **For non-root user:**

```
sudo ./ReportManagerServiceSetup.bin -i silent -f silent.properties
```

The installation begins. It uses the property values from the response file. An empty prompt in the command line indicates that the `srm-rib` Linux daemon (the NetOps Report Manager Service) is installed. The installer creates and starts the NetOps Report Manager Service after the installation completes.

## **Verify the Installation**

After you have installed the NetOps Report Manager Service, [verify the installation](#).

## **Set up to Run NetOps Business Reports**

Before you can manage and run NetOps business reports, set up to run them.

(23.3.2 and higher)

### **IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

Use the following process to set up to run NetOps business reports:

1. [Allow Users to Run NetOps Business Reports](#)

## 2. [Enable and Configure the NetOps Report Manager Service](#)

### **Allow Users to Run NetOps Business Reports**

Grant users who will run NetOps business reports the Run NetOps Business Reports role right. By default, the Administrator user account role includes this role right. You can assign this user account role to the user who will run NetOps business reports, or you can select this role right for a custom user account role, and then assign this custom user account role to the user who will run NetOps business reports.

For more information:

- About this role right, see [Role Rights](#).
- About how to assign user account roles to a user account, see [Manage User Accounts](#).
- About how to select role rights for a custom user account role, see [Manage User Account Roles](#).

### **Enable and Configure the NetOps Report Manager Service**

#### **Follow these steps:**

1. Log into NetOps Portal with a user account that has the Administer Data Sources role right.
2. Hover over **Administration**, **Data Sources**, and then click **Data Sources**.  
The **Manage Data Sources** page appears.
3. From the list, select a DX NetOps Spectrum (Spectrum) data source that points to the Spectrum OneClick instance with Report Manager (SRM) enabled, and then click **Edit**.  
The **Edit Data Source** dialog opens. The following image shows an example of this dialog:

Figure 10: Edit Data Source Dialog

## Edit Data Source



Source Type

Spectrum Infrastructure Mar

Status

Enabled

— Data Source —

Host Name \*

Port \*

8080

Protocol

☒ http☐ https

Display Name \*

Spectrum Infrastructure Manager@☒ Contribute inventory to Data Aggregator☒ Synchronize device life cycle state from SpectrumWeb Console ☒ Same as data source☒ Enable Report Manager Service

— Report Manager Service —

Host Name \*

Port \*

8080

Protocol

☒ http☐ https

— Authentication —

Username \*

spectrum

Password \*

\*\*\*\*\*

Confirm Password \*

\*\*\*\*\*

4. Select the **Enable Report Manager Service** checkbox:

**NOTE**

You can enable and configure the NetOps Report Manager Service for only *one* registered Spectrum data source in NetOps Portal. If you add, or register, another Spectrum data source in NetOps Portal and the NetOps Report Manager Service is enabled for an existing registered Spectrum data source, this checkbox is cleared and you cannot select it.

By default, this checkbox is cleared.

The **Report Manager Service** section expands.

5. In the **Report Manager Service** section, complete the following fields:

- **Host Name**

Specifies the host name of the system where the NetOps Report Manager Service has been installed.

**Example:** `srm-rib.domain.com`

- **Port**

Specifies the port for communication between the NetOps Report Manager Service and the Spectrum OneClick server for authentication purposes.

**Default:** 8080

- **Protocol**

Specifies the protocol for the NetOps Report Manager Service web server.

**Options:** `http` or `https`

**Default:** `http`

6. (Optional) Click **Test** to validate the connectivity between NetOps Portal and the NetOps Report Manager Service. If the data that you entered is correct, the `Data Source Test Successful` message appears, otherwise, check the information for correctness.

7. Click **Save**.

The NetOps Report Manager Service is enabled and configured in NetOps Portal.

### **Next Step**

After you have set up to run NetOps business reports, you can:

- [Enable HTTPS for the NetOps Report Manager Service](#)
- [Manage and Run NetOps Business Reports](#)

## **Install a Low-Scale System**

For smaller deployments of 150,000 polled items or less, you can run the data aggregator and the data collector on a *shared* single node (a low-scale system install). Install a low-scale system by deploying DX NetOps Performance Management with the data aggregator and the data collector on a single node and installed to run as the same user.

For information about the sizing requirements, see the [DX NetOps Sizing Tool](#).

Use the following process to deploy DX NetOps Performance Management with the data aggregator and data collector on a single node:

1. [Review the installation requirements and considerations.](#)
2. [Install NetOps Portal.](#)
3. [Install the data repository.](#)

**NOTE**

In low-scale systems, the data repository is deployed on a single node. Therefore, skip the steps that refer to cluster installation.

4. [Install the data aggregator.](#)

5. [Install and configure the data collector on the same host and owned/managed by the same user as the data aggregator.](#)
6. [Complete the post-installation configuration.](#)

### **Limitations**

The following options are not supported in a low-scale system:

- Calculated metrics beyond 150,000 metrics per second
- Fast (1-minute) polling
- Integration with DX NetOps Mediation Manager or DX NetOps Virtual Network Assurance
- Multiple data collectors
- Multiple IP domains
- Multiple tenants
- Polling beyond 150,000 items
- The data aggregator and data collector installed to run as different users.

## **Install the Data Collector on the Same Host as the Data Aggregator**

You can install the data collector on the same host and owned/managed by the same user as the data aggregator.

Use the following process to install the data collector on the same host in a low-scale system:

1. [Install the Data Collector](#)
2. [Configure the Data Collector](#)
3. [Validate the Deployment](#)

### **Install the Data Collector**

**Follow these steps:**

1. Log in to the shared host.
2. Verify that the Data Aggregator service is running by issuing the following commands:
 

```
systemctl status dadaemon
```
3. If the service is not running, do *one* of the following steps:
  - Start the Data Aggregator service by issuing the following command:
 

```
systemctl start dadaemon
```
  - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:
 

```
<installation_directory>/IMDataAggregator/scripts/dadaemon activate
```

    - ***installation\_directory***  
The installation directory for the data aggregator.  
**Default:** /opt
4. Verify that the Apache ActiveMQ service is running by issuing the following command:
 

```
systemctl status activemq
```

If the service is not running, start the ActiveMQ service by issuing the following command:

```
systemctl start activemq
```
5. Download the data collector installer by issuing the following command:
 

```
wget -nv http://<da_host>:8581/dcm/InstData/Linux/VM/install.bin
```
6. Make the `install.bin` file executable by issuing the following command:
 

```
chmod a+x install.bin
```



7. Do one of the following steps:
  - Stop the data aggregator and ActiveMQ services by issuing the following commands:
 

```
systemctl stop dadaemon
systemctl stop activemq
```
  - (Fault-tolerant environment) If the local data aggregator is running, shut it down and prevent it from restarting until maintenance is complete by issuing one the following commands:
 

```
<installation_directory>/IMDataAggregator/scripts/dadaemon maintenance
```

    - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** /opt
8. Start the data collector installation by issuing the following command:
 

```
./install.bin -i console
```
9. Follow the instructions in the console.  
The installer prompts for the hostname or IP address of the data aggregator host. Supply the hostname of the shared host. Specify the same user as when you installed the data aggregator.  
The data collector is installed and the data collector and ActiveMQ services start. After a short time, the data collector service fails because the data aggregator service is not running.
10. Uninstall the ActiveMQ service on the data aggregator by issuing the following command:
 

```
<installation_directory>/IMDataAggregator/scripts/activemq uninstall
```

  - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** /opt
11. Install the ActiveMQ service on the data collector by issuing the following command:
 

```
<DC_installation_directory>/IMDataCollector/scripts/activemq install
```

  - **DC\_installation\_directory**  
The installation directory for the data collector.  
**Default:** /opt

The data collector is installed on a shared host and owned/managed by the same user as the data aggregator.

## **Configure the Data Collector**

Configure the data collector to use the shared host. On a shared host, the data aggregator and data collector use a single ActiveMQ broker.

### **Follow these steps:**

1. Stop the Data Collector and ActiveMQ services by issuing the following commands:
 

```
systemctl stop activemq
systemctl stop dcmd
systemctl stop icmpd
```

The data collector, the ICMP daemon, and ActiveMQ stop.
2. Modify the data aggregator startup script (the `dadaemon` file) by completing the following steps:
  - a. Edit the `<installation_directory>/IMDataAggregator/scripts/dadaemon` file.
    - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** /opt
  - b. Modify the `ACTIVEMQ_HOME` variable to point to the ActiveMQ broker on the data collector by issuing the following command:
 

```
export eval `grep ACTIVEMQ_HOME /opt/DCM.cfg`
```

**TIP**

To find the correct version, go to the `<DC_installation_directory>/broker` directory.

- c. Save your changes to the file.
3. Disable the data aggregator ActiveMQ broker heartbeat:
  - a. Create the `<installation_directory>/IMDataAggregator/apache-karaf/etc/com.ca.im.dm.core.amq.cfg` file as the user running the data aggregator. This ensures that the correct permissions are set on the file.
    - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt`
  - NOTE**  
Create this file as the user running the data aggregator to ensure the correct permissions are set on the new file.
  - b. Insert the following property configuration into the file:
 

```
jmsbroker-heartbeat-disabled=true
```
  - c. Save your changes to the file.
  - d. Disable log messages that are related to the data aggregator heartbeat manager by editing the `<installation_directory>/IMDataAggregator/apache-karaf/etc/org.ops4j.pax.logging.cfg` file.
    - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt`
  - e. Insert the following property configuration into the file:
 

```
# Disable ActiveMQ Health Monitoring log messages in shared DA/data collector mode
log4j2.logger.JmsHeartbeatManager.name=com.ca.im.core.jms.heartbeat.JmsHeartbeatManager
log4j2.logger.JmsHeartbeatManager.level=OFF
```
  - f. Save your changes to the file.  
The data aggregator ActiveMQ broker heartbeat is disabled.
4. Modify the Karaf JMX management properties by completing the following steps:
  - a. Edit the `<DC_installation_directory>/IMDataCollector/apache-karaf/etc/org.apache.karaf.management.cfg` file.
    - **DC\_installation\_directory**  
The installation directory for the data collector.  
**Default:** `/opt`
  - b. Set the following properties:
 

```
rmiRegistryPort = 1199
rmiServerPort = 44445
```
5. (Optional) To support the ability to debug the data aggregator and data collector components simultaneously, modify the Karaf debug port in the `karaf` file:
  - a. Edit the `<DC_installation_directory>/IMDataCollector/apache-karaf/bin/karaf` file.
    - **DC\_installation\_directory**  
The installation directory for the data collector.  
**Default:** `/opt`
  - b. Set the following property:
 

```
DEFAULT_JAVA_DEBUG_PORT=5006
```
  - c. Save your changes to the file.
6. Do *one* of the following steps:
  - Start the Data Aggregator service by issuing the following command:

```
systemctl start dadaemon
```

- (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:

```
<installation_directory>/IMDataAggregator/scripts/dadaemon activate
```

- **installation\_directory**

The installation directory for the data aggregator.

**Default:** /opt

7. Start the Data Collector service by issuing the following command:

```
systemctl start dcmd
```

The Data Aggregator and Data Collector services start on the single host.

The data collector is configured.

## Validate the Deployment

Validate the deployment by verifying that the required services are functioning as expected.

### Follow these steps:

1. Validate that the Data Aggregator service is running by issuing the following command:

```
systemctl status dadaemon
```

2. Validate that the Data Collector service is running by issuing the following command:

```
systemctl status dcmd
```

3. Validate that the ActiveMQ service is running by issuing the following command:

```
systemctl status activemq
```

4. Validate that the ActiveMQ Broker being run is the data collector broker by issuing the following command:

```
systemctl status activemq
```

The returned message must include the following information:

```
Status for ActiveMQ: INFO: Loading '<DC_installation_directory>/IMDataCollector/broker/<apache-activemq-
*>/bin/env'
```

### Example:

```
Status for ActiveMQ: INFO: Loading '/opt/IMDataCollector/broker/apache-activemq-5.18.3/bin/env'
```

- **DC\_installation\_directory**

The installation directory for the data collector.

**Default:** /opt

- **apache-activemq-\***

The installation directory for Apache ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

If the message does not show the data collector installation directory, modify the data aggregator startup script (the `dadaemon` file) by repeating that step in [the "Configure the Data Collector" section](#).

## Uninstall Performance Management

Use the following process to uninstall DX NetOps Performance Management:

1. [Uninstall NetOps Portal.](#)
2. [Uninstall the data aggregator.](#)
3. [Uninstall the data collectors.](#)
4. [Uninstall the data repository.](#)

## Uninstall NetOps Portal

Uninstalling NetOps Portal is one step in the process of uninstalling DX NetOps Performance Management.

You can uninstall NetOps Portal using the command line by running the uninstaller program from a command prompt. Uninstalling NetOps Portal removes NetOps Portal—including the NetOps Portal files, folders, registry entries, and shortcuts—and the unwanted components from the server.

The following directories are preserved:

- `/opt/CA/MySQL/data`  
The MySQL data directory.
- `/opt/CA/PerformanceCenter/InstallLogs`  
This directory contains the uninstallation log.
- `/opt/CA/jre`  
This directory is preserved for other products that use the JRE.

### Follow these steps:

1. Log in to the server as root or as non-root using `sudo`.
2. Navigate to the `<installation_directory>/Uninstall_PerformanceCenter` the uninstaller program.
  - **`installation_directory`**  
The default installation directory for NetOps Portal.  
**Default:** `/opt/CA/PerformanceCenter`
3. Start the uninstaller program by issuing the following command:  
`./Uninstall_PerformanceCenter -i console`
4. Press the **Enter/Return** key on your keyboard to continue.  
The uninstaller shows the progress as it removes the NetOps Portal files, folders, registry entries, and shortcuts.

## Uninstall the Data Aggregator

Use the following process to uninstall the data aggregator:

**Prerequisite:** (If you do not have root access to run the data aggregator) [Configure the sudo user account for the data aggregator](#).

1. Uninstall the data aggregator using *one* of the following options:
  - [Uninstall from the command line](#)
  - [Uninstall in silent mode](#)
  - [Uninstall using the installation wizard](#)
2. (If the data aggregator is configured for fault tolerance) [Complete the next steps](#)

### Uninstall the Data Aggregator from the Command Line

You can uninstall the data aggregator that is installed on a computer in your networking environment from the command line.

### Follow these steps:

1. From a command prompt, issue the following command to log in to the computer where you want to uninstall the data aggregator as the root user:  
`su - root`
2. Access the uninstallation directory by issuing the following command:
  - **`installation_directory`**

The default installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

3. Run the uninstaller by issuing the following command:

```
./Uninstall
```

You are prompted to uninstall.

4. Select option 1 for a complete uninstallation, and then press the **Enter** key on your keyboard.

The data aggregator is uninstalled.

### **Uninstall the Data Aggregator in Silent Mode**

You can uninstall the data aggregator silently. You cannot discover new devices or manage existing polled devices whose IP addresses fall within the IP domain that is associated with the uninstalled data aggregator.

#### **Follow these steps:**

1. From a command prompt, issue the following command to log in to the computer where you want to uninstall the data aggregator as the root user:

```
su - root
```

2. Access the uninstallation directory by issuing the following command:

```
cd <caimda> <installation_directory>/Uninstall
```

#### **– *installation\_directory***

The default installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

3. Run the uninstaller by issuing the following command:

```
./Uninstall
```

The data aggregator is uninstalled.

### **Uninstall the Data Aggregator using the Installation Wizard**

You can uninstall the data aggregator using the installation wizard.

#### **Follow these steps:**

1. From a command prompt, log in to the computer where you want to uninstall the data aggregator as the root user by issuing the following command:

```
su - root
```

2. Access the uninstallation directory by issuing the following command:

```
cd <caimda> <installation_directory>/Uninstall
```

#### **– *installation\_directory***

The default installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

3. Run the uninstaller by issuing the following command:

```
./Uninstall
```

You are prompted to uninstall.

4. Click **Next**.

The **Uninstall Options** dialog opens.

5. Select **Complete Uninstall**, and then click **Next**.

The data aggregator is uninstalled.

## Next Steps

In a fault-tolerant environment, a shared data directory (example: `/DASharedRepo`) is defined to help limit data loss. If the data aggregator is configured for fault tolerance, after uninstalling the data aggregator, remove the items within this shared data directory. These files can include any vendor certifications, metric families, and other files that were custom designed or altered. Removing the items in this shared directory completely uninstalls the fault-tolerant data aggregator pair.

For more information, see [Fault Tolerance](#).

## Uninstall the Data Collectors

Uninstalling the data collectors is part of the process of uninstalling DX NetOps Performance Management.

You can uninstall the data collectors that are installed in your networking environment.

**Prerequisite:** (If you do not have root access to log in or run the data collectors) [Configure the sudo user account for the data collector](#).

### Follow these steps:

1. Log in to the data collector host as the root user or sudo user.
2. Open a command prompt.
3. Access the uninstallation directory by issuing the following command:

```
cd DC_installation_directory/Uninstall
```

4. Run the uninstaller:

```
./Uninstall
```

You are prompted to uninstall.

5. Select option 1 for a complete uninstallation, and then press the **Enter/Return** key on your keyboard.
6. Remove the `/var/.com.zerog.registry.xml` file.

The data collector is uninstalled. Repeat this step for each data collector that you want to uninstall.

## Uninstall the Data Repository

After you uninstall NetOps Portal, the data collector, and the data aggregator components, uninstall the data repository component.

Uninstall the components using the process described in [Uninstall Performance Management](#).

### TIP

Back up your data repository for later use.

### Follow these steps:

1. From an open console, become the Linux user account for the database administrator user by issuing the following command:

```
su - Linux user account for the database administrator user
```

### Example:

```
su - dradmin
```

2. Open the Vertica Administration Tools utility, `adminTools`, from the `/opt/vertica/bin` directory.

### NOTE

If the database is running in a cluster, you can launch the utility from any node in the cluster.

3. Select **(4) Stop Database**, and then select **OK**.

If the data repository does *not* stop, select **(7) Advanced Menu**, then **(2) Stop Vertica on Host**. If the data repository still does not stop, select **(3) Kill Vertica Process on Host** in the Advanced menu.

**NOTE**

If the data repository is running in a cluster, you might need to select **(3) Kill Vertica Process on Host** on more than one host in the cluster.

Data aggregator stops automatically.

4. Note the location of the data repository data directory as follows:

- a. Select **(6) Configuration Menu**.
- b. Select **(3) View Database**.

The database directory for the database is the parent of the catalog directory that is shown in the output.

5. Drop the database as follows:

- a. Select **(6) Configuration Menu**.
- b. Select **(7) Drop Database**.

6. Exit as the database admin user (dradmin), and then log in (su ) as the root/sudo user.

7. Find the name of the data repository package that is installed by issuing the following command. If the data repository is running in a cluster, repeat this step for each host participating in the cluster:

```
rpm -qa | grep vertica
```

8. Remove the data repository package by issuing the following command. If the data repository is running in a cluster, repeat this step for each host participating in the cluster:

```
rpm -e <data_repository_RPM_Package_Manager_installation_package>
```

9. Complete the following steps. If the data repository is running in a cluster, repeat these steps for each host participating in the cluster:

- a. Delete the /opt/vertica/ directory by issuing the following command:

```
rm -rf /opt/vertica/
```

- b. Display the database directory by issuing the following command:

```
ls <data_repository_directory>
```

- **data\_repository\_directory**

Specify the installation directory of your data repository.

**Default:** /opt/CA/IMDataRepository\_verticaVersion/

- c. Verify that the database directory that you specified is correct.

- d. Delete the database directory by issuing the following command:

```
rm -rf <data_repository_directory>
```

- **data\_repository\_directory**

Specify the installation directory of your data repository.

**Default:** /opt/CA/IMDataRepository\_verticaVersion/

The data repository component is uninstalled.

## Uninstall NetOps Kafka

If you deployed NetOps Kafka for use with your installed DX NetOps solutions, you can uninstall it.

Uninstall NetOps Kafka by running the <installation\_directory>/scripts/uninstall-netops-kafka.sh script.

- **installation\_directory**

The installation directory for NetOps Kafka.

**Default:** /opt/CA/netops-kafka

This script removes the following from the system:

- The contents that are under the top-level installation directory for NetOps Kafka (default is `/opt/CA/netops-kafka`).
- The `/etc/netops-kafka.cfg` file.

**NOTE**

The NetOps Kafka installer creates and uses this configuration file.

- All services that NetOps Kafka requires.

## Uninstall the NetOps Report Manager Service

If you deployed the NetOps Report Manager Service for use with your installed DX NetOps solutions, you can uninstall it.

### (23.3.2 and higher)

**IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

You uninstall the NetOps Report Manager Service using the uninstall script. This script removes the following from the system:

- The contents that are under the top-level installation directory for the NetOps Report Manager Service (default is `/opt/SRM`).
- The `/etc/systemd/system/multi-user.target.wants/srm-rib.service` file.

### Follow these steps:

1. Change to the `srm-rib` directory by issuing the following command:

```
cd <installation_directory>/srm-rib
```

**Example:**

```
cd /opt/SRM/srm-rib
```

- **installation\_directory**

The installation directory for the NetOps Report Manager Service.

**Default:** `/opt/SRM`

2. Uninstall based on your version:

- **(23.3.6 and higher)**

Run the uninstall script by issuing the following command:

```
./uninstall-netops-report-manager-service
```

The following output appears:

```
Stopping srm-rib service
Removing srm-rib service
Removed /etc/systemd/system/multi-user.target.wants/srm-rib.service.
Removing Report Manager Service directories
Deleting: /opt/SRM
Uninstall complete
```

- **(23.3.2 through 23.3.5)**

Do the following:

- a. Run the uninstall script by issuing the following command with the `-i` option:

```
./Uninstall_srm-rib -i console
```

The following output appears:

```
-----
```



```
Preparing CONSOLE Mode Uninstallation ...
```

```
=====
```

```
Uninstall Report Manager Service
```

```
-----
```

```
About to uninstall ...
```

```
Report Manager Service
```

```
This will remove features installed by InstallAnywhere. It will not remove files and folders
created after the installation.
```

```
PRESS <ENTER> TO CONTINUE:
```

- b. Press the **ENTER/Return** key on your keyboard.

The uninstall begins.

The following output appears:

```
=====
```

```
Uninstalling ...
```

```
-----
```

```
... *
```

```
*
```

```
*****
```

```
*****
```

```
*****
```

```
*****
```

```
... *
```

After the NetOps Report Manager Service is uninstalled successfully, a message similar to the following appears:

```
=====
```

```
Uninstall Complete
```

```
-----
```

```
Some items could not be removed.
```

The NetOps Report Manager Service is uninstalled.

## Upgrading

---

Upgrade DX NetOps Performance Management to the current version.

For the latest upgrade information, see the new unified [upgrade section](#) in the [DX NetOps documentation](#).

The following is a list of topics that you might find helpful while upgrading DX NetOps Performance Management:

- [Upgrade Requirements and Considerations](#)
- [Upgrade Prerequisites](#)
- [Upgrade Path](#)
- [Download Artifacts](#)
- [Upgrade Performance Monitoring](#)

# Migrating

Migrate DX NetOps Performance Management, either when moving to a new operating system (OS) on the same host, or to another (new) host.

Migration involves moving the DX NetOps Performance Management components—including configurations, customizations, and data—from one system to another system.

The following situations require that you migrate:

- You are moving to new hosts for an operating system (OS) install (for example, from Red Hat Enterprise Linux (RHEL) 6.9 to RHEL 8.1 or SUSE Linux Enterprise Server (SLES) 12 SP2).
- The current database hardware no longer meets the sizing requirements.
- You are moving from virtual machines to physical hardware for the database.

However, if you only need to rehydrate your cloud environment machines, see [Rehydrate Data in a Cloud Environment](#).

Migrate to new hosts using *one* of the following options:

## TIP

You can identify *when* to migrate by [reviewing the upgrade scenarios graphic](#).

- [Migrate from a supported OS.](#)
- [Migrate from an unsupported OS.](#)

## Migrate to New Hosts from a Supported Operating System

Use the following process if you are moving DX NetOps Performance Management components to new systems with new IP addresses and hostnames from a supported OS:

1. [Upgrade the components using their supported upgrade path.](#)

### NOTE

Upgrade the existing system to the DX NetOps Performance Management version to which you are migrating *before* migrating.

2. Migrate the components in the following order:

- a. [Migrate the data repository.](#)

### NOTE

Depending on the amount of data, migrating the data repository can take a significant amount of time. To minimize downtime, perform incremental backups after your initial backup.

- b. [Migrate NetOps Portal.](#)
- c. [Migrate the data aggregator.](#)
- d. [Migrate the data collectors.](#)

The DX NetOps Performance Management components are migrated to new hosts from a supported OS.

## Migrate to New Hosts from an Unsupported Operating System

Use the following process to moving DX NetOps Performance Management components to new systems with new IP addresses and hostnames from an unsupported OS:

1. [Upgrade each component using its supported upgrade path until you reach a version from which you can upgrade directly to this release.](#)
2. Back up the components to another server in the following order:
  - a. [Back up the data repository.](#)

**NOTE**

Depending on the amount of data, backing up the data repository can take a significant amount of time. To minimize downtime, perform incremental backups after your initial backup.

- b. [Back up NetOps Portal.](#)
- c. [Back up the data aggregator.](#)
3. Install the OS on each host.
4. [Install the same DX NetOps Performance Management version from your recent upgrade on each host.](#)
5. Restore the components in the following order:
  - a. [Restore the data repository.](#)
  - b. [Restore NetOps Portal.](#)
  - c. [Restore the data aggregator.](#)
6. [Upgrade the components to this version.](#)
7. Migrate the components in the following order:
  - a. [Migrate the data repository.](#)

**NOTE**

Depending on the amount of data, migrating the data repository can take a significant amount of time. To minimize downtime, perform incremental backups after your initial backup.

- b. [Migrate NetOps Portal.](#)
- c. [Migrate the data aggregator.](#)
- d. [Migrate the data collectors.](#)

The DX NetOps Performance Management components are migrated to new hosts from an unsupported OS.

## Migrate NetOps Portal

You can migrate NetOps Portal, either when moving to a new operating system (OS) on the same host, or to another (new) host.

Use the following process to migrate NetOps Portal:

1. [Back Up the NetOps Portal Database](#)
2. [Install NetOps Portal](#)
3. [Stop the NetOps Portal Services](#)
4. [Restore NetOps Portal](#)
5. (If the new NetOps Portal is going to run on a new IP address and hostname) [Update the IP Address and Hostname to Reflect the New Environment](#)
6. [Start the NetOps Portal Services](#)

### **Back Up the NetOps Portal Database**

For more information, see [Back up NetOps Portal](#).

### **Install NetOps Portal**

Prepare the host and install NetOps Portal on the new host.

**IMPORTANT**

The new host inherits the LDAP integration and HTTPS from the originating host. Do not configure these on the new host.

For more information, see [Prepare to Install NetOps Portal](#) and [Install NetOps Portal](#).

## **Stop the NetOps Portal Services**

Stop the NetOps Portal services *except* MySQL.

For more information, see [Restart NetOps Portal](#).

## **Restore NetOps Portal**

For more information, see [Restore NetOps Portal](#).

## **Update the IP Address and Hostname to Reflect the New Environment**

(If the new NetOps Portal is going to run on a new IP address and hostname) Update NetOps Portal and the event manager to reflect the new environment (the new host) by editing and running the NetOps Portal migration script. While editing the script, if you are changing the data aggregator IP and hostnames, edit those values. Run the script on the new NetOps Portal host.

### **Follow these steps:**

1. On the new NetOps Portal host, open the `<installation_directory>/PerformanceCenter/Tools/bin/update_pc_da_database_references.sh` NetOps Portal migration script, and update the following bold sections of the script to match your system:

```
...
#####
# UPDATE THE FOLLOWING PC/DA VARIABLES TO REFLECT NEW ENVIRONMENT
#####
NEW_PC_IP_ADDRESS="<New NetOps Portal IP Address>"
NEW_PC_HOSTNAME="<New NetOps Portal Hostname>"
NEW_PC_EVENT_PRODUCER_PORT=8181
NEW_PC_EVENT_PRODUCER_PROTOCOL="http" # change to "https" if using SSL
NEW_DA_IP_ADDRESS="<Existing/New Data Aggregator IP Address>"
NEW_DA_HOSTNAME="<Existing/New Data Aggregator Hostname>"
NEW_DA_PORT_NUMBER=8581
...
```

#### **– *installation\_directory***

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

2. Run the script.

The NetOps Portal and data aggregator data sources are updated with the new IP and hostname.

## **Start the NetOps Portal Services**

For more information, see [Restart NetOps Portal](#).

# **Migrate the Data Repository**

You can migrate the data repository to a new host, for example, if the current host does not meet the requirements for the existing or new data repository version.

Migrating the data repository copies the data repository data from the primary (source) cluster, which is the original data repository cluster containing the data, to a target cluster. Migrate the data repository in the following situations:

- The current host does not meet the requirements for the existing or new data repository version.
- You need to move to a new set of systems (new host) for Vertica, either to switch from virtual to physical or the other way around.
- You are upgrading the operating system (OS), and the OS requires a new machine.
- You must move to new machines.

Migrate the cluster by migrating the existing source data. The migration uses the `copycluster` command.

Use the following process to migrate the data repository to a new host:

1. [Prepare the environment.](#)
2. [Install Vertica on the new host.](#)
3. [Verify the new host.](#)
4. [Configure data transfer for the data repository.](#)
5. [Start the database on the target cluster.](#)
6. [Stop the data aggregator.](#)
7. [Copy recent data.](#)
8. [Verify the copy of the recent data.](#)
9. (If you are migrating only the data repository, which enables communication between the data aggregator and the new data repository cluster) [Update the database connection information.](#) Otherwise, if you are migrating the data repository and the data aggregator, [migrate the data aggregator](#), which includes steps that update the database connection information.

After you have migrated the data repository to a new host, complete the [next steps](#).

## **Prepare the Environment**

Before migrating the data repository to a new host, ensure that the target cluster meets the following requirements:

- It is accessible from the primary (source) cluster.
- It is a companion or equivalent cluster of the primary (source) cluster on a different system, with the same number of nodes as the primary (source) cluster.
- For each node in the primary (source) cluster, you have verified that the Database Administrator user account (default: `dradmin`) has been set up with passwordless SSH. The `copycluster` command requires that it be set up.

### **TIP**

If you have not yet, you can verify that passwordless SSH is working by issuing the following command from any node in the primary (source) cluster:

```
ssh dradmin@<paired-hostname> '/opt/vertica/bin/admintools -t list_allnodes'
```

This command executes on the paired-hostname of a node on the target cluster.

- ***paired-hostname***

The hostname of a node in the target cluster.

If the Database Administrator user is prompted for a password, passwordless SSH is not set up.

For more information, see the ["Configure Passwordless SSH for the Database Administrator User"](#) section.

- You have enabled TCP forwarding by setting `AllowTcpForwarding = Yes` in the `/etc/ssh/sshd_config` file on the Vertica hosts, on both the primary (source) and destination (target) systems. Backing up the data repository and the `copycluster` command within Vertica using the `vbr` utility requires that TCP forwarding be enabled. This allows the utility to forward connections from database hosts to backup hosts. For more information, see [the Vertica documentation](#).
- Port 50000 is open between the primary (source) data repository nodes and the target hosts.

## **Install Vertica on the New Host**

The installation process is the same as a normal data repository installation.

For more information, see [Install the Data Repository](#).

During the installation, use the same configuration for the new cluster as for the primary (source) cluster. For example, the Vertica version, node count, database name, administrator, user, catalog directory, and data directory must be the same as the originating data repository.

## **Verify the New Host**

Verify that the new host (the target cluster) has the same settings as the primary (source) cluster for the following settings:

- Database version
- Node names

### **TIP**

To get the node name, issue the following command on each node:

```
/opt/vertica/bin/admintools -t list_allnodes
```

This command also returns the installed Vertica version and the database name.

- Database name
- Database administrator
- Database user
- Catalog directory

### **TIP**

To get the catalog directory configuration, issue the following command on each node:

```
/opt/vertica/bin/admintools -t list_db -d <database name>
```

- Data directory

## **Configure Data Transfer for the Data Repository**

For the data repository, duplicate the primary (source) database to the recovery (target) database using the `vbr` script with the `--copycluster` option. `copycluster` is an incremental backup that copies all updates to the database. Because this data transfer is the longest transfer, the backup frequency to the recovery system is limited by the runtime of `copycluster`. Issue the command multiple times before you schedule a regular transfer to verify the runtime. Ensure that the backup frequency is at least twice the runtime of `copycluster`.

### **NOTE**

For existing large databases, `copycluster` takes as long as a full backup to complete. To minimize the performance impact to the system, restore a backup of the primary (source) system to the recovery system, then configure and issue the `vbr` script with the `--copycluster` option.

### **TIP**

For a large database, an incremental `copycluster` for one day takes about one hour. Run an incremental `copycluster` at least daily.

Use the following process to configure data transfer for the data repository:

1. [Configure passwordless SSH for the Database Administrator user.](#)
2. [Create the configuration files for copy cluster.](#)
3. [Stop the database on the target cluster.](#)
4. [Copy the historical data.](#)
5. [Verify the copy of the historical data.](#)

## Configure Passwordless SSH for the Database Administrator User

The `copycluster` command requires that passwordless SSH is enabled for the Database Administration user account between the nodes on the primary (source) cluster and the target cluster.

For more information, see [Configure Passwordless SSH](#).

## Create the Configuration Files for Copy Cluster

Create the configuration files for the `copycluster` command.

### Follow these steps:

1. Log in to the data repository host as the Database Administrator user.
2. Create a password file, for example `/opt/vertica/config/password.txt`.

#### NOTE

You can choose a different location for the password file.

```
[Passwords]
; Specified password for db admin account
dbPassword = DBpassword
; Specifies password for rsync user account - if different than DB admin
; serviceAccessPass = rsyncpwd
; Specifies password for the dest_dbuser Vertica account. Used only for restoring to
; alternate cluster.
; dest_dbPassword = DestinationPwd
```

3. Create a `copycluster.ini` configuration file, specifying the *full* names of the nodes in the source system (source node names) and the host names of the nodes in the target cluster (target host names) in the `[Mapping]` section, as separate lines, as shown in the following example. Create this file in any location on the source system, and give it any name with a `.ini` extension.

#### TIP

You can find the full node names in the `/opt/vertica/config/admintools.conf` file.

### Example:

The following example configuration file copies a database on a three node cluster (`v_drdata_node0001`, `v_drdata_node0002`, and `v_drdata_node0003`) to another cluster consisting of three nodes (`test-host01`, `test-host02`, and `test-host03`):

```
[Misc]
snapshotName = Copydrdata
restorePointLimit = 5
objectRestoreMode = createOrReplace
tempDir = /tmp/vbr
retryCount = 5
retryDelay = 1
passwordFile = <backupdir>/password.txt

[Database]
dbName = <drdata>
dbUser = <dradmin>
dbPromptForPassword = False

[Transmission]
encrypt = False
```



```
checksum = False
port_rsync = 50000
```

```
[Mapping]
```

```
; backupDir is not used for cluster copy
```

```
v_drdata_node0001= test-host01
```

```
v_drdata_node0002= test-host02
```

```
v_drdata_node0003= test-host03
```

- **tempDir**

The directory for the `vbr` utility. Ensure that this directory has read and write permissions.

**Example:** `/tmp/vbr`

- **backupdir**

The location of the `password.txt` file.

- **drdata**

The name of the database to copy.

**Case-sensitive:** yes

- **dradmin**

The database administrator account name.

The configuration files are created.

### **Stop the Database on the Target Cluster**

Before you start the migration, shut down the database on the target cluster.

#### **Follow these steps:**

1. Log in to the target cluster as the Database Administrator user.
2. Start the Vertica Administration Tools utility (`adminTools`) by issuing the following command:  
`/opt/vertica/bin/adminTools`
3. Select **(4) Stop Database**.  
Wait for the shutdown to complete.

### **Copy the Historical Data**

After you install the database on the target cluster, copy the data from the primary (source) database using the `copycluster` command. This command simultaneously backs up the existing source database and restores the data to the target cluster. It configures, starts, and runs the copy cluster on the primary (source) system to point to the destination (target) system.

Issuing the `copycluster` command copies the data in the data repository from *before* you run the command. Because DX NetOps Performance Management continues to collect data while the command is running, the process requires that you issue the command multiple times.

#### **NOTE**

The `copycluster` command requires that passwordless SSH is enabled for the Database Administrator user account. You enabled this when you set up the data repository.

#### **Follow these steps:**

1. Log in to the primary (source) cluster with the Database Administrator account.
2. Ensure that passwordless SSH is enabled (the passwordless SSH key is set up) for the Database Administrator user account.  
For more information, see [Configure passwordless SSH for the Database Administrator user](#).

3. Issue the following commands:

```
chown dradmin /opt/vertica/config/password.txt
chmod 600 /opt/vertica/config/password.txt
```

4. From any node on the primary (source) system, copy the historical data for the source database by issuing the following command:

```
vbr.py --task copycluster --config-file <CopyClusterConfigurationFile>.ini
```

**Example:**

```
vbr.py --task copycluster --config-file copycluster.ini
```

The following message is displayed:

```
Preparing...
Copying...
1871652633 out of 1871652633, 100%
All child processes terminated successfully.
copycluster done!
```

The historical data is copied over to all nodes on the destination (target) system.

### **Verify the Copy of the Historical Data**

After the historical data is copied, ensure the integrity of the data.

**Follow these steps:**

1. [Start the database on the target cluster](#)
2. From any node in the cluster, open the Vertica SQL prompt by issuing the following command:  

```
/opt/vertica/bin/vsql -U dauser
```
3. Verify the timestamp of these key database tables by issuing the following queries:  

```
SELECT to_timestamp(max(tstamp)) from dauser.reach_rate;
SELECT to_timestamp(max(tstamp)) from dauser.ifstats_rate;
```

The date and time must correspond to the time when you started the copy.

The copy of the historical data is verified.

### **Start the Database on the Target Cluster**

Start the database on the target cluster.

**Follow these steps:**

1. Log in to the target cluster as the Database Administrator user.
2. Start the Vertica Administration Tools utility (`adminTools`) by issuing the following command:  

```
/opt/vertica/bin/adminTools
```
3. Select **(3) Start Database**.

The target database is started.

### **Stop the Data Aggregator**

To maintain integrity of your data, stop the data aggregator before the final copy of the data repository. Polling continues on the data collector if it is running and polling when the data aggregator is stopped. During this phase, the data collector continues to receive and queue poll responses from the managed devices. After copy cluster is done and you have restarted the data aggregator and unblocked the data collector firewall rules, the data collector delivers these polled responses to the data repository.

**Follow these steps:**

1. Log in to the data aggregator host as the root user or a sudo user.

**NOTE**

If you installed the data aggregator as the sudo user, you set up a sudo command alias for the service `dadaemon` command. Use the sudo commands.

2. Use firewall rules to block traffic from the data collectors. Issue the following command for each data collector:

```
iptables -A INPUT -s <DC_IP> -j DROP
```

- **DC\_IP**

The IP of the data collector.

3. Do one of the following steps:

- Stop the data aggregator service by issuing the following command:

```
systemctl stop dadaemon
```

- (Fault-tolerant environments) If the local data aggregator is running, shut it down and prevent it from restarting until maintenance is complete by issuing the following commands

```
<installation_directory>/scripts/dadaemon maintenance
```

- **installation\_directory**

The installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator stops.

**Copy Recent Data**

Copy recent data in the source cluster by issuing the `copycluster` command again. The command copies only new data that has arrived after the initial copy.

**Follow these steps:**

1. Log in to the source cluster with the database administrator account.
2. Issue the following command:

```
vbr.py --task copycluster --config-file <CopyConfigurationFile>.ini
```

**Example:**

```
vbr.py --task copycluster --config-file copycluster.ini
```

The recent data in the source cluster is copied.

**Verify the Copy of the Recent Data**

To ensure the integrity of your data, verify the data of the source cluster that you copied.

**Follow these steps:**

1. [Start the database on the target cluster.](#)
2. From any node in the cluster, open the Vertica SQL prompt by issuing the following command:

```
/opt/vertica/bin/vsql -U dauser
```

3. Verify the timestamp of these key database tables by running the following queries:

```
SELECT to_timestamp(max(tstamp)) from dauser.reach_rate;
```

```
SELECT to_timestamp(max(tstamp)) from dauser.ifstats_rate;
```

The date and time must correspond to the time when you started the copy.

The copy of the recent data is verified.

### **Update the Database Connection Information**

Prepare to point the data aggregator to the target database. Complete this procedure if you are migrating *only* the data repository, which enables communication between the data aggregator and the new data repository cluster. Otherwise, if you are migrating both the data repository and the data aggregator, skip this procedure and [migrate the data aggregator](#), which includes steps that update the database connection information.

#### **Follow these steps:**

1. Log in to the data aggregator host.
2. Open the `dbconnection.cfg` file by issuing the following command:
 

```
vi <installation_directory>/apache-karaf/etc/dbconnection.cfg
```

  - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
3. Update the following parameter with the hostnames of the target cluster:
 

```
dbHostNames=<hostname1,hostname2,hostname3>
```

The data repository migration is complete.

### **Next Steps**

After you have migrated the data repository to a new host, complete the following procedures:

1. [Start the data aggregator.](#)
2. [Monitor the data aggregator restart process.](#)
3. [Validate the migration.](#)

### **Start the Data Aggregator**

#### **Follow these steps:**

1. Log in to the data aggregator host as the root user or a sudo user.
 

**NOTE**  
If you installed the data aggregator as the sudo user, you set up a sudo command alias for the service `dadaemon` command. Use the sudo commands.
2. Do one of the following steps:
  - Start the data aggregator service by issuing the following command:
 

```
systemctl start dadaemon
```
  - (Fault-tolerant environments) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:
 

```
<installation_directory>/scripts/dadaemon activate
```

    - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`

The data aggregator starts and synchronizes with NetOps Portal and the data repository. When the `iptables` entries are removed, the data collector sends any queued poll responses to the data aggregator.

If the queued data exceeds the disk space limit, then the data collector discards the oldest data. This can result in a data gap.

### **Monitor the Data Aggregator Restart Process**

#### **Follow these steps:**

1. After waiting for few minutes to let the data aggregator restart, log in to the data aggregator host and navigate to the `/opt/IMDataAggregator/data/performance-spool` directory.
2. Verify that there are no data transfer object (DTO) files with a size greater than zero.
3. Enable traffic from the Simple Network Management Protocol (SNMP) data collector with largest number of polled items by issuing the following command:

```
iptables -D INPUT -s <DC_IP> -j DROP
```

The data aggregator starts schema validation and processing of cached and new polled data from this data collector.

4. After the data aggregator system utilization decreases, do the following:
  - a. Enable traffic from the remaining SNMP data collectors.
  - b. Enable traffic from the DX NetOps Mediation Manager data collectors.

The data aggregator restart process is monitored.

### **Validate the Migration**

Ensure that the system is healthy and then validate that the migrated data is present.

#### **Follow these steps:**

1. Log in to NetOps Portal
2. Verify that the system is healthy. Do the following steps:
  - a. For the data aggregator, do the following:
    - a. Navigate to the data aggregator data source list.
    - b. Verify that the system status is good.
    - c. Verify that the data aggregator data source is available.
    - d. Verify the **Last Polled On** data and time.
  - b. For the data collector, do the following:
    - a. Navigate to the data collector list.
    - b. Verify that the data collectors are up and collecting data.

For more information:

- About how to view system status, see [View System Status](#).
  - About how to view system health, see [Monitor System Health](#).
3. Validate that the migrated data is present. Go to the **Infrastructure Overview** dashboard, and then verify that data is available for the following time ranges:
    - **Last hour**  
The system is receiving SNMP poll responses when poll data exists for the last hour.
    - **Last 7 days**  
Data was properly migrated when poll data exists for the last seven days.
 For more information about this dashboard, see [Out-of-the-Box Dashboards](#).

The migration is validated.

## Reinstall or Migrate the Data Aggregator

You can reinstall or migrate the data aggregator on the same host or migrate the data aggregator to another (new) host.

Migrate the data collectors in the following situations:

- The current host does not meet the requirements for the existing or new data aggregator version.
- You need to move to a new host, either to switch from virtual to physical or the other way around.
- You are upgrading the operating system (OS), and the OS requires a new machine.
- You need to move to new machines.

Use the following process to migrate the data aggregator on a new host:

1. (If the existing DX NetOps Performance Management version does not support the OS for the new host) [Upgrade the existing system to the DX NetOps Performance Management version to which you are migrating.](#)
2. [Back up the data aggregator.](#)
3. [Install the data aggregator on a new machine.](#)
4. [Complete the reinstallation/migration to the new host.](#)
5. (If you are migrating to a new host with a new IP address and hostname) [Update the IP address and hostname to reflect the new environment.](#)
6. [Update the data collector configurations.](#)
7. [Update the data aggregator data source in NetOps Portal.](#)

The data aggregator is migrated to a new host.

### Upgrade the Existing System

(If the existing DX NetOps Performance Management version does not support the OS for the new host) [upgrade the existing system to the DX NetOps Performance Management version to which you are migrating.](#)

For more information about the OS versions that the currently installed DX NetOps Performance Management version supports *for the version from which you are migrating*, see [Installation Requirements and Considerations](#).

### Back up the Data Aggregator

Use the following process to back up the data aggregator:

1. [Back up the files on the data aggregator on the originating \(current\) host.](#)
2. Copy the `DA.tar.gz` file (the data aggregator backup file) by issuing the following command:
 

```
scp -r <backup_dir>/DA.tar.gz <user>@<backuphost>:<directory_location>
```

  - **backup\_dir**  
The directory that includes the tar file.
  - **user**  
The user who owns the data aggregator process (root or sudo user).
  - **backuphost**  
The host to which you are migrating or a temporary host when reinstalling the existing system.
  - **directory\_location**  
The location to save the `DA.tar.gz` data aggregator backup file.  
If `/tmp` is not a preferred location, specify another directory location.

### Install the Data Aggregator on a New Machine

For more information, see [Install the Data Aggregator](#).

## Complete the Reinstallation/Migration to the New Host

Use the following process to complete the reinstallation/migration to the new host:

1. [Stop the data aggregator.](#)
2. [Restore the data aggregator.](#)
3. On the data aggregator host, ensure that the `dbconnection.cfg` file points to the correct data repository hosts, either the new migrated one or the previous one if you are not migrating:
  - a. Open the `<installation_directory>/apache-karaf/etc/dbconnection.cfg` file on the data aggregator host.
    - **`installation_directory`**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
  - b. Modify the following line in the file. Modify the line to reference the hostname or IP address of each data repository host:  
`dbHostNames=dbNode1Hostname,dbNode2Hostname,dbNode3Hostname`
  - c. Save the file.

The `dbconnection.cfg` file now points to the data repository hosts.

## Update the IP Address and Hostname to Reflect the New Environment

(If the new data aggregator will run on a new IP address and hostname) Update data aggregator database to reflect the new environment (the new host) by editing and running the migration script. While editing the script, if you are changing the data aggregator IP and hostnames, edit those values. Run the script on the data repository host.

Follow these steps:

1. On the data repository host, update the bold sections of the `<data_repository_directory>/update_da_dc_database_references.sh` migration script to match your system:

– **`data_repository_directory`**

The installation directory for the data repository.

**Default:** `/opt/CA/IMDataRepository_verticaVersion`

```
#####
# UPDATE DAUSER/DAPASS BELOW TO REFLECT THE NON-ADMIN
# VERTICA USERNAME/PASSWORD FOR THIS SYSTEM
#####
DAUSER=dauser
DAPASS=dapass
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S IP ADDRESS BELOW
#####
RECOVERY_DA_IP_ADDRESS="<New IP Address for the Data Aggregator>"
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S HOSTNAME BELOW
#####
SOURCE_DA_HOSTNAME="<Original Hostname for the Data Aggregator>"
RECOVERY_DA_HOSTNAME="<New Hostname for the Data Aggregator>"
#####
# UPDATE THE FOLLOWING ARRAYS TO REFLECT THE SOURCE DATA
# COLLECTOR HOSTNAMES, NEW HOSTNAMES, AND NEW
```

```
# IP ADDRESSES RESPECTIVELY.
# IMPORTANT: THE ORDER OF THE ENTRIES BELOW IS CRITICAL FOR
# MAPPING PURPOSES. IN ADDITION, IF MULTIPLE VALUES
# ARE REQUIRED, SEPARATE VALUES WITH A SINGLE SPACE.
#####
declare -a SOURCE_DC_HOSTNAMES=(<remove the default text>)
declare -a RECOVERY_DC_HOSTNAMES=(<remove the default text>)
declare -a RECOVERY_DC_IP_ADDRESSES=(<remove the default text>)
```

**Example:**

```
#####
# UPDATE DAUSER/DAPASS BELOW TO REFLECT THE NON-ADMIN
# VERTICA USERNAME/PASSWORD FOR THIS SYSTEM
#####
DAUSER=dauser
DAPASS=dapass
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S IP ADDRESS BELOW
#####
RECOVERY_DA_IP_ADDRESS="192.168.2.200"
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S HOSTNAME BELOW
#####
SOURCE_DA_HOSTNAME="oldhostname"
RECOVERY_DA_HOSTNAME="newhostname"
#####
# UPDATE THE FOLLOWING ARRAYS TO REFLECT THE SOURCE DATA
# COLLECTOR HOSTNAMES, NEW HOSTNAMES, AND NEW
# IP ADDRESSES RESPECTIVELY.
# IMPORTANT: THE ORDER OF THE ENTRIES BELOW IS CRITICAL FOR
# MAPPING PURPOSES. IN ADDITION, IF MULTIPLE VALUES
# ARE REQUIRED, SEPARATE VALUES WITH A SINGLE SPACE.
#####
declare -a SOURCE_DC_HOSTNAMES=()
declare -a RECOVERY_DC_HOSTNAMES=()
declare -a RECOVERY_DC_IP_ADDRESSES=()
```

**2. Run the script.**

The data collectors are updated with the new IP and hostname.

**Update Data Collector Configurations**

Update the data collectors to point to the new data aggregator host.

For more information, see [Configure Data Collector When the Data Aggregator IP Address Changes](#).

**Follow these steps:**

1. On each data collector host, edit the data aggregator host information in following files:



**NOTE**

If you are upgrading or migrating the data collectors, do this procedure *after* the upgrade or migration.

- /opt/IMDataCollector/apache-karaf/etc/com.ca.im.dm.core.collector.cfg
- /opt/IMDataCollector/broker/<apache-activemq-\*/>/conf/activemq.xml

- **apache-activemq-\***

The installation directory for Apache ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

2. Restart the Data Collector Karaf and Data Collector ActiveMQ services by issuing the following commands:

```
systemctl stop activemq
systemctl start activemq
systemctl stop dcmd
systemctl start dcmd
```

### **Update the Data Aggregator Data Source in NetOps Portal**

Update the data aggregator data source in NetOps Portal to point to the new data aggregator host.

#### **Follow these steps:**

1. Hover over **Administration**, **Data Sources**, and then click **Data Sources**.  
The **Manage Data Sources** page appears.
2. Select the data aggregator data source that you want to update, and then click **Edit**.  
The **Edit Data Source** dialog opens.
3. Specify the following fields, and then click **Save**:
  - **Host Name**  
The IP address or host name of the new data aggregator host.
  - **Protocol**  
If you have set up HTTPS, select **https** as the communication protocol.
  - **Synchronize component items that are not currently present on the monitored device**  
When the data aggregator finds a device component that is no longer present in the environment, that status of the component is set to **Not Present**. By default, the data aggregator does not synchronize these items because data can no longer be collected for the item. Historical data that has not reached the data retention limit is still available for these items.

#### **IMPORTANT**

If the properties of an active component match the identifying properties of the not present component, the components are indistinguishable. Data from the previous component might contribute to group based dashboards instead of data from the active component. Enable this feature only under the following circumstances:

- You want to report on historical data for items that are no longer present in the environment.
- You can ensure that the not present item does not conflict with an actively monitored item.
- You can identify the not present items so that you can exclude the items from groups.

- **Discover devices from other data sources**

This option specifies whether synchronization includes devices that are reported by other data sources. Automatic synchronization includes only devices that are discovered after you select this option. To discover previously discovered devices, perform a full synchronization of the data aggregator.

DX NetOps Performance Management creates a discovery profile that includes the IP addresses of the devices.

This discovery profile attempts discovery once per day.

4. If the URL to access the data aggregator is a selected authorized URL (it is listed in the **Select Authorized URLs** field on the **Security Settings** page in NetOps Portal), [configure login URL protection by adding the new data aggregator URL to the list](#).

If the data aggregator data source is enabled, data appears in NetOps Portal after the next synchronization.

## Reinstall or Migrate the Data Collectors

You can migrate, or move, the data collector to another system without having to rediscover network devices and components or lose historical data.

The data collector polls 500,000 devices and components, and you do not want to lose data or perform rediscovery.

Migrate the data collectors in the following situations:

- The current host does not meet the requirements for the existing or new data collector version.
- You need to move to a new host, either to switch from virtual to physical or the other way around.
- You are upgrading the operating system (OS), and the OS requires a new machine.
- You need to move to new machines.

Use the following process to *move the data collector to another system*:

1. [Verify the prerequisites.](#)
2. [Determine the unique identifier for the data collector.](#)
3. [Stop the data collector.](#)
4. [Reinstall the data collector on a clean host.](#)
5. (If you are migrating to a new host with a new IP address and hostname) [Update the IP address and hostname to reflect the new environment.](#)
6. (Optional) [Verify data collection on the new host.](#)

### NOTE

Backup and restore is not required for the data collector.

### Verify the Prerequisites

Ensure that you have met the following prerequisites prior to moving the data collector:

- If your installation uses DX NetOps Mediation Manager, you have migrated the device packs.
- You have noted the following considerations:
  - The amount of data loss is equal to the amount of time that has elapsed from the time that you stop the old data collector to the time that you deploy (install) the new data collector.
  - If the old data collector happens to start accidentally, the Simple Network Management Protocol (SNMP) data is polled twice. A warning appears in the data aggregator karaf log:

```
WARN | Session Task-810 | 2013-01-02 13:52:09,062 | DCHeartBeatLog |
ore.collector.interfaces |
| HeartBeat message not received. Expected: 93, Received: 255
```

To fix this problem, stop or uninstall the old data collector.

### Determine the Unique Identifier for the Data Collector

If you do not know the unique identifier for the data collector, retrieve it.

**Complete one of the following steps:**

- Log in to NetOps Portal as a user with the Administrator role, and do the following steps:
  - a. Hover over **Administration, Data Source Settings**, and then select a data aggregator data source from the menu. The Data Aggregator Admin UI opens. The **Monitored Devices** page appears.
  - b. Select **Data Collectors** below the **System Status**. The **Data Collectors** page appears.

- c. Find the data collector component that you want to move, and notate its ID.  
The format of the data collector component ID is `HOSTNAME:UUID`.
- Open a web browser and issue the following web service call:  
`http://<da_host>:<port>/rest/dcms`
  - **da\_host**  
Specifies the data aggregator host name.
  - **port**  
Specifies the data aggregator required port number.  
**Default port:** 8581
- Find the `<DataCollectionMgrInfo>` section whose `HostName` and `IPAddress` match the one you want to move.  
Note the value for `<DcmID>`.

### Stop the Data Collector

Stop the data collector services on the current host.

#### Follow these steps:

1. (If you have installed device packs for this data collector) Complete the following steps:
  - a. In NetOps Portal, on the **Monitored Devices** page, select **EMS Integration Profiles** from the **Monitoring Configuration** menu.
  - b. Right-click a profile that is associated with this data collector host, and then select **Stop**. Do this step for every EMS profile that is related to this data collector host.
  - c. Archive the DX NetOps Mediation Manager artifacts by issuing the following command:
 

```
tar -
zcvf <filename>
<DC_installation_directory>/apache-karaf/MediationCenter
```

**Example:**

```
tar -zcvf filename /opt/IMDataCollector/apache-karaf/MediationCenter
```

    - **filename**  
Specifies the name of the archive file, which is moved to the new data collector host.
    - **DC\_installation\_directory**  
The default installation directory for the data collector.  
**Default:** /opt/IMDataCollector
2. Stop the Data Collector, the ICMP daemon, and ActiveMQ services on the data collector host by issuing the following command:
 

```
systemctl stop dcmd
systemctl stop icmpd
systemctl stop activemq
```

or

```
sudo systemctl stop dcmd
sudo systemctl stop icmpd
sudo systemctl stop activemq
```
3. Verify that the data collector has stopped by completing the following steps:
  - a. Log in to NetOps Portal as a user with the Administrator role.
  - b. Hover over **Administration, Data Source Settings**, and then select a data aggregator data source from the menu.
  - c. Select **System Status, Data Collectors** from the menu.

The Data Aggregator Admin UI opens.

- d. Verify that the data collector status is "Not Connected".

### **Reinstall the Data Collector on a Clean Host**

Reinstall the data collector on a clean host using the same data collector unique identifier (DCM\_ID). This allows the new data collector to retrieve the original (old) data collector configuration from the data aggregator when the new data collector process starts.

#### **Follow these steps:**

1. Verify that DX NetOps Mediation Manager and DX NetOps Virtual Network Assurance point to the correct data collectors:
  - For DX NetOps Mediation Manager, see [the DX NetOps Mediation Manager documentation](#).
  - For DX NetOps Virtual Network Assurance, see [the DX NetOps Virtual Network Assurance documentation](#).
2. Complete one of the following based on your installation:
  - (DX NetOps Mediation Manager only) Migrate your device packs. On the old data collector host, issue the `$CMM_HOME/tools/migratePMtoCMM` script with the `-t` option on a data collector server where a Local Controller is installed. The DX NetOps Mediation Manager Console is running on another server. For more information, see [the DX NetOps Mediation Manager documentation](#).
  - (DX NetOps Virtual Network Assurance only) See [the DX NetOps Virtual Network Assurance documentation](#).
3. Log in to the new host system and open a command shell session.
4. Set an environment variable with the ID of the data collector by issuing the following command:
 

```
export DCM_ID=<data_collector_id>
```

  - **data\_collector\_id**  
The DcmID that you noted for the data collector.
5. Install the data collector from the same session by running the `install.bin` file.  
For more information, see [Install the Data Collectors](#).
6. Install the DX NetOps Mediation Manager Local Controller on the same server.
7. If you have previously-installed device packs for this data collector, complete these additional steps:
  - a. Copy the zip files that you created previously with the migration script to local directories on this host.
  - b. Using the DX NetOps Mediation Manager web console, deploy these device packs and start them.

#### **NOTE**

It is not required that you redeploy the certification packs to the data aggregator host.

### **Update the IP Address and Hostname to Reflect the New Environment**

(If the new data collector is going to run on a new IP address and hostname) Update data aggregator database to reflect the new environment (the new host) by editing and running the migration script. While editing the script, if you are changing the data collector IP and hostnames, edit those values. Run the script on the data repository host.

#### **Follow these steps:**

1. On the data repository host, update the bold sections of the `<data_repository_directory>/update_da_dc_database_references.sh` migration script to match your system:

- **data\_repository\_directory**

The installation directory for the data repository.

**Default:** `/opt/CA/IMDataRepository_verticaVersion`

Ensure that the order of the data collectors for the original system and the new system is the same. The script uses the order of the list to map the original system components to the new system.

```
#####
# UPDATE DAUSER/DAPASS BELOW TO REFLECT THE NON-ADMIN
```

```

# VERTICA USERNAME/PASSWORD FOR THIS SYSTEM
#####
DAUSER=dauser
DAPASS=dapass
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S IP ADDRESS BELOW
#####
RECOVERY_DA_IP_ADDRESS="<remove the default text>"
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S HOSTNAME BELOW
#####
SOURCE_DA_HOSTNAME="<remove the default text>"
RECOVERY_DA_HOSTNAME="<remove the default text>"
#####
# UPDATE THE FOLLOWING ARRAYS TO REFLECT THE SOURCE DATA
# COLLECTOR HOSTNAMES, NEW HOSTNAMES, AND NEW
# IP ADDRESSES RESPECTIVELY.
# IMPORTANT: THE ORDER OF THE ENTRIES BELOW IS CRITICAL FOR
# MAPPING PURPOSES. IN ADDITION, IF MULTIPLE VALUES
# ARE REQUIRED, SEPARATE VALUES WITH A SINGLE SPACE.
#####
declare -a SOURCE_DC_HOSTNAMES=("<Source/Original DC Hostname 1>" "<Source/Original
DC Hostname 2>")
declare -a RECOVERY_DC_HOSTNAMES=("<New DC Hostname 1>" "<New DC Hostname 2>")
declare -a RECOVERY_DC_IP_ADDRESSES=("<New DC Hostname 1 IP Address>" "<New DC
Hostname 2 IP Address>")

```

**Example:**

```

#####
# UPDATE DAUSER/DAPASS BELOW TO REFLECT THE NON-ADMIN
# VERTICA USERNAME/PASSWORD FOR THIS SYSTEM
#####
DAUSER=dauser
DAPASS=dapass
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S IP ADDRESS BELOW
#####
RECOVERY_DA_IP_ADDRESS=""
#####
# UPDATE TO REFLECT THE NEW DATA AGGREGATOR'S HOSTNAME BELOW
#####
SOURCE_DA_HOSTNAME=""
RECOVERY_DA_HOSTNAME=""
#####
# UPDATE THE FOLLOWING ARRAYS TO REFLECT THE SOURCE DATA
# COLLECTOR HOSTNAMES, NEW HOSTNAMES, AND NEW

```

```
# IP ADDRESSES RESPECTIVELY.
# IMPORTANT: THE ORDER OF THE ENTRIES BELOW IS CRITICAL FOR
# MAPPING PURPOSES. IN ADDITION, IF MULTIPLE VALUES
# ARE REQUIRED, SEPARATE VALUES WITH A SINGLE SPACE.
#####
declare -a SOURCE_DC_HOSTNAMES=("oldDChost1" "oldDChost2")
declare -a RECOVERY_DC_HOSTNAMES=("newDChost1" "newDChost2")
declare -a RECOVERY_DC_IP_ADDRESSES=("192.168.2.100" "192.168.2.105")
```

## 2. Run the script.

The data collectors are updated with the new IP and hostname.

### **(Optional) Verify Data Collection on the New Host**

After several polling cycles, verify that the data collector is collecting data on the new host. Uninstall the old data collector, and delete any associated EMS profiles.

## Rehydrate Data in a Cloud Environment

If you have set up DX NetOps Performance Management in a cloud environment, you can patch operating systems from a common image instead of patching each operating system individually.

Rehydrating data migrates existing operating systems (OS) to new OSs for all components.

Use the following process to rehydrate the DX NetOps Performance Management nodes with minimal data loss:

1. [Verify the prerequisites.](#)
2. [Rehydrate the data collectors.](#)
3. [Rehydrate the Vertica nodes.](#)
4. [Verify the Vertica nodes.](#)
5. [Rehydrate the data aggregator.](#)
6. [Rehydrate NetOps Portal.](#)
7. [Rehydrate DX NetOps Virtual Network Assurance \(VNA\).](#)
8. (If you rehydrate the DX NetOps Performance Management environment, and want to reconnect it to an existing SpectroSERVER) [Reconnect the DX NetOps Spectrum \(Spectrum\) data source.](#)
9. [Rehydrate DX NetOps Network Flow Analysis \(NFA\).](#)

### **Verify the Prerequisites**

Before rehydrating the data, ensure that your environment is in a good state.

#### **Follow these steps:**

1. Verify the following details:
  - The data aggregator and data repository are connected.
  - The data aggregator is up and running.
  - Backed up or cached poll data does not exist.

#### **NOTE**

If the PRQ queue is not empty, the data collectors need to send the rest of the polled data to the data aggregator. The queue fills up when an outage occurs, which causes data to be cached on the data collectors. View the **System Status** page to verify that all the data collectors have a green status. The

**Polling Status** column shows whether cached values exist on the data collector.

Verify these details from the **System Status** page.

**TIP**

To view this page, hover over **Administration**, **Data Sources**, and then click **System Status**.

2. If you have VNA in your environment, set the VNA Gateway **Administrative State** to "Down" by doing the following steps:
  - a. Hover over **Administration**, **Monitored Items Management**, and then click **VNA Administration**.  
The **VNA Administration** page appears.
  - b. In the **Gateway Configuration** section, set **Administrative State** to "Down", and then save your changes.  
The VNA Gateway **Administrative State** is "Down".

### **Rehydrate the Data Collectors**

Rehydrate each data collector one at a time. Ensure each data collector recovers and starts polling before you rehydrate each Vertica node and the data aggregator. During this process, some polls are cached on the data collectors for a time.

#### **Follow these steps:**

1. Build the new operating system on the data collector container or virtual machine.
2. Copy the `DCM_ID` by issuing the following command:
 

```
grep "manager\-id\=" <DC_installation_directory>/apache-karaf/etc/com.ca.im.dm.core.collector.cfg
```

  - ***DC\_installation\_directory***  
The installation directory for the data collector.  
**Default:** `/opt`
3. Bring down the old container or virtual machine.
4. Give the new container or virtual machine a new IP address or name and bring it online.
5. Reinstall the data collector with the `DCM_ID` of the original data collector by issuing the following commands:
 

```
export DCM_ID="Original_DC_Host:DCM_ID"
cd /tmp;
rm -rf install.bin;
wget http://DA_Host:Port/dcm/InstData/Linux/VM/install.bin;
chmod a+x install.bin;
./install -i silent
```

The data collector installs, reconnects to the data aggregator, and then starts polling.

### **Rehydrate the Vertica Nodes**

Rehydrate each Vertica node one at a time.

**Prerequisite:** You have verified that all nodes are up and running by issuing the following command using the Vertica Administration Tools utility (adminTools):

```
/opt/vertica/bin/admintools -t list_allnodes
```

#### **Follow these steps:**

1. Bring down the Vertica node by issuing the following command:
 

```
/opt/vertica/bin/admintools -t stop_node -s IP_Address
```
2. Unmount the `data` directory and the `catalog` directory.
3. [Create a new node with the same IP address and name.](#)
4. Mount the `data` and `catalog` directory to the new node.
5. Validate the system settings by running the `dr_validate.sh` validation script, and then reviewing and resolving any errors or warnings. Issue the following command:

```
./dr_validate.sh -n -p drinstall.properties
```

You can run this script multiple times to verify that all system configuration options are set properly. The validation script might prompt you to reboot.

6. Install Vertica from an up-and-running node by issuing the following command:

```
/opt/vertica/sbin/install_vertica -u dradmin -l /export/dradmin -d /export/data -L ./resources/vlicense.dat -Y -r ./resources/vertica-<version>.rpm
```

#### NOTE

Values should match those in the properties files for `dr_install.sh`, and should point to the same resources.

7. Start the node and verify that it is up and running by issuing the following command:

```
/opt/vertica/bin/admintools -t restart_node -s Host_Name -d DB_Name
```

#### NOTE

The state starts as DOWN, then changes to REBUILDING until it changes to UP.

### Verify the Vertica Nodes

Verify that all nodes are back up and running as the `dradmin` user by issuing the following command:

```
/opt/vertica/bin/admintools -t list_allnode
```

### Rehydrate the Data Aggregator

During this Vertica refresh, the data aggregator collects data from the data collectors. The data aggregator pushes data to Vertica the entire time on the up and running nodes. The speed of the ingestion is sometimes cut in half during this time. After you rehydrate the data collectors and Vertica, you can rehydrate the data aggregator.

#### Follow these steps:

1. [Install the data aggregator.](#)
2. Do one of the following steps:
  - Stop the Data Aggregator service by issuing the following command:

```
systemctl stop dadaemon
```

- (Fault-tolerant environment) If the local data aggregator is running, shut it down and prevent it from restarting until maintenance is complete by issuing the following command:

```
<installation_directory>/scripts/dadaemon maintenance
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

The data aggregator completes processing and the service stops.

3. [Move the configuration files to the new node.](#)

#### NOTE

In fault-tolerant data aggregator environments, use the shared data directory, and then reattach it.

For more information, see [Fault Tolerance](#).

4. Do one of the following steps:
  - Start the Data Aggregator service by issuing the following command:

```
systemctl start dadaemon
```

- (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:

```
<installation_directory>/scripts/dadaemon activate
```

- **installation\_directory**



The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator consumes the queued polls and pushes them to Vertica.

**NOTE**

Depending on the total outage time, all cached data is consumed and ready for reporting in approximately two times the outage time. When ActiveMQ consumption returns to normal, this indicates there is no longer a backlog. Verify that all the data collectors have a green status and that the system is receiving approximately the same number of polls as it was before the process started from the **System Status** page.

The data aggregator is rehydrated.

## **Rehydrate NetOps Portal**

**Follow these steps:**

1. [Install NetOps Portal](#).
2. Stop the NetOps Portal services by issuing the following commands:
 

```
systemctl stop caperfcenter_console
systemctl stop caperfcenter_devicemanager
systemctl stop caperfcenter_eventmanager
systemctl stop caperfcenter_sso
```
3. [Move the database to the new node](#).
4. Start NetOps Portal by doing the following steps:
  - a. Start the SSO service by issuing the following command:
 

```
systemctl start caperfcenter_sso
```
  - b. Wait one minute, then start the event manager and device manager by issuing the following commands:
 

```
systemctl start caperfcenter_eventmanager
systemctl start caperfcenter_devicemanager
```
  - c. Wait one minute, then start the console service by issuing the following commands:
 

```
systemctl start caperfcenter_console
```

NetOps Portal is rehydrated.

## **Rehydrate VNA**

If you have VNA in your environment, rehydrate it now.

**Follow these steps:**

1. Retrieve the engine ID by issuing a GET request to the following REST URL:
 

```
http://<VNA_host>:<port>/vna/rest/v1/admin/engines
```

  - **VNA\_host**  
Specifies the VNA host name.
  - **port**  
Specifies the VNA required port number.  
**Default:** 8080  
For more information about the ports that are required for communication between DX NetOps Performance Management and VNA, see [Installation Requirements and Considerations](#).
2. Retrieve the plug-in configuration by issuing a GET request to the following REST URL:
 

```
http://<VNA_host>:<port>/vna/rest/v1/admin/engines/Engine_ID/config
```

  - **VNA\_host**  
Specifies the VNA host name.
  - **port**

Specifies the VNA required port number.

**Default:** 8080

3. Stop the application server by issuing the following command:

```
systemctl stop wildfly
```

4. Back up the existing database to a specified directory by issuing the following command:

```
<VNA_home>/VNA/tools/bin/db_backup.sh <backup_directory/backup_filename>
```

- **VNA\_home**

The installation directory for VNA.

**Default:** /opt/CA

- **backup\_directory/backup\_filename**

The location of the backup directory and file. Use any secure location with sufficient space for the backup directory.

**Example:** /tmp/vna\_db.sql

5. [Install VNA on the new server](#), and then restore the database from the backup by issuing the following command:

```
<VNA_home>/tools/bin/db_restore.sh <backup_directory/backup_filename>
```

- **VNA\_home**

The installation directory for VNA.

**Default:** /opt/CA

- **backup\_directory/backup\_filename**

The location of the backup directory and file. Use any secure location with sufficient space for the backup directory.

**Example:** /tmp/vna\_db.sql

6. Reconfigure the plug-ins using the information from the original query.

**NOTE**

Use the same Domain ID from the original configuration.

7. Set the VNA Gateway **Administrative State** to "Up" by doing the following steps:

- a. In NetOps Portal, hover over **Administration, Monitored Items Management**, and then click **VNA Gateways**.

The **VNA Gateways** page appears.

- b. Change the new VNA server ID.

- c. Set the **Administrative State** to "Up".

The VNA Gateway **Administrative State** is "Up".

VNA is rehydrated.

### **Reconnect an Existing Spectrum Data Source**

If you rehydrate the DX NetOps Performance Management environment, and want to reconnect it to an existing Spectrum server, complete the following procedure.

**Follow these steps:**

1. Disable the Spectrum data source by doing the following steps:

- a. In NetOps Portal, hover over **Administration, Data Sources**, and then click **Data Sources**.

The **Manage Data Sources** page appears.

- b. Select the Spectrum data source, and then click **Edit**.

The **Edit Data Source** dialog opens.

- c. Change the **Status** to "Disabled", and then click **Save**.

The Spectrum data source is disabled.

2. Remove the NetOps Portal entries by issuing the following commands:

```
cd <installation_directory>/vnmsh
./connect
./show models | grep CAPC
```

– **installation\_directory**

The installation directory for Spectrum.

3. Issue the following command for every CAPCIPDomain and CAPCTenant model found:

```
./destroy model mh=0XXXXXX
```

4. Remove the NetOps Portal integration database from Spectrum by issuing the following command:

```
bash -login
cd mysql
cd bin
./mysqladmin --defaults-file=../my-spectrum.cnf -u netqos -p password drop
netqos_integ
```

5. Restart Spectrum Tomcat by issuing the following command:

```
cd <installation_directory>/tomcat/bin
./stopTomcat.sh
./startTomcat.sh
```

– **installation\_directory**

The installation directory for Spectrum.

6. Enable the Spectrum data source by doing the following steps:

- a. In NetOps Portal, hover over **Administration, Data Sources**, and then click **Data Sources**.  
The **Manage Data Sources** page appears.
- b. Select the Spectrum data source, and then click **Edit**.  
The **Edit Data Source** dialog opens.
- c. Change the **Status** to "Enabled", and then click **Save**.  
The Spectrum data source is enabled.

The existing Spectrum data source is reconnected.

## **Rehydrate NFA**

If you have NFA in your environment, rehydrate it now.

### **Follow these steps:**

1. Disable the NFA data source by doing the following steps:
  - a. In NetOps Portal, hover over **Administration, Data Sources, Data Sources**.  
The **Manage Data Sources** page appears.
  - b. Select the NFA data source, and then click **Edit**.  
The **Edit Data Source** dialog opens.
  - c. Change the **Status** to "Disabled", and click **Save**.  
The NFA data source is disabled.
2. Determine the database files to backup:
  - Customized data\_retention database (Stand-alone or Harvester server): `data_retention`
  - Harvester database (Stand-alone or Harvester server): `harvester`
  - Reporter database (Stand-alone or NFA console): `reporter`
3. Copy each of the target directories or files to a remote location.
4. Back up the database files to a remote location by issuing the following command:

### **IMPORTANT**

Back up the `reporter` database last, regardless of the deployment architecture.

```
mysqldump --routines --events -u root DB_Name --skip-lock-tables > dbbackupname.sql
```

5. (Optional) Verify that the backup was successful by checking that the size of the backup is over 1 KB.
6. Restore each of the target directories or files from its remote location to its original location.

7. Restore each of the database files by issuing the following commands:

**IMPORTANT**

- Restore the `reporter` database last, regardless of the deployment architecture.
- For best results, restore to a clean installation.

```
mysql -e "drop database DB_Name;"
mysql -e "create database DB_Name;"
mysql -u root dbbackupname.sql
mysql -u root mysql > proc.sql
```

8. Enable the NFA data source by doing the following steps:
- a. In NetOps Portal, hover over **Administration**, **Data Sources**, and then click **Data Sources**.  
The **Manage Data Sources** page appears.
  - b. Select the NFA data source, and then click **Edit**.  
The **Edit Data Source** dialog opens.
  - c. Change the **Status** to "Enabled", and then click **Save**.  
The NFA data source is enabled.

NFA is rehydrated.

## Migrate Virtual Disk Usage Data for SDN Devices

You can migrate virtual usage data for Software-Defined Networking (SDN) devices using the `migrate_sdn_device_metrics.sh` migration script.

If you are upgrading from 21.2.x to 23.3.x and you require contiguous reports (which includes data from before the upgrade) for aggregate virtual disk usage data for SDN devices or if you have applied extensions to the Virtual Disk metric family, including baselines, migrate the data using this procedure.

**IMPORTANT**

Complete this procedure at least 24 hours after upgrading the data aggregator and DX NetOps Virtual Network Assurance (VNA) from 21.2.x to 23.3.x. A successful run of the migration does not print any output.

**Follow these steps:**

1. From a command prompt, export the data aggregator user password by issuing the following command:  
`export DB_PASS=<da_user_password>`  
– **da\_user\_password**  
The data aggregator user password.
2. Validate that the database is in a state that allows the *inactive* table partitions to be migrated by issuing the following command:  
`./migrate_sdn_device_metrics.sh <da_db_name> <da_db_user> <mode> <partition type>  
<data type>`

**Example:**

```
./migrate_sdn_device_metrics.sh <da_db_name> <da_db_user> test inactive all
```

- **da\_db\_name**  
The data aggregator database name.
- **da\_db\_user**  
The data aggregator database user name.
- **mode**  
Defines the action (test, commit, or revert) for the script to take on the virtual disk usage data for SDN devices in the Virtual Disk metric family.

**Options:**

- **test:** Validate that the database is in a state that allows virtual disk usage data for SDN devices to be migrated.
- **commit:** Copy the virtual disk usage data for SDN devices from the Virtual Disk metric family to the SDN Devices Metrics metric family for the table partitions.
- **revert:** Revert the copied virtual disk usage data for SDN devices from the SDN Devices Metrics metric family for the table partitions.

– **partition type**

Defines the type of table partitions to test, commit, or revert.

**Options:**

- **all:** All database files to which the data aggregator writes with incoming virtual disk usage data for SDN devices.
- **active:** The set of database files to which the data aggregator writes with incoming virtual disk usage data for SDN devices.
- **inactive:** The set of database files to which the data aggregator no longer writes with incoming virtual disk usage data for SDN devices.

– **data type**

Defines the type of virtual disk usage data for SDN devices to test, commit, or revert.

**Options:**

- **all:** Test, commit, or revert all virtual disk usage data for SDN devices.
- **aggregate:** Test, commit, or revert only aggregated virtual disk usage data for SDN devices.
- **polled:** Test, commit, or revert only polled virtual disk usage data for SDN devices.

The results show that data can be migrated (it indicates whether the script finishes successfully or not).

**NOTE**

To map columns, the migration script requires that the metric family includes a metric family extension. If the migration script cannot map a column, apply the missing metric family extensions to the metric family, and then run the migration script again.

For more information about how to apply metric family extensions, see [Create or Extend Metric Families](#).

3. Copy the existing virtual disk usage data for SDN devices from the Virtual Disk metric family to the SDN Devices Metrics metric family for all *inactive* table partitions for those columns that can be mapped by issuing the following command:

```
./migrate_sdn_device_metrics.sh <da_db_name> <da_db_user>commit inactive all
```

– **da\_db\_name**

The data aggregator database name.

– **da\_db\_user**

The data aggregator database user name.

**NOTE**

The time that this command takes is comparable to the previous validate the database command (`test`).

4. Validate that the database is in a state that allows *active* aggregate table partitions to be migrated by issuing the following command:

```
./migrate_sdn_device_metrics.sh <da_db_name> <da_db_user> test active aggregate
```

– **da\_db\_name**

The data aggregator database name.

– **da\_db\_user**

The data aggregator database user name.

The results show that the data can be migrated.

5. [Stop the data aggregator](#).

6. Copy the existing aggregate virtual disk usage data for SDN devices from the Virtual Disk metric family to the SDN Devices Metrics metric family for *active* aggregate table partitions by issuing the following command:

```
./migrate_sdn_device_metrics.sh <da_db_name> <da_db_user> commit active aggregate
```

– **da\_db\_name**

The data aggregator database name.

– ***da\_db\_user***

The data aggregator database user name.

**NOTE**

The time that this command takes is comparable to the previous validate the database command (`test`).

7. [Restart the data aggregator](#), and then wait for it to come back online.

**NOTE**

This can take a bit of time for large systems.

8. Validate that those dashboards displaying virtual disk usage data for SDN devices from the SDN Devices Metrics metric family now show historic virtual disk usage from before the upgrade.

9. Clean up existing migration backup tables by issuing the following command:

```
./migrate_sdn_device_metrics.sh <da_db_name> <da_db_user> clean
```

– ***da\_db\_name***

The data aggregator database name.

– ***da\_db\_user***

The data aggregator database user name.

The virtual disk usage data for SDN devices is migrated.

---

# Building

---

Configure your system to collect performance data.

DX NetOps Performance Management provides data for dashboards and views by collecting data from devices in your network. This section includes information about device management, discovery, and self-certification, and other information related to how DX NetOps Performance Management collects data.

## Self-Certification

DX NetOps Performance Management uses metric families and vendor certifications to support devices. These components determine how DX NetOps Performance Management collects configuration and operational metrics for a device.

Out of the box, DX NetOps Performance Management supports the common vendors, metrics, and components in your network infrastructure. For unsupported devices, extend monitoring capabilities using self-certification.

### TIP

You can view a list of the out-of-the-box certifications by the data aggregator version, vendor certification, and metric families using the [DX NetOps Certification Portal](#).

### IMPORTANT

Changes to metric families and vendor certifications apply to all tenants.

In this article:

- [Types of Self-Certification](#)
- [When to Self-Certify](#)
- [Self-Certification Prerequisites](#)
- [The Data Model](#)
- [Example: Support for a Router Device](#)
- [Self-Certification Workflow](#)

### Types of Self-Certification

The following are the available self-certifications methods:

- **Custom Certification**  
Create a custom vendor certification, metric family, or component.
- **Extend**  
Modify an out-of-the-box vendor certification or metric family. Changes to extended vendor certifications or metric families are maintained when you update the original vendor certification or metric family. Extending a vendor certification creates an XML file, but keeps a backup of the out-of-the-box vendor certification.
- **Update**  
Apply changes to a custom certification. Updating a vendor certification completely replaces the existing XML file.

### NOTE

Some DX NetOps Performance Management features, such as configuring percentiles or projections, automatically extend or update the vendor certification.

### TIP

When possible, use an extension instead of a custom vendor certification or metric family. Updating a vendor certification or metric family also updates the extended vendor certifications.

## **When to Self-Certify**

Self-certification applies in the following situations:

- To support an existing metric family on a device that DX NetOps Performance Management does not yet support, create a custom vendor certification.
- To support a new vendor or device, create a custom vendor certification and a metric family.
- To support a new technology in DX NetOps Performance Management, create a custom vendor certification, metric family, or custom component.
- To add a metric to an out-of-box vendor certification, extend the the vendor certification or metric family.
- To change the name of a metric, the calculation of a metric for an out-of-box vendor certification, or to extend the metric family.
- To add custom discovery filtering to an out-of-box vendor certification, extend the vendor certification.
- To change the OID that a device uses, extend the vendor certification.
- To change polled and baseline configuration for metrics in an out-of-box metric family, extend the metric family.

## **Self-Certification Prerequisites**

To ensure a successful self-certification process, verify that you have met the following prerequisites:

- You have access to the management information base (MIB) and objectID (OID) from the device.
- You have a MIB browser, such as the free version of iReasoning.
- You have a test environment for the certifications.

### **NOTE**

You cannot delete metric families and vendor certifications from your system. Test custom certifications before you implement the changes in your production environment.

## **The Data Model**

Understanding the data collection model helps you understand what is required in the self-certification process.

You can configure the following:

- **Discovery profiles**  
Determines which devices, interfaces, and component the data aggregator discovers in your environment, typically based on a range of IP addresses. The discovery process identifies the "type" for each item that it finds.
- **Device collections**  
Organizes your inventory into groups of related items. Items are automatically added to a device collection based on the item type and IP address.
- **Monitoring profiles**  
Controls the polling rate for a device collection, and determines which metric families to poll. Monitoring profiles can poll one or more metric families. To poll the same metric family at different rates, and on different groups, add the same metric family to more than one monitoring profile.

### **NOTE**

To ensure that the system is not overloaded with polling traffic, use monitoring profiles to adjust the polling rate for different sets of metrics.

- **Metric families**  
Controls which metrics are gathered for a monitoring profile. Metric families are associated with one or more vendor certifications, which are listed in priority order.
- **Vendor certifications**  
Maps attributes from a vendor MIB to the metrics in a metric family. Also determines how metrics that are collected from an item are formatted. Metrics that are provided for an item can vary, depending on the item vendor. Mapping these values ensures that the metric values are reported consistently, regardless of the vendor. You can associate multiple vendor certifications with a single metric family. In such cases, the data aggregator maps metric values using



a ranked list of vendor certifications. The data aggregator calculates a metric value using the highest-priority vendor certification that matches the polled item.

### **Example: Support for a Router Device**

When running your discovery profile, the data aggregator finds and identifies an item as a router. The router managed item is automatically added to the factory **All Routers** device collection. This device collection is associated with the router's monitoring profile. This profile uses the CPU and Memory metric families to discover the CPU and Memory components on the device. These metric families also determine the vendor certification to use when calculating the metric values for these components. Based on this monitoring profile, the data aggregator polls the router every 5 minutes. The vendor certifications that are associated with a metric family determine how to calculate and format the raw metric data. DX NetOps Performance Management stores the collected metric data for your router, and uses the data in dashboards and reports.

### **Self-Certification Workflow**

For more information about self-certification workflows, including an overview of common self-certification scenarios, see [Self-Certification Workflows](#).

Complete a self-certification using the following process:

1. (If you do not have existing components) [Create components](#).
2. (If you do not have existing metric families) [Create or extend metric families](#).
3. [Create or extend vendor certifications](#).
4. [Manage the vendor certification priorities](#).

## **Manage Components**

To support a new technology, create a component.

You manage components by creating them and updating them. You create components using the Data Aggregator REST web services. You can create an entry that lists all instances of your component in the **Inventory** menu, and create a context page by creating a component. Use a component to create a REST endpoint or to create attributes that are shared between metric families.

### **NOTE**

You cannot extend out-of-the-box components.

**Prerequisite:** To avoid possible data loss, you have backed up the certification directory.

### **Create a Component XML Template**

Create the component XML template using an existing component.

### **TIP**

Retrieve a list of existing components from the following REST URL:

```
http://da_hostname:8581/typcatalog/components
```

### **Follow these steps:**

1. Set up a REST client with a connection to the data aggregator server.
2. Select a component that is similar to the component that you require.
3. Retrieve the template component by entering the following URL:  

```
http://da_hostname:8581/typcatalog/components/<component_name>
```

  - **component\_name**  
Specifies the name of the template component.
4. On the **Method** tab, select **GET**, and then run the method.

The REST client returns the XML information for the component. Use the XML file as a template to create the component.

### **Edit the Component XML Template**

To apply the necessary changes to the component, edit the XML file (the component XML template). Make changes to the `<ItemSyncDefinition>` section to create an entry that lists all instances of your component in the Inventory menu and to create a Context page for your component.

For more information about the XML structure, see [Component XML Structure](#).

### **Create or Update a Component using the Component XML Template**

**Follow these steps:**

1. Using a REST client, specify `http://da_hostname:8581/typelog/components` as the URL.
2. On the **Method** tab, do one of the following:
  - To create a component, select `POST`.
  - To update a component, select `PUT`.
3. In the **Body** settings, set `application/xml` as the **Body Content-type**.
4. Copy the component XML template into the **Body** tab.
5. Run the method.

The component is created or updated. If no errors occur, the Status field in the HTTP Response section displays the following result:

```
HTTP/1.1 200 OK
```

## **Create or Extend Metric Families**

DX NetOps Performance Management includes metric families. You can create new ones or extend the ones that are included.

If an existing metric family is close to meeting your monitoring needs, you can extend it. If you require something different, create a custom metric family.

For an example, see [Add Metrics to Existing Metric Families](#).

Create or extend a metric family using the following process:

1. [Verify the prerequisites](#).
2. [Create a metric family](#).
3. [Import a metric family](#).
4. [Update the metric family properties](#).
5. [Trigger rediscovery](#).
6. [Verify the metric family results](#).

### **Verify the Prerequisites**

Before you create custom metric families, verify that you have completed the following prerequisites:

- You have downloaded and reviewed the related schema XSD files. The schema is required to validate your XML files. For more information, see [Certification Schema Files and Examples](#).
- To prevent possible data loss, you have backed up the certification directory.

## Create a Metric Family

Create a custom metric family or extended an existing metric family using the metric family XML as a template. Create the XML template using an existing metric family. For extensions, get the XML for the target metric family.

### NOTE

When you create a metric family, DX NetOps Performance Management adds the metrics to the OpenAPI schema.

For more information about this schema, see [OpenAPI](#).

## Create a Metric Family XML Template

### Follow these steps:

1. Set up a REST client with a connection to the data aggregator server.
2. Retrieve the metric family XML template by entering the following URL:

### TIP

You can retrieve the name of the metric family using the user interface (UI). Click the downward arrow on any column, hover over **Columns**, and click **Internal Name**, and then go to the **Metric Families** page.

- **New Metric Family:** `http://da_hostname:8581/typcatalog/metricfamilies/mf_name`
  - **Extension:** `http://da_hostname:8581/typcatalog/metricfamilies/mf_name`  
`mf_name` specifies the name of the metric family template.
3. Select GET in the **Method** tab, and run the method.

The REST client returns the metric family XML.

## Edit the Metric Family XML

You edit metric families by editing the metric family XML file.

For more information about the XML structure, see [Metric Family XML Structure](#).

### NOTE

The units label on reports is **Units** if the `Units` attribute is not defined in the XML.

### NOTE

When extending an existing metric family:

- Do not edit restricted tabs or attributes.  
For more information, see [Restricted XML Tags](#).
- Include only the XML nodes that require changes. When adding metrics to metric families, check the existing collected metrics for those monitoring profiles that are configured to select collected metrics (the profile has assigned metric families), and add the new metrics.  
For more information, see [Add Metrics to Existing Metric Families](#).

## Import a Metric Family

Select one of the following options to import a metric family:

- [Import Using a REST Client](#)
- [Import Using NetOps Portal](#)

### Import Using a REST Client

#### Follow these steps:

1. Specify one of the following URLs:

- **Custom Metric Family:** `http://da_hostname:8581/typecatalog/metricfamilies/`
  - **Extend Metric Family:** `http://da_hostname:8581/typecatalog/metricfamilies/extension/mf_name`
2. On the **Method** tab, Select **POST** for a new custom metric family, or select **PUT** to update or extend a metric family.
  3. In the **Body** settings, select **application/xml** as the Body Content-type.
  4. Copy the metric family XML into the **Body** tab.
  5. Run the method.

The custom metric family is imported. If no errors occur, the Status field in the HTTP Response section displays the following result:

```
HTTP/1.1 200 OK
```

## Import using NetOps Portal

### Follow these steps:

1. Hover over **Administration**, **Data Sources**, and then click the data aggregator data source.
2. Under **Monitoring Configuration**, click **Metric Families**.

#### **TIP**

Locate specific information that is related to a pane using Search. Alternatively, navigate between pages in a pane using the arrows.

3. Click **Import**.
4. Click **Browse**, and then select the metric family file.

#### **NOTE**

You can import ZIP files. For example, you can import the downloaded ZIP file for a certification from the [On-Demand Certification support page](#).

5. Click **Open**, and then click **Import**.

The custom metric family is imported.

## Update the Metric Family Properties

To change attributes, such as the display name, update the metric family properties.

### Follow these steps:

1. Specify the following URL:
  - **Extend Metric Family:** `http://da_hostname:8581/typecatalog/metricfamilies/extension/mf_name`
2. To update or extend a metric family, on the **Method** tab, select **PUT**.
3. In the **Body** settings, select **application/properties** as the **Body Content-type**.
4. Specify updates with one line per property in the **Body** tab as illustrated in the following template and example.

#### **Template:**

#### **NOTE**

The `metricfamilyname` and `attributename` variables in the following example are case-sensitive and are lowercase. These variables are a different case from their values in the XML. For example, `NormalizedPortInfo` from the XML is `normalizedportinfo` here and `PctDiscardsIn` in the XML is `pctdiscardsin` here.

```
im.ca.com.normalizer.metricfamilyname.displayname=DisplayName
im.ca.com.normalizer.metricfamilyname.documentation=Documentation text
im.ca.com.normalizer.metricfamilyname.attribute.attributename.attributedisplayname=Attribute
```

```
im.ca.com.normalizer.metricfamilyname.attribute.attributename.documentation=Documentation
text
```

**Example:**

```
im.ca.com.normalizer.normalizedportinfo.displayname=Interface
im.ca.com.normalizer.normalizedportinfo.documentation=Defines the identification
information, configuration information, and polled metrics for interfaces.
im.ca.com.normalizer.normalizedportinfo.attribute.pctdiscardsin.attributedisplayname=Percent
Discards In
im.ca.com.normalizer.normalizedportinfo.attribute.pctdiscardsin.documentation=The
percentage of the frames (packets) received by the interface that were discarded.
```

**5. Run the method.**

The metric family properties are updated.

**Trigger Rediscovery**

After you extend a metric family, the changes occur during the nightly automatic rediscovery. If the changes do not apply automatically, trigger the update manually.

**TIP**

Prevent severe impacts on performance by triggering rediscoveries *after* normal business hours.

**Follow these steps:**

1. In NetOps Portal, from the list of metric families, select the metric family.
2. Click **Update Metric Family**.

The data aggregator rediscovers components on all devices that support the selected metric family.

**Verify the Metric Family Results**

To ensure successful operation, verify the results of the import. In NetOps Portal, from the list of metric families, verify that the metric family is in the list and that the **Last Modified** time has been updated.

## Create or Extend Vendor Certifications

Create the vendor certifications for the vendors or devices that you want to support. To change the way that a metric is calculated, or to add custom discovery filtering for an out-of-the-box vendor certification, extend the vendor certification.

You can also remove vendor certifications.

**IMPORTANT**

You cannot delete metric families and vendor certifications.

Use the following process to create or extend vendor certifications:

- [Verify the Prerequisites](#)
- [Get a Vendor Certification XML Template](#)
- [Extend the Vendor Certification](#)
- [Verify the Vendor Certification](#)
- [Import the Custom Vendor Certification](#)
- [Verify the Vendor Certification Results](#)

## Verify the Prerequisites

Before you create or extend vendor certifications, verify that you have completed the following prerequisites:

- You have downloaded and reviewed the related schema XSD files. The schema is required to validate your XML files. For more information, see [Certification Schema Files and Examples](#).
- To prevent possible data loss, you have backed up the certification directory.

## Get a Vendor Certification XML Template

You can use an XML as a template to create the custom, or extended, vendor certification.

### Follow these steps:

1. Set up a REST client with a connection to the data aggregator server.
2. Retrieve the template vendor certification by entering the following URL:

#### – New vendor certification

Create an XML template using an existing vendor certification for a similar device.

#### TIP

To find a suitable vendor certification to use as a template, look at the metric family that you want to support with the new certification, and then pick one that is similar to your device.

```
http://da_hostname:8581/typecatalog/certifications/snmp/<cert_name>
```

#### – Extension

Get the XML for the target vendor certification.

```
http://da_hostname:8581/typecatalog/certifications/snmp/extension/<cert_name>
```

#### NOTE

The **<cert\_name>** specifies the name of the vendor certification, which is an attribute of the `FacetType` tag.

3. Select GET in the **Method** tab, and then run the method.

The REST client returns the XML information that you can use as a template to extend the vendor certification.

## Extend the Vendor Certification

Extend the vendor certification by editing the XML file or using NetOps Portal. You can also edit existing custom vendor certifications when you want to collect additional data for reporting. Include only the XML nodes that require changes.

For more information about the XML structure, see [Vendor Certification XML Structure](#).

### IMPORTANT

Do not edit restricted XML tags or attributes.

For more information, see [Restricted XML Tags](#).

### TIP

Use NetOps Portal to edit vendor certifications if you have limited changes that you need immediately. For example, you can edit the expression that maps to the normalized metric family variables.

For more information, see [Edit Custom Vendor Certifications Using NetOps Portal](#).

You can enable the `Bytes`, `Bytes In`, and `Bytes Out` metrics for individual vendor certifications by extending the vendor certification. The example XML to enable these metrics is included in the `/opt/IMDataAggregator/examples/vendorCertification/` directory.

### WARNING

Enabling these metrics for the Interface metric family can affect system performance.

For more information about this metric family, see [Add Metrics to Existing Metric Families](#).

Merge the example XML with your extension XML.

## Verify the Vendor Certification

Verify the vendor certification in a test environment before you import it to your production environment. You cannot delete vendor certifications.

## Import the Custom Vendor Certification

Import the vendor certification to the system.

Use *one* of the following methods to import the vendor certification:

- [Import the Custom Vendor Certification using a REST Client](#)
- [Import the Custom Vendor Certification using NetOps Portal](#)

## Import the Custom Vendor Certification using a REST Client

Follow these steps:

1. Specify the following URL:  
`http://da-hostname:8581/typecatalog/certifications/snmp`
2. On the **Method** tab, select **POST**.
3. In the **Body** settings, select **application/xml** as the **Body Content-type**.

### NOTE

Set the Content-type to avoid 415 errors.

4. Copy the vendor certification XML into the **Body** tab.
5. Run the method.

Your vendor certification is imported or updated. If no errors occur, the **Status** field in the **HTTP Response** section displays the following result:

```
HTTP/1.1 200 OK
```

### NOTE

To import an unchanged vendor certification extension template, use a file that is similar to the following template:

```
<?xml version="1.0" encoding="UTF-8"?>
<DataModel namespace="http://im.ca.com/certifications/snmp" xmlns:xsi="http://www.w3.org/2001/
XMLSchema-instance" xsi:noNamespaceSchemaLocation="SNMPCertificationFacet.xsd">
  <Author>CA</Author>
  <Version>2.02</Version>
  <FacetType name="IfXTableMib"
descriptorClass="com.ca.im.core.datamodel.certs.CertificationFacetDescriptorImpl">
    <FacetOf namespace="http://im.ca.com/core" name="Item" />
  </FacetType>
</DataModel>
```

## Import the Custom Vendor Certification using NetOps Portal

After you create a custom vendor certification, you can share it with other data aggregator users who want to gather metrics for the same vendor by importing it using NetOps Portal. You can also remove a vendor certification extension by importing an *unchanged* vendor certification extension template.

You can share custom vendor certifications between data aggregator users. Users can gather metrics for vendor devices when a factory certification is not yet available using custom vendor certifications.

**Prerequisite:** The metric family that is associated with the vendor certification is available in NetOps Portal.

**Follow these steps:**

1. Log in to NetOps Portal.
2. Hover over **Administration**, **Data Sources**, and then click the data aggregator data source.  
The **Monitored Devices** page appears.
3. In the **Monitoring Configuration** section, click **Vendor Certifications**.  
The **Vendor Certifications** page appears.
4. Click **Import**.  
The **Import Vendor Certification** window opens.

**NOTE**

- To use a shared custom vendor certification, import it in XML format into your installation of data aggregator. You are not required to import the MIB.
- You can import ZIP files. For example, you can import the downloaded ZIP file for a certification from the [On-Demand Certification support page](#).

5. Browse for, select the custom vendor certification file that you want to import, and then click **Open**.
6. Click **Import**.

The custom vendor certification is imported. The data aggregator immediately begins collecting metrics for the metric families that are associated with the newly imported custom vendor certification.

**Verify the Vendor Certification Results**

To ensure successful operation, verify the results of the import. Test custom certifications before you implement the changes in your production environment.

**Follow these steps:**

1. On the **Vendor Certifications** page, verify that the vendor certification appears in the list, and that the **Last Modified** time has been updated.
2. From the **Monitoring Configuration** section, click **Metric Families**.  
The **Metric Families** page appears.
3. Verify that the metric family appears in the list.
4. Click the **Vendor Certification Priorities** tab on the right side of the page.  
The vendor certifications that you create (custom vendor certifications) are automatically added to the bottom of the priority list for the specified metric family.
5. If necessary, modify the priority list and move the vendor certification to a higher priority.

**Edit Custom Vendor Certifications Using NetOps Portal**

Edit existing custom vendor certifications when you want to collect additional data for reporting. You can edit these certifications by editing the XML or by using NetOps Portal. This article explains how to edit vendor certifications using NetOps Portal. Edit the certification using NetOps Portal if you have limited changes that you need immediately. For example, you can edit the expression that maps to the normalized metric family variables.

For more information about how to edit vendor certifications by editing the XML, see [Create or Extend Vendor Certifications](#).

Common edits to vendor certification expressions include the following:

- Change a value that is assigned to the expression, such as an average.
- Add multiple vendor certification variables to the expression.
- Remove an expression from a metric family variable.



**NOTE**

The `Names` and `Indexes` metric names require an expression. For custom vendor certifications, the data aggregator automatically provides the `Indexes` metric expression for the selected MIB table.

The edits that you make to vendor certifications applies only to the following:

- Those certifications that you created or edited using NetOps Portal.
- Devices with a single MIB and a single MIB table.
- Some aspects of the vendor certification.

**IMPORTANT**

To avoid possible data loss, back up the certifications directory each time that you create or update a vendor certification, metric family, or component.

**Follow these steps:**

1. As an Administrator, hover over **Administration**, **Data Sources**, and then select the data aggregator data source. The **Monitored Devices** page appears.
2. Click **Vendor Certifications** from the **Monitoring Configuration** menu for a data aggregator data source. The **Vendor Certifications** page appears.
3. Select the custom vendor certification that you want to edit from the list, select the metric family that you want to edit, and then click **Edit**.
4. Manually edit the expression, and then click **Accept Expression**. The expression is mapped, and the top table is populated with the updated values.
5. Click **Save**.

The vendor certification details grid is updated with the changes to the metric family variables.

## Manage Vendor Certification Priorities

You can modify the priority list and move vendor certifications to a higher priority, and group vendor certification priorities.

Metric families for devices can support one or more vendor certifications. If the metric family supports more than one vendor certification, the vendor certification with the highest priority is selected as the backing vendor certification. As you add custom vendor certifications, they are automatically added to the bottom of the priority list for the specified metric family. If necessary, you can modify the priority list and move the vendor certification to a higher priority. To use more than one vendor certification for the same device, group the vendor certification priorities.

### Prioritize the Vendor Certification Within the Metric Family

Changing the priority of the vendor certifications for a metric family updates the metric family on all affected devices. The change generates an event on the metric family, indicating that the vendor certification priority on the metric family has changed. If the backing vendor certification changes on a device, an event is generated on the device. The event indicates that a vendor certification has changed. A second event is also generated, indicating the specific changes.

**NOTE**

To take advantage of new vendor certifications that are part of an installation upgrade, manually change the vendor certification priorities. For example, F5 CPU vendor certifications are modeled as normal CPUs but are not discovered because F5 also supports Host Resources. After an upgrade, the Host Resources CPU priority entry will be higher than the F5 entries appended to the end of the priority list. To discover F5 CPU devices and components, update the vendor certification priority for the CPU metric family. A fresh installation does not have this issue.

**Follow these steps:**

1. Navigate to the data aggregator data source in NetOps Portal. The **Monitored Devices** page appears.

2. Under **Monitoring Configuration**, click **Metric Families**.  
The **Metric Families** page appears.
3. Select a metric family, and then click the **Vendor Certification Priorities** tab in the pane on the right.  
The list of prioritized vendor certifications appears.
4. Click **Manage**.  
The **Manage Vendor Certification Priority** window opens.
5. Arrange the priority order as necessary using the **Move Up** and **Move Down** arrows, and then click **Save**.  
The metric family uses the new priority to determine which vendor certification to use for monitored devices.

### **Group Vendor Certification Priorities**

You can use more than one vendor certification for the same device by grouping the vendor certification priorities. A single vendor certification priority can belong to as many priority groups as necessary for your application. When a device is discovered, the device supports the vendor certifications in the priority group.

#### **IMPORTANT**

To prevent disrupting automatic rediscovery, use separate REST calls to update the `<PriorityGroup>` tag and to change the order of the vendor certifications (`<CertificationOrder>`).

For more information about this issue, see [Automatic Rediscovery Does Not Run After Updating Vendor Group Priority](#).

You group vendor certification priorities using only the data aggregator REST web services.

#### **Follow these steps:**

1. Run a GET operation on the following REST URL:  
`http://da_host:8581/rest/vendorpriorities/`  
The REST call returns a list of all the vendor certification priorities.
2. Determine the metric family ID in the list of vendor certification priorities.
3. Retrieve the vendor certification priority (the XML) that you want to modify by running a GET REST call on the following REST URL:  
`http://da_host:8581/rest/vendorpriorities/ID`  
The XML file is retrieved.
4. Add one or more vendor certifications to a priority group using the `<PriorityGroup>` tag in the vendor certification priority XML, placing them in the priority order:

#### **Example:**

In this example, the `CiscoAirespaceWirelessRADIUSServer` priority group is used to group the `CiscoAirespaceWirelessRADIUSAccServerMib` and `CiscoAirespaceWirelessRADIUSAuthServerMib` vendor certifications:

```
<CertificationOrder>
  <CollectionID>3232</CollectionID>
  <VendorCertID>{http://im.ca.com/certifications/
snmp}CiscoAirespaceWirelessRADIUSAccServerMib</VendorCertID>
  <PriorityGroup>CiscoAirespaceWirelessRADIUSServer</PriorityGroup>
</CertificationOrder>
<CertificationOrder>
  <CollectionID>3232</CollectionID>
  <VendorCertID>{http://im.ca.com/certifications/
snmp}CiscoAirespaceWirelessRADIUSAuthServerMib</VendorCertID>
  <PriorityGroup>CiscoAirespaceWirelessRADIUSServer</PriorityGroup>
</CertificationOrder>
```

– **VendorCertID**

The ID for the vendor certification.

**IMPORTANT**

To prevent vendor certification priority failures, the XML for the `<VendorCertID>` tag must be written on a single line, and must not include spaces and carriage returns.

– **PriorityGroup**

The priority group to which you are adding the vendor certification. (Optional) To add a single vendor certification priority to multiple groups, separate the group names using commas.

**Example:**

```
<PriorityGroup>F5, Huawei</PriorityGroup>
```

**TIP**

You can enable recommended priority groups by replacing the `<RecommendedPriorityGroup>` tag for the `<PriorityGroup>` tag.

5. Remove the `<ID>` and `<MetricFamilyID>` tags from the vendor certification priority XML.

6. Save the XML file.

7. Import the updated XML file by running a PUT REST call on the following REST URL:

```
http://da_host:8581/rest/vendorpriorities/ID
```

– **ID**

Specifies the ID of the vendor certification priority.

DX NetOps Performance Management rediscovers all devices that support the metric family.

## Verify Vendor Certification Priority Grouping

Verify that the vendor certification priorities are grouped as required.

**Follow these steps:**

1. Navigate to the data aggregator data source in NetOps Portal.

The **Monitored Devices** page appears.

2. Select a device from the tree view, and then click the **Polled Metric Families** tab in the pane on the right.

The vendor certifications that you grouped under the corresponding metric families appear on multiple rows in the **Vendor Cert** column.

## Create or Edit Vendor Certification Expressions

To modify the mapping of normalized metric family variables, edit the vendor certification expressions.

Consider the following information when you edit expressions:

- Use delimiters to separate vendor certification variables.
- MVFLEX Expression Language (MVEL) functions and custom functions are valid.
- The Names and Indexes metric names require an expression. The remaining metric names are optional.

**NOTE**

For vendor certifications that are created using the Vendor Certification wizard, DX NetOps Performance Management automatically provides the Indexes metric expression for your selected MIB table.

### Example: Change the Value for Averages

Cisco router CPU statistics are mapped to the normalized variable 'CPU Utilization', as shown in the following expression:

```
cpmCPULoadAvg5min+cpmCPUUseravg5min
```

You can change the 5-minute average to a 1-minute average by editing the expression as follows:

```
cpmCPULoadAvg1min+cpmCPUUseravg1min
```

### Example: Use an Advanced Expression

The following expression verifies whether `hrStorageSize < 0`, and returns the value of `hrStorageSize` converted to an Unsigned value, and multiplied by 100. Otherwise, the expression returns the following value: `hrStorageUsed/hrStorageSize * 100`.

```
(hrStorageSize < 0) ? (hrStorageUsed/
convertSignedIntToUnsignedDecimal(hrStorageSize)) * 100 : hrStorageUsed/hrStorageSize *
100
```

## Functions and Global Variables

MVEL is the language that vendor certifications use to manipulate data from monitored devices. In vendor certification expressions, use MVEL to normalize data and perform calculations.

MVEL is a publicly available Expression Language for Java environments that can be embedded. MVEL supports expressions similar to Java expressions. You can build expressions using operators, use braces to control precedence, and terminate statements using semi-colons. For a detailed reference of the MVEL language, see the [MVEL 2.0 Wiki](#).

### Performing Calculations with MVEL

If your expression contains a calculation without at least one attribute variable, you might see unexpected results. Include at least one attribute in your expression that does not affect the calculation result. For example, the following expression might return unexpected results:

```
<Expression destAttr="metric1">15 * 15</Expression>
```

However, the following expression does not affect the result and does return the expected result:

```
<Expression destAttr="metric1">Index; 15 * 15</Expression>
```

- **Index**

Specify an attribute variable that is defined in the **Attribute** or **AttributeGroup**.

The following certification shows the proper usage of MVEL functions for vendor certifications: [MVEL Test Certification](#). The **assert** tag is used for internal testing purposes only, and is not used in creating vendor certifications.

#### IMPORTANT

To ensure that you do not lose all the data from an expression that uses an assert tag, add a space before the semicolon at the end of the expression.

The following custom functions and global variables are available for use in vendor certification expressions:

### availabilityWithSysUptime Function

This function calculates availability as a percentage using `sysUptime` and the poll duration, granting a grace period.

#### Syntax

This function has the following format:

```
Object availabilityWithSysUptime (Long sysUpTime, Long duration)
```

**Parameters**

- **sysUpTime**  
The time (in centiseconds) since the network management portion of the system was last reinitialized.
- **duration**  
The poll duration time in seconds. Use the global variable `_rspDuration`. See the Advanced example for more information.

**Return Values**

Returns the availability as a percentage (0 - 100), or returns "null" when invalid data is passed.

**Examples**

The following expression produces the following result for a sysUpTime of 30000 and a poll duration of 300:

**Expression:**

```
availabilityWithSysUptime (sysUpTime, duration)
```

**Result:**

```
100
```

The same expression produces the following result for a sysUpTime of 6000 and a poll duration of 300:

**Result:**

```
20
```

The same expression produces the following result for a sysUpTime of 30005 and a poll duration of 300:

**Result:**

```
100
```

**Advanced Example**

The following expression is taken from "System Statistics" Vendor Certification:

```
Availability=availabilityWithSysUptime(sysUpTime,_rspDuration)
```

**mapModel Function**

This function uses the value of an objectID (sysObjectID) and maps the system OID to a model name string. Use this function to certify devices.

**Syntax**

This function has the following format:

```
String mapModel ( ObjectID sysObjectID )
```

**Parameters**

- **sysObjectID**  
The object ID value to parse.

**Return Values**

Returns the string containing the mapped model name.

**Examples**

The following expression produces the following result for an OID value of 1.3.6.1.4.1.9.1.223:

**Expression:**

```
mapModel (oid )
```

**Result:**

```
Cisco7204VXR
```

The same expression produces the following result for an OID value of 1.3.6.1.4.1:

**Result:**

```
Unknown 1.3.6.1.4.1
```

**Advanced Example**

The following expression is taken from “System Statistics” Vendor Certification:

```
Model=mapModel (sysObjectID)
```

**mapVendor Function**

This function uses the value of an objectID (sysObjectID) and maps the system OID to a vendor name string. Use this function to find the vendor of a device.

**Syntax**

This function has the following format:

```
String mapVendor( ObjectID sysObjectID )
```

**Parameters**

- **sysObjectID**  
The object ID value to parse.

**Return Values**

Returns the string containing the mapped vendor name. If a vendor is not found, returns "".

**Examples**

The following expression produces the following result for an OID value of 1.3.6.2.1.2.2636.0:

**Expression:**

```
mapVendor (oid )
```

**Result:**

```
Juniper
```

The same expression produces the following result for an OID value of 1.3.6.2.1.2.1234567.0:

**Result:**

```
Unknown
```

**Advanced Example**

The following expression is taken from “System Statistics” Vendor Certification:

```
Model=mapVendor (sysObjectID)
```

### **mvelInfo Function**

This function populates the INFO level of the karaf.log file with the input parameters. Use this function to log the polled values from a report that returns a result that you believe is incorrect.

The poll log of the Data Collector only shows the values of the polled attributes, while the report only shows the result of the calculation. The mvelInfo function allows you to view the input poll values to determine where the calculation went wrong.

#### **Syntax**

This function has the following format:

```
String mvelInfo (Array objects)
```

#### **Parameters**

- **objects**  
The object array is logged under the INFO level in the karaf.log file of the Data Collector.

#### **Return Values**

Null

#### **Examples**

The following expression logs cpmCPUTotal5minRev in the karaf.log file.

#### **Expression:**

```
mvelInfo({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev})
```

#### **Result:**

Null

#### **Result (karaf.log):**

MVEL info: cpmCPUTotal5minRev=15

#### **Advanced Example**

```
mvelInfo({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev, " cpmCPUTotal10minRev=",  
cpmCPUTotal10minRev}); cpmCPUTotal10minRev;
```

#### **Result:**

12

#### **Result (karaf.log):**

MVEL info: cpmCPUTotal5minRev=15 cpmCPUTotal10minRev=12

### **mvelWarn Function**

This function populates the WARN level of the karaf.log file with the input parameters. Use this function to log the polled values from a report that returns a result that you believe is incorrect.

The poll log of the Data Collector only shows the values of the polled attributes, while the report only shows the result of the calculation. The mvelWarn function allows you to view the input poll values to determine where the calculation went wrong.

#### **Syntax**

This function has the following format:

---

```
String mvelWarn (Array objects)
```

### Parameters

- **objects**  
The object array is logged under the WARN level in the karaf.log file of the Data Collector.

### Return Values

Null

### Examples

The following expression logs cpmCPUTotal5minRev in the karaf.log file.

#### Expression:

```
mvelWarn({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev})
```

#### Result:

Null

#### Result (karaf.log):

MVEL warn: cpmCPUTotal5minRev=15

#### Advanced Example

```
mvelWarn({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev, " cpmCPUTotal10minRev=",  
cpmCPUTotal10minRev}); cpmCPUTotal10minRev;
```

#### Result:

12

#### Result (karaf.log):

MVEL warn: cpmCPUTotal5minRev=15 cpmCPUTotal10minRev=12

### mvelError Function

This function populates the ERROR level of the karaf.log file with the input parameters. Use this function to log the polled values from a report that returns a result that you believe is incorrect.

The poll log of the Data Collector only shows the values of the polled attributes, while the report only shows the result of the calculation. The mvelError function allows you to view the input poll values to determine where the calculation went wrong.

### Syntax

This function has the following format:

```
String mvelError (Array objects)
```

### Parameters

- **objects**  
The object array is logged under the ERROR level in the karaf.log file of the Data Collector.

### Return Values

Null

### Examples



The following expression logs cpmCPUTotal5minRev in the karaf.log file.

**Expression:**

```
mvelError({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev})
```

**Result:**

Null

**Result (karaf.log):**

MVEL error: cpmCPUTotal5minRev=15

**Advanced Example**

```
mvelError({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev, " cpmCPUTotal10minRev=",  
cpmCPUTotal10minRev}); cpmCPUTotal10minRev;
```

**Result:**

12

**Result (karaf.log):**

MVEL error: cpmCPUTotal5minRev=15 cpmCPUTotal10minRev=12

**mvelDebug Function**

This function populates the DEBUG level of the karaf.log file with the input parameters. Use this function to log the polled values from a report that returns a result that you believe is incorrect.

The poll log of the Data Collector only shows the values of the polled attributes, while the report only shows the result of the calculation. The mvelDebug function allows you to view the input poll values to determine where the calculation went wrong.

**Syntax**

This function has the following format:

```
String mvelDebug (Array objects)
```

**Parameters**

- **objects**  
The object array is logged under the DEBUG level in the karaf.log file of the Data Collector.

**Return Values**

Null

**Examples**

The following expression logs cpmCPUTotal5minRev in the karaf.log file.

**Expression:**

```
mvelDebug({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev})
```

**Result:**

Null

**Result (karaf.log):**

MVEL debug: cpmCPUTotal5minRev=15

## Advanced Example

```
mvelDebug({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev, " cpmCPUTotal10minRev=",
cpmCPUTotal10minRev}); cpmCPUTotal10minRev;
```

### Result:

12

### Result (karaf.log):

MVEL debug: cpmCPUTotal5minRev=15 cpmCPUTotal10minRev=12

## mvelTrace Function

This function populates the TRACE level of the karaf.log file with the input parameters. Use this function to log the polled values from a report that returns a result that you believe is incorrect.

The poll log of the Data Collector only shows the values of the polled attributes, while the report only shows the result of the calculation. The mvelTrace function allows you to view the input poll values to determine where the calculation went wrong.

### Syntax

This function has the following format:

```
String mvelTrace (Array objects)
```

### Parameters

- **objects**  
The object array is logged under the TRACE level in the karaf.log file of the Data Collector.

### Return Values

Null

### Examples

The following expression logs cpmCPUTotal5minRev in the karaf.log file.

### Expression:

```
mvelTrace({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev})
```

### Result:

Null

### Result (karaf.log):

MVEL trace: cpmCPUTotal5minRev=15

## Advanced Example

```
mvelTrace({"cpmCPUTotal5minRev=", cpmCPUTotal5minRev, " cpmCPUTotal10minRev=",
cpmCPUTotal10minRev}); cpmCPUTotal10minRev;
```

### Result:

12

### Result (karaf.log):

MVEL trace: cpmCPUTotal5minRev=15 cpmCPUTotal10minRev=12

### **snmpConstArrayMap Function**

This function maps a value (index) to a set of constant values (array). If necessary, this function rounds the input value to nearest integer value. Then, it uses the integer value as an index to the set of constant values (array) that are shown as c0, c1, up to cn-1. The c values must be integers. This function checks these values when the expression is parsed and returns cx. If the value is not in the domain from 0 to n-1 (inclusive), the result is 0 (without an error message). Use this function to certify devices.

#### **Syntax**

This function has the following format:

```
Integer snmpConstArrayMap(Double index, Integer[] array)
```

#### **Parameters**

- **index**  
A Double value, which is used as an index into the array.
- **array**  
Any range of integer values.

#### **Return Values**

Returns an integer value from the array. An index value of null returns null.

#### **Examples**

The following expression produces the following result for an index of 2 and an array of {5, 6, 7, 8, 9, 4}:

##### **Expression:**

```
snmpConstArrayMap (index, array)
```

##### **Result:**

```
7
```

The following expression produces the following result for an index of 4.88 and an array of {5, 6, 7, 8, 9, 4}:

##### **Expression:**

```
snmpConstArrayMap (value, array)
```

##### **Result:**

```
4
```

#### **Advanced Example**

The following expression is taken from “Generic Modem” Vendor Certification:

```
SpeedOut=snmpConstArrayMap(mdmCsFinalTxLinkRate,
{0,110,300,600,1200,2400,4800,7200,9600,12000,14000,16000,19000,38000,75,450,0,57000,21000,24000})
```

### **snmpCounter64 Function**

This function evaluates two 32-bit numeric values and returns a value containing the 64-bit representation. Use this function to certify devices. The hiVal is shifted 32 bits left and the lowVal is added and the result is placed in a 64-bit return variable.

#### **Syntax**

This function has the following format:

---

```
Object snmpCounter64 (Long hiVal, Long lowVal)
```

### Parameters

- **hiVal**  
The 32-bit numeric value representing the high-order bits.
- **lowVal**  
The 32-bit numeric value representing the low-order bits.

### Return Values

Returns the 64-bit representation of the two 32-bit numeric values, or returns "null" when either 32-bit value input is null.

### Examples

The following expression produces the following result for a hiVal of 88 and a lowVal of 558.

#### Expression:

```
snmpCounter64 (hiVal, lowVal)
```

#### Result:

```
377957122606
```

### Advanced Example

The following expression is taken from "Cisco CBQos ClassMap" Vendor Certification. This certification contains many snmpMax examples:

```
PrePolicyPackets=snmpMax(0, snmpCounter64 (cbQosCMPrePolicyPktOverflow, cbQosCMPrePolicyPkt))
```

### snmpGet

This function gets the value of an ObjectID on the polled device.

#### NOTE

This function issues extra SNMP requests. To avoid a negative performance impact, certifications that include this expression are moved to a potentially slower processing queue.

### Syntax

This function has the following format:

```
<GetResponse> snmpGet (OID<string>)
```

### Parameters

- **OID** The ObjectID of the polled device.

### Functions

- **getError()** Returns the response error code. **Example:** SUCCESS
- **getIp()** Returns the IP address of the device.
- **getOid()** Returns the requested ObjectID.
- **getResult()** Returns the SNMP response object.
- **getValue()** Returns the value from the SNMP response object.
- **getType()** Returns the type of the SNMP response object, such as COUNTER32 or GAUGE.

### Return Values

Returns the response object of an SNMP GET request.

### Example

The following expression includes sample results for each function:

```
response = snmpGet('1.3.6.1.2.1.3.0');
response.getError() --- SUCCESS
response.getIp() --- 10.42.96.5
response.getOid() --- 1.3.6.1.2.1.3.0
value = response.getResult();
value.getType() --- TIMETICKS
value.getValue() --- 31528546
```

The following expression does a null check (?.) to avoid breaking the logic when the result is null:

```
response = snmpGet('1.3.6.1.2.1.3.0');
response.error --- SUCCESS
response.ip --- 10.42.96.5
response.oid --- 1.3.6.1.2.1.3.0
response.?result.type ---- TIMETICKS
response.?result.value ---- 31528546
```

### snmpGetTable

This function gets the values of all ObjectIDs that share the prefix ObjectID.

#### NOTE

This function issues extra SNMP requests. To avoid a negative performance impact, certifications that include this expression are moved to a potentially slower processing queue.

### Syntax

This function has the following format:

```
<GetBulkResponse> snmpGetTable (OID<string>)
```

### Parameters

- **OID**  
The ObjectID of the polled device.

### Functions

- **getError()** Returns the response error code. **Example:** SUCCESS
- **getIp()** Returns the IP address of the device.
- **getOid()** Returns the requested ObjectID.
- **getResult()** Returns a list of SNMP response objects.
- **getValue()** Returns the value from the SNMP response object.
- **getType()** Returns the type of the SNMP response object, such as COUNTER32 or GAUGE.
- **getIndex()** Returns the SNMP instance index of the SNMP response object.

### Return Values

Returns a table that contains a list of SNMP response objects.

### Examples

The following expression includes sample results for each function:

```
response = snmpGetTable('1.3.6.1.2.1.2.2.1.10');
response.getError() --- SUCCESS
response.getIp() --- 10.42.96.5
response.getOid() --- 1.3.6.1.2.1.2.2.1.10
list = response.getResult();
list.size() --- 200
list.get(0).getType() --- COUNTER
list.get(0).getValue() --- 1849
list.get(0).getIndex() --- 1
item = list.get(199);
item.getType() --- COUNTER
item.getResult() --- 1855
item.getIndex() --- 200
for(int i =0; i < list.size(); i++){
    instance = list.get(i);
    instanceType = instance.getType();
    instanceValue = instance.getValue();
    instanceIndex = instance.getIndex();
}
```

#### NOTE

The `instanceType = list.get(i).getType()` method chain is unsupported. The result of `list.get(i)` must be stored in a variable first. Use the following expression instead:

```
instance = list.get(i);
instanceType = instance.getType();
```

### **snmpGetUpSinceTime Function**

This function returns the time that the system was turned on based on the number of seconds since the current epoch.

#### **Syntax**

This function has the following format:

```
snmpGetUpSinceTime(Long upTime)
```

#### **Parameters**

- **upTime**  
The number of seconds since the beginning of the current epoch. You can get the system uptime from the following OID: 1.3.6.1.2.1.1.3.0. Convert it into centiseconds before you pass it in.

#### **Return Values**

Returns the time that the device was powered on in the form of total seconds since the current epoch.

### **snmpMax Function**

This function returns the larger of two 64-bit values. Use this function to certify devices.

#### **Syntax**

This function has the following format:

```
Object snmpMax(BigInteger val1, BigInteger val2)
```

#### Parameters

- **val1**  
The first 64-bit BigInteger value.
- **val2**  
The second 64-bit BigInteger value.

#### Return Values

Returns the maximum of the two BigInteger values that are passed in, or returns "null" when either BigInteger value input is null.

#### Examples

The following expression produces the following result for a val1 of 2^32 and a val2 of 10:

#### Expression:

```
snmpMax (val1, val2)
```

#### Result:

```
2^32
```

The same expression produces the following result for a val1 of 5864 and a val2 of 134556890:

#### Result:

```
134556890
```

#### Advanced Example

The following expression is taken from "Cisco CBQos ClassMap" Vendor Certification. This certification contains many snmpMax examples:

```
PrePolicyPackets=snmpMax(0, snmpCounter64(cbQosCMPrePolicyPktOverflow, cbQosCMPrePolicyPkt))
```

#### snmpObjectIDToASCIIString Function

This function converts an SNMP OID value to its string representation. Any leading or trailing spaces are removed.

#### Syntax

This function has the following format:

```
snmpObjectIDToASCIIString( Object Id oid )
```

#### Parameters

- **oid**  
The object ID to convert to a string.

#### snmpOIDParser Function

This function uses the value of an objectID (OID) and parses out a subset of the OID based on the startIndex and endIndex values. The indexes are 1 based. If the endIndex is -1, then we go to the end of the OID. Use this function to certify devices.

#### Syntax

This function has the following format:

```
ObjectID snmpOIDParser( ObjectID OID, Integer startIndex, Integer endIndex )
```

#### Parameters

- **OID**  
The object ID (OID) value to parse.
- **startIndex**  
An integer value of the index at which to begin parsing.
- **endIndex**  
An integer value of the index at which to stop parsing.

#### Return Values

Returns the parsed subset ObjectID (OID).

#### Examples

The following expression produces the following result for an OID value of 1.2.3.4.5.6.7.8.9.10, a startIndex value of 1, and an endIndex value of 5:

#### Expression:

```
snmpOIDParser(oid, startIndex, endIndex )
```

#### Result:

```
1.2.3.4.5
```

The same expression produces the following result for an OID value of 1.2.3.4.5.6.7.8.9.10, a startIndex value of 6, and an endIndex value of -1:

#### Result:

```
6.7.8.9.10
```

#### Advanced Example

The following expression is taken from “Cisco CBQos ClassMap” Vendor Certification:

```
ItemUniqueIDs=snmpOIDParser(cbQosConfigIndex, 2, 2)
```

#### snmpOctetStringFloat Function

This function converts an SNMP octet string to a floating-point value. Use this function to certify devices. An SNMP octet string is a seven-bit ASCII string.

#### Syntax

This function has the following format:

```
Object snmpOctetStringFloat(byte[] octetString)
```

#### Parameters

- **octetString**  
The SNMP octet string.

#### Return Values

Returns the converted string value, or returns "null" when the function cannot convert the string.



## Examples

The following expression produces the following result for an `octetString` of {0x33, 0x33, 0x2E, 0x33, 0x33}:

### Expression:

```
snmpOctetStringFloat (octetString)
```

### Result:

```
33.33
```

The same expression produces the following result for an `octetString` of {0x36, 0x36, 0x36}:

### Result:

```
666.0
```

## snmpProtectedDiv Function

This function divides two `Double` values and returns the result of the division as a `Double`. If the dividend or divisor is null or 0.0 the return value is 0.0. Use this function to protect the expression from dividing with null or 0. Data Repository can contain null or zero values, such as when a poll fails. In this case, you avoid a divide-by-zero exception by using this function.

### Syntax

This function has the following format:

```
Double snmpProtectedDiv(Double val1, Double val2)
```

### Parameters

- **val1**  
The dividend, which is a `Double` value (floating number) to be divided by `val2`. (`Double` is a Java data type.)
- **val2**  
The divisor, which is a `Double` value (floating number). (`Double` is a Java data type.)

### Return Values

Returns the result of the division as a `Double` or 0.0 if the dividend or divisor is null or 0.0 (*Double* is a Java data type).

## Examples

The following expression produces the following result for a `val1` of 7.2 and `val2` of 2:

### Expression:

```
snmpProtectedDiv(val1, val2)
```

### Result:

```
3.6
```

The following expression produces the following result for a `val1` of 7.2 and `val2` of null or 0.0:

### Result:

```
0.0
```

## Advanced Example

The following expression is taken from Vendor Certification:

```
Utilization=snmpProtectedDiv((cpuStatsUser + cpuStatsSys),(cpuStatsUser + cpuStatsSys +
(isdef(cpuStatsWait)?cpuStatsWait:0) + cpuStatsIdle))*100
```

### **snmpRound Function**

This function rounds a numeric value to the nearest integer value.

#### **Syntax**

This function has the following format:

```
Long snmpRound(Double dNumber)
```

#### **Parameters**

- **dNumber**  
A Double value (floating number) to be rounded (*Double* is a Java data type).

#### **Return Values**

Returns a Long value, which is the nearest integer value to the value provided in dNumber (*Long* is a Java data type).

#### **Examples**

The following expression produces the following result for a dNumber of 3.5:

#### **Expression:**

```
snmpRound(dNumber)
```

#### **Result:**

4

The same expression produces the following result for a dNumber of 3.4:

#### **Result:**

3

### **Advanced Example**

The following expression is taken from “Cisco IPSLA Jitter Precision Statistics” Vendor Certification:

```
PathAvailability=snmpRound(rttMonJitterStatsNumOfRTT / (rttMonJitterStatsNumOfRTT
+ rttMonJitterStatsPacketLossSD + rttMonJitterStatsPacketLossDS +
rttMonJitterStatsPacketOutOfSequence + rttMonJitterStatsPacketMIA +
rttMonJitterStatsPacketLateArrival + rttMonJitterStatsError + rttMonJitterStatsBusies +
1/100) * 100)
```

### **snmpStringParser Function**

This function was written for internal use only. The function parses IP addresses that are received from a CA Application Insight Module (AIM). CA has only tested this function with an internal class. Another type of class may not be supported.

#### **Syntax**

This function has the following format:

```
snmpStringParser(Delimiter, Type to convert to, String to parse 1, String to parse 2)
```

#### **Parameters**

- **Type to convert to**

The type of class to which to parse the supplied strings.

- **Strings to parse 1 and 2**

The IP address to parse. Two strings enable you to provide addresses in IPv4 and IPv6 format.

## Return Values

Returns the converted value of the string, or returns "null" when the function cannot parse the string.

## snmpSvcs Function

This function takes the values from sysObjectOID, sysService, and ipForwarding MIB variables of an agent and determines what services the SNMP agent supports. For example, Router/Switch/Repeater/Host could be a supported service, as defined in the SNMP MIB RFC 1213.

The return from the function is evaluated as follows because custom device types have high precedence over system ones:

- If sysObjectID OID is mapped in the DeviceTypes file, then the return services is from the file.
- If sysObjectID OID is not mapped in the DeviceTypes file, then the sysServices and ipForwarding is used to return the supported services.

## Syntax

This function has the following format:

```
DeviceService[] snmpSvcs(ObjectID sysObjectID, Integer sysServices,
Integer ipForwarding)
```

## Parameters

- **sysObjectID**  
The object ID value to parse.
- **sysServices**  
An integer where each bit represents a different service, such as switch/repeater/host.
- **ipForwarding**  
An integer indicating whether this entity is acting as an IP gateway or IP host regarding the forwarding of datagrams. This entity receives the forwarded datagrams, but the forwarded datagrams are not addressed to this entity. Possible values are 1 (Forwarding) and 2 (notForwarding).

## Return Values

Returns a list of one or more of the following device services:

- ROUTER
- REPEATER
- SWITCH
- HOST
- UNKNOWN\_TYPE

## Example

The following expression produces the following result for a sysServices value of 8, an ipForwarding value of 0, and sysObjectID not found in the DeviceTypes file:

### Expression:

```
snmpSvcs(sysObjectOID, sysServices, ipForwarding)
```

### Result:

```
DeviceService[HOST]
```

## Advanced Example

The following expression is taken from the "System Statistics" Vendor Certification:

```
Services=sntpSvc(sysObjectID, isdef(sysServices)?sysServices:0, isdef(ipForwarding)?  
ipForwarding:0)
```

## storePortReconfig Function

This function returns a string containing XML representing the values of ifNumber, ifTableLastChange, ifStackLastChange. You can use XML to track device changes and to rediscover interfaces on a device, when needed.

### Syntax

This function has the following format:

```
String storePortReconfig ( Integer ifNumber, Long ifTableLastChange,  
Long ifStackLastChange )
```

### Parameters

- **ifNumber**  
The number of ports on the device.
- **ifTableLastChange**  
The date and time of the latest port table change in milliseconds, calculated from a start date and time of January 1, 1970 GMT.
- **ifStackLastChange**  
The date and time of the latest port stack change in milliseconds, calculated from a start date and time of January 1, 1970 GMT.

### Return Values

Returns a string containing XML in the form that is shown in the following example.

### Example

The following expression produces the following result for an ifNumber of 5, ifTableLastChange of 123456, and ifStackLastChange of 234567:

### Expression:

```
storePortReconfig ( ifNumber, ifTableLastChange, ifStackLastChange )
```

### Result:

```
<ReconfigData>  
  <ReconfigValue name="ifNumber" value="5"/>  
  <ReconfigValue name="ifTableLastChange" value="123456"/>  
  <ReconfigValue name="ifStackLastChange" value="234567"/>  
</ReconfigData>
```

## Global Variables

Data Aggregator supports the following global variables:

- **\_rspDuration**  
A Long (Java data type) containing the duration of the current poll cycle in seconds.

**NOTE**

The "System Statistics" Vendor Certification contains an example of the use of `_rspDuration`.

- **`_rspTimestamp`**

A Long (Java data type) containing the timestamp at which the current poll cycle started in milliseconds since January 1, 1970 GMT.

- **`_context`**

Java class that contains the details of the polled item, such as the Device Item ID.

**WARNING**

`_context` is a reserved keyword and global variable that is intended only for system use. Do not use `_context` under any circumstance.

## Vendor Certification Expression Operators

Use MVEL syntax in vendor certification expressions. You can build expressions using operators, use braces to control precedence, and terminate statements by semi-colons. The MVEL language has vendor certification utility functions and vendor certification global variables that you can also use in vendor certification expressions.

MVEL is a publicly available embeddable Expression Language for Java environments that has a syntax close to Java. MVEL supports expressions similar to Java expressions. For a detailed reference of the MVEL language, see [https://en.wikibooks.org/wiki/Transwiki:MVEL\\_Language\\_Guide](https://en.wikibooks.org/wiki/Transwiki:MVEL_Language_Guide).

To study the use of functions, operators, and global variables, you can use the Vendor Certification tab in NetOps Portal.

The following table summarizes the available operators:

**NOTE**

In XML documents, use the XML Named Entities (XNE) presentation.

Operator	XML Named Entities	Description	Example
=	N/A	Assign	a = 1
==	N/A	Equals	"fred" == "fred"
!=	N/A	Not Equals	"fred" != "tom"
>	&gt;	Greater Than	1 > 0 is true
<	&lt;	Less Than	0 < 1 is true
>=	N/A	Greater Than or Equal	1 >= 0 is true
<=	N/A	Less Than or Equal	1 <= 1 is true
contains	N/A	Verify if the value on the left side contains the value on the right	"tomcat" contains "cat"
isdef	N/A	Tests whether a variable is defined	isdef a
+	N/A	Add	1 + 1
+	N/A	Concatenate	"one " + "two"
-	N/A	Minus	2 - 1
*	N/A	Multiply	2 * 2
/	N/A	Divide	4 / 2
%	N/A	Modulus	5 % 2
&&	&amp;&amp;	Logical AND	(x>-1) && (x<1)
	N/A	Logical OR	(x<-1)    (x>1)

&	&amp;	AND bit operation	17 & 0xF
	N/A	OR bit operation	4   1
^	N/A	Exclusive OR bit operation	5 ^ 1
!	N/A	Negate	! True
?	N/A	Ternary operator	age > 17 ? "allow" : "deny"

## Self-Certification XML

Vendor certifications, metric families, and components are defined in XML.

You can create, extend, or update certifications by updating the associated XML files.

The articles in this section provide details about the self-certification XML attributes and structure.

## Certification Schema Files and Examples

The schema files that are provided with data aggregator include detailed information about element types, occurrence, and allowed lengths. The files also contain annotations that provide more information, such as allowed characters and naming conventions. These files are required when validating your XML files during vendor certification creation and extension.

For more information about how to create or extend a vendor certification, see [Create or Extend Vendor Certifications](#).

### Download the Schema and Example XML Files

Enter the following URLs in your web browser address bar:

*Hostname* is the data aggregator host. The default *port* is 8581.

- <http://hostname:port/resource/xsd/IMDBCertificationFacet.xsd>
- <http://hostname:port/resource/xsd/ComponentFacet.xsd>
- <http://hostname:port/resource/xsd/ItemSyncDefinition.xsd>
- <http://hostname:port/resource/xsd/SNMPCertificationFacet.xsd>
- <http://hostname:port/resource/xsd/CertificationFacet.xsd>
- <http://hostname:port/resource/xsd/webservices.xsd>
- <http://hostname:port/resource/xsd/baseweb services.xsd>
- <http://hostname:port/resource/xsd/datamodel.xsd>

The example XML files are located in the `/opt/IMDataAggregator/examples` directory.

## Restricted XML Tags

You cannot use the following XML tags in custom and extended vendor certifications:

- `AggregateToDevice`
- `SupportsDeviceAggregation`

When extending a vendor certification or metric family, you cannot modify the following XML tags:

- AlarmRules
- Author
- ComponentFacets
- descriptorClass
- DeviceType
- IsDynamicDiscoveryAttribute
- ItemFacets
- ItemRelationshipMappings
- Normalized
- ParentNodeList
- Protocol
- RollupExpression
- TableName
- type
- UsesDynamicIndex
- Variance
- VCSupportExpression
- WriteOnPoll

## Vendor Certification XML Structure

Vendor certifications map vendor- and device-specific data to performance metrics and configuration data that are defined in a metric family using XML. Mapping this data from various sources to the `normalized` metric family values helps the data aggregator uniformly report on this data, regardless of the device vendor.

### IMPORTANT

The properties included in the XML example and listed in the following descriptions are presented in a recommended order. List them in the XML in that order.

### Example:

The following example is an example of a custom vendor certification XML that supports Frame-Relay PVC. The example `frPVCInfo` custom metric family is included in the `ExpressionGroup` section:

### NOTE

**Best Practice:** When managing the vendor certification XML, use a REST client. In web browsers, certain tags are hidden.

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- Auto-generated by the type catalog local manager.-->
  <DataModel namespace="http://im.ca.com/certifications/
snmp" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="SNMPCertificationFacet.xsd">
    <Author>Custom</Author>
    <Version>1.0</Version>
    <FacetType name="frPVCInfoCustom"
descriptorClass="com.ca.im.core.datamodel.certs.CertificationFacetDescriptorImpl">
        <Documentation>Frame Relay PVC Vendor Certification</Documentation>
        <FacetOf namespace="http://im.ca.com/core" name="Item" />
        <DisplayName>Frame Relay PVC Certification</DisplayName>
        <MIB>RFC1315-MIB</MIB>
```

```
<Protocol>SNMP</Protocol>
<AttributeGroup name="AttributeGroup" external="true" list="true">
  <Documentation />
  <Attribute name="INDEX" type="ObjectID">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.4</Source>
    <IsIndex>true</IsIndex>
    <IsKey>false</IsKey>
    <NeedsDelta>false</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitReceivedBECNs" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.5</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitSentFrames" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.6</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitSentOctets" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.6</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitReceivedFrames" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.8</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitReceivedOctets" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.9</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
</AttributeGroup>
```



```

    <Expressions>
    <ExpressionGroup destCert="{http://im.ca.com/normalizer}frPVCInfo"
name="frPVCInfoDS">
    <Expression destAttr="Indexes">INDEX</Expression>
    <Expression destAttr="Names">"Frame Relay " + INDEX</Expression>
    <Expression destAttr="FECNIn">frCircuitReceivedFECNs</Expression>
    <Expression destAttr="BECNIn">frCircuitReceivedBECNs</Expression>
    <Expression destAttr="FramesIn">frCircuitReceivedFrames</Expression>
    <Expression destAttr="FramesOut">frCircuitSentFrames</Expression>
    <Expression destAttr="BytesIn">frCircuitReceivedOctets</Expression>
    <Expression destAttr="BytesOut">frCircuitSentOctets</Expression>
    </ExpressionGroup>
    </Expressions>
    </FacetType>
</DataModel>

```

In this article:

- [Custom Vendor Certification Basic Properties](#)
- [AttributeGroup](#)
- [IndexTagList](#)
- [ExpressionGroup](#)
- [HierarchyList](#)

### **Custom Vendor Certification Basic Properties**

The custom vendor certification basic properties help to distinguish it from other custom vendor certifications that you create. They also indicate from which vendor MIB you are collecting metric data.

Consider the following restrictions when you determine the basic properties:

- The `FacetType/name` and `FacetType/DisplayName` properties must be unique for each vendor certification.
- The Protocol tag is either SNMP or EMS. The `/typecatalog/certifications/snmp` property supports only SNMP certifications. In this case, the only value that is supported is SNMP. The `/typecatalog/certifications/camm` property supports only DX NetOps Mediation Manager (DX NetOps MM) certifications. In this case, the only value that is supported is EMS.
- Set the `FacetType/descriptorClass` property and all `DataModel` and `FacetOf` properties as shown in the example XML in the previous illustration.

The following list details basic vendor certification properties:

- **FacetType/name**

Uniquely identifies a vendor certification.

**Best practice:** Conform to `<MibName><TableName>Mib`.

**Can be updated:** No

**Possible values:** Alphanumeric and underscore. Dot and dash are not permitted.

The `FacetType` section manifests a particular vendor certification. The same XML document can contain multiple `FacetType` sections when those vendor certifications expose various aspects of the vendor-specific device such as TCP and UDP statistics from a MIB-2 implementation.

The `FacetType` section contains some basic properties. For example, this section contains the name of the vendor MIB, followed by one or more `AttributeGroup` sections. These `AttributeGroup` sections define which attributes this certification uses from the MIB. One or more `ExpressionGroup` sections map attributes from the `AttributeGroup` sections to the metrics specified in a metric family.

**NOTE**

You can update the following items. They support plain text.

- **FacetType/Documentation**

Describes what is certified with the vendor certification.

**Best practice:** Include the details about the vendor, MIB name, and table name.

**Effect of updating:** None

**NOTE**

List this property first under the FacetType/name .

- **FacetType/DisplayName**

Specifies the name of the vendor certification as it displays in NetOps Portal.

**Best practice:** Start with the vendor name and include the MIB and functionality information.

**Effect of updating:** A change to the name in the Administrator user interface (UI).

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** Refresh the user interface.

**IMPORTANT**

Ensure that the DisplayName property is unique to the vendor certification.

- **FacetType/MIB**

Specifies the name of the MIB, which the DEFINITIONS clause defines in the ASN.1 file.

**Best practice:** Conform to <MibName> .

**Effect of updating:** Change to the SNMP MIB Name column in the **Vendor Certification** tab of the Administrator UI.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** Refresh the user interface.

- **FacetType/Protocol** The Protocol tag is either SNMP or EMS. The /typecatalog/certifications/snmp property supports only SNMP certifications. In this case, the only value that is supported is SNMP. The /typecatalog/certifications/camm property supports only DX NetOps MM certifications. In this case, the only value that is supported is EMS.

**AttributeGroup**

The following example illustrates the AttributeGroup section of your custom vendor certification. This section identifies the attributes (variable OIDs) of a particular table in the vendor MIB that are used to map raw device data. This data is mapped to the performance metrics and the configuration data that is defined in a metric family.

You set the AttributeGroup/list and AttributeGroup/external properties to true , as shown in the example XML in the previous illustration. These properties specify that each attribute represents a list of values that is obtained from an external source (a MIB table). The following information summarizes the XML elements to customize.

**NOTE**

You can update the following items. They support plain text. The update does not affect performance.

- **AttributeGroup/name**

Specifies the attribute group name.

**Best practice:** Conform to <FacetType/name>Group .

- **Documentation**

(Optional) Specifies the description for the attribute group.

**NOTE**

List this property first under the AttributeGroup/name .

- **UseIndex**

Specifies the name of the attribute to be used as the index for this attribute group for joining multiple MIB tables.

**NOTE**

When using MultiMIB Table Support, the AttributeGroup order must match the IndexTagList order.

**Best Practice:** Set to the value of the AttributeGroup/name property.

**NOTE**

In the attributes list, the attributes used for calculating indexes and names should be listed first.

**General Attributes**

The general attributes for all vendor certifications are as follows:

**NOTE**

Unless specified otherwise, the update takes effect immediately and no actions are required to trigger the update. You can update the entries in this list.

- **Attribute/name**  
Specifies the attribute name.  
**Best practice:** Set to the MIB variable name, which the `OBJECT-TYPE` clause defines in the `ASN.1` file.  
**Possible values:** Alphanumeric and underscore. Dot and dash are not permitted.  
**Effect of updating:** Update any expressions that reference this attribute.
- **Attribute/type**  
Specifies the data type of the attribute.  
**Best practice:** Use the attribute type that best matches the variable type that the `SYNTAX` clause defines in the `ASN.1` file.  
**Possible values:** Boolean, Int, Long, Double, BigInteger, String, DateTime, IPAddress, MACaddress, IPSubnet, OctetString, ObjectID  
**Effect of updating:** Polled SNMP data is converted to this type.  
**When does the update take effect:** At the next poll.
- **Documentation**  
(Optional) Specifies the description for the attribute, which documents the semantics (such as the unit) of the MIB variable.  
**Best practice:** Use the descriptions that are taken from the MIB `ASN.1` file.  
**Possible values:** Plain text  
**Effect of updating:** None  
**NOTE**  
List this property first under the `Attribute/name` and `Attribute/type`.
- **IsKey**  
(Optional) Indicates whether the MIB variable is key for determining support for a table. When you specify multiple fields as keys, together all are considered as a compound key.  
**Default:** false  
**Best practice:** Set to `true` if it is a key MIB object for component discovery. If the contents of the MIB attributes are necessary to determine support, use `VCSupportExpression` instead.  
**Possible values:** true, false  
**Effect of updating:** Components could change to a new vendor certification.  
**When does the update take effect:** After component rediscovery.  
**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.
- **IsIndex**  
(Optional) Uses a flag to indicate whether this variable is an index to the MIB table.  
**Default:** false  
**Best practice:** Set to `true` for an index attribute.  
**Possible values:** true, false  
**Effect of updating:** Component indexing could change.  
**When does the update take effect:** After component rediscovery.  
**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.
- **NeedsDelta**  
(Optional) Uses a flag to indicate whether to delta (that is, store the difference between current and last poll for Counters) the MIB variable.

**Default:** false

**Best practice:** Set to `true` if the variable is defined as a Counter, Counter32, Counter64, or TimeTicks quantity in the MIB.

**Possible values:** true or false.

**Effect of updating:** The polled data changes.

**When does the update take effect:** At the next poll.

#### NOTE

You can use this attribute only during the polling phase. For example, do not use this attribute as part of the Name or Description.

- **Source**

Specifies the ObjectID of the attribute.

**Best practice:** Set to the fully qualified MIB variable OID that the OBJECT-TYPE defines.

**Possible values:** Dot-separated numbers (for example, 1.3.6.1.4.1...).

**Effect of updating:** Data is polled from the specified OID.

**When does the update take effect:** At the next poll.

By default, the `Source` attribute specifies an OID to poll from the device. This default behavior is defined with `src='polled'`. You can set `src='mvel'` to process an MVEL expression using any polled attribute instead of an OID. For example, you could use the `src='mvel'` parameter to combine two 32-bit OIDs into a single 64-bit counter. You could also use the `src='mvel'` parameter to poll a counter that is not stored as a numeric type.

**Example:** The following example uses the `src='mvel'` parameter to combine two 32-bit OIDs into a single 64-bit counter:

```
<Attribute name="memberbitsout" type="Long">
  <NeedsDelta>true</NeedsDelta>
  <Source src='mvel'>snmpCounter64(memberbitsoutHi32,memberbitsoutLo32)</Source>
</Attribute>
```

You can use the `src='mvel'` parameter only during polling. You cannot use it during the discovery phases. For example, you cannot use it as part of the Name or Description.

- **Version**

Specifies the version of the vendor certification. Update this attribute when you update the certification. You cannot decrease the Version value.

**Possible values:** Floating point or decimal number Example: 1.0 or 1.01

- **Author**

Specifies the creator of the vendor certification.

**Default:** "Custom"

**Possible values:** Any alphanumeric string.

**Effect of updating:** The author attribute is updated.

- **UsesDynamicIndex**

(Optional) Enables dynamic ObjectID values for IPSLA polling.

**Default:** False

**Possible values:** true, false

- **IsDynamicDiscoveryAttribute**

(Optional) Specifies whether the attribute should be used to discover the dynamic index.

#### IMPORTANT

If you do not specify this, the data collector chooses the first dynamic OID it encounters in the attribute group. The device might not support the OID that the data collector chooses, which can result in poll failure.

The list of attributes specifies the set of data that a metric family collects when supported by this vendor certification. Typically, this data falls into two categories:

- Configuration data of the device component (such as name or indexes) that is collected only at discovery time.
- Performance data that is collected every poll cycle.

## Configuration Data Attributes

An attribute with the name `INDEX` and type `ObjectID` is mapped to the `Indexes` attribute of the target metric family. You can set the value for the `Source` tag to any variable `OID`. However, you typically use one of the variables that are listed in the `INDEX` clause of the table. For example, consider `ifIndex` in the `interfaces` table of MIB-2. This variable serves as the index for the other variables in the same MIB table. In addition, the `IsIndex` tag (and typically also the `IsKey` tag) for this attribute is set to `true`.

In this example, attributes such as `ifDesc` or `ifType` provide more configuration information about an interface. Therefore, these attributes are useful for the `Names` and `Description` attributes of the target metric family.

## Performance Data Attributes

These attributes provide the raw data for performance metrics in the target metric family. Consider the following points:

- You can directly map one of these attributes to a metric family performance metric, or
- You can use the attribute in an expression with other attributes to compute a value for the metric.

## IndexTagList

To poll attributes from multiple MIB tables, we need an attribute group per MIB table containing these attributes. The index tag list provides a mechanism to relate two attribute groups (or MIB tables) with different indexes. The groups are related such that one item (row) of one table is linked to a corresponding row in a second table.

### NOTE

For these items, the following criteria apply to all:

- You can update the entries in this list.
  - The update changes indexing.
  - The update takes effect after component rediscovery.
  - For the updates to take effect, update the metric family or change the vendor certification priority.
- **PrimaryTag**  
References the primary Attribute group (that is, the group that defines an index attribute with the `ObjectID` type). The value of this element must equal the `UseIndex` tag of the attribute group for the primary group.  
**Possible values:** The `UseIndex` tag of the attribute group corresponding to the primary attribute group.
  - **IndexTag**  
Defines how to relate rows of the primary group (or MIB table) to rows in the secondary group. This element relates the rows by specifying attributes of both groups that must match.
  - **IndexTag/Name**  
References the secondary group (or MIB table). The value of this element must equal the `UseIndex` tag of the secondary attribute group that you are trying to relate with the primary one.  
**Possible values:** The `UseIndex` tag of the secondary attribute group.
  - **IndexTag/PrimaryKeyExpression**  
Specifies an MVEL expression containing attributes of the primary attribute group or an attribute group corresponding to any of the previously defined `IndexTag`. The calculated value is matched up with the `ThisTagKeyExpression`. If there is a match, the rows of both attribute groups (or MIB tables) are linked. Then, these attributes can be used together in an `Expression` backing a `destAttr` (or `Metric`).  
**Possible values:** A valid MVEL expression.
  - **IndexTag/ThisTagKeyExpression**  
Specifies an MVEL expression containing attributes of the secondary attribute group. The calculated value is matched up with the `PrimaryKeyExpression`. If there is a match, the rows of both groups (or MIB tables) are linked. Then, you can use these attributes together in an `Expression` backing a `destAttr` (or `Metric`).  
**Possible values:** A valid MVEL expression.

## ExpressionGroup

The `ExpressionGroup` maps attributes as follows:

- From the `AttributeGroup` (that defines how to get a metric from an SNMP MIB)
- To the metrics specified in a metric family (that defines how an attribute is stored in the database)

### NOTE

List the `Filter`, `VariableGroup`, `VCSupportExpression`, `Expression`, and `SetExpression` properties in this order in the XML.

You can store a MIB value in the database as it is received from the device or after some normalization operations are performed. For example, normalization operations include dividing or multiplying with 1024 to transform to/from kilobytes.

### NOTE

Unless specified otherwise, the update takes effect immediately and no actions are required to trigger the update. You can update the entries in this list.

- **ExpressionGroup/destCert**  
Specifies the metric family that contains the `destAttrs` to populate.  
**Possible values:** Any valid metric family  
**Effect of updating:** Changes the `destAttr` permissible expression.
- **ExpressionGroup/name**  
Specifies the expression group name.  
**Possible values:** Plain text  
**Effect of updating:** None
- **ExpressionGroup/Filter**  
(Optional) Specifies which components are discovered. Using the `Filter` reduces the number of components that are managed.

### NOTE

The expression group filter *does not* exclude the specified components. The filter selects the specified components and excludes components that do not match the criteria.

**Possible values:** Boolean MVEL expression using available Attributes

**Effect of updating:** Changes which components are discovered.

**When does the update take effect:** After component rediscovery.

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

- **VariableGroup**  
Defines variables that are used in the `ExpressionGroup`.

### NOTE

Within a `VariableGroup`, variables are processed in the order listed.

**Possible Values:** Calculated vendor certification values.

**Effect of updating:** Changes which components are discovered.

**When does the update take effect:** After component rediscovery.

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

### NOTE

The Juniper and Cisco/Standard High Speed Interface vendor certifications include the `UtilizationMaxPercent` variable. This variable defines the percentage at which to drop the data for the utilization metric. Dropped data preserves the integrity of rollout data for the interface and results in a gap in views and reports.

- **VCSupportExpression**  
(Optional) Extracts and calculates the MIB attribute values to determine whether the VC is supported.

### NOTE

To ensure vendor certifications can discover, use `IsKey` and `VCSupportExpression` in separate vendor certifications.

**Possible values:** Boolean MVEL expression using available Attributes

**Effect of updating:** Changes which components are discovered.

**When does the update take effect:** After component rediscovery.

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

- **Expression**

Converts vendor certification attribute values to normalized attribute values.

**Possible values:** Normalized attribute value

- **SetExpression**

Converts normalized attribute values to vendor certification attribute values.

**Possible values:** Vendor certification attribute value

**Effect of updating:** Changes which components are discovered.

**When does the update take effect:** After component rediscovery.

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

## **ExpressionGroup Filter Examples**

Filters limit the device components that a vendor certification discovers and polls. Discovery occurs only if all the criteria in the filter are true. If any of the specified attributes in the filter criteria cannot be evaluated for a given component, the component is discovered because the complete filter criteria could not be evaluated. The most common reason an attribute cannot be evaluated is that the component has no value for that attribute.

### **Example 1:**

```
<Filter>(ifType!=24) & & (ifType!=1)</Filter>
```

DX NetOps Performance Management does not poll the device component if the `ifType` value is 1 or 24. Interfaces with no value for `ifType` are discovered and polled.

### **Example 2:**

#### **NOTE**

Use the `.toString()` method when comparing `OctetString` attributes, as `OctetStrings` are not `Strings`.

```
<Filter> hrStorageType.toString() == "1.3.6.1.2.1.25.2.1.4" & &
    hrStorageSize != 0
</Filter>
```

DX NetOps Performance Management discovers and polls the device component if the the `StorageType` is `hrStorageFixedDisk` (1.3.6.1.2.1.25.2.1.4) and the size is not 0. However, if a component has no value for the `hrStorageSize`, that component is discovered. If you do not intend this behavior, extend the filter to use the **isdef** function to verify that the attribute has a valid value.

```
<Filter> hrStorageType.toString() == "1.3.6.1.2.1.25.2.1.4" & &
    isdef (hrStorageSize) & & hrStorageSize != 0
</Filter>
```

DX NetOps Performance Management discovers and polls the device component if the `StorageType` is `hrStorageFixedDisk` and the size is not 0. Only components with a specified value for `hrStorageSize` are discovered and polled.

### **Example 3:**

#### **NOTE**

Use the `.toString()` method when comparing `OctetString` attributes, as `OctetStrings` are not `Strings`.

```
<Filter> (rttMonCtrlAdminRttType==9) & &
    ( !(rttMonCtrlAdminOwner.toString() contains "Network Health") )
</Filter>
```



DX NetOps Performance Management discovers and polls the device component if the `rttMonCtrlAdminRttType` is 9 and the `rttMonCtrlAdminOwner.toString` does not contain Network Health.

### **Expression/destAttr Metrics**

The following information describes the Expression/destAttr metrics. You can update all of these metrics:

#### **NOTE**

Unless specified otherwise, the update takes effect immediately and no actions are required to trigger the update.

- **Indexes**

Specifies to use the vendor certification attributes of the ObjectID to define the MVEL expression to provide the value for the Indexes metric family attribute.

**Best practice:** Set to `INDEX`.

**Possible values:** Any attribute that has `<IsIndex>true</IsIndex>`.

**Effect of updating:** Component indexing could change.

**When does the update take effect:** After component rediscovery

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

- **Names**

Specifies to use the vendor certification attributes to collect configuration data. This configuration data helps define the MVEL expression to provide the value for the Names metric family attribute.

**Best practice:** Include as much information as necessary to identify an instance uniquely.

**Possible values:** String MVEL expression using available attributes.

**Effect of updating:** Component name change

**When does the update take effect:** After component rediscovery

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

- **Descriptions**

(Optional) Specifies to use the vendor certification attributes to collect configuration data. This configuration data helps define the MVEL expression to provide the value for the Descriptions metric family. Not all metric families support a `Descriptions` attribute.

**Best practice:** Include as much information as is available to describe an instance.

**Possible values:** String MVEL expression using available attributes.

**Effect of updating:** Component description change

**When does the update take effect:** Component rediscovery

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

- **Other Metrics**

Specifies to use the vendor certification attributes to collect configuration or performance data. This data is used to define the MVEL expression to provide the value for the metric family attribute.

**Can be added:** Yes, if the `destAttr` permissible expression exists in the metric family.

**Possible values:** MVEL expression using available Attributes, producing a value that matches the type of the destination attribute.

**Effect of updating:** Polled value changes.

**When does the update take effect:** At the next poll.

The metric family exposes URIs (such as, `{http://im.ca.com/normalizer}FamilyName.AttributeName`), which are separately referred to in the ExpressionGroup. The `ExpressionGroup/destCert` property is set to the URI (for example, `{http://im.ca.com/normalizer}FamilyName`), and the `Expression/destAttr` expression is set to `AttributeName`.



## Speed Override

Vendor certifications for interfaces include variables in the `ExpressionGroup` that can override the `SpeedIn` and `SpeedOut` values. The `SpeedInOverride` and `SpeedOutOverride` variables let you override the `SpeedIn` and `SpeedOut` values in the UI. The following example shows how to use the override variables:

```
<VariableGroup>
    <Variable name="SpeedInOverride" providedBy="override"/>
    <Variable name="SpeedOutOverride" providedBy="override"/>
    <Variable name="RawIfSpeed">ifSpeed</Variable>
    <Variable name="CalculatedSpeedIn">
        isdef(SpeedInOverride) ? SpeedInOverride : RawIfSpeed
    </Variable><Variable name="CalculatedSpeedOut">
        isdef(SpeedOutOverride) ? SpeedOutOverride : RawIfSpeed
    </Variable>
    <Variable name="CalculatedIfInOctets">ifInOctets <= 786432000 ? ifInOctets :
null</Variable>
    <Variable name="CalculatedIfOutOctets">ifOutOctets <= 786432000 ? ifOutOctets :
null</Variable>
</VariableGroup>
```

## HierarchyList

The following list defines the hierarchy behavior:

### NOTE

For these items, the following criteria apply to all:

- All entries in this list can be updated.
  - The update changes the hierarchy construction.
  - The update takes effect after component rediscovery.
  - For the updates to take effect, update the metric family or change the vendor certification priority.
- **Hierarchy/ParentFacet**  
Specifies the QName of the facet that is used to find the candidate parent items.  
**Possible values:** Any valid facet.
  - **Hierarchy/ParentAttribute**  
Specifies the QName of the attribute that is used to identify the specific parent item.  
**Possible values:** Any valid attribute QName.
  - **Hierarchy/ChildAttribute**  
Specifies the QName of the attribute on the child item that is used to match the `ParentAttribute` on the parent item.  
**Possible values:** Any valid attribute QName.

## Multi-MIB Table Support

Some situations exist where you must collect the raw data for a particular metric family from two or more MIB tables. Custom certification includes support for multiple MIB tables. The XML structure is similar to a standard vendor certification, and uses a common key (index) to join data that is collected from multiple tables.

### Example:

In this example, the Frame Relay PVCs can be named using a combination of the ifName MIB object in the ifXTable and the frCircuitDlci object that provides the data link connection identifier (DLCI) for this PVC. This kind of naming convention is useful for determining which Frame Relay interface the PVC is layered onto.

To support both MIB tables, add the following information to the XML:

- Add an attribute to the existing AttributeGroup to represent the frCircuitDlci MIB object.
- The ifName MIB object comes from a MIB that is not included in your custom vendor certification. Add an AttributeGroup (in this case, ifXTable), and then add the new attribute (ifName).

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- Auto-generated by the type catalog local manager.-->
<DataModel namespace="http://im.ca.com/certifications/snmp" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="SNMPCertificationFacet.xsd">
  <FacetType name="frPVCInfoCustom"
    descriptorClass="com.ca.im.core.datamodel.certs.CertificationFacetDescriptorImpl">
    <Documentation>Frame Relay PVC Vendor Certification</Documentation>
    <FacetOf namespace="http://im.ca.com/core" name="Item" />
    <AttributeGroup name="ifXTableGroup" external="true" list="true">
<Documentation>This pulls data from the ifXTable so that the ifName corresponding to the
PVC can be referenced</Documentation>
<UseIndex>ifXIndexTag</UseIndex>
      <Attribute name="ifXTableIndex" type="ObjectID">
<Documentation />
<IsKey>false</IsKey>
<IsIndex>true</IsIndex>
<Source>1.3.6.1.2.1.31.1.1.1.1</Source>
<Polled>false</Polled>
</Attribute>
<Attribute name="ifName" type="OctetString">
<Documentation />
<IsKey>false</IsKey>
<IsIndex>false</IsIndex>
<Source>1.3.6.1.2.1.31.1.1.1.1</Source>
<Polled>false</Polled>
</Attribute>
</AttributeGroup>
      <IndexTagList>
<PrimaryTag>PVCIndexTag</PrimaryTag>
<IndexTag>
<Name>ifXIndexTag</Name>
<PrimaryKeyExpression>snmpOIDParser (INDEX,1,1)</PrimaryKeyExpression>
<ThisTagKeyExpression>ifXTableIndex</ThisTagKeyExpression>
</IndexTag>
</IndexTagList>
      <AttributeGroup name="AttributeGroup" external="true" list="true">
        <Documentation />

```

```
<UseIndex>PVCIndexTag</UseIndex>
<Attribute name="INDEX" type="ObjectID">
  <Documentation />
  <Source>1.3.6.1.2.1.10.32.2.1.4</Source>
  <IsIndex>true</IsIndex>
  <IsKey>false</IsKey>
  <NeedsDelta>false</NeedsDelta>
</Attribute>
<Attribute name="frCircuitReceivedFECNs" type="Long">
<Documentation />
<Source>1.3.6.1.2.1.10.32.2.1.4</Source>
<IsIndex>false</IsIndex>
<IsKey>true</IsKey>
<NeedsDelta>true</NeedsDelta>
</Attribute>
  <Attribute name="frCircuitReceivedBECNs" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.5</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitSentFrames" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.6</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitSentOctets" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.6</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitReceivedFrames" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.8</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
  </Attribute>
  <Attribute name="frCircuitReceivedOctets" type="Long">
    <Documentation />
    <Source>1.3.6.1.2.1.10.32.2.1.9</Source>
```

```

        <IsIndex>false</IsIndex>
        <IsKey>true</IsKey>
        <NeedsDelta>true</NeedsDelta>
    </Attribute>
    <Attribute name="frCircuitState" type="int">
<Documentation />
<Source>1.3.6.1.2.1.10.32.2.1.3</Source>
<IsIndex>false</IsIndex>
<IsKey>false</IsKey>
<NeedsDelta>false</NeedsDelta>
</Attribute>
    <Attribute name="frCircuitDlci" type="int">
<Documentation />
<Source>1.3.6.1.2.1.10.32.2.1.2</Source>
<IsIndex>false</IsIndex>
<IsKey>false</IsKey>
<NeedsDelta>false</NeedsDelta>
</Attribute>
</AttributeGroup>
<Protocol>SNMP</Protocol>
<DisplayName>Frame Relay PVC Certification</DisplayName>
<Expressions>
    <ExpressionGroup destCert="{http://im.ca.com/normalizer}frPVCInfo"
name="frPVCInfoDS">
        <Filter>(frCircuitState==2)</Filter>
        <Expression destAttr="Indexes">INDEX</Expression>
        <Expression destAttr="Names">isdef(ifName)? (isdef(frCircuitDlci) ? ifName + "
DCLI:" + frCircuitDlci : "Frame Relay " + INDEX) : "Frame Relay " + INDEX</Expression>
        <Expression destAttr="FECNIn">frCircuitReceivedFECNs</Expression>
        <Expression destAttr="BECNIn">frCircuitReceivedBECNs</Expression>
        <Expression destAttr="FramesIn">frCircuitReceivedFrames</Expression>
        <Expression destAttr="FramesOut">frCircuitSentFrames</Expression>
        <Expression destAttr="BytesIn">frCircuitReceivedOctets</Expression>
        <Expression destAttr="BytesOut">frCircuitSentOctets</Expression>
        <Expression destAttr="BitsIn">frCircuitReceivedOctets*8</Expression>
        <Expression destAttr="BitsOut">frCircuitSentOctets*8</Expression>
    </ExpressionGroup>
</Expressions>
<MIB>RFC1315-MIB</MIB>
</FacetType>
</DataModel>

```

### AttributeGroup

Each table must go into its own AttributeGroup section. Each Attribute on that table is added as a child of that AttributeGroup.

Refer to these sections for the following information:

- The **AttributeGroup** information  
Details about the XML elements that are used to define primary and secondary table attributes.
- The **UseIndex** and **IndexTagList** information  
Details about the XML elements that are used to join the primary and secondary attribute groups.

In the example, the primary attribute group represents the table that you want to “extend” with more information. The secondary attribute group contains the “extension” information for the primary one.

The primary **AttributeGroup** contains an **Attribute** identifying the MIB table variable serving as the common “key” into the secondary **AttributeGroup**.

The secondary **AttributeGroup** includes the **Attribute** definitions for all MIB table variables carrying the “extension” information for the primary table. In addition, there is an **Attribute** identifying the variable matching the common “key” from the primary **AttributeGroup**.

### **UseIndex**

Each **AttributeGroup** is given a **UseIndex** tag. The **UseIndex** tag lets you group OIDs under a common name. This common name is associated with a given variable serving as the common key (index) into the respective MIB table.

The following information summarizes the XML elements to customize:

- **AttributeGroup/UseIndex**  
Uniquely identifies the primary and secondary tag name (respectively) that is used in the **IndexTagList** section.  
**Recommendation:** Set to the value of the **AttributeGroup/name** property.

### **IndexTagList**

The **IndexTagList** section is a mechanism to relate two attribute groups (or MIB tables) with different indexes. When the groups are related, one item has multiple index IDs from multiple tables.

The **IndexTagList** section contains all the join information, including an **IndexTag** section for every secondary attribute group.

- **IndexTagList/PrimaryTag**  
Defines the primary attribute group (or MIB table). Set to the value of the **UseIndex** property of the primary **AttributeGroup**.
- **IndexTag/Name**  
Defines the secondary attribute group. Set to the value of the **UseIndex** property of the secondary **AttributeGroup**.
- **IndexTag/PrimaryKeyExpression**  
Specifies the expression to generate the common key in the primary table. Consider using the MVEL functions to derive a common key from the designated primary table index **Attribute**.
- **IndexTag/ThisTagKeyExpression**  
Specifies the expression to generate the common key in the secondary table. Consider using the MVEL functions to derive a common key from the designated secondary table index **Attribute**.

The multitable approach supports the chaining of more than two tables. Two types of relationships exist in multiple table joins:

- **Primary - > Secondary #1, Primary - > Secondary #2**  
Ordering of the secondary tables does not matter in an index tag list.
- **Primary - > Secondary #1 - > Secondary #2**  
List secondary table #1 before secondary table #2 because of the way tables are merged.

One or more rows in the primary table can legally map to the same row in the secondary table. Keys on the secondary table are searched in order, and the first match wins.

## Metric Family XML Structure

DX NetOps Performance Management normalizes these metrics so that reporting is uniform regardless of the vendor (data source). Metrics are "null" when the vendor does not provide a value. Any report views based on the null metrics are empty.

A metric family also defines attributes that are captured during discovery, like the item name and index. There can also be discovery rules defined that reconcile component matching. You include a metric family in a monitoring profile. The set of metric families in a monitoring profile determines which metrics to collect for the devices in each device collection that is associated with the profile.

### NOTE

List some properties in the XML in a particular order. The properties included in the XML example and listed in the following descriptions are presented in the recommended order.

In this article:

- [Example Metric Family](#)
- [Basic Properties](#)
- [ComponentFacets](#)
- [ItemFacets](#)
- [AttributeGroup \(Metric Family\)](#)
- [Expressions](#)
- [Hierarchy](#)
- [BaselineDefinitions](#)
- [ComponentReconciliation](#)
- [Tags in Custom and Extended Metric Families](#)
- [MatchAlgorithms](#)
- [ReconfigDetectionAttr](#)
- [Tags in Custom and Extended Metric Families](#)

### Example Metric Family

This example metric family supports the vendor certification for Frame-Relay PVC:

### NOTE

If you view the metric family XML in a browser, certain tags are hidden. For this reason, copy and paste the metric family XML only from a REST client.

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- Auto-generated by the type catalog local manager. -->
<DataModel xmlns:xsi=http://www.w3.org/2001/XMLSchema-instance namespace="http://
im.ca.com/normalizer" xsi:noNamespaceSchemaLocation="IMDBCertificationFacet.xsd">
  <Author>Custom</Author>
  <Version>1.0</Version>
  <FacetType name="frPVCInfo"
    descriptorClass="com.ca.im.core.datamodel.certs.NormalizedFacetDescriptorImpl">
    <Documentation>Frame Relay Permanent Virtual Circuit</Documentation>
    <FacetOf namespace="http://im.ca.com/core" name="Item" />
    <DisplayName>Frame Relay PVC</DisplayName>
    <TableName>FR_PVC_INFO</TableName>
    <Protocol>IMDB</Protocol>
    <Normalized>true</Normalized>
```

```

<ComponentFacets>
<Facet>{http://im.ca.com/inventory}frPVC</Facet>
</ComponentFacets>
<AttributeGroup name="AttributeGroup" external="true" list="true">
<Documentation />
<Attribute name="Indexes" type="ObjectID[]">
<Documentation />
<AttributeDisplayName />
<Polled>>false</Polled>
<Baseline>>false</Baseline>
<IsDbColumn>>false</IsDbColumn>
<Variance>>false</Variance>
<StandardDeviation>>false</StandardDeviation>
<Minimum>>false</Minimum>
<Maximum>>false</Maximum>
<WriteOnPoll>>false</WriteOnPoll>
<RollupStrategy />
<Percentile>0</Percentile></Attribute><Attribute name="Names" type="String">
<Documentation>The name of the frame relay circuit</Documentation>
<AttributeDisplayName />
<Polled>>false</Polled><Baseline>>false</Baseline>
<IsDbColumn>>false</IsDbColumn>
<Variance>>false</Variance>
<StandardDeviation>>false</StandardDeviation>
<Minimum>>false</Minimum>
<Maximum>>false</Maximum>
<WriteOnPoll>>false</WriteOnPoll>
<RollupStrategy />
<Percentile>0</Percentile>
</Attribute>
<Attribute name="Description" type="String">
<Documentation>A description for the frame relay circuit</Documentation>
<AttributeDisplayName />
<Polled>>false</Polled>
<Baseline>>false</Baseline>
<IsDbColumn>>false</IsDbColumn>
<Variance>>false</Variance>
<StandardDeviation>>false</StandardDeviation>
<Minimum>>false</Minimum>
<Maximum>>false</Maximum>
<WriteOnPoll>>false</WriteOnPoll>
<RollupStrategy />
<Percentile>0</Percentile>
</Attribute>
<Attribute name="BECNIn" type="Double">
<Documentation>Backward congestion since the virtual circuit was created</Documentation>

```

```
<AttributeDisplayName />
<Polled>true</Polled>
<Baseline>false</Baseline>
<IsDbColumn>true</IsDbColumn>
<Variance>false</Variance>
<StandardDeviation>false</StandardDeviation>
<Minimum>false</Minimum>
<Maximum>false</Maximum>
<WriteOnPoll>false</WriteOnPoll>
<RollupStrategy>Sum</RollupStrategy>
<Percentile>0</Percentile>
</Attribute>
<Attribute name="FECNIn" type="Double">
<Documentation>Forward congestion since the virtual circuit was created</Documentation>
<AttributeDisplayName />
<Polled>true</Polled>
<Baseline>false</Baseline>
<IsDbColumn>true</IsDbColumn>
<Variance>false</Variance>
<StandardDeviation>false</StandardDeviation>
<Minimum>false</Minimum>
<Maximum>false</Maximum>
<WriteOnPoll>false</WriteOnPoll>
<RollupStrategy>Sum</RollupStrategy>
<Percentile>0</Percentile>
</Attribute>
<Attribute name="FramesIn" type="Double">
<Documentation>Frames received since the virtual circuit was created</Documentation>
<AttributeDisplayName />
<Polled>true</Polled>
<Baseline>false</Baseline>
<IsDbColumn>true</IsDbColumn>
<Variance>false</Variance>
<StandardDeviation>false</StandardDeviation>
<Minimum>false</Minimum>
<Maximum>false</Maximum>
<WriteOnPoll>false</WriteOnPoll>
<RollupStrategy>Sum</RollupStrategy>
<Percentile>0</Percentile>
</Attribute>
<Attribute name="FramesOut" type="Double">
<Documentation>Frames sent since the virtual circuit was created</Documentation>
<AttributeDisplayName />
<Polled>true</Polled>
<Baseline>false</Baseline>
<IsDbColumn>true</IsDbColumn>
```



```

<Variance>false</Variance>
<StandardDeviation>false</StandardDeviation>
<Minimum>false</Minimum>
<Maximum>false</Maximum>
<WriteOnPoll>false</WriteOnPoll>
<RollupStrategy>Sum</RollupStrategy>
<Percentile>0</Percentile>
</Attribute>
<Attribute name="BytesIn" type="Double">
<Documentation>Bytes received since the virtual circuit was created</Documentation>
<AttributeDisplayName />
<Polled>true</Polled>
<Baseline>false</Baseline>
<IsDbColumn>true</IsDbColumn>
<Variance>false</Variance>
<StandardDeviation>false</StandardDeviation>
<Minimum>false</Minimum>
<Maximum>false</Maximum>
<WriteOnPoll>false</WriteOnPoll>
<RollupStrategy>Sum</RollupStrategy>
<Percentile>0</Percentile>
</Attribute>
<Attribute name="BytesOut" type="Double">
<Documentation>Bytes sent since the virtual circuit was created</Documentation>
<AttributeDisplayName />
<Polled>true</Polled>
<Baseline>false</Baseline>
<IsDbColumn>true</IsDbColumn>
<Variance>false</Variance>
<StandardDeviation>false</StandardDeviation>
<Minimum>false</Minimum>
<Maximum>false</Maximum>
<WriteOnPoll>false</WriteOnPoll>
<RollupStrategy>Sum</RollupStrategy>
<Percentile>0</Percentile>
</Attribute>
</AttributeGroup>
<Attribute name="SourceFacetTypes" cached="true" list="true" persistent="true"
  type="QName">
<Documentation />
</Attribute>
<Expressions>
<ExpressionGroup destCert="{http://im.ca.com/core}Item">
<Expression destAttr="Name">Names</Expression>
</ExpressionGroup>
<ExpressionGroup destCert="{http://im.ca.com/inventory}DeviceComponent">

```

```

<Expression destAttr="IndexList">Indexes</Expression>
</ExpressionGroup>
</Expressions>
</FacetType>
</DataModel>

```

## Basic Properties

The basic properties of your custom metric family help to distinguish it from other custom metric families you create.

Consider the following restrictions when you determine basic properties:

- The FacetType/name, DisplayName, and TableName properties must be unique for each metric family.
- The Protocol tag is always IMDB.
- The Normalized tag is always true.
- Set the FacetType/descriptorClass property and all DataModel and FacetOf properties.

The following list details basic metric properties:

### NOTE

Unless stated otherwise, you can update all entries in this list.

- **FacetType/name**

Specifies the metric family name. For each metric family, the name must be a unique name that identifies it internally within the system. Carefully select a name with a minimal possibility of naming conflicts with future similar metric families. For example, define a naming scheme that ensures that these metric family names are unique.

### NOTE

DX NetOps Performance Management does not expose this name externally. To display a metric family name in NetOps Portal, use the `DisplayName` element.

### Updatable: No

**Possible values:** Alphanumeric and underscore. Dot (.) and dash (-) are not permitted. The value must be unique across all metric families.

- **Documentation**

Specifies the external description for the metric family. To make these comments useful, describe why and when you added or changed the metric family.

**Possible values:** Plain text

**Effect of updating:** None

### NOTE

List this property first under the `FacetType/name`.

- **FacetOf**

Asserts that this metric family is an item.

**Possible values:** `namespace="http://im.ca.com/core" name="Item"`

- **DisplayName**

Specifies the metric family name that displays in NetOps Portal.

**Possible values:** Plain text

### IMPORTANT

Ensure that this property is unique to the metric family.

**Effect of updating:** Change to the name in the Administrator section of NetOps Portal.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** Refresh NetOps Portal.

### NOTE

List this property before the `AttributeGroup` property.

- **TableName**

Specifies the database table name that is used to store the metrics that the metric family collects.

**Possible values:** Uppercase alphanumeric and underscore. The value must begin with a letter. The value must be unique across all metric families.

**Example:** `PROCESS_STATS`

**Effect of updating:** Poll data is stored in a new set of database tables.

#### IMPORTANT

If you update this property, the old poll data is lost. Old report views are broken.

**When does the update take effect:** Immediately. Before you can create new views, there is a delay of up to 5 minutes while DX NetOps Performance Management loads the new MIB files.

**Required actions for updates to take effect:** Views must be recreated.

- **Protocol**  
The Protocol tag is always IMDB.
- **Normalized**  
The Normalized tag is always true.
- **SupportsDeviceAggregation**  
Supports thresholding at the device level for some metrics.

#### NOTE

You cannot use this property for custom or extended certification.

### ComponentFacets

The ComponentsFacets section lists the facets that are created during discovery. Discovery identifies items as device components or creates a hierarchy relationship between items.

- **Facet**  
Specifies a facet that is attached to the *component* item during component discovery.  
**Updatable:** Yes  
**Possible values:** QName of the facet  
**Effect of updating:** If the component facet is synchronized to NetOps Portal, the component is visible in NetOps Portal.  
**When does the update take effect:** Rediscover  
**Required actions for updates to take effect:** Delete the device and rediscover.

### ItemFacets

The ItemFacets section lists the facets that are created during discovery that identify items as devices.

#### IMPORTANT

This section supports complex metric family structures. Do not use this section.

- **Facet**  
Specifies a facet that is attached to the *item* during discovery.  
**Updatable:** Yes  
**Possible values:** QName of the facet.  
**Effect of updating:** Component is visible on the REST service for the specified facet. If the component facet is synchronized to NetOps Portal, the component is visible in NetOps Portal.  
**When does the update take effect:** Rediscover  
**Required actions for updates to take effect:** Delete the device and rediscover.

#### Example:

```
<ItemFacets>
  <Facet>{http://im.ca.com/inventory}Host</Facet>
  <Facet>{http://im.ca.com/inventory}Device</Facet>
```

```

    <Facet>{http://im.ca.com/inventory}ConsolidatedAndDiscoveredMetricFamilyHistory</Facet>
    <Facet>{http://im.ca.com/core}Syncable</Facet>
    <!-- The IPDomainID attribute will be filled in by discovery -->
    <Facet>{http://im.ca.com/core}IPDomainMember</Facet>
</ItemFacets>

```

### **AttributeGroup (Metric Family)**

An `AttributeGroup` is a collection of item discovery attributes and metric attributes. The item discovery attributes are set during discovery, like item descriptions. The metric attributes are collected during polling. The following information describes the elements that you use in the `AttributeGroup` section.

Set the `AttributeGroup/list` and `AttributeGroup/external` properties to true. These properties specify that each attribute represents a list of values that is obtained from an external source. Customize the following XML elements:

#### **NOTE**

You can update the entries in this list using plain text. The update has no behavioral impact.

- **AttributeGroup/name**  
Specifies the attribute group name. Conform to the "<FacetType/name>Group" naming scheme.
- **Documentation**  
(Optional) Specifies the description for the attribute group.

This includes:

- [General Attributes \(Metric Family\)](#)
- [Discovery Attributes](#)
- [Polled and Baseline Attributes](#)
- [SourceFacetTypes Attribute](#)

### **General Attributes (Metric Family)**

The general attributes for all metric families are as follows:

#### **NOTE**

You can update the entries in this list. Unless specified otherwise, the update takes effect immediately and no actions are required to trigger the update.

- **Attribute/name**  
Specifies the unique, internal name. For metrics, this name is also used for naming the database column.

#### **NOTE**

DX NetOps Performance Management does not expose this name externally. To display an attribute name in NetOps Portal, use the `AttributeDisplayName` element. To change this attribute, [update the metric family properties](#).

**Possible values:** Alphanumeric and underscore.

**Effect of updating:** For metrics, the values for this attribute are stored in a new database column corresponding to the updated name. The user loses the historical data that is collected for this metric (with the older name). The custom reports reporting on this metric fails.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **Attribute/type**  
Indicates the data type of this attribute. The most frequently used data types are Int, Long, Double, String, or ObjectID. The database stores metric attributes as a float. Therefore, these attributes must use a numeric type (use a Double). Other types are used for item attributes.

**Possible values:** Boolean, Int, Long, Double (floating-point), BigInteger, String, DateTime, IPAddress, MACAddress, IPSubnet, OctetString (hex representation), ObjectID, ItemID, QName (Qualified Name)

**NOTE**

The type names are case insensitive, for example, "boolean" is the same as Boolean.

**Effect of updating:** For metrics, none. All metrics are stored in the database as a float. For item attributes, the device must be deleted and rediscovered.

**When does the update take effect:** For metrics, next poll. For item attributes, on rediscover.

**Required actions for updates to take effect:** For metrics, none. For item attributes, delete the device and rediscover.

- **Documentation**

Displays the attribute description in NetOps Portal. The documentation is also displayed in tool tips when you hover the cursor over the attribute name.

**Possible values:** Plain text

**Effect of updating:** Hovering the cursor over the attribute name shows the updated documentation.

**NOTE**

List this property first under the `Attribute/name` and `Attribute/type`. This property and the `AttributeDisplayName` property must be listed before the other attribute properties in the XML.

- **AttributeDisplayName**

Specifies the name of the attribute in the UI. To change this attribute, see [update the metric family properties](#).

**Possible values:** Alphanumeric, space, and underscore.

**Effect of updating:** The metric reflects the updated `AttributeDisplayName` in the Metric Families UI and custom reports.

**NOTE**

This property and the `Documentation` property must be listed before the other attribute properties in the XML.

- **AttributeAbbreviation**

This parameter is not supported.

- **Polled**

Indicates whether the attribute is polled. If it is set to false, it is only accessed during discovery.

**Possible values:** `true`, `false`

**Effect of updating:** If set to false, the OIDs corresponding to this attribute/metric are not polled when no other polled attribute/metric is using that OID in its expression. If set to true, the OIDs corresponding to this attribute/metric are polled.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **IsDbColumn**

Stores its value in the database table. `IsDbColumn` is used for metric attributes. Set the `IsDbColumn` value to true when `Polled` is set to true.

**Possible values:** `true`, `false`

**Effect of updating:** If set to false, the data for this attribute/metric is not stored in the database. If set to true, the data for this attribute/metric is stored in the database.

## Discovery Attributes

For many attributes only the value that is retrieved during discovery is stored in the database. No further polling or processing, such as an evaluation of a baseline, is performed.

The Indexes and Names attributes must exist for all metric families. The Descriptions attribute is optional.

```
<Attribute name="Indexes" type="ObjectID[]" />
<Attribute name="Names" type="String" />
<Attribute name="Descriptions" type="String" />
```

The metric families supporting Hierarchy must include these attributes:

```
<Attribute name="ItemUniqueIDs" type="String" />
<Attribute name="ParentUniqueIDs" type="String" />
```

## Polled and Baseline Attributes

The following information describes the polled and baseline attribute elements:

### NOTE

You can update the entries in this list.

- **Baseline**

Indicates whether to calculate a mean value for this attribute. If you set this attribute to true, you must define a corresponding `BaselineList` definition.

### NOTE

This attribute requires that the `StandardDeviation` attribute be set to true.

**Possible values:** `true`, `false`

**Effect of updating:** Baseline values are calculated when true.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **Maximum**

Indicates whether to calculate the maximum of this attribute during the rollup. Creates a 'max\_' column in the database table. If `RollupStrategy` is defined, this attribute must also be defined.

**Possible values:** `true`, `false`

**Effect of updating:** True provides a calculation of, and a reporting field for maximum.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **Minimum**

Indicates whether to calculate the minimum of this attribute during the rollup. Creates a 'min\_' column in the database table. If the `RollupStrategy` attribute is defined, you must also define this attribute.

**Possible values:** `true`, `false`

**Effect of updating:** True provides a calculation of, and a reporting field for minimum.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **StandardDeviation**

Indicates whether to calculate the standard deviation of this attribute during the rollup. Creates a 'std\_' column in the database table. If `RollupStrategy` is defined, this attribute must also be defined.

**Possible values:** `true`, `false`

**Effect of updating:** True provides a calculation of, and a reporting field for standard deviation.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **DeviationFromBaseline**

Requires that the `Baseline` attribute is set to true. Provides the extra **Average Baseline** and **Percent Deviation** reporting fields, calculated using baseline data. These fields are not available for building custom views. No changes are made to the database table.

**Possible values:** `true`, `false`

**Effect of updating:** True provides the **Average Baseline** and **Percent Deviation** fields for the internal report development.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** None

- **Percentile**

Indicates whether to calculate the 95th percentile of this attribute during the rollup. Creates a 'pct\_' column in the database table. If `RollupStrategy` is defined, this attribute must also be defined.

**Possible values:** 0, 95

**Effect of updating:** A value of 95 provides a calculation of, and the **95th Percentile** reporting field. Zero specifies that no calculation is performed, and the reporting field is not available.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **Percentile2** (Optional) Specifies the value of a user-configurable percentile.

**Possible values:** 0-99

**Effect of updating:** A non-zero value specifies the percentile to calculate. Zero specifies that no calculation is performed, and the reporting field is not available.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **Percentile3** (Optional)

Specifies the value of a user-configurable percentile.

**Possible values:** 0-99

**Effect of updating:** A non-zero value specifies the percentile to calculate. A zero value specifies that no calculation is performed, and the reporting field is not available.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

#### NOTE

You cannot set two percentiles to the same non-zero value. You cannot set `Percentile2` or `Percentile3` to 95.

#### IMPORTANT

Changes to `Percentile2` and `Percentile3` attributes can cause discontinuity in the trend view depending on the time range.

- **ProjectionPercentile**

(Optional) Specifies the percentile to calculate for metric projection.

**Possible values:** 0-99

**Effect of updating:** Changes the percentile to use to calculate projections. Zero specifies that no calculation is performed.

#### IMPORTANT

Changes to this attribute can cause inaccurate projections for up to 90 days. When you change the value of this attribute, the percentile values for days before the change are not recalculated.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **AggregateToDevice**

Supports thresholding at the device level for some metrics.

#### NOTE

You cannot use this attribute for custom or extended certification.

- **RollupStrategy**

Specifies the operation that is performed every cycle during the rollup of the individually polled values. When `Polled` and `IsDbColumn` are set to true, this element is required.

**Possible values:** Sum (a summation for counters), Avg (an average for gauges)

**Effect of updating:** The specified strategy is used to perform rollup calculations.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

- **Rate**

Provides the extra **Average Rate** reporting field, calculated as AVG (metric value / time). No changes are made to the database table.

#### NOTE

The Rate is available for reporting but not for use when monitoring the profile event rules.

**Possible values:** `true`, `false`

**Effect of updating:** Provides the **Average Rate** reporting field.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** None

- **Units**

Specifies the name of the units label used in reports. The actual label that is displayed is translated according to the language setting of the report.

**NOTE**

If you do not define this attribute, the units labels on reports is 'Units'.

**Possible values:** Percent, Packets, PacketsPerSecond, DiscardedPackets, ErroredPackets, Bits, BitsPerSecond, Bytes, BytesPerSecond, Seconds, Microseconds, Milliseconds, UnixTime, Observations, FramesPerSecond, Frames, RequestsPerSecond, Requests

**Effect of updating:** The specified units label is displayed in reports.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** None

**Example: Using polled and baseline attribute elements**

```
<Attribute name="Utilization" type="double">
  <AttributeDisplayName>Utilization</AttributeDisplayName>
  <IsDbColumn>true</IsDbColumn>
  <Baseline>true</Baseline>
  <Minimum>true</Minimum>
  <Maximum>true</Maximum>
  <RollupStrategy>Avg</RollupStrategy>
  <StandardDeviation>true</StandardDeviation>
  <DeviationFromBaseline>true</DeviationFromBaseline>
  <Percentile>95</Percentile>
  <Polled>true</Polled>
  <Units>Percent</Units>
</Attribute>
```

### **SourceFacetTypes Attribute**

A `SourceFacetTypes` attribute is required for discovery. You must define this attribute.

Use these required values:

- **Name:** SourceFacetTypes
- **Type:** QName
- **Cached:** true
- **Persistent:** true
- **List:** true

**Example:**

```
<Attribute name="SourceFacetTypes" type="QName" cached="true" persistent="true" list="true" />
```

**Updatable:** No

### **Expressions**

The `Expressions` section is composed of `ExpressionGroup` tags that are used for component discovery. During the component discovery, the values for the component item properties (such as the `IndexList`, `Name`, and `Description`) are calculated. The vendor certification expressions supporting the metric family expressions are used for this calculation.



**NOTE**

The metric family and vendor certification `ExpressionGroup` tags are different than these tags.

The `ExpressionGroup` tags for the following `DestCert` URIs must exist:

DestCert	DestAttr
{http://im.ca.com/core}Item	Name
{http://im.ca.com/core}Item	Description
{http://im.ca.com/inventory} DeviceComponent	IndexList

**NOTE**

You can update the entries in this list.

- **ExpressionGroup/name**  
(Optional) Specifies the expression group name.  
**Possible values:** Plain text  
**Effect of updating:** None
- **ExpressionGroup/destCert**  
Specifies the component facet that contains the `destAttr` to populate. The facet name typically comes from the `ComponentFacets` section, except the `Item` and `DeviceComponent` facets.  
**Possible values:** Facets that are defined in `ComponentFacets`, or the `Item`, `DeviceComponent` facet.  
**Effect of updating:** Changes permissible expression `destAttr`  
**When does the update take effect:** Component Rediscovery  
**Required actions for updates to take effect:** None
- **ExpressionGroup/Expression**  
Specifies the expression for the component facet attribute.  
**Possible values:** Any valid metric  
**Effect of updating:** Changes permissible expression `destAttr`  
**When does the update take effect:** Component Rediscovery  
**Required actions for updates to take effect:** None
- **ExpressionGroup/Expression/destAttr**  
Specifies the component facet attribute name.  
**Possible values:** Any valid attribute from that component facet.  
**Effect of updating:** Changes the attribute name  
**When does the update take effect:** Component Rediscovery  
**Required actions for updates to take effect:** None

**Hierarchy**

You can define a hierarchy, or parent-child relationship, between items of different metric families, for example, `Interface` and `CBQoS Classmap`. In the metric family definitions, the `Hierarchy` must be specified in the child metric family with:

- The `Hierarchy QName` in the `ComponentFacets`
- The `ItemUniqueID` and `ParentUniqueID destAttr` values in the `Hierarchy ExpressionGroup`
- The `ItemUniqueIDs` and `ParentUniqueIDs` attributes in the `AttributeGroup`

The supporting expressions are defined in the vendor certifications.

The `Hierarchy ExpressionGroup` tags for the following `DestCert` URIs must exist:

DestCert	DestAttr
{http://im.ca.com/inventory}Hierarchy	ItemUniqueID

{http://im.ca.com/inventory}Hierarchy	ParentUniqueID
---------------------------------------	----------------

**Example:**

```
<ComponentFacets>
  <Facet>{http://im.ca.com/inventory}QoSClassMap</Facet>
  <Facet>{http://im.ca.com/inventory}Hierarchy</Facet>
</ComponentFacets>
<ExpressionGroup name="Hierarchy" destCert="{http://im.ca.com/inventory}Hierarchy">
  <Expression destAttr="ItemUniqueID">ItemUniqueIDs</Expression>
  <Expression destAttr="ParentUniqueID">ParentUniqueIDs</Expression>
</ExpressionGroup>
<AttributeGroup name="QosCosGroup" list="true" external="true">
  <Attribute name="ItemUniqueIDs" type="String" />
  <Attribute name="ParentUniqueIDs" type="String" />
  ...
</AttributeGroup>
```

**BaselineDefinitions**

The **BaselineDefinitions** section contains the baseline definitions to calculate for this metric family. You must specify a baseline definition for each metric in the **AttributeGroup** section whose **Baseline** property is set to true.

You can define an Hourly (required) and Daily (optional) baseline. Hourly baselines are used both for event processing and for displaying baselines in reports. Daily baselines are used for displaying baselines in reports with a time frame of one month or greater.

The following information describes the baseline elements used:

- **Name**  
Specifies the type of baseline definition for a metric. The type is either hourly or daily.  
**Updatable:** No  
**Possible values:** HourlyBaseline , DailyBaseline
- **ID**  
Specifies a value that is no longer used. However, you must specify the field as a positive integer and it must be unique across all hourly and daily baseline definitions within this metric family.  
**Updatable:** Yes  
**Possible values:** Any unique and positive integer.  
**Effect of updating:** None
- **PerformanceMetric**  
Specifies the name (case sensitive) of the metric for which the baseline is calculated. Set the Polled and Baseline properties for the metric attribute to true.  
**Updatable:** Yes  
**Possible values:** A valid metric name (case sensitive).  
**Effect of updating:** Baseline calculations are performed for the metric.  
**When does the update take effect:** Next baseline calculation, either hourly or daily.  
**Required actions for updates to take effect:** None
- **Period**  
Specifies the type of baseline calculation, hourly or daily. Specify the value "1 Hour" for an HourlyBaseline Name or "1 Day" for a DailyBaseline Name.  
**Updatable:** Yes  
**Possible values:** 1 Hour, 1 Day  
**Effect of updating:** Baseline calculations are performed either hourly or daily.

**When does the update take effect:** Next baseline calculation, either hourly or daily

**Required actions for updates to take effect:** None

- **ProjectionInterval**

(Optional) Specifies the projection period in days for metric projection. Specify this expression in `DailyBaseline`.

**Updatable:** Yes

**Possible values:** 0 or any positive integer

**Effect of updating:** Changes the period that the system calculates projection values for. Zero specifies that no calculation is performed.

**When does the update take effect:** Next baseline calculation.

**Required actions for updates to take effect:** `PredictionInterval` requires that `ProjectionPercentile` is set to a non-zero value.

- **ProjectionInterval2**

(Optional) Defines a second projection interval. For details, see **ProjectionInterval**.

- **ProjectionInterval3** (Optional) Defines a third projection interval. For details, see **ProjectionInterval**.

- **Window**

Specifies a value that is no longer used.

**Updatable:** No

**Possible values:**

- **30 Days:** Specify for hourly baselines.
- **90 Days:** Specify for daily baselines.

- **StartDate, EndDate, DaysOfWeek**

Specifies more values that are not used, but they must be specified as 0 (zero).

**Updatable:** No

**Possible values:** 0

## **ComponentReconciliation**

The following information defines the component reconciliation logic that is used in component discovery. First, this information determines whether the system has already discovered this component or not. Then, the reconciliation logic determines whether to update an existing component or create a new one.

Metric families can include multiple ordered algorithms. If a metric family does not define a reconciliation algorithm, a default one with match attribute `Item.Name` is applied.

## **ItemReconciliation**

The following information defines the item reconciliation logic that is used in item discovery. The logic determines whether the system has already discovered an item or not. Based on this determination, an existing item is updated or a new item is created. Item reconciliation is similar to component reconciliation. However, item reconciliation is used for items that are not components, such as virtual hosts. The `ItemFacets` are added to any new items or to any matching items (if the facets do not exist).

### **Example:**

```
<ItemReconciliation>
<SourceAgentScopedReconciliation>
<MatchAlgorithms>
<ExactMatch>
<MatchAttribute name="{http://im.ca.com/inventory}SourceAgentInfo.SourceAgentIndexes" />
</ExactMatch>
</MatchAlgorithms>
</SourceAgentScopedReconciliation>
<GlobalScopedReconciliation matchDevices="true" />
```

```
</ItemReconciliation>
```

- **SourceAgentScopedReconciliation**

Defines the match algorithms that are used to reconcile items.

**Updatable:** Yes

**Effect of updating:** Changes the item reconciliation logic.

**When does the update take effect:** Rediscovery

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

- **GlobalScopedReconciliation**

Defines the match algorithms that are used when items could not be reconciled for the source agent. The `GlobalScopedReconciliation` algorithms are used to locate items that have been created for other agents but match the potential new items. If the `matchDevices` property is set to true, the system default (built-in, not visible in XML) `ComponentReconciliationInfo` match algorithm is used. The match algorithm is based on the device primary IP address and host name.

**Updatable:** Yes

**Effect of updating:** Changes the item reconciliation logic.

**When does the update take effect:** Rediscovery

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

## MatchAlgorithms

Component reconciliation and item reconciliation use match algorithms. You can use the following algorithms:

- **ExactMatch**

All specified attributes must be matched to say the item matches with the new data.

- **BestOfMatch**

Users must specify the least number of attributes to be matched by using the "leastMatchCount" value. Also, each attribute has a "required" key property. If the required property is set to true, that attribute must be matched to be considered a match.

The algorithm has a match precedence when multiple algorithms are provided for a metric family. The order of the algorithms determines the precedence. The algorithm at the top has the highest precedence. The bottom one has the lowest precedence.

Each algorithm must have at least one matching attribute. When data matches to multiple items with the same algorithm, the item with the most matched attributes wins. When multiple matched items have the same number of matched attributes, the winner is picked at random from these items.

## Examples: How the reconciliation works

Two match algorithms are provided for a metric family: alg1 and alg2. Alg1 has higher precedence than alg2. The metric family has three existing component items: 1, 2 and 3. Rediscovering the metric family finds three entries: A, B, and C. Now, we apply the two algorithms to determine which entry is new, changed, and unchanged.

Reconciliation Meta Data	New Data	Existing Components
<ComponentReconciliation>	A	1
<MatchAlgorithms>	B	2
<MatchAlgorithm1>	C	3
<MatchAttribute name="attr2"/>		
</MatchAlgorithm1>		
<MatchAlgorithm2>		
<MatchAttribute name="attr1"/>		
<MatchAttribute name="attr3"/>		

```

<MatchAttribute name="attr4"/>
</MatchAlgorithm2>

</MatchAlgorithms>
</ComponentReconciliation>

```

MatchAlgorithm1 and MatchAlgorithm2 can be either ExactMatch or BestOfMatch . The order of the two match algorithms tells us that MatchAlgorithm1 has a higher precedence than MatchAlgorithm2 .

- **Case 1: Unique 1-to-1 Match**

Entry A matches to item 1, and item 1 does not have any other match.

```
A -----> 1
```

This example is the simplest case. This match is unique, so it does not matter if it also matches alg1 or alg2. Entry A matches item 1.

A good match algorithm produces more unique matches.

- **Case 2: One entry has multiple matches**

Entry A matches to item 1 by alg1 and also matches to 2 by alg2.

```
---> 1 (alg1)  (1 wins)
```

```
/
```

```
A                Since alg1 has higher precedence, item 1 wins the match.
```

```
\
```

```
---> 2 (alg2)
```

- **Case 3: Multiple entries match to the same item with different algorithms**

Entry A matches to 1 by alg1 and entry B also matches to item 1 by alg2.

```
A -----> 1 (alg1)  (A wins)
```

```
B -----> 1 (alg2)
```

Since alg1 has higher precedence, entry A wins.

- **Case 4: Multiple entries match to the same item with same algorithm but different numbers of matched attributes**

Both A and B match to 1 by alg1.

```
A -----> 1 (alg1, # of matched attrs: 2)  (A wins)
```

```
B -----> 1 (alg1, # of matched attrs: 1)
```

Because A has more matched attributes, A wins.

If the number of match attributes is the same, the winner is randomly picked and a warning is generated.

- **Case 5: Mixed match 1**

```
alg1
```

```
A -----> 1
```

```
/  alg2(match attr count: 3)
```

```
B
```

```
\  alg2(match attr count: 2)
```

```
-----> 2
```

A matches to 1 because it matches with a higher precedence algorithm.

B matches to 2 because 1 has matched to A.

- **Case 6: Mixed match 2**

```
-----> 3
```

```
/  alg1
```

```
A                ==> A wins 3 because alg1 has a higher matching precedence
```

```
\  alg2
```

```

-----> 1
/  alg2
B           ==> B wins 2 because alg1 has a higher matching precedence
\  alg1
-----> 2
/  alg2
C           ==> C has no match because 2 is matched to B and 3 is matched to
A
\  alg2
-----> 3

```

Entry C is treated as a new component. 1 is considered as an unmatched item.  
The more match case 1 (unique match), the better the match algorithm is.

**Updatable:** Yes

**Effect of updating:** Changes the component reconciliation logic.

**When does the update take effect:** Rediscovery

**Required actions for updates to take effect:** Update metric family or change vendor certification priority.

### **ReconfigDetectionAttr**

The `ReconfigDetectionAttr` element defines a metric family attribute that is used for change detection. You can enable the change detection on a monitoring profile. Data Aggregator polls that attribute only to verify whether the target device has changed, instead of doing a complete rediscovery. This feature helps performance and helps reduce the network traffic.

- **Updatable:** Yes
- **Possible values:** The full name of the metric family attribute. The specified metric family attribute flag must have the cached, persistent, and external flags set to true.
- **Effect of updating:** Change to component reconfiguration detection.
- **When does the update take effect:** After rediscovery
- **Required actions for updates to take effect:** Update metric family or change vendor certification priority.

### **Tags in Custom and Extended Metric Families**

The following tags have restricted use in custom and extended metric families:

- **Variance**  
The value of this tag must be `false`.
- **RollupExpression**  
Remove this tag or the value must be blank.
- The **Normalized** tag is always `true`.

## **Component XML Structure**

A device component uses XML to define a class of component items that are associated with a device. DX NetOps Performance Management provides out-of-the-box components, but you typically define a custom component for a custom metric family. Device components can define an optional `ItemSyncDefinition` attribute, which synchronizes component items to NetOps Portal. You can then view the components in Inventory lists, groups, and context pages.

### **NOTE**

You must list some properties in the XML in a particular order. The properties included in the XML example and listed in the following descriptions are presented in the required order.

### **Example:**

This example shows a custom component named `frPVC` :

#### NOTE

If you view the component XML in a browser, certain tags are hidden. For this reason, copy and paste the component XML only from a REST client.

```
<?xml version="1.0" encoding="UTF-8"?>
  <!--Auto-generated by the type catalog local manager.-->
  <DataModel namespace="http://im.ca.com/inventory" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance" xsi:noNamespaceSchemaLocation="ComponentFacet.xsd">
    <Author>Custom</Author>
    <Version>1.0</Version>
    <FacetType name="frPVC">
      <Documentation>A Frame Relay PVC</Documentation>
      <FacetOf namespace="http://im.ca.com/core" name="Item" />
      <Component>true</Component>
      <ItemSyncDefinition itemTypeName="component" itemSubtypeName="frpvc"
itemTypeLabel="FrameRelayPVC" itemTypeLabelPlural="FrameRelayPVCs" categorize="false"
groupBy="false" context="true">
        </ItemSyncDefinition>
      </FacetType>
    </DataModel>
```

You can navigate to this page from the `frPVC` device component where it appears in the **Inventory, Device Components** list.

In this article:

- [Basic Properties](#)
- [The ItemSyncDefinition Property](#)
- [Remove an ItemSyncDefinition](#)
- [Unsupported Properties for Custom Components](#)

### **Basic Properties**

The following basic properties of your custom component help NetOps Portal to distinguish it from other custom components that you create:

#### NOTE

From the following list, you can update only the `documentation` property.

- **FacetType name**

Specifies the component name. Each component must have a unique name that identifies it internally within the system. Carefully choose a name with a minimal possibility of naming conflicts with future similar components. For example, define a naming scheme that ensures that these component names are unique.

#### NOTE

DX NetOps Performance Management does not expose this name externally. To have NetOps Portal display a component name, use the `ItemSyncDefinition` attribute's `itemTypeLabel` and `itemTypeLabelPlural` properties.

**Possible values:** Alphanumeric and underscore (dot and dash are not permitted)

- **Documentation**

Specifies internal comments for the component, in plain text. To make these comments useful, describe why and when you added or changed the component.

**Effect of updating:** None

**When does the update take effect:** Immediately

**Required Actions for updates to take effect:** None

**NOTE**

List this property first under the `FacetType/name` in the XML. List this property and the `FacetOf` property before the `Component` and `ItemSyncDefinition` properties in the XML.

- **FacetOf**

Asserts that this component is an item.

**Possible values:** `namespace="http://im.ca.com/core" name="Item"`

**NOTE**

List this property and the `Documentation` property *before* the `Component` and `ItemSyncDefinition` properties in the XML.

- **Component**

Asserts that this item is a component.

**Possible values:** `true`

### **The ItemSyncDefinition Property**

The `ItemSyncDefinition` attribute is optional. This attribute specifies how NetOps Portal synchronizes and displays the component items. If you do not specify the component items, NetOps Portal does not display them in Inventory lists (for example, Device Components). But, NetOps Portal can still report them in custom views.

**NOTE**

Unless otherwise noted, there are no required actions for the updates to take effect.

- **ItemSyncDefinition/itemTypeName**

Specifies the item type. For custom components, this value must be set to `Component`.

**Updatable:** No

**Possible values:** `Component`

- **ItemSyncDefinition/itemSubtypeName**

Specifies the internal name of the component in NetOps Portal. This value must be unique for all components.

Use a naming convention that avoids conflicts with out-of-the-box and custom components, such as using a prefix representing your organization, for example, `acmeFan`.

**Updatable:** No

**Possible values:** Alphanumeric, unique for all components

- **ItemSyncDefinition/itemTypeLabel**

Specifies the label that NetOps Portal uses when displaying a single component of this type. For example, NetOps Portal uses this value in the **Inventory** (menu), Device Components, **Type** column.

**Updatable:** Yes

**Possible values:** Plain text, unique for all components

**Effect of updating:** Label is displayed in NetOps Portal Inventory user interfaces.

**When does the update take effect:** Allow for resync to occur and up to 15 minutes to complete updates.

- **ItemSyncDefinition/itemTypeLabelPlural**

Specifies the label that NetOps Portal uses when displaying multiple components of this type. NetOps Portal uses this value for the **Inventory** (menu) (see `groupBy`) and the Group name (see `categorize`).

**Updatable:** Yes

**Possible values:** Plain text, unique for all components

**Effect of updating:** NetOps Portal displays this label in **Inventory**.

**When does the update take effect:** Allow for resync to occur and up to 15 minutes to complete updates.

- **ItemSyncDefinition/categorize**

Specifies whether NetOps Portal creates an inventory group under **Inventory** (menu), **All Items** to contain all items of this component type. The group is named `{itemTypeLabelPlural}`.



**NOTE**

Do not use this inventory group for reporting dashboards. This group is for inventory purposes only. If you use this group for reporting, then no data displays. Otherwise, you can use device-based inventory groups (under **Inventory** (menu), **All Items**) for reporting, such as Routers and Servers. However, you cannot use component-based inventory groups, such as Device Components.

**Updatable:** Yes

**Possible values:**

- **true:** NetOps Portal creates an inventory group under **Inventory** (menu), **All Items**. For NetOps Portal group administrators, on the **Manage Groups** page, NetOps Portal creates the group under the **Inventory** (menu), All Items, *{itemTypeLabelPlural}*.
- **false:** NetOps Portal removes the inventory group from **Inventory** (menu), **All Items**.

**When does the update take effect:** After a resync, and up to 30 minutes to complete updates.

**Required actions for updates to take effect:** Items typically appear in the group within 30 minutes. If they do not, [manually resynchronize the data aggregator datasource](#).

- **ItemSyncDefinition/groupBy**

Specifies whether NetOps Portal creates the inventory menu item (under **Inventory** (menu)) for viewing all items of this component item type, named *{itemTypeLabelPlural}*. The component type shows up in the **Context Type** drop-down list when setting a view context. The `groupBy` property does not create a group (see **categorize**).

**Updatable:** Yes

**Possible values:**

- **true:** NetOps Portal creates the inventory menu item (under **Inventory** (menu)) for viewing all items of this component item type.
- **false:** NetOps Portal removes the inventory menu item. The components are listed in the **Inventory** (menu), **Device Components** table with the type of *{itemTypeLabel}*.

**When does the update take effect:** Allow for resynchronization to occur and allow up to 15 minutes to complete updates.

- **ItemSyncDefinition/context**

Specifies whether NetOps Portal makes each component item name a context hyperlink in the Inventory component views that, when clicked, navigates the user to the individual component context page.

**Updatable:** Yes/No

**Possible values:**

- **true:** NetOps Portal creates links to custom context pages. You can also select the metric family as a “context,” which makes the custom metric family available in Dynamic Trend charts.
- **false:** NetOps Portal removes the links from custom context pages.

**When does the update take effect:** Allow for resynchronization to occur and allow up to 15 minutes to complete updates.

## **Remove an ItemSyncDefinition**

You can completely remove the `ItemSyncDefinition` section from the XML.

**Follow these steps:**

1. Remove the `ItemSyncDefinition` section, including the `ItemSyncDefinition` property's start and end tags.
2. Apply the component change to the data aggregator.
3. [Synchronize the data aggregator data source](#).

Allow between 15-30 minutes for the changes to synchronize. When this process completes, all NetOps Portal behaviors that the `ItemSyncDefinition` property defined for the component are removed.

## Unsupported Properties for Custom Components

You cannot use the following properties for custom components in your XML:

- Attribute
- WebService
- ItemSyncDefinition/isDeviceComponent
- ItemSyncDefinition/mapped
- ItemSyncDefinition/ItemProperty

## Self-Certification Workflows

Self-certification workflows are examples of real use cases that you might encounter.

Use the following workflows as models to help you with self-certification:

- [Create a Component](#)
- [Create a Metric Family](#)
- [Create a Vendor Certification](#)

### Create a Component

You can introduce a managed item type using the components XML. The components XML defines how a discovered component is synced over to NetOps Portal. The XML creates a component in your inventory and is required only when a component does not exist already.

**Follow these steps:**

1. Customize the following XML:

```
<?xml version="1.0" encoding="UTF-8"?>
<!--Auto-generated by the type catalog local manager.-->
<DataModel namespace="http://im.ca.com/inventory"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="ComponentFacet.xsd">
<Author>Custom</Author><Version>Version</
Version>
<FacetType
name="Name">
<Documentation>Clear_Text_Description</Documentation>
<FacetOf namespace="http://im.ca.com/core" name="Item" />
<Component>true</Component>
<ItemSyncDefinition
itemTypeName="component" itemSubtypeName="Internal_Name"
itemTypeLabel="External_Name"
itemTypeLabelPlural="External_Names_Plural"
categorize="false"
groupBy="false"
context="true">
</ItemSyncDefinition>
</FacetType>
</DataModel>
```

- **Version**

Specifies the sequential number of the version definition.

Example: 1.0

- **Name**

Specifies the component name. Each component must have a unique name that identifies it internally within the system. Choose a name with a minimal possibility of naming conflicts with future similar components. For example, define a naming scheme that ensures that these component names are unique. This string is used in the Component Facet Definition in the metric family, and cannot be updated.

**Possible values:** Alphanumeric and underscore. Dot (.) and dash (-) are not permitted.

**NOTE**

This name is never exposed externally. To display a component name in the user interface, use the `ItemSyncDefinition`, `itemTypeLabel`, and `itemTypeLabelPlural` elements.

- **Clear\_Text\_Description**

Specifies internal comments for the component. To make these comments useful, describe why and when you added or changed the component.

**Possible values:** Plain text

**Effect of updating:** None

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** None

- **Internal\_Name**

Specifies the internal name of the component in NetOps Portal. This value must be unique for all components.

Use a naming convention that avoids conflicts with out-of-the-box and custom components, such as using a prefix representing your organization, `acmeFan`.

**Can be updated:** No

**Possible values:** Alphanumeric, unique for all components

- **External\_Name**

Specifies the user interface label that is used when displaying a single component of this type. For example, this value is used in the Inventory, Device Components UI “Type” column.

**Can be updated:** Yes

**Possible values:** Plain text, unique for all components

**Effect of updating:** Label is displayed in NetOps Portal Inventory user interfaces.

**When does the update take effect:** Allow for resynchronization to occur and up to 15 minutes to complete updates.

- **External\_Name\_Plural**

Specifies the user interface label that is used when displaying multiple components of this type. Used by the Inventory menu (see `groupBy`) and the Group name (see `categorize`).

**Can be updated:** Yes

**Possible values:** Plain text, unique for all components

**Effect of updating:** Label is displayed in NetOps Portal Inventory UIs.

**When does the update take effect:** Allow for resynchronization to occur and up to 15 minutes to complete updates.

- **categorize**

Instructs NetOps Portal to create an inventory group under Inventory, All Items. This group contains all items of this component type. The group is named `{itemTypeLabelPlural}`.

**NOTE**

The inventory group that is created cannot be used for reporting dashboards. Instead, this group is for inventory purposes only. If the group is selected for reporting, then no data displays. Other, device-based inventory groups (under Inventory, All Items) can be used for reporting, such as Routers and Servers. However, component-based inventory groups cannot, such as Device Components.

**Can be updated:** Yes

**Possible values:** true, false

**Effect of updating:** Creates or removes an Inventory group in NetOps Portal. For NetOps Portal group administrators, on the Manage Groups page, the group is created under Inventory, All Items, `{itemTypeLabelPlural}`.

**When does the update take effect:** Allow for resynchronization to occur and up to 30 minutes to complete updates.

**Required actions for updates to take effect:** Items typically appear in the group within 30 minutes. If they do not, manually resync the data aggregator data source. Select **Perform a Full Resynchronization**, and then click **Resync confirmation**.

– **groupBy**

Instructs NetOps Portal to create an inventory menu item (under Inventory) for viewing all items of this component item type. The menu is named `{itemTypeLabelPlural}`. This attribute also causes the component type show up in the Context Type drop-down list when setting a view context. When false, the components are listed in the Inventory, Device Components table with the type of `{itemTypeLabel}`. The `groupBy` property does not create a group (see categorize).

**Can be updated:** Yes

**Possible values:** true, false

**Effect of updating:** The menu item is created when true, or removed when false.

**When does the update take effect:** Allow for resynchronization to occur and allow up to 15 minutes to complete updates.

– **context**

Makes each component item name a context hyperlink in the Inventory component views that, when clicked, navigate to the individual component context page.

**Can be updated:** Yes/No

**Possible values:** Plain text

**Effect of updating:** Makes each component item name a context hyperlink in the Inventory component views.

**When does the update take effect:** Allow for resynchronization to occur and allow up to 15 minutes to complete updates.

**NOTE**

You can create a component without the `ItemSyncDefinition` tag. In this case, a component discovered in Performance Management for this managed item type is not synced over to NetOps Portal (like CPU, or memory). Although the component is not synced over, you can still report on these managed items at the device level.

2. Set up a REST client with a connection to the data aggregator server.

3. Use the following format for the URL in the REST client:

`http://DA_host:8581/typcatalog/components`

**NOTE**

To update or extend a vendor component, use the following format for the URL in the REST client:

`http://DA_host:8581/typcatalog/components/component_name`

• **component\_name**

Specifies the name of the existing component.

4. Select **POST** for the **HTTP Method**.

**NOTE**

To update an existing component, select **PUT**.

5. Select **application/xml** as the **Body Content-type**.

6. Add the customized XML within the "Body" text section.

7. Run the method.

## Create a Metric Family

A metric family describes the metrics available for collection for a component. Different vendors can provide information for all or part of the metrics in a metric family. For each metric, a metric family defines the behavior of the system. For

example, a metric family can define whether to calculate baselines, percentiles, minimum, maximum, or projections. They also define how to calculate rollups. Only some fields in the metric family XML require changes. The `Indexes`, `Names`, and `Description` attributes are required. You should add one or more metric attributes to reflect the number and type of metrics in the metric family.

### Follow these steps:

#### 1. Customize the following XML:

```
<?xml version="1.0" encoding="UTF-8" ?>
    <!-- Auto-generated by the type catalog local manager. -->
    <DataModel xmlns:xsi=http://www.w3.org/2001/
XMLSchema-instance namespace="http://im.ca.com/normalizer"
    xsi:noNamespaceSchemaLocation="IMDBCertificationFacet.xsd">
        <Author>Custom</Author>
        <Version>1.0</Version>
        <FacetType name="Internal_Metric_Family_Name"

descriptorClass="com.ca.im.core.datamodel.certs.NormalizedFacetDescriptorImpl">
    <Documentation>Metric_Family_Description</Documentation>
    <FacetOf namespace="http://im.ca.com/core" name="Item" />
    <AttributeGroup name="AttributeGroup" external="true" list="true">
    <Documentation />
    <Attribute name="Indexes" type="ObjectID[]">
    <Documentation />
    <Polled>false</Polled>
    <Baseline>false</Baseline>
    <IsDbColumn>false</IsDbColumn>
    <Variance>false</Variance>
    <StandardDeviation>false</StandardDeviation>
    <Minimum>false</Minimum>
    <Maximum>false</Maximum>
    <WriteOnPoll>false</WriteOnPoll>
    <RollupStrategy />
    <AttributeDisplayName />
    <Percentile>0</Percentile>
    </Attribute>
    <Attribute name="Names" type="String">
        <Documentation>Element/Component_Name</Documentation>
    <Polled>false</Polled>
    <Baseline>false</Baseline>
    <IsDbColumn>false</IsDbColumn>
    <Variance>false</Variance>
    <StandardDeviation>false</StandardDeviation>
    <Minimum>false</Minimum>
    <Maximum>false</Maximum>
    <WriteOnPoll>false</WriteOnPoll>
    <RollupStrategy />
    <AttributeDisplayName />
```

```

    <Percentile>0</Percentile>
  </Attribute>
  <Attribute name="Description" type="String">
    <Documentation>Element/Component_Description</Documentation>
    <Polled>>false</Polled>
    <Baseline>>false</Baseline>
    <IsDbColumn>>false</IsDbColumn>
    <Variance>>false</Variance>
    <StandardDeviation>>false</StandardDeviation>
    <Minimum>>false</Minimum>
    <Maximum>>false</Maximum>
    <WriteOnPoll>>false</WriteOnPoll>
    <RollupStrategy />
    <AttributeDisplayName />
    <Percentile>0</Percentile>
  </Attribute>
  <Attribute name="Metric_Name" type="MIB_Metric_Type">
    <Documentation>Metric_Description</Documentation>
    <Polled>>true</Polled>
    <Baseline>>false</Baseline>
    <IsDbColumn>>true</IsDbColumn>
    <Variance>>false</Variance>
    <StandardDeviation>>false</StandardDeviation>
    <Minimum>>false</Minimum>
    <Maximum>>false</Maximum>
    <WriteOnPoll>>false</WriteOnPoll>
    <RollupStrategy>Sum</RollupStrategy>
    <AttributeDisplayName />
    <Percentile>0</Percentile>
  </Attribute>
</AttributeGroup>
<Attribute name="SourceFacetTypes" cached="true" list="true"
persistent="true" type="QName">
  <Documentation />
</Attribute>
<DisplayName>External_Metric_Family_Name</DisplayName>
<Expressions>
  <ExpressionGroup destCert="{http://im.ca.com/core}Item">
    <Expression destAttr="Name">Names</Expression>
  </ExpressionGroup>
  <ExpressionGroup destCert="{http://im.ca.com/inventory}DeviceComponent">
    <Expression destAttr="IndexList">Indexes</Expression>
  </ExpressionGroup>
</Expressions>
<TableName>DB_Table_Metric_Family_Name</TableName>
<ComponentFacets>

```

```

<Facet>{http://im.ca.com/inventory}Component_Name</Facet>
</ComponentFacets>
<Protocol>IMDB</Protocol>
<Normalized>true</Normalized>
</FacetType>
</DataModel>

```

#### – Internal\_Metric\_Family\_Name

Specifies the metric family name. For each metric family, the name must be a unique name that identifies it internally within the system. Carefully select a name with a minimal possibility of naming conflicts with future similar metric families. For example, define a naming scheme that ensures that these metric family names are unique.

##### NOTE

This name is never exposed externally. To display a metric family name in the user interface, use the `DisplayName` element.

**Can be updated:** No

**Possible values:** Alphanumeric and underscore. Dot and dash are not permitted. The value must be unique across all metric families.

#### – Metric\_Family\_Description

Specifies the external description for the metric family. To make these comments useful, describe why and when you added or changed the metric family.

**Possible values:** Plain text

**Effect of updating:** None

#### – Element/Component\_Name

The documentation is also displayed in tool tips when you hover the cursor over the attribute name.

#### – Element/Component\_Description

Displays the attribute description in the user interface. The documentation is also displayed in tool tips when you hover the cursor over the attribute name.

**Possible values:** Plain text

**Effect of updating:** Hovering the cursor over the attribute name shows the updated documentation.

#### – Metric\_Name

Specifies the unique, internal name. For metrics, this name is also used for naming the database column.

##### NOTE

This name is never exposed externally. To display an attribute name in the user interface, use the `AttributeDisplayName` element. To change the `AttributeDisplayName`, see [Create or Extend Metric Families](#) and update the metric family properties.

**Possible values:** Alphanumeric and underscore.

**Effect of updating:** For metrics, the values for this attribute are stored in a new database column corresponding to the updated name. The user loses the historical data that is collected for this metric (with the older name). The custom reports that report on this metric fails.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

#### – MIB\_Metric\_Type

Indicates the data type of this attribute. The most frequently used data types are `Int`, `Long`, `Double`, `String`, or `ObjectID`. The database stores metric attributes as a float. Therefore, these attributes must use a numeric type (we recommend a `Double`). Other types are used for item attributes.

**Possible values:** `Boolean`, `Int`, `Long`, `Double` (floating-point), `BigInteger`, `String`, `DateTime`, `IPAddress`, `MACAddress`, `IPSubnet`, `OctetString` (hex representation), `ObjectID`, `ItemID`, `QName` (Qualified Name)

##### NOTE

The type names are case insensitive, for example, "boolean" is the same as `Boolean`.

**Effect of updating:** For metrics, none. All metrics are stored in the database as a float. For item attributes, the device must be deleted and rediscovered.

**When does the update take effect:** For metrics, next poll. For item attributes, on rediscover.

**Required actions for updates to take effect:** For metrics, none. For item attributes, delete the device and rediscover.

– **Metric\_Description**

Displays the attribute description in the user interface. The documentation is also displayed in tool tips when you hover the cursor over the attribute name.

**Possible values:** Plain text

**Effect of updating:** Hovering the cursor over the attribute name shows the updated documentation.

– **Baseline**

Indicates whether to calculate a mean value for this attribute. If it is set to true, a corresponding `BaselineList` definition must be defined.

**NOTE**

The `Baseline` attribute requires that the `StandardDeviation` attribute is set to true.

**Possible values:** `true`, `false`

**Effect of updating:** Baseline values are calculated when true.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

– **Variance** This tag has restricted use in custom and extended metric families. The value of this tag must be `false`.

– **Standard Deviation**

Indicates whether to calculate the standard deviation of this attribute during the rollup. Creates a 'std\_' column in the database table. If `RollupStrategy` is defined, this attribute must also be defined.

**Possible values:** `true`, `false`

**Effect of updating:** True provides a calculation of, and a reporting field for, `Standard Deviation`.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

– **Minimum**

Indicates whether to calculate the minimum of this attribute during the rollup. Creates a 'min\_' column in the database table. If `RollupStrategy` is defined, this attribute must also be defined.

**Possible values:** `true`, `false`

**Effect of updating:** True provides a calculation of, and a reporting field for, `Minimum`.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

– **Maximum**

Indicates whether to calculate the maximum of this attribute during the rollup. Creates a 'max\_' column in the database table. If `RollupStrategy` is defined, this attribute must also be defined.

**Possible values:** `true`, `false`

**Effect of updating:** True provides a calculation of, and a reporting field for, `Maximum`.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

– **RollupStrategy**

Specifies the operation that is performed every cycle during the rollup of the individually polled values. When `Polled` and `IsDbColumn` are set to true, this element is required.

**Possible values:** Sum (a summation for counters), Avg (an average for gauges)

**Effect of updating:** The specified strategy is used to perform rollup calculations.

**When does the update take effect:** Next poll

**Required actions for updates to take effect:** None

– **Percentile**

Indicates whether to calculate the 95th percentile of this attribute during the rollup. Creates a 'pct\_' column in the database table. If `RollupStrategy` is defined, this attribute must also be defined.

**Possible values:** 0, 95

**Effect of updating:** A value of 95 provides a calculation of, and a reporting field for, "95th Percentile." Zero specifies that no calculation is performed, and the reporting field is not available.

**When does the update take effect:** Next poll



**Required actions for updates to take effect:** None

– **External\_Metric\_Family\_Name**

Specifies the metric family name that displays in NetOps Portal.

**Possible values:** Plain text

**IMPORTANT**

Ensure that the `DisplayName` property is unique to the metric family.

**Effect of updating:** Change to the name in the administrator UI.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** Refresh the UI.

– **DB\_Table\_Metric\_Family\_Name**

Specifies the database table name that is used to store the metrics that the metric family collects.

**Possible values:** Uppercase alphanumeric and underscore. The value must begin with a letter. The value must be unique across all metric families.

**Example:** `PROCESS_STATS`

**Effect of updating:** Poll data is stored in a new set of database tables.

**IMPORTANT**

Updating the `TableName` property erases the old poll data. Old report views are broken.

**When does the update take effect:** Immediately. Before you can create views, there is a delay of up to 5 minutes while DX NetOps Performance Management loads the new MIB files.

**Required actions for updates to take effect:** Recreate the views.

– **Component\_Name**

Specifies a facet that is attached to the component item during component discovery.

**Can be updated:** Yes

**Possible values:** QName of the facet

**Effect of updating:** If the component facet is synchronized to NetOps Portal, the component is visible in NetOps Portal.

**When does the update take effect:** Rediscover

**Required actions for updates to take effect:** Delete the device and rediscover.

2. Set up a REST client with a connection to the data aggregator server.

3. Use the following format for the URL in the REST client:

`http://DA_host:8581/typecatalog/metricfamilies`

**NOTE**

To update or extend a metric family, use the following format for the URL in the REST client:

`http://DA_host:8581/typecatalog/metricfamilies/metric_family_name`

• ***metric\_family\_name***

Specifies the name of the existing metric family.

4. Select `POST` as the **HTTP Method**.

**NOTE**

To update an existing component, select `PUT`.

5. Select `application/xml` as the **Body Content-type**.

6. Add the customized XML within the "Body" text section.

7. Run the method.

## Create a Vendor Certification

A vendor certification provides a link between the vendor MIB and the metrics in the metric family. A vendor certification specifies which MIB variables are used for the relevant metrics. A vendor certification also specifies the formulas or expressions that are used to calculate the relevant values. Some variables directly match a metric (for example, rate, errors, discards, and volume). Others variables require calculations (for example, utilization or free disk space).

**NOTE**

If scalar OIDs need to be polled in the vendor certification, the definition of attributes changes slightly. If you are grouping scalar OIDs for better organization within an `AttributeGroup`, set the list variable to false (list="false"). Or, omit the variable in the `AttributeGroup` tag definition. list="true" in a vendor certification appends the index to the OID when polling. When you define the attribute source, append a '.' to the end of the OID. You can also have scalar OIDs defined outside of an `AttributeGroup`.

**Follow these steps:**

## 1. Customize the following XML:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- Auto-generated by the type catalog local manager.-->
<DataModel namespace="http://im.ca.com/certifications/snmp"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="SNMPCertificationFacet.xsd">
  <Author>Custom</Author><Version>Version</
Version>
  <FacetType name="Internal_Vendor_Cert_Name"

descriptorClass="com.ca.im.core.datamodel.certs.CertificationFacetDescriptorImpl">
    <Documentation>Vendor_Cert_Description</Documentation>
    <FacetOf namespace="http://im.ca.com/core" name="Item" />
    <AttributeGroup name="AttributeGroup" external="true" list="true">
    <Documentation />
    <Attribute name="INDEX" type="ObjectID">
    <Documentation />
        <Source>Index_OID</Source>
    <IsIndex>true</IsIndex>
    <IsKey>false</IsKey>
    <NeedsDelta>false</NeedsDelta>
    </Attribute>
        <Attribute name="Internal_Metric_Name" type="Data_Type_from_MIB">
    <Documentation />
        <Source>MIB_Variable_OID</Source>
    <IsIndex>false</IsIndex>
    <IsKey>true</IsKey>
    <NeedsDelta>true</NeedsDelta>
    </Attribute>
    </AttributeGroup>
    <Protocol>SNMP</Protocol>
    <DisplayName>Vendor_Cert_Name</DisplayName>
    <Expressions>
        <ExpressionGroup destCert="{http://im.ca.com/
normalizer}Metric_Family_Name" name="Metric_Family_Name">
    <Expression destAttr="Indexes">INDEX</Expression>
    <Expression destAttr="Names">"Element_Name" + INDEX</Expression>
    <Expression
destAttr="Metric_Name">Internal_Metric_Name_or_MVEL_Expression
```

**<MIB>MIB\_Name</MIB>**

</FacetType>

</DataModel>

– **Version**

Specifies the vendor certificate version.

Example: 1.0

– **Internal\_Vendor\_Cert\_Name**

Uniquely identifies a vendor certification.

**IMPORTANT**

Conform to <MibName><TableName>Mib .

**Can be updated:** No

**Possible values:** Alphanumeric and underscore. Dot and dash are not permitted.

– **Vendor\_Cert\_Description**

Describes what is certified with the vendor certification.

**TIP**

Include the details about the vendor, MIB name, and table name.

**Effect of updating:** None

– **Index\_OID**

Specifies the ObjectID of the attribute.

**TIP**

Set to the fully qualified MIB variable OID that the OBJECT-TYPE defines.

**Possible values:** Dot (.)-separated numbers (for example, 1.3.6.1.4.1...)

**Effect of updating:** Data is polled from the specified OID.

**When does the update take effect:** Next poll

**NOTE**

By default, the `Source` attribute specifies an OID to poll from the device. This default behavior is defined with `src='polled'`. You can set `src='mvel'` to process an MVEL expression using any polled attribute instead of an OID. For example, you could use the `src='mvel'` parameter to combine two 32-bit OIDs into a single 64-bit counter. You could also use the `src='mvel'` parameter to poll a counter that is not stored as a numeric type.

**Example:** The following example uses the `src='mvel'` parameter to combine two 32-bit OIDs into a single 64-bit counter:

```
<Attribute name="memberbitsout" type="Long">
    <NeedsDelta>true</NeedsDelta>
```

```
    <Source src='mvel'>snmpCounter64(memberbitsoutHi32,memberbitsoutLo32)</Source>
</Attribute>
```

**NOTE**

The `src='mvel'` parameter can only be used during polling. This parameter cannot be used during the discovery phases. For example, this parameter cannot be used as part of the Name or Description.

– **Internal\_Metric\_Name**

Specifies the attribute name.

**TIP**

Set to the MIB variable name, which the OBJECT-TYPE clause defines in the ASN.1 file.

**Possible values:** Alphanumeric and underscore. Dot (.) and dash (-) are not permitted.

**Effect of updating:** Update any expressions that reference this attribute.

– **Data\_Type\_from\_MIB**

Specifies the data type of the attribute.

**TIP**

Use the attribute type that best matches the variable type that the SYNTAX clause defines in the ASN.1 file.

**Possible values:** Boolean, Int, Long, Double, BigInteger, String, DateTime, IPAddress, MACAddress, IPSubnet, OctetString, ObjectID

**Effect of updating:** Polled SNMP data is converted to this type.

**When does the update take effect:** Next poll

- **MIB\_Variable\_OID**

Specifies the ObjectID of the attribute.

**TIP**

Set to the fully qualified MIB variable OID that the OBJECT-TYPE defines.

**Possible values:** Dot-separated numbers (for example, 1.3.6.1.4.1...)

**Effect of updating:** Data is polled from the specified OID.

**When does the update take effect:** Next poll

**NOTE**

By default, the `Source` attribute specifies an OID to poll from the device. This default behavior is defined with `src='polled'`. You can set `src='mvel'` to process an MVEL expression using any polled attribute instead of an OID. For example, you could use the `src='mvel'` parameter to combine two 32-bit OIDs into a single 64-bit counter. You could also use the `src='mvel'` parameter to poll a counter that is not stored as a numeric type.

**Example:** The following example uses the `src='mvel'` parameter to combine two 32-bit OIDs into a single 64-bit counter:

```
<Attribute name="memberbitsout" type="Long">
    <NeedsDelta>true</NeedsDelta>

    <Source src='mvel'>snmpCounter64(memberbitsoutHi32,memberbitsoutLo32)</Source>
</Attribute>
```

**NOTE**

You can use the `src='mvel'` parameter only during polling. You cannot use the parameter during the discovery phases. For example, you cannot use this parameter as part of the Name or Description.

- **Vendor\_Cert\_Name**

Specifies the name of the vendor certification as it displays in NetOps Portal.

**TIP**

Start with the vendor name and include the MIB and functionality information.

**Effect of updating:** A change to the name in the Administrator UI.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** Refresh the UI.

**IMPORTANT**

Ensure that the `DisplayName` property is unique to the vendor certification.

- **Metric\_Family\_Name**

Specifies the metric family that contains the `destAttrs` to populate.

**Possible values:** Any valid metric family

**Effect of updating:** Changes the permissible expression `destAttr`.

- **Element\_Name**

Specifies to use the vendor certification attributes to collect configuration data. This configuration data helps define the MVEL expression to provide the value for the Names metric family attribute.

**TIP**

Include as much information as necessary to identify an instance uniquely.

**Possible values:** String MVEL expression using available Attributes

**Effect of updating:** Component name change

**When does the update take effect:** After component rediscovery

**Required actions for updates to take effect:** Update the metric family or change the vendor certification priority.

- **Metric\_Name**  
Specifies the metric name as defined in the metric family.
- **Internal\_Metric\_Name\_or\_MVEL\_Expression**  
Specifies the internal metric name as defined in the attribute section. Or, specifies the relevant formula that is used to calculate the metric from the MIB variable represented by the internal metric name.
- **>MIB\_Name**  
Specifies the name of the MIB, which the DEFINITIONS clause defines in the ASN.1 file.

**IMPORTANT**

Conform to "<MibName>"

**Effect of updating:** Change to the **SNMP MIB Name** column in the Vendor Certification tab of the Administrator user interface.

**When does the update take effect:** Immediately

**Required actions for updates to take effect:** Refresh the UI.

2. Set up a REST client with a connection to the data aggregator server.
3. Use the following format for the URL in the REST client:

```
http://DA_host:8581/typecatalog/certifications/snmp
```

**NOTE**

To update or extend a vendor certification, use the following format for the URL in the REST client:

```
http://DA_host:8581/typecatalog/certifications/snmp/cert_name
```

- **cert\_name**  
Specifies the name of the existing certification.

4. Select **POST** for the **HTTP Method**.

**NOTE**

To update or extend a vendor certification, select **PUT**.

5. Select **application/xml** as the **Body Content-type**.
6. Add the customized XML within the "Body" text section.
7. Run the method.

## Add Metrics to Existing Metric Families

In this workflow, you add a BitsRateCustom metric that is based on collected MIB objects to the Interface metric family:

**TIP**

You can report data on the metrics that you add by creating a view for them.

For more information, see [Manage Dashboards](#).

1. [Update the Metric Family with the Metric](#)
2. [Update the Vendor Certification with the Metric](#)
3. (Optional) [Add the Metric to Monitoring Profiles Configured to Select Collected Metrics](#)
4. [Validate the Metric](#)

## Update the Metric Family with the Metric

**Follow these steps:**

1. Export existing extensions for the Interface metric family by making the following GET REST call:

```
http://da_hostname:8581/typecatalog/metricfamilies/extension/NormalizedPortInfo
```

The current extensions are returned. If you have not defined any extensions, the following XML is returned:

```
<DataModel xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" namespace="http://im.ca.com/normalizer"
xsi:noNamespaceSchemaLocation="IMDBCertificationFacet.xsd">
```

```

<Author>CA</Author>
<Version>2.03</Version>
<FacetType name="NormalizedPortInfo"
descriptorClass="com.ca.im.core.datamodel.certs.NormalizedFacetDescriptorImpl">
  <FacetOf namespace="http://im.ca.com/core" name="Item"/>
</FacetType>
</DataModel>

```

2. Define the BitsRateCustom attribute in the extension XML file:

```

<DataModel xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" namespace="http://
im.ca.com/normalizer" xsi:noNamespaceSchemaLocation="IMDBCertificationFacet.xsd">
  <Author>CA</Author>
  <Version>2.03</Version>
  <FacetType name="NormalizedPortInfo"
descriptorClass="com.ca.im.core.datamodel.certs.NormalizedFacetDescriptorImpl">
    <FacetOf namespace="http://im.ca.com/core" name="Item"/>
    <AttributeGroup name="PortInfoPollable" list="true" external="true">
      <Attribute name="BitsRateCustom" type="Double">
        <Documentation/>
        <IsDbColumn>true</IsDbColumn>
        <Baseline>false</Baseline>
        <Minimum>true</Minimum>
        <Maximum>true</Maximum>
        <RollupStrategy>Sum</RollupStrategy>
        <StandardDeviation>false</StandardDeviation>
        <Variance>false</Variance>
        <Percentile>95</Percentile>
        <Percentile2>0</Percentile2>
        <Percentile3>0</Percentile3>
        <ProjectionPercentile>0</ProjectionPercentile>
        <Polled>true</Polled>
        <Units>BitsPerSecond</Units>
      </Attribute>
    </AttributeGroup>
  </FacetType>
</DataModel>

```

3. Import the extensions for the Interface metric family by making the following PUT REST call:

```
http://da_hostname:8581/typecatalog/metricfamilies/extension/NormalizedPortInfo
```

The Interface metric family is updated with the BitsRateCustom metric.

### Update the Vendor Certification with the Metric

#### Follow these steps:

1. Export the extensions for the vendor certification by making the following GET REST call:

**TIP**

You can retrieve the name of the vendor certification using DX NetOps Performance Management. From the **Monitoring Configuration** menu for your data aggregator data source, click **Vendor Certifications**, click the downward arrow on any column, hover over **Columns**, and then click **Internal Name**.

```
http://da_hostname:8581/typecatalog/certifications/snmp/extension/cert_name
```

2. Add an expression for calculating the BitsRateCustom attribute:

```
<?xml version="1.0" encoding="UTF-8"?>
<!--Auto-generated by the type catalog local manager.-->
<DataModel namespace="http://im.ca.com/certifications/snmp"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="SNMPCertificationFacet.xsd">
  <Author>CA</Author>
  <Version>2.02</Version>
  <FacetType name="IfTableMib"
    descriptorClass="com.ca.im.core.datamodel.certs.CertificationFacetDescriptorImpl">
    <FacetOf namespace="http://im.ca.com/core" name="Item" />
    <Expressions>
      <ExpressionGroup destCert="{http://im.ca.com/
normalizer}NormalizedPortInfo" name="PortNRMDS">
        <Expression destAttr="BitsRateCustom">((ifInOctets+ifOutOctets)*8)/
_rspDuration</Expression>
      </ExpressionGroup>
    </Expressions>
  </FacetType>
</DataModel>
```

3. Import the extensions for the vendor certification by making the following PUT REST call:

```
http://da-hostname:8581/typecatalog/certifications/snmp/extension/cert_name
```

The vendor certification is updated. DX NetOps Performance Management resumes polling and uses this extended vendor certification.

**NOTE**

The BitsRateCustom metric becomes available 10 minutes after you import the extensions for the vendor certification.

**(Optional) Add the Metric to Monitoring Profiles Configured to Select Collected Metrics**

Add the metric to those monitoring profiles that are configured to select collected metrics (the monitoring profile has assigned metric families).

For more information:

- About how to add a metric to existing collected metrics, see [Configure Metric Filtering](#).
- About how to assign a metric family to a monitoring profile, see [Manage Monitoring Profiles](#).

**Validate the Metric**

Data that DX NetOps Performance Management polls for the BitsRateCustom metric appears in views after 3 poll cycles (15 minutes).

**Follow these steps:**

1. Hover over **Administration, Monitored Items Management**, and then click **Monitored Devices**.  
The **Monitored Devices** page opens.
2. Under **Monitoring Configuration**, click **Metric Families**.  
The **Metric Families** tab appears.
3. Verify that the `BitsRateCustom` metric appears.  
For more information about how to view data aggregator health monitoring information, see [View Health Monitoring Information](#).

**Add a Filter to a Vendor Certification**

This self-certification workflow is an example of a real use case that you may encounter. Use this workflow as a model to help you with self-certification.

In this example, you add a filter to the `IfTableMib` vendor certification to prevent items with `ifType 53 (propVirtual)` from being discovered.

After the filter is applied, you can expect the following results:

- In NetOps Portal, if you go to **Administration, Monitored Devices, Polled Metric Families**, the interface appears as **Not Available** or **Not Present** and **Not Polled**.
- In NetOps Portal, if you go to **Administration, Monitored Devices, Filter Report**, the filter report is unaffected. The report shows only the filters from the monitoring profile. The report excludes the filter from the vendor certification.
- The interface is still available through REST on the Data Aggregator.
- The interface is unavailable through OpenAPI.

After the next sync, you can expect the following results in NetOps Portal:

- If you go to **Inventory, Interfaces**, the interface is removed.
- The interface is removed from all dashboards and context pages.

**Follow these steps:**

1. Make a GET REST call to export existing extensions for the vendor certification:

```
http://da-hostname:8581/typcatalog/certifications/snmp/extension/Vendor_Certification_Name
```

For example:

```
http://da-hostname:8581/typcatalog/certifications/snmp/extension/IfTableMib
```

**TIP**

To retrieve the name of the vendor certification in the UI, go to the Vendor Certification page. Click the downward arrow on any column, hover over **Columns**, and click **Internal Name**.

2. Find the metric definition in the XML file.

```
<ExpressionGroup name="PortNRMDS" destCert="{http://im.ca.com/normalizer}NormalizedPortInfo">
```

3. Enter the metric filtering expression inside a Filter tag:

```
<Filter>(ifType!=24) &amp;&amp; (ifType!=1) &amp;&amp; (ifType!=53)</Filter>
```

4. Make a PUT REST call to import the new XML file:

```
http://da-hostname:8581/typcatalog/certifications/snmp/extension/IfTableMib
```

After you import the vendor certification, the Data Aggregator resumes polling using the new extended vendor certification.

5. Click **Update Metric Family** on the Metric Families page.



The vendor certification no longer monitors interfaces of ifType 53.

6. To verify that the filter is added, click **Admin, Monitored Devices**.

7. Click **Polled Metric Families**.

In the list of Components, the Status of the filtered interfaces is changed to **Not Available** or **Not Present**, and the **SNMP Poll Rate** is changed to **Not Polled**.

### Example

The following example shows how to define the metric filtering expression in the IfTableMib vendor certification:

```
<?xml version="1.0" encoding="UTF-8"?>
<!--Auto-generated by the type catalog local manager.-->
<DataModel namespace="http://im.ca.com/certifications/
snmp" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="SNMPCertificationFacet.xsd">
  <Author>CA</Author>
  <Version>2.02</Version>
  <FacetType name="IfTableMib"
descriptorClass="com.ca.im.core.datamodel.certs.CertificationFacetDescriptorImpl">
    <FacetOf namespace="http://im.ca.com/core" name="Item" />
    <Expressions>
      <ExpressionGroup destCert="{http://im.ca.com/normalizer}NormalizedPortInfo"
name="PortNRMDS">
        <Filter>(ifType!=24) & & (ifType!=1) & & (ifType!=53)</Filter>
      </ExpressionGroup>
    </Expressions>
  </FacetType>
</DataModel>
```

## Change the Calculation Method for an Existing Metric

You want to change how the Discards metric is calculated for the IfTableMIB vendor certification. Discards is an existing out-of-the-box metric.

1. Make a GET REST call to export the full IfTableMIB vendor certification.

```
http://da_hostname:8581/typecatalog/certifications/snmp/IfTableMIB
```

2. Make a GET REST call to export existing extensions for the IfTableMIB vendor certification.

3. Find the metric expression that you want to change.

```
<Expression destAttr="Discards"></Expression>
```

4. Enter the new calculation in the vendor certification extension file.

```
<Expression destAttr="Discards">(ifInDiscards+ifOutDiscards+ifInErrors+ifOutErrors)</
Expression>
```

## 5. Import the new extended vendor certification.

```
http://da-hostname:8581/typecatalog/certifications/snmp/extension/IfTableMIB
```

The DA resumes polling using the new extended IfTableMIB vendor certification. The new calculation for the Discards metric is applied.

6. After 3 poll cycles (15 minutes), verify the new calculation in a CAPC view.
  - a. Create a view using the new extended IfTableMIB vendor certification.
  - b. Verify in CA NetOps Portal that the new calculation for Discards is used.

### Example

The following example shows how to define the new calculation in the vendor certification extension XML file:

```
<?xml version="1.0" encoding="UTF-8"?>

<DataModel xsi:noNamespaceSchemaLocation="SNMPCertificationFacet.xsd" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance" namespace="http://im.ca.com/certifications/snmp">

  <Author>CA</Author>

  <Version>100.0</Version>

  <FacetType
    descriptorClass="com.ca.im.core.datamodel.certs.CertificationFacetDescriptorImpl"
    name="IfTableMib">

    <Documentation>Supports interfaces new metrics</Documentation>

    <FacetOf namespace="http://im.ca.com/core" name="Item"/>

    <Expressions>

      <ExpressionGroup name="PortNRMDS" destCert="{http://im.ca.com/
normalizer}NormalizedPortInfo">

        <Expression destAttr="Discards">(ifInDiscards+ifOutDiscards+ifInErrors
+ifOutErrors)</Expression>

      </ExpressionGroup>

    </Expressions>

  </FacetType>

</DataModel>
```

## Deploy a New Metric Family or Vendor Certification on the Production System

You have created a new metric family or vendor certification on your test environment. You want to deploy the new metric family or vendor certification to your production system. Do not experiment on the production system. Always develop metric families and vendor certifications on the test environment, and then deploy them to the production system.

1. Export the metric family or vendor certification from the test environment type catalog.

### Export a Metric Family

```
GET http://da_testhostname:8581/typecatalog/metricfamilies/mf_name
```

### Export a Vendor Certification

```
GET http://da_testhostname:8581/typecatalog/certifications/snmp/vc_name
```

2. Import the XML file that contains the metric family or vendor certification to the production system.

### Import a Metric Family

```
POST http://da_hostname:8581/typecatalog/metricfamilies
```

### Import a Vendor Certification

```
POST http://da_hostname:8581/typecatalog/certifications/snmp
```

## Manage SNMP Profiles

You can determine what credentials NetOps Portal uses when it accesses a device by discovering devices using *SNMP profiles*.

SNMP profiles provide SNMP parameters to data sources and ensure data security. SNMP profiles provide access credentials that allow discovery to perform secure queries of device Management Information Bases (MIBs). The default profile provides NetOps Portal access to those devices that use SNMPv1/v2C with public credentials. Create an SNMP profile for devices that do not use SNMPv1/v2C with public credentials. NetOps Portal encrypts community strings and credentials.

During discovery, NetOps Portal tries each profile for device access, and uses the SNMP profile with the highest rank (priority order) that can access a device.

### NOTE

- NetOps Portal can quickly discover SNMP devices, without the need for SNMP profiles or discovery profiles. For more information, see [Quickly Discover SNMP Devices](#).
- You can limit the SNMP profiles that NetOps Portal uses during discovery using a specific list of assigned SNMP profiles. For more information, see [Discovery Profiles](#).

You can manage SNMP profiles by creating them, editing existing SNMP profiles, or deleting them. Users with the **Administer SNMP Profiles** role right can manage (create, edit, and delete) SNMP profiles.

CA Application Delivery Analysis, DX NetOps Network Flow Analysis, and CA Unified Communications Monitor query the MIBs of managed items for performance information using SNMP profiles. When you register one of these data sources, the profiles that were created in that data source are added to NetOps Portal. NetOps Portal resolves naming conflicts automatically. During synchronization, NetOps Portal propagates the changes that you make to an SNMP profile in NetOps Portal to these data sources.

In this article:

- [View a List of SNMP Profiles](#)
- [Create or Edit an SNMP Profile](#)
- [Change the Priority Order of the SNMP Profiles](#)
- [SNMP Profile Changes](#)
- [Modify the Timeout and Retries Parameters](#)

### **View a List of SNMP Profiles**

NetOps Portal displays only the SNMP profiles that were created in the tenant space. Only users in that tenant space can access the SNMP profile. The Default Tenant Administrator can access the SNMP profiles that are associated with the Default Tenant. In multi-tenant environments, NetOps Portal displays only the SNMP profiles for that tenant.

Hover over **Administration**, **Configuration Settings**, and then click **SNMP Profiles**. The list includes high-level information about the contents of each profile.

### **Create or Edit an SNMP Profile**

NetOps Portal can query devices through SNMP using SNMP profiles that include communication credentials.

The following video shows the SNMP profile creation process:

#### **Follow these steps:**

1. Log in as an administrator with **Administer SNMP Profiles** role rights.
2. Hover over **Administration**, **Configuration Settings**, and then click **SNMP Profiles**. The **Manage SNMP Profiles** page appears.
3. Complete one of the following:
  - To create an SNMP profile, click **New**. The **Add SNMP Profile** dialog opens.
  - To edit an SNMP profile, select the SNMP profile that you want to edit, and then click **Edit**. The **Edit SNMP Profile** dialog opens.

#### **NOTE**

By default, the data in these dialogs is encrypted. To enable an administrator to troubleshoot issues with SNMP polling, allow that administrator to view secure data in clear text by assigning the **SNMP Clear Text** role right to the Administrator role.

For more information, see [Show Secure SNMP Data in Clear Text](#).

4. Complete the following fields, and change any default settings:

#### **NOTE**

The fields that apply only to SNMPv3 are noted.

- **Profile Name**  
Specifies the name for this SNMP profile. This name must be unique, cannot be duplicated across SNMP versions, and are not case-sensitive.  
**Required:** Yes
- **SNMP Version**  
Specifies the protocol that NetOps Portal uses to discover SNMP devices using this SNMP profile.  
**Options:**

- **SNMPv1/SNMPv2C:** NetOps Portal uses the SNMPv1/SNMPv2C protocol to discover SNMP devices using this SNMP profile.

**NOTE**

If a device that uses an SNMPv1/SNMPv2C profile responds to SNMPv1 and SNMPv2C, NetOps Portal polls devices using SNMPv2C.

- **SNMPv3:** NetOps Portal uses the SNMPv3 protocol to discover SNMP devices using this SNMP profile.

**IMPORTANT**

Electing for NetOps Portal to discover devices that use the SNMPv3 protocol adds an extra load of about 30 percent to the CPU of data collectors.

**Default:** SNMPv1/SNMPv2C

**Required:** Yes

- **Port**

Specifies the port that NetOps Portal uses to make SNMP connections to SNMP devices associated with this SNMP profile.

**Default:** 161

**Required:** Yes

**NOTE**

NetOps Portal can also send SNMP traps to trap receivers associated with this profile through notifications using this port. In this scenario, use port 162.

For more information, see [Configure Notifications](#).

- **Community Name**

Defines a secure string that lets the data source query the MIB of the associated device. The community name that you enter must provide read access to the device MIB.

**NOTE**

The community name for the default SNMP profile is 'public'.

**Required:** Yes

- **Verify Community Name**

**Required:** Yes

- **User Name**

Specifies the name of the user for this profile associated to the secret keys that NetOps Portal uses to authenticate and encrypt the SNMP packets.

**Required:** Yes

- **Context Name**

Specifies the collection of management information SNMP entities can access. The context name is an octet octet string that NetOps Portal requires to provide end-to-end identification and to retrieve data from SNMP agents.

**NOTE**

Only NetOps Portal communicates with the device using the context name on SNMPv3 profiles.

- **Authentication Protocol (SNMPv3)**

Specifies the authentication protocol for NetOps Portal to use when contacting devices associated with this SNMP profile. DX NetOps Performance Management supports the following algorithms for authenticating SNMPv3 packets:

- None
- MD5 (Message Digest 5)
- SHA (Secure Hash Algorithm)
- SHA 256
- SHA 512

**Default:** None

- **Authentication Password**

Specifies the password that NetOps Portal uses to authenticate using SNMPv3 and the selected authentication protocol.

**NOTE**

This option is enabled if authentication is enabled for the SNMP profile (you have chosen an option for **Authentication Protocol** other than **None**).

**Required:** Yes

– **Verify Authentication Password**

**Required:** Yes

– **Privacy Protocol** (SNMPv3)

Specifies the encryption protocol for NetOps Portal to use when contacting devices associated with this profile for data flows sent to those devices or servers.

**NOTE**

This option is enabled if authentication is enabled for the SNMP profile (you have chosen an option for **Authentication Protocol** other than **None**).

**Options:**

- None
- DES
- AES 128
- Triple DES
- AES 256 with 3DES key

**Default:** None

– **Privacy Password**

Defines the password that NetOps Portal uses when exchanging encryption keys.

**NOTE**

This option is enabled if you have chosen an encryption protocol for contacting devices for data flows sent to those devices or servers (you have chosen an option for **Privacy Protocol** other than **None**).

**Required:** Yes

– **Verify Privacy Password**

**Required:** Yes

– **Use by default for new devices**

Specifies whether CA Application Delivery Analysis, DX NetOps Network Flow Analysis, and CA Unified Communications Monitor contact new items (discovery) using this SNMP profile.

**Options:**

- **Enabled:** The data sources contact new items (discovery) using this SNMP profile.
- **Disabled:** The data sources do not use this SNMP profile for discovery by default when contacting new items (discovery) .

**Default:** Enabled

**Required:** Yes

– **Use for SNMP SET**

Specifies whether the SNMP profile provides write credentials on the devices that NetOps Portal has discovered. For more information, see [Configure Round Trip Time \(RTT\) Tests](#).

**Options:**

- **Yes:** The SNMP profile provides write credentials on the devices that NetOps Portal has discovered.
- **No:** The SNMP profile does not provide write credentials on the devices that NetOps Portal has discovered.

**Default:** No

**Required:** Yes

5. Click **Save**.

The SNMP profile is added to NetOps Portal, and NetOps Portal uses it for discovery and polling. NetOps Portal sends the profile information to registered data sources (global synchronization).

### **Change the Priority Order of the SNMP Profiles**

The **Order** parameter determines which SNMP profile discovery uses and which NetOps Portal uses for polling when a device responds to multiple SNMP profiles. NetOps Portal uses this order for unreachable devices. Administrator users can change this (priority) order.

#### **NOTE**

Changes to the order do not affect existing polled devices. NetOps Portal continues to use the associated SNMP profile to poll those devices.

The new order takes effect in the following situations:

- Discovery discovers a new device.
- An existing device becomes unreachable through SNMP for at least two poll cycles.
- You delete the SNMP profile for a device.

On the **Manage SNMP Profiles** page, click and then drag or click **Move Up** and **Move Down**.

### **SNMP Profile Changes**

When the SNMP credentials on a polled device change, add the new SNMP profile information to NetOps Portal. When the device becomes unreachable with the deprecated SNMP profile for two poll cycles, NetOps Portal attempts to contact the device with other profiles. When DX NetOps Performance Management successfully contacts the device with an SNMP profile, that profile is assigned to the device for future polling.

#### **TIP**

You can view the SNMP profile that NetOps Portal uses to poll the device from the **Monitored Devices** page for the device.

For more information, see [Manage Monitored Devices](#).

### **Modify the Timeout and Retries Parameters**

If SNMP requests go across a WAN or across a slow network connection, they might time out. The timeouts can cause missing polled data or device discovery failure. You can modify the **Timeout** and **Retries** parameters for a profile.

For example, by default, an SNMP request is given the following amount of time:

3 seconds x (first attempt + 2 retries) = 3 seconds x 3 tries = 9 seconds to respond before it times out

#### **IMPORTANT**

To prevent unintended consequences such as resource starvation (CPU/Memory) and unnecessary traffic on the data collectors, modify these parameters with careful consideration and only if you have a basic understanding of SNMP communication.

### **Follow these steps:**

1. Set up a REST client with a connection to the data aggregator server.
2. Specify the following URL:

```
http://da_hostname:8581/rest/profiles/profile_item_id
```

3. PUT the XML for modifying the parameters:

```
<?xml version="1.0" encoding="UTF-8"?>
  <CommunicationProfile version="1.0.0">
    <CommunicationFailurePolicy version="1.0.0">
```

```

    <Timeout>3000</Timeout>
    <Retries>2</Retries>
  </CommunicationFailurePolicy>
</CommunicationProfile>

```

- **Timeout**  
The amount of time a device is given to respond to an SNMP request per tryDefault: 3000 milliseconds
- **Retries**  
The number of times an SNMP query is reissued before it times out  
**Default:** 2 retries

The `Timeout` and `Retries` parameters are modified.

## Show Secure SNMP Data in Clear Text

Allow users to troubleshoot SNMP profiles and view secure SNMP data that is typically masked in clear text.

The data in the **Add SNMP Profile** and **Edit SNMP Profile** dialogs is encrypted. You can enable Administrators to troubleshoot issues by allowing the Administrator to view secure data in clear text instead of encrypted. By default, the **SNMP Clear Text** role right is not assigned to any roles.

Only the predefined Administrator role can have this role right. By default, the predefined Administrator role is assigned only to the global administrator. To allow another user to view secure SNMP data, assign the Administrator role to another user account.

For more information about this role right, see [Role Rights](#).

### Follow these steps:

1. Log in as an Administrator user.
2. Hover over **Administration, User Settings**, and then click **Roles**.  
The **Manage Roles** page appears.
3. Select the **Administrator** role, and then click **Edit**.
4. Select **NetOps Portal**, and then click **Edit**.  
The **Edit Role** page appears.
5. Select **DX NetOps**, and then click **Edit**.  
The **Edit Role Rights** dialog opens.
6. In the **Available Rights** column, select the **SNMP Clear Text** role right, and then click the right arrow.  
The role right is added to the **Selected Rights** list.
7. Click **Save**.

Users with the Administrator role can now view secure SNMP data in clear text.

## IP Domains

IP domains resolve IP address conflicts by separating monitored devices from different networks. For example, you monitor two networks and each network includes a separate device with the same IP address.

As a service provider, use IP domains to monitor multiple discrete networks that belong to different customers. Each customer, or tenant, contains one or more IP domains. The data collectors associate managed items and data with the assigned IP domain. Customer user accounts see only the relevant IP domains. Service provider administrators see the data from all IP domains.

Administrators and Designers can create custom dashboards to monitor activity on a specific domain or group of domains.

DX NetOps Network Flow Analysis, CA Application Delivery Analysis, and CA Unified Communications Monitor data sources can use IP domains. To apply domains to the data sources, register the data source with NetOps Portal. When



you create IP domains in NetOps Portal, the domain identifiers become visible in the registered data sources after synchronization.

## Manage IP Domains

To monitor multiple tenants or environments with overlapping IP addresses, create an IP domain. Data sources report a domain identifier. This identifier associates monitored items with specific IP domains and tenants. Items that are not assigned to a custom IP domain in the data source are associated with the Default Domain.

### NOTE

The Default Domain is created automatically. This domain includes the items that are not assigned to a custom domain in the data source.

In this article:

- [View the List of IP Domains](#)
- [Add or Edit an IP Domain](#)
- [Assign Items to an IP Domain](#)
- [Delete an IP Domain](#)

### View the List of IP Domains

You can view a list of IP domains from the **Manage IP Domains** page. To view this page, hover over **Administration, Configuration Settings**, and then click **IP Domains**.

### Add or Edit an IP Domain

### NOTE

- Updates to device alias names and interface descriptions can take up to 5 minutes to appear.
- Editing an IP domain does not change the historical data for monitored items.

The procedure requires the Administer IP Domains role right.

### Follow these steps:

1. From the **Manage IP Domains** page, click **New** or select the IP domain that you want to edit, and then click **Edit**. The **IP Domains Administration** dialog opens.
2. Specify a domain name and description.
3. (Optional) Configure a device alias name or interface description override. Click **Browse**, select the file, and then click **Open**.

### NOTE

The CSV or TXT file must include the following format:

```
IP_Address,name,description,alias/alternate_description
```

Use the primary IP address of the associated devices. To find the primary IP address, look at the Address column of the Inventory Devices list.

You can use the same alternate interface descriptions for more than one interface.

### Example:

```
172.24.36.107,ethernet_7,vmxnet3 Ethernet Adapter,Connection to Dallas
```

To change or remove an alternate description or device name alias, import a new file. To remove the alias or alternate description, leave the final column blank. When you remove the alias or alternate description, the original name or description reappears. If you change or remove alternate descriptions or device name aliases from a file using a spreadsheet program, include a column heading for that column.

**TIP**

The preferred method to update device alias names and alternate interface descriptions is through REST. For more information, see [Set Device Alias Names Using the Devices Web Service](#).

To update multiple device alias names, use the `update_alias_name.sh` script that is included with DX NetOps Performance Management.

For more information about this script, see [Set Alias Names Using a Script](#).

4. (Optional) To assign a primary and secondary DNS address for the domain in DX NetOps Network Flow Analysis (NFA), select **DNS Settings**, and then specify the required values.

**NOTE**

Only NFA uses the DNS settings. DX NetOps Performance Management does not directly use the DNS settings.

5. Click **Save**.

The IP domain displays in the list. The changes apply to the data sources in the next synchronization.

**Assign Items to an IP Domain**

You can assign the following managed item types with an IP domain:

- Devices
- Interfaces and interface addresses
- Networks
- VoIP Locations

Data collectors associate managed items with a single IP domain. To enable multi-tenant deployments, assign an IP domain to each data collector.

**NOTE**

You cannot move items from one IP domain to another.

To assign items from other data sources to an IP domain, see the documentation for that data source or see the following pages:

- [Assign Network Flow Analysis Items to an IP Domain](#)
- [Assign Application Delivery Analysis Items to an IP Domain](#)
- [Assign CA Unified Communications Monitor Items to an IP Domain](#)

**TIP**

DX NetOps Spectrum (Spectrum) determines which models are synchronized with NetOps Portal using IP domains.

For more information, see the [DX NetOps Spectrum documentation](#).

Use the following process to assign managed items to an IP domain:

1. [Assign the data collector with the IP domain](#).
2. [Create a discovery profile that uses that IP domain](#).
3. [Run discovery](#).

**Delete an IP Domain**

Deleting an IP domain does the following:

- Deletes the monitored items in that IP domain, including historical data. This also deletes the items from registered data sources that use IP domains, such as NFA, CA Application Delivery Analysis, and DX NetOps Mediation

Manager. Spectrum devices in the IP domain are removed from NetOps Portal, but Spectrum continues to monitor the devices.

- Deactivates the data collectors and discovery profiles that are associated to the IP domain.

**TIP**

You can continue polling or discovery by assigning these data collectors and discovery profiles to another IP domain.

**IMPORTANT**

To retain associated data, do not delete an IP domain. You cannot recover deleted data.

**Follow these steps:**

- From the **Manage IP Domains** page, select the IP domain that you want to delete, and then click **Delete**.
- At the prompt, confirm the deletion by clicking **Yes**.

The IP domain is deleted.

## Assign Network Flow Analysis Items to an IP Domain

You can edit the settings for an item (router, interfaces, and CVIs) to associate them with a tenant and IP domain.

Interfaces and CVIs inherit their initial tenant-domain setting from the parent router and Harvester when the parent Harvester is added and the router and interfaces first become active. If the Harvester is not associated with a custom domain, the routers and interfaces are assigned to the Default Domain as they become active.

You can edit the settings for an item at any time. The items do not have to match the parent router or Harvester.

Changing the settings can affect which operators have access to interface data. The settings do not affect which Simple Network Management Protocol (SNMP) profiles you are using for polling. The router tenant determines the set of SNMP profiles for polling.

**NOTE**

You can also change the tenant-domain setting for Harvesters and routers.

**Follow these steps:**

- From the DX NetOps Network Flow Analysis console menu, select **Administration**.
- Select **Interfaces, Physical & Virtual**.  
The **Active Interfaces** page appears.
- Select the interfaces that you want to associate with a tenant and domain.

**TIP**

- To search for parent routers, interfaces, or CVIs, enter all or part of a router IP address, a router or interface name, or an interface description in the **Search** field, and then click **Search**. Expand the router details.
- To navigate to an interface or CVI manually, go to the page that contains the parent router, and then click the arrow next to the router name. The router details expand to show the interfaces and CVIs.

- Click **Edit**.

The **Edit Router** or **Edit Interface** dialog opens for editing the selected item or items. The Domain selection list is included in the dialog only if multiple domains exist.

- Select a tenant/domain option from the Domain list, and then click **Save**.  
The dialog closes. The changes are shown on the **Active Interfaces** page.

## Assign Application Delivery Analysis Items to an IP Domain

CA Application Delivery Analysis (ADA) can observe duplicate Internet Protocol (IP) traffic, which occurs in a managed service provider environment. The provider can host an application on a single server for multiple customers whose environments contain overlapping client IP addresses.

You enable ADA to identify separate IP traffic during data collection setup. As you verify and modify data collection parameters, assign the same IP domain to the appropriate:

- Monitor feeds
- Client networks
- Servers or server subnets

With the same IP domain assignments for these feeds, ADA reports on the application traffic between a client and a server by domain.

Applications are domain-independent. Therefore, you are not required to define the same application twice, such as Exchange Company A and Exchange Company B, to enable ADA to report on application performance across domains. However, to set different thresholds for application performance, performance operational level agreements (OLAs), and availability OLAs, create an application for each IP domain.

If you do not need to separate duplicate IP traffic, you can use the DNS settings in the Default Domain to query Domain Name System (DNS) and resolve the hostname of an ADA server. Otherwise, ADA uses the monitor feed that is assigned to the server to resolve the hostname.

You can do the following in ADA:

- **View a list of domains**  
You can view a list of domain definitions and current domain associations in the Administration section of the ADA management console.
- **Assign a domain to a monitor feed**  
You can instruct each Standard Monitor to associate the items that it monitors with a custom domain as part of ADA collection device setup.
- **Assign a domain to a client network**  
After you add a client network, you cannot change its IP domain association. If you need to change the assigned IP domain, first delete the network, and then add it to the correct domain.
- **Assign a domain to a server or server subnet**  
After you add a server or server subnet, you cannot change its IP domain. If you add a server or server subnet to the wrong IP domain, first delete it, and then add it to the correct domain.

For more information, see [the CA Application Delivery Analysis documentation](#).

## Assign CA Unified Communications Monitor Items to an IP Domain

In the CA Unified Communications Monitor management console, you can instruct collectors to associate the items that they discover with custom domains in NetOps Portal. The act of creating a single custom domain in NetOps Portal enables domain associations for locations, voice gateways, and call servers in any registered data sources.

Items appear with domain designations when they are discovered from call traffic. Previously discovered items do not receive retroactive associations.

Locations are automatically associated with IP domains by the subnets that they contain. To preserve the flow of data collection and the appropriate association of data with IP domains, take care when moving Locations to new IP domains. Follow the procedure in the CA Unified Communications Monitor online Help to change IP domain assignments.

Instruct collectors to associate items with custom IP domains.

**Follow these steps:**

1. Click Administration, Data Collection, Collectors.
2. Edit each collector to select its domain for the IP Domain parameter.
3. Reload the collectors to send them the domain information.  
Domains are populated with managed items after the next product synchronization.

## Discovery

Discovery is the process that NetOps Portal uses to identify the devices on your network. NetOps Portal identifies devices using the IP domain, IP addresses, IP ranges, and hostnames defined in discovery profiles.

NetOps Portal attempts to discover devices through Internet Control Message Protocol (ICMP) and Simple Network Management Protocol (SNMP) protocols. If a device does not respond to SNMP but responds to ICMP, NetOps Portal creates a pingable device.

In this article:

- [Device Discovery](#)
- [Metric Family Discovery](#)

### Device Discovery

Device discovery identifies the following information about devices:

- Which protocols the device responds to, SNMP or ICMP
- The classification of the device, such as router or switch
- The device vendor, such as Cisco or Juniper
- The device type, such as 7700 or 8200

Running device discovery produces device discovery events, which include events for successful and failed discoveries, as well as events for unreachable and pingable devices. Discovered devices and monitored components take up to five minutes to begin synchronizing with NetOps Portal. The data aggregator automatically adds the discovered devices to device collections depending on the rules that control each device collection membership. During the first synchronization between the data aggregator and NetOps Portal *after* NetOps Portal discovers devices, it adds the devices to collections.

You can discover devices using the following methods:

- [Run device discovery using NetOps Portal.](#)
- [Discover devices in the data aggregator from other data sources, such as DX NetOps Spectrum \(Spectrum\),](#) which creates a discovery profile.
- [By configuring a DX NetOps Virtual Network Assurance \(VNA\) plug-in.](#)
- [Quickly discover SNMP devices.](#)

The following workflow offers a best practice to use as a quick reference when performing a discovery of your inventory. Perform this process as either a user with the Administrator role or as the tenant administrator:

1. To enable the data collector to perform queries of device MIB tables that use SNMP, [create SNMP profiles](#).

**NOTE**

For the discovery of SNMP devices, NetOps Portal can quickly discover these devices without the need for SNMP profiles.

2. [Create discovery profiles](#).

**NOTE**

For the discovery of virtual network devices or SNMP devices, NetOps Portal can quickly discover these devices without the need for discovery profiles.

3. [Run device discovery, and then view the results.](#)

The following videos examines how to create, run, and view the results of a discovery profile in NetOps Portal, using an IP address or IP address ranges, to define which devices to discover on the network for performance monitoring:

### **Metric Family Discovery**

Metric family discovery, also known as component discovery, is a separate process that determines whether a specific metric family is supported for a device. A metric family defines the set of metrics to collect and report on for a given technology. NetOps Portal normalizes the values of the metrics so that reporting is uniform regardless of the data source.

Polling begins automatically after metric family discovery completes. NetOps Portal collects and retains operational metrics at regular polling intervals for reporting. Examples of operational metrics include error rate, utilization, and bytes in. Configuration data represents or identifies a component or the component configuration, such as port type or component index.

For more information, see [Manage Monitoring Profiles](#).

## **Manage Discovery Profiles**

*Discovery profiles* specify how inventory discovery operates in your environment. They determine the IP domain, IP addresses, IP address ranges, and host names for discovery.

Discovery profiles specify a list of SNMP profiles that NetOps Portal uses during discovery. You can manage discovery profiles by creating them, editing existing discovery profiles, or deleting them. Administrators can manage discovery profiles using NetOps Portal or by way of the `profiles` Data Aggregator REST web service or use this API in your scripts for managing discovery profiles. This topic describes how to manage them using NetOps Portal.

For more information about how to automate discovery profile management by way of the web service, see [Manage Discovery Using REST](#).

You can optimize discovery and reduce Simple Network Management Protocol (SNMP) traffic by setting up granular discovery profiles. Create separate discovery profiles for the following:

- For each group of devices that share an SNMP profile.  
For more information about SNMP profiles, see [Manage SNMP Profiles](#).
- For groups of devices that require different rediscovery schedules.

### **NOTE**

NetOps Portal can quickly discover SNMP devices, without the need for SNMP profiles or discovery profiles.  
For more information, see [Quickly Discover SNMP Devices](#).

The following video walks through a single device discovery:

For more information about how to run discovery (run a discovery profile), including the methods that you can use to run discovery, how to show discovery history for a specific discovery profile, and how to show the devices that were discovered during a specific discovery instance, see [Run Discovery](#).

### **View a List of Discovery Profiles**

Only users in the tenant space where the discovery profile was created can access the discovery profile. Users that are assigned to the Default Tenant can access discovery profiles in the Default Tenant space.

### **Follow these steps:**

1. Log in as a tenant administrator.
2. Hover over **Administration, Monitored Items Management**, and then click **Discovery Profiles**.  
The **Discovery Profiles** page appears.

## Create or Edit a Discovery Profile

### Follow these steps:

1. From the **Discovery Profiles** page, complete one of the following:

- To create a discovery profile, click **New**.
- To edit a discovery profile, select the discovery profile that you want to edit, and then click **Edit**.

The **Discovery Profile** dialog opens.

2. Complete the following required fields:

- **Name**

The name for the discovery profile.

The following characters are not permitted: single quotes, double quotes, backward slashes, forward slashes, ampersands.

- **IP Domain**

The IP domain for the discovery profile.

#### IMPORTANT

Newly-discovered devices are created in this IP domain. When you have multiple data collectors in an IP domain, each issues a discovery request to each device as specified by the discovery profile. When more than one data collector can contact a device, NetOps Portal randomly selects one of the data collectors to monitor the device. Review the capacity of the data collectors and rebalance as required.

For more information, see [Rebalance the Load on the Data Collectors](#).

3. On the **IPs/Hosts** tab, complete one or more of the following actions:

- Specify IP address ranges, individual IP addresses, and host names that you want to discover for IPv4 addresses, and then click **Add**. You can add comma-delimited values. IPv4 address ranges can contain the following characters:

- Wildcards (\*). A wildcard represents a full range for an IP octet: 0-255.
- Hyphens (-). A hyphen can exist between the lower IP address and upper IP address. A hyphen can also be in the IP octets in the lower IP address.

#### NOTE

If an IP range includes multiple IP addresses, and one of the IP addresses maps to the hostname, discovery always uses the hostname IP as the primary IP address.

For more information about how to specify an IP address range, including viewing examples, see the "[Discovery Profile IP Ranges](#)" section.

- (Optional) Navigate to and import a CSV file of IP addresses. The CSV file can contain a comma-separated list of IPv4 addresses, IPv6 addresses, IPv4 address ranges, and hostnames. Click **Import**, browse to select the file, and then click **Open**.

#### NOTE

To apply Chinese characters to the alias name, save the CSV file in UTF-8 format.

4. (Optional) To have NetOps Portal regularly update information for discovered devices and discover new devices, select the **Schedule** tab, and then configure a schedule for the discovery profile by completing the following fields:

- **Scheduling interval**

Defines the interval to discover devices.

**Options:** None, Daily, or Weekly

**Default:** None

- **On Days**

Defines the days to discover devices.

**Options:** Sun, Mon, Tue, Wed, Thu, Fri, or Sat

**Default:** Mon - Fri

- **Run profile starting at**

Defines the time to discover devices.

**Default:** 12:00 AM

5. (Optional) To define the type of protocol to use when discovering devices, select the **Protocols** tab, and then complete the following steps:

The following image shows an example of this tab:

## Discovery Profile ✕

Name: \*

IP Domain: \*

IPs/Hosts
Schedule
Protocols
Advanced

---

Discovery Protocols: \* SNMP and ICMP ▼

SNMP Discovery

☐ Use specific list of assigned SNMP profiles

Available SNMP profiles

Assigned SNMP profiles

☒ Create Pingables
 ☐ Save As Default

Save
Cancel

- a. Choose the type of protocol to use when discovering devices by completing the following field:
- **Discovery Protocols**  
Defines the type of protocol to use when discovering devices.  
**Options:**
    - **SNMP:** Discover devices using SNMP, management information bases (MIBs), and object identifiers (OIDs).
    - **ICMP:** Discover devices using Internet Control Message Protocol (ICMP) to perform ping sweep and device reachability discovery.
    - **SNMP and ICMP:** Discover devices using SNMP and ICMP.
  - Default:** SNMP and ICMP
- b. Complete the following:



- If you have chosen to discover devices using SNMP or SNMP and ICMP, in the **SNMP Discovery** section, define whether to discover devices using a specific list of assigned SNMP profiles by selecting the **Use specific list of assigned SNMP profiles** checkbox, and then moving one or more SNMP profiles from the **Available SNMP profiles** list to the **Assigned SNMP profiles** list.

**NOTE**

Using a subset of SNMP profiles reduces SNMP traffic.

**Default:** Cleared

- If you have chosen to discover devices using ICMP or SNMP and ICMP, in the **ICMP Discovery** section, edit the following fields:
  - **Create Pingables**  
Defines whether NetOps Portal discovers pingable devices using ICMP.
  - **Save As Default**  
Have NetOps Portal use the ICMP discovery settings defined in the discovery profile as the default.

**Default:** Cleared

6. (Optional) Select the **Advanced** tab, and then configure the following advanced options:

The following image shows an example of this tab:

## Discovery Profile ✕

Name: \*

IP Domain: \*

IPs/Hosts   Schedule   Protocols   **Advanced**

### Naming Order and Reconciliation Options

- ↕ 1 Host Name
- ↕ 2 System Name
- ↕ 3 IP Address

↑   ↓

☐ Exclude Host Name   ☐ Save As Default

Save

Cancel

- In the **Naming Order and Reconciliation Options** section, define the priority NetOps Portal uses to name the discovered devices. NetOps Portal names the device items that the discovery profile creates using the highest available naming convention. To change the priority, select a name from the list and then use the **Move Up** or **Move Down** arrows. If a higher priority attribute is unavailable for the device, NetOps Portal uses the next highest priority attribute. For virtual machines, NetOps Portal ignores the naming order and uses the names from vCenter.

#### NOTE

If you use host name to name devices, NetOps Portal updates the device name automatically when the hostname changes. If you use another attribute, such as system name, the change to the device name occurs when the discovery profile runs again.

- Exclude Host Name**  
In some configurations, the network might not have unique DNS host names. To reconcile devices by the IP address and system name only, select this checkbox.  
**Default:** Cleared
- Save As Default**

Have NetOps Portal use the priority naming order defined in the discovery profile to name the device items that the discovery profile creates.

**Default:** Cleared

7. Click **Save**.

The discovery profile is created and is displayed in the **Discovery Profiles** list on the **Discovery Profiles** page. If the profile has a schedule, discovery runs at the scheduled time. The discovery profile is in the "Scheduled" state, and the next scheduled run time appears in the **Next Run Time** column of this page.

### **Discovery Profile IP Ranges**

In a discovery profile, you can specify the IP address ranges that you want to discover for only IPv4.

#### **Examples: Valid IP Ranges**

- The following examples attempt to discover devices at every IP address from 10.25.1.0 to 10.25.1.190:

10.25.1.0-10.25.1.190

OR

10.25.1.0-190

- The following examples attempt to discover devices at every IP address from 10.25.0.0 to 10.25.255.255:

10.25.\*.\*

OR

10.25.0.0 - 10.25.255.255

- The following examples attempt to discover devices at every IP address from 10.25.0.3 to 10.25.0.40 and from 10.25.1.3 to 10.25.1.40:

10.25.0-1.3-40

OR

10.25.0.3 - 10.25.0.40, 10.25.1.3 - 10.25.1.40

- The following examples attempt to discover devices at every IP address from 10.25.0.0 to 10.25.0.5, from 10.25.1.0 to 10.25.1.5, and so on, up to 10.25.255.0 to 10.25.255.5:

10.25.\*.0-5

OR

10.25.0.0 - 10.25.0.5, 10.25.1.0 - 10.25.1.5 ... 10.25.255.0 - 10.25.255.5

#### **Examples: Invalid IP Ranges**

- The following example is invalid because the upper IP address is incomplete:

10.25.1.0 - 10.23

- The following example is invalid because when a hyphen (-) is used in an octet in the lower IP address, the upper IP address cannot be present:

10.25.1.0-190 - 10.25.1.255

- The following example is invalid because when a wildcard (\*) is used in an octet in the lower IP address, the upper IP address cannot be present:

10.25.\*.0 - 10.25.255.255

- The following example is invalid because it is unclear whether the wildcard octet (1\*) implies 10.25.10-19.0 or 10.25.10-199.0:

10.25.1\*.0

## Run Device Discovery

Device discovery is the process of the data aggregator discovering devices as specified by a discovery profile.

Device discovery produces device discovery events, which include events for successful and failed device discoveries, as well as events for unreachable and pingable devices.

For more information about device discovery events, see [Event Types](#).

In this article:

- [Run Device Discovery](#)
- [View Discovery Profile History Results](#)
- [Export Discovery Profile History Results](#)
- [Run Device Discovery and Generate a Discovery Log for Debugging Purposes](#)

### Run Device Discovery

You can run device discovery using the following methods:

- [Configure a discovery profile to run on a daily or weekly schedule.](#)
- [Automate running discovery by way of the `discoveryprofiles` Data Aggregator REST web service or use this API in your scripts for running discovery.](#)
- [Use NetOps Portal.](#)

This article describes how to run device discovery using NetOps Portal.

### Run Device Discovery from the Discovery Profile using NetOps Portal

When you run device discovery using NetOps Portal, you run it from (as specified by) the discovery profile.

#### **Prerequisites:**

- (To run a device discovery as an Administrator) You have configured the data collector in the desired IP domain.
- The discovery profile is created and is in "Ready" or "Scheduled" state.

For more information about how to create discovery profiles, see [Manage Discovery Profiles](#).

**Follow these steps:**

1. As an Administrator or a Tenant administrator, hover over **Administration**, **Monitored Items Management**, and then **Discovery Profiles**.  
The **Discovery Profiles** page appears.
2. From the **Discovery Profiles** list, select the discovery profile that you want to run device discovery, and then click **Run**.  
The **Run Discovery Profile** dialog opens.
3. Confirm your selection by clicking **Yes**.

The data aggregator adds discovered devices to device collections, which initiates device monitoring and polling. The state of the discovery profile changes to "Running".

**NOTE**

- Discovered devices and monitored components can take up to five minutes to synchronize with NetOps Portal. A discovery is considered to be hanging when NetOps Portal does not discover new devices within 10 minutes and the state of the discovery profile has not changed within 10 minutes. If discovery hangs, NetOps Portal aborts device discovery, and then generates an audit event on the data aggregator.
- You can get a list of the unreachable devices [by way of the `discoveryinstances` Data Aggregator REST web service](#).
- During discovery, NetOps Portal updates only those devices that are in the "Active" life cycle state.
- If discovery fails for any of your devices (devices are not found), the state of the discovery profile indicates "FAILURE". If discovery finds at least one device, the state indicates "PARTIAL\_FAILURE".
- If device discovery fails to discover *all* devices, or it does not discover the device as expected, you can [generate a detailed discovery log of what the discovery process finds when it communicates with a device for debugging purposes](#).

**View Discovery Profile History Results**

The discovery profile history results show the devices that were discovered during a specific discovery instance. You can view specific details about these discovered devices, including the IP address, model, type, vendor name, location, and protocols.

**NOTE**

You can also automate reviewing (and exporting) discovery profile history by way of the `discoveryprofiles` and `discoveryinstances` Data Aggregator REST web services or use these APIs in your scripts for managing discovery profiles.

For more information about this API, see [Manage Discovery Using REST](#).

**Follow these steps:**

1. On the **Discovery Profiles** page, select the discovery profile for which you want to view discovery results, and then click **History**.  
The **Discovery History** dialog opens.
2. Select the discovery history instance for which you want to view results from the **Select Discovery Instance** list.
3. (Optional) In the **Discovered Devices** section, filter the **Discovered Devices** table from the following options, and then click **Apply**:
  - **by device type**  
Choose the device type to filter from the **Show Devices of Type** options. If discovery does not find a single IP address or hostname that was specified in the discovery profile, the device type indicates "Inaccessible". NetOps Portal reports only accessible devices from IP ranges.

**Options:**

- **Manageable:** The data aggregator classifies the device type as manageable when the following criteria are true:

- It can associate the device with a Firewall, Load Balancer, wireless LAN controller (WLC), or wireless access point (AP) context type.
- It cannot identify the device type as Router, Switch, or Server.
- **Pingable:** The data aggregator classifies the device type as pingable when it cannot classify the device type as manageable. For example, the device does not respond to SNMP requests and the data aggregator cannot determine its device type.

- **Inaccessible**

For more information about device type classifications, see [Override Device Types](#).

– **by device state**

Choose the device state to filter from the **Show Devices With State** options.

**Options:**

- **New**

Indicates a device that was discovered for the first time during the discovery run.

- **Changed**

Indicates that a device type has changed from a previous discovery. For example, a previously discovered pingable device is now discovered as a manageable device. Or a previously manageable device with the device type of Switch has now changed to the device type of Router. Devices with only attribute changes, such as hostname, or system description are not classified as "Changed".

- **Deleted**

Indicates that the device has been deleted from the data aggregator since the discovery ran.

- **Unchanged**

Indicates that existing devices have not changed. Existing devices with only attribute changes are also classified as Unchanged. Existing devices that show different IP addresses were discovered and are being monitored with a different IP address. Many devices can respond to multiple IP addresses. DX NetOps Performance Management maintains the full set of IP addresses for each device.

The discovery results appear in the **Discovered Devices** table.

In the **Discovered Devices** table, the **SNMP Profile** column shows the highest ranked SNMP profile to which the device responded. The **State** column indicates the device state.

4. Click **OK** to close the **Discovery History** dialog.

### **Export Discovery Profile History Results**

You can export the discovery profile history results (discovery instance) from discovery runs. On the **Discovery History** dialog, select the discovery instance that you want to download, and then click the CSV icon. If only one discovery instance exists, it is selected by default. If there are more than one, the latest discovery instance is selected. Exported discovery profile history results are downloaded as .CSV files.

### **Run Device Discovery and Generate a Discovery Log for Debugging Purposes**

Generate a discovery log when directed by Broadcom Support. This information is highly technical and might require assistance from Support to understand what is happening.

**Follow these steps:**

1. On the **Discovery Profiles** page, from the **Discovery Profiles** list, select the discovery profile that you want to run device discovery with logging, and then click **Run with Logging**.  
The **Run Discovery Profile with Logging** dialog opens.
2. Confirm your selection by clicking **Yes**.  
Logging with history is run. The data aggregator adds discovered devices to device collections, which initiates component monitoring and polling. The state of the discovery profile changes to "Running".  
If discovery fails for any of your devices (devices are not found), the state of the discovery profile indicates "FAILURE".  
If discovery finds at least one device, the state indicates "PARTIAL\_FAILURE".

**TIP**

You can get a list of the unreachable devices [by way of the discoveryinstances Data Aggregator REST web service](#).

3. Click **History**.

The **Discovery History** dialog for the discovery profile opens. The log is included in the CSV download of the discovery history of the selected discovery instance.

In the **Discovery Summary** list, the following line appears:

Discovery debug log available in the download.

4. To download the debug log, [export the results](#).

## Quickly Discover SNMP Devices

You can have NetOps Portal quickly discover SNMP devices, without the need for SNMP profiles or discovery profiles.

In a quick discovery of Simple Network Management Protocol (SNMP) devices, NetOps Portal attempts to discover them through Internet Control Message Protocol (ICMP) and SNMP protocols. If a device does not respond to SNMP but responds to ICMP, the data aggregator creates a pingable device.

You can have NetOps Portal quickly discover SNMP devices using one of the following protocols:

- [Use the SNMPv1/SNMPv2C Protocol](#)
- [Use the SNMPv3 Protocol](#)

### Use the SNMPv1/SNMPv2C Protocol

If a device that uses an SNMPv1/SNMPv2C profile responds to SNMPv1 and SNMPv2C, NetOps Portal polls devices using SNMPv2C.

#### Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Quick Device Discovery**. The **Quick Device Discovery** page appears.
2. In the **What Devices To Discover** section, for **IP Domain**, specify the IP domain that you want NetOps Portal to use to quickly discover SNMP devices.  
**Default:** Default Domain  
**Required:** Yes
3. In the **Which Devices To Discover** section, complete the following fields, and change any default settings:
  - **SNMP protocol to use**  
 Specify **SNMPv1/SNMPv2C** as the protocol that NetOps Portal uses to quickly discover SNMP devices.  
**Default:** SNMPv3
  - **IP Addresses or Host Names**  
 Specifies the ranges for for Internet Protocol version 4 (IPv4) addresses that NetOps Portal uses to discover SNMP devices. Specify the addresses by entering a comma or space-separated list of IP addresses, host names, or IP address ranges, for example, 191.36.94.1-10 .  
**Required:** Yes
  - **Port**  
 Specifies the port that NetOps Portal uses to make SNMP connections to SNMP devices.  
**Default:** 161  
**Required:** Yes

**NOTE**

NetOps Portal can also send SNMP traps to trap receivers through notifications using this port. In this scenario, use port 162 .

For more information, see [Configure Notifications](#).

- **Community Name**  
Specifies the community name that NetOps Portal uses to gain access to the device.  
**Required:** Yes
- **Re-enter Community Name**  
**Required:** Yes

4. Click **Discover**.

NetOps Portal discovers the devices within the specified IP addresses and host names, and then the data aggregator adds the devices to the system inventory.

### Use the SNMPv3 Protocol

#### **IMPORTANT**

Electing for NetOps Portal to discover devices that use the SNMPv3 protocol adds an extra load of about 30 percent to the CPU of data collectors.

#### **Follow these steps:**

1. From the **Quick Device Discovery** page, in the **What Devices To Discover** section, for **IP Domain**, specify the IP domain that you want NetOps Portal to use to quickly discover SNMP devices.  
**Default:** Default Domain  
**Required:** Yes
2. In the **Which Devices To Discover** section, complete the following fields, and change any default settings:
  - **SNMP protocol to use**  
Specify **SNMPv3** as the protocol that NetOps Portal uses to quickly discover SNMP devices.  
**Default:** SNMPv3
  - **IP Addresses or Host Names**  
Specifies the ranges for for Internet Protocol version 4 (IPv4) addresses that NetOps Portal uses to discover SNMP devices. Specify the addresses by entering a comma or space-separated list of IP addresses, host names, or IP address ranges, for example, 191.36.94.1-10 .  
**Required:** Yes
  - **Port**  
Specifies the port that NetOps Portal uses to make SNMP connections to SNMP devices.  
**Default:** 161  

#### **NOTE**

NetOps Portal can also send SNMP traps to trap receivers through notifications using this port. In this scenario, use port 162 .

For more information, see [Configure Notifications](#).

**Required:** Yes
  - **User Name**  
Specifies the name of the user associated to the secret keys that NetOps Portal uses to authenticate and encrypt SNMP packets.  
**Required:** Yes
  - **Context Name**  
Specifies the collection of management information that SNMP entities can access. The context name is an octet octet string that NetOps Portal requires to provide end-to-end identification and to retrieve data from SNMP agents.
  - **Authentication Protocol**  
Specifies the authentication protocol for NetOps Portal to use when contacting devices to authenticate SNMPv3 packets. DX NetOps Performance Management supports the following algorithms for authenticating SNMPv3 packets:



**Options:**

- None
- MD5 (Message Digest 5)
- SHA (Secure Hash Algorithm)
- SHA2-256
- SHA2-512

**Default:** None**Required:** Yes– **Authentication Password**

Specifies the password that NetOps Portal uses to authenticate using SNMPv3 and the selected authentication protocol.

**NOTE**

This option is enabled if authentication is enabled for the SNMP profile (you have chosen an option for **Authentication Protocol** other than **None**).

**Required:** Yes– **Re-enter Authentication Password****Required:** Yes– **Privacy Protocol**

Specifies the encryption protocol for NetOps Portal to use when contacting devices for data flows sent to those devices or servers.

**NOTE**

This option is enabled if authentication is enabled for the SNMP profile (you have chosen an option for **Authentication Protocol** other than **None**).

**Options:**

- None
- DES
- AES 128
- Triple DES
- AES 256 with 3DES key

**Default:** None**Required:** Yes– **Privacy Password**

Defines the password that NetOps Portal uses when exchanging encryption keys.

**NOTE**

This option is enabled if you have chosen an encryption protocol for contacting devices associated with this profile for data flows sent to those devices or servers (you have chosen an option for **Privacy Protocol** other than **None**).

**Required:** Yes– **Re-enter Privacy Password****Required:** Yes3. Click **Discover**.

NetOps Portal discovers the devices within the specified IP addresses and host names, and then the data aggregator adds the devices to the system inventory.

## Run Device Rediscovery

Running device rediscovery for a monitored device updates the device and runs metric family rediscovery for the device.

**NOTE**

You can also automate device rediscovery [by way of the profiles Data Aggregator REST web service](#).

In this article:

- [Run Device Rediscovery](#)
- [Run Device Rediscovery and Generate a Discovery Log for Debugging Purposes](#)

**Run Device Rediscovery**

**Prerequisite:** The device is in "Active" state. For more information about device life cycle states, see [Manage Device Life Cycles](#).

**Follow these steps:**

1. Log in to NetOps Portal as an Administrator or a Tenant administrator.
2. Hover over **Administration**, **Monitored Items Management**, and then click **Monitored Devices**.  
The **Monitored Devices** page appears. The **Monitored Devices** list displays a list of the devices that NetOps Portal has discovered.
3. From the **Monitored Devices** list, with the **Tree View** tab selected, sort by **Device by Collection**, expand a collection, locate and select the active device that you want to rediscover (the **Life Cycle State** is "Active"), and then click **Rediscover** at the bottom of the page.  
The **Rediscover Device** dialog opens.
4. Confirm your selection by clicking **Yes**.

NetOps Portal rediscovers the device and metric families. Changes in the device attributes can take up to 5 minutes to show in inventory or dashboards views.

**NOTE**

If device discovery fails to discover *all* devices, or it does not discover the device as expected, you can [generate a detailed discovery log of what the rediscovery process finds when it communicates with a device for debugging purposes](#).

**Run Device Rediscovery and Generate a Discovery Log for Debugging Purposes****IMPORTANT**

Generate a discovery log when directed by Broadcom Support. This information is highly technical and might require assistance from Support to understand what is happening.

**Follow these steps:**

1. On the **Monitored Devices** page, from the **Monitored Devices** list, with the **Tree View** tab selected, sort by **Device by Collection**, expand the **All Manageable Devices** collection, locate and select the device that you want to rediscover and generate a discovery log, and then click **Rediscover with Logging** at the bottom of the page.  
The **Rediscover Device with Logging** dialog opens.
2. Confirm your selection by clicking **Yes**.

NetOps Portal rediscovers the device and generates a discovery log. The discovery log appears in the **Discovery History Log** section of the **Details** tab. Only the most recent log is available. The following image shows an example of this log:

**Figure 11: Discovery History Log**

**NOTE**

If the discovery log does not display, click **Refresh**.

## Discover Logical Systems Through SNMP Context

You can discover logical systems on physical devices using the SNMP context name.

The method you use is dependent on the target device. Discover logical systems using one of the following methods:

- [Automatically](#)  
Use this method to automatically discover logical systems on physical devices using the SNMP context name.
- [Manually](#)  
Use this method if the device *does not* provide the logical system context name through SNMP.

### Automatic Discovery of Logical Systems

Out of the box, DX NetOps Performance Management discovers the Check Point Virtual Firewalls logical system. You can add vendor certifications for other logical systems.

#### **Follow these steps:**

1. Create an SNMPv3 profile to discover the physical parent devices.  
For more information, see [SNMP Profiles](#).
2. Associate the **Context System** metric family to a monitoring profile that is attached to the device that hosts the logical systems.  
For more information, see [Manage Monitoring Profiles](#).
3. Wait for discovery to complete.  
The logical system are automatically discovered.

#### **NOTE**

Logical systems have the same IP address as the parent device. However, each logical system has its own context name. The data collector polls the data for the logical system (such as the CPU, Memory, and Interfaces) under the context name.

### Manual Discovery of Logical Systems

You can manually discover logical systems on physical devices using the SNMP profile's context name using REST. For Juniper SRX devices with logical systems, use this method.

Use the following process to manually discover logical systems:

1. [Verify the Prerequisites](#)
2. [Create an Item for a Logical System using a REST Call](#)

#### Verify the Prerequisites

Before creating an item for a logical system using a REST call, verify that you have the following information:

- The tenant item ID (`tenant-item-id`).

#### **TIP**

You can get the tenant item ID by retrieving information about the tenant. Issue the following GET call from a REST client editor or HTTP tool with a connection to the data aggregator:

```
http://da-host:8581/rest/tenants
```

- The item ID of the parent device (`parent-item-id`). This ID corresponds to where you want to discover devices.

#### **TIP**

You can get the item ID for the parent device from the **Item ID** field that is on the **Details** context tab of the **Monitored Devices** page for the device.

For more information about how to view this page, see [Manage Monitored Devices](#).

- The context name of the logical system that is associated with (at the same IP as) the parent device.

**TIP**

You can get the context name for the logical system from the parent system configuration.

**Create an Item for a Logical System using a REST Call**

You can have the data aggregator create an item for the logical system that has a context name associated with the parent device using a REST call. The data aggregator discovers and monitors the logical system as a manageable device at the same IP of the parent device using this context name.

**Follow these steps:**

1. From a REST client, issue the following call with the **POST** method:

```
http://<da-host>/rest/tenant/<tenant-item-id>/devices/subsystems
```

- **da-host**  
The hostname of the data aggregator host.
- **tenant-item-id**  
The ID for the tenant.

**REST Body:**

```
<SubSystemCommInfo version="1.0.0">
  <ParentDeviceItemID><parent-item-id></ParentDeviceItemID>
  <ContextName><context-name></ContextName>
  <ContextNameUsage><context-name-usage></ContextNameUsage>
</SubSystemCommInfo>
```

**Example REST Body:**

In this example, the community string for the parent device's SNMP v2c profile is public, and the logical system's context name (the `ContextName`) is `LSYS1/`. This examples uses `LSYS1/public` as the full community string for the logical system.

```
<SubSystemCommInfo version="1.0.0">
  <ParentDeviceItemID>7776</ParentDeviceItemID>
  <ContextName>LSYS1</ContextName>
  <ContextNameUsage>PREFIX</ContextNameUsage>
</SubSystemCommInfo>
```

- **parent-item-id**  
The item ID of the parent device for the logical system that you want to create.
- **context-name**  
The name of the logical system that is associated with (at the same IP as) the parent device.

**Example:** `LSYS1/`

- (Parent devices that use SNMP v1 or SNMP v2c profiles only) **context-name-usage**  
Defines the usage for the context name.

**Values:**

- **PREFIX:** If the community string for the parent device's SNMP v2c profile is public, the full community string for the logical system (the `ContextNameUsage`) is `<context-name>public`, for example, `LSYS1/public`.
- **SUFFIX:** The community string for the parent device's SNMP v2c profile is public, the full community string for the logical system (the `ContextNameUsage`) is `public<context-name>`, for example, `publicLSYS1/`.

2. Wait approximately one minute.

The item for the logical system is created. The data aggregator discovers and monitors the logical system as a manageable device at the same IP of the parent device using this context name.

**Discovery From Other Data Sources**

You can configure whether new inventory from a data source contributes to the data aggregator.

**Follow these steps:**

1. Hover over **Administration**, **Data Sources**, and then click **Data Sources**.
2. Click the data source that you want contribute to the data aggregator, and then click **Edit**.  
The **Edit Data Source** dialog appears.
3. For **Contribute inventory to the Data Aggregator**, define whether new inventory contributes to the data aggregator:  
**Options:**
  - **Selected:** Contribute new data source inventory to the data aggregator.

**NOTE**  
Automatic synchronization includes only devices that NetOps Portal discovers after you select this option. To discover previously discovered devices, perform a full synchronization of the data aggregator. DX NetOps Performance Management creates a discovery profile that includes the IP addresses of the devices. This discovery profile attempts discovery 1 minute after new IP addresses are pushed into the IP domain discovery profile for the data aggregator. Otherwise, the discovery profile attempts discovery once per day.  
For more information, see [Configure a Data Source](#) and [Synchronize Data Sources](#).

  - **Cleared:** The existing inventory remains the same, and the new data source inventory is not contributed to the data aggregator.

**Default:** Cleared  
**Required:** No
4. Save your changes.

## Discovery and Polling in VMware Environments

DX NetOps Performance Management can discover and monitor your VMware virtual machines and ESX hosts.

The discovery and monitoring process for these devices and components differs to accommodate the collection of data from vCenter. DX NetOps Performance Management can discover the VMs and ESX hosts directly with SNMP, and can collect vCenter data through the vCenter Server Application Insight Module (VCAIM).

DX NetOps Performance Management discovers devices (ESX hosts and virtual machines) in the following ways:

- Through ICMP
- Through SNMP, if the servers have an SNMP agent deployed
- [Through discovery of a server running SystemEDGE with the VCAIM](#)

**NOTE**

If the lifecycle state of a virtual device that DX NetOps Performance Management discovers with SystemEDGE is "Retired", DX NetOps Performance Management does not poll for its vCenter statistics.

DX NetOps Performance Management creates only one device item for each ESX host or virtual machine. By default, the change detection rate for virtual machines is 15 minutes. The change detection rate for ESX hosts is 24 hours.

### SystemEDGE Discovery

When discovery runs for a server running SystemEDGE with the VCAIM, DX NetOps Performance Management uses the following logic during discovery of the vCenter Server:

- The VMware ESX host monitoring profile discovers the ESX hosts, and DX NetOps Performance Management uses the following rules to create devices:
  - DX NetOps Performance Management creates SNMP devices for each device with an IP address that is used by another discovered device. The device type is ESX Host and Server.
  - DX NetOps Performance Management creates pingable devices with the ESX Host device type in the following situations:

- A device does not have an IP address.
- The IP address of a device is not used by another device.
- The VMware virtual machine monitoring profile discovers the virtual machines, and DX NetOps Performance Management uses the following rules to create devices:
  - DX NetOps Performance Management creates SNMP devices for each device with an IP address that is used by another discovered device. The device type is Virtual Machine and Server.
  - DX NetOps Performance Management creates pingable devices with the Virtual Machine device type in the following situations:
    - A device does not have an IP address.
    - The IP address of a device is not used by another device.

## SystemEDGE System Response Path Test Metrics

You can view the overall status for devices that support service level agreements (SLA) using the response path test metrics for SystemEDGE that DX NetOps Performance Management includes.

You can view the overall status using various protocols and get a broad sense of overall SLA performance for the enterprise or a specific device. For example, as a network infrastructure manager, you can identify which protocols are experiencing the highest latency and slowest response times from the dashboard. You can then look at the device-level views to determine which devices are experiencing the problems. This information can help you determine if increased activity, a device that is down, or something else caused an issue.

DX NetOps Performance Management includes the following response path test metrics for SystemEDGE for the Service Availability Response Path metric family:

- Attempts
- Avg. DNS Lookup Time
- Avg. Response Time
- Avg. TCP Connect Time
- Avg. Transaction Time
- Bytes Received Total
- Descriptions
- Failed Transactions
- Indexes
- Max. DNS Lookup Time
- Maximum Response
- Max. TCP Connect Time
- Min. Transaction Time
- Names
- Percent Failed Attempts
- Percent Successful Attempts
- Response/Limit
- Response Throughput
- Successful Attempts
- Successful Transactions
- Total Errors

The SystemEDGE Response Path vendor certification includes this metric family.

# Groups

Groups determine the data that you see in dashboards when you log in. The group that is applied as a filter to the current dashboard is the group context for that dashboard.

When you log in to NetOps Portal, the pages reflect the context of your default permission group. You can change the default group for your user account to view data from another group in the dashboards.

*Groups* are organized into a hierarchical tree structure. In NetOps Portal, use the **Groups** tree to define relationships, policies, and dependencies among services, devices, applications, locations, and users within your organization. Organize your group structure according to business and reporting needs. To create a regional structure that represents regions, countries, and locations, use site groups. Use custom groups for other types of organizations, such as customers, services, or technologies.

For more information:

- About how to use site groups, see [Organize Group Items Geographically](#).
- About how to use custom groups, see [Manage Groups](#).

*Threshold profiles* apply threshold rules to all items in a group. The group hierarchy requirements for thresholding are probably different from the requirements for reporting. Create separate groups that address both sets of requirements. Consider the different layers of the network and how to create thresholds for components in those layers. For example, you might threshold on CPU, memory, and interface metrics on the core network differently to the distribution layer. Create multiple groups to apply threshold rules appropriately.

For more information:

- About how to assign groups to a threshold profile, see [Configure Threshold Profiles](#).
- About how to add managed items, such as devices, switches, routers, interfaces, device components, virtual interfaces, aggregated components, service set identifiers (SSIDs), and tunnels to groups assigned to a threshold profile, see [Manage Groups](#).

*Tenants* include special types of system groups to maintain separation among deployments. Tenants can also contain custom grouping structures.

## NOTE

The lock icon that appears on a group icon, such as system groups, indicates that you cannot edit it.

## Group Types

NetOps Portal includes the following group types:

- [System Groups](#)
- [Custom Collections Group](#)
- [Inventory Group](#)
- [VNA Domains Group](#)
- [Custom Groups](#)
- [Site Groups](#)
- [Groups for Multi-Tenant Deployments](#)

## System Groups

System groups are read-only groups that are automatically created based on information from data sources. You can view system groups, apply them as permission groups to user accounts, or copy them to custom or site groups.

## Custom Collections Group

The **Custom Collections** group represents the collections of devices. Collections are groupings of devices that are monitored using the rules that are specified in monitoring profiles. The out-of-the-box collections are not visible in the Groups tree. Collections include only devices.

This group lets you create custom collections. Any subgroup that you add to the Collections group is synchronized to the data aggregator as a collection.

## Inventory Group

The **Inventory** system group includes all managed items that are discovered by all registered data sources. This group also organizes data sources, IP domains, and managed items in subgroups. The Inventory group contains its own system subgroups to organize managed items by their type.

### IMPORTANT

NetOps Portal does not synch `<DataSource>@Hostname` groups, such as `DataAggregator@mydahost.ca.com`, to data sources. To avoid synchronization errors, do not use these groups for reporting or as default user groups.

The following system groups appear under the **Inventory** node:

- **All Items**

This system group includes subgroups of managed items, which are categorized by type. Expand the All Items system group to display the following subgroups:

- **Application Delivery Analysis Networks**

This subgroup includes all networks that CA Application Delivery Analysis (ADA) has observed. An ADA network consists of an IP address and mask.

- **ESX Hosts**

This subgroup includes all VMware servers that host virtual machines.

- **Interfaces**

This subgroup includes router and switch interfaces from all data sources.

- **Pingable Devices**

This subgroup includes all discovered devices that cannot be contacted using SNMP.

- **Routers**

This subgroup includes all routers from all data sources.

- **Servers**

This subgroup includes all servers from all data sources.

- **Switches**

This subgroup includes switches from all data sources.

- **Virtual Machines**

This subgroup includes all virtual machines running on all ESX servers.

- **Data Sources**

This system group includes groups of registered data sources. Some data sources include subgroups, which appear when you expand the data source group.

### NOTE

Some data sources, such as data aggregator and event manager data sources, do not have groups. These groups are not reporting groups and might not provide data when selected as the context.

- **IP Domains**

This system group includes the custom IP domains that the administrator created. This system group also includes the Default Domain subgroup. This subgroup contains the items that are not explicitly assigned to a custom domain. For more information, see [IP Domains](#).



## **VNA Domains Group**

The **VNA Domains** system group includes the groups for the DX NetOps Virtual Network Assurance (VNA) plug-ins, include the Default Domain. These groups contain hierarchically-organized VNA devices. DX NetOps Performance Management uses VNA domains as a reporting group. NetOps Portal automatically organizes VNA data by plug-in under this system group. Use these groups for reporting. So that you can view VNA data in the context of a NetOps Portal user, NetOps Portal includes predefined dashboards that you can filter using these groups to only show only VNA data in the dashboards.

You can verify that the data from the plug-ins are integrating as expected using the groups in this system group. For more information, see the following video:

### **TIP**

You can configure the highest-level groups in the VNA Domains system group in VNA.

For more information about how to filter dashboards using these groups, see [Manage Dashboards](#).

## **Custom Groups**

Custom groups create hierarchical levels and organize items into logical relationships within the Groups tree. Custom groups at the top level of the Groups tree represent geographical, topological, or functional divisions within your organization. Lower-level custom groups, or subgroups, typically represent one of the following options:

- Applications
- Devices
- Job functions of IT staff
- Services

Users with the Administer Groups Owned by You role right can create and edit custom groups, which filter the data in dashboards and views. The group context for a dashboard or view determines the data that is presented.

Monitor and manage your system, such as organizing data and assigning operator permissions to access data, using custom groups.

You can use group rules to add items to groups automatically as they are discovered. Setting up rules makes it easier to populate and maintain groups. You can also populate custom groups by manually adding specific items, such as routers or interfaces that are logically or geographically related.

## **Site Groups**

Site groups are special custom groups that are based on sites, such as branch offices, or on physical locations, such as regions or cities. Site groups let you create navigation functions within dashboards to present views across all sites. They include a Time Zone and a Business Hours parameter to let you see prioritized data from business-critical times of day.

Site groups also provide a granular context to apply to dashboards. For example, after you create a site group for each of your sites, a single dashboard can individually report on each site. Create a site group for each data center within your enterprise, and for other major infrastructure locations.

## **Groups for Multi-Tenant Deployments**

When the administrator for the Default Tenant creates at least one tenant, features to support multi-tenancy are enabled. Multi-tenant deployments consist of multiple discrete enterprises with IP addresses that might overlap. More groups appear in the Groups tree to let the administrator organize tenant inventories and allocate permissions:

- **Tenants Group**  
The Tenants group includes all tenants. Tenants are used with IP domains to monitor separate customer environments with a single NetOps Portal instance. Each tenant can contain multiple subgroups of items that are not shared among tenants.

Tenant administrators can create custom groups within their tenant. For the global administrator, tenant groups appear under the Tenant node in the Groups tree.

- **Global Tenant Groups Group**

The Global Tenant Groups group contain groups of items that help the global administrator manage tenant environments. These groups let the administrator visualize and organize shared items, which are not explicitly associated with a tenant IP domain. The groups that allocate access to data from shared items appear under each tenant.

When you expand the top-level Inventory group, the following group appears in a multi-tenant deployment:

- **IP Domains**

Includes all custom IP domains that are used to associate managed items with tenants. Also includes the Default Domain, which contains all items that are not explicitly assigned to a custom domain. For more information, see IP Domains.

In a multi-tenant deployment, each tenant has its own groups. Tenant users cannot see items outside of the tenant group unless the global administrator grants access with Service Provider Defined groups:

- **Groups (Tenant)**

Allows global administrators or tenant administrators to create, or add, custom groups.

- **Inventory (Tenant)**

Includes all managed items that are associated with the tenant IP domains. Items from all registered data sources might appear in this group.

Each tenant also has the following system subgroups in its Inventory group:

- **IP Domains**

The IP Domains subgroup represents the IP domains that are associated with this tenant. Any managed items that have been discovered are associated with this tenant through its IP domains. To see the managed items of the tenant, click a tenant IP domain in the Groups tree.

- **Global Tenant Groups**

The Global Tenant Groups include groups that the global administrator has populated with shared items that this tenant can access. Use these groups to grant access to data from shared devices to selected tenant user accounts. For example, a router that the service provider owns handles traffic from multiple tenant domains. Using Service Provider Defined groups, the global administrator can allocate tenant access to data from that router. The tenant can then independently monitor and verify system performance.

- **Global Tenant Items**

Items that are not explicitly associated with a tenant IP domain are automatically placed in the Service Provider Items group. Global administrators can place these items into Service Provider Defined groups to allocate tenant access to data from shared items.

## Device Collections

Collections are groups of devices with common monitoring configuration. Custom (device) collections provide granular control over polling. Most deployment require custom device collections. NetOps Portal includes out-of-the-box device collections for use in a lab or in a demo setting. In a production environment, the best practice is to design and configure custom device collections to have granular control over what is being polled.

The following video examines how to create a custom collection in NetOps Portal to organize network items with common monitoring configurations:

In this article:

- [Custom Device Collections](#)
- [Out-of-the-Box Device Collections](#)
- [View a List of Device Collections](#)

## **Custom Device Collections**

In a production environment, as a best practice, design and configure custom device collections to have granular control over what NetOps Portal polls. For example, disable polling of a device by disassociating the device from any other device collection that has monitoring profiles that are associated to it. If you are associating monitoring profiles to the out-of-the-box device collections (such as All Routers), then you cannot stop a single device from being polled. Devices cannot be removed from out-of-the-box device collections, so you disassociate the monitoring profiles instead to disable polling. You then create custom device collections that contain devices to which you want to apply the same polling policy. Associate monitoring profiles (or custom monitoring profiles) to those custom device collections to begin polling.

Create custom device collections in NetOps Portal, then either synchronize them immediately with Data Aggregator, or wait for the automatic synchronization. Upon synchronization, Data Aggregator creates the corresponding device collections for use in monitoring devices.

Access the Monitoring Configuration menu to see a list of device collections and to see the monitoring profiles that are applied to each. Administrators can view the device collections for the tenant they are administering. A tenant administrator can view its own list of device collections.

### **NOTE**

Only users with the Administrator role can create or edit collections because collections control polling. Collections can significantly impact system load and performance.

Consider the following best practices for organizing devices into collections for monitoring:

- Create custom collections that match the monitoring requirements in the environment:
  - Consider the different layers of the network, access, distribution, and core. Devices in different layers might require different levels of monitoring.
  - Consider which technologies and metric families are required. Metric families that would be applied to all devices, such as CPU and memory, apply to broad collections. Targeted monitoring, such as QoS and IPSLA, apply to limited collections.
- Create collections that enable the flexibility to break out monitoring:
  - Some devices are included in multiple collections so that specific metric families are polled at different rates.
  - Devices in different collections have different filtering criteria.
  - Different monitoring requirements depending on importance of device

## **Out-of-the-Box Device Collections**

NetOps Portal uses these device collections to get data into NetOps Portal quickly, and to test NetOps Portal. NetOps Portal associates the devices that it detects during discovery to these device collections depending on their type. For example, NetOps Portal associates routers to the All Routers device collection. Upon synchronization, NetOps Portal associates these monitored devices to the corresponding device collections.

NetOps Portal then automatically applies the out-of-the-box monitoring profiles to the out-of-the-box device collections, allowing NetOps Portal to collect data immediately without any intervention on your part. After NetOps Portal collects this data, you can run reports on the data to gain a better understanding of your network.

NetOps Portal includes the following out-of-the-box device collections:

- [All Devices](#)
- [All Routers](#)
- [All Servers](#)
- [All Switches](#)
- [All Manageable Devices](#)
- [All ESX Hosts](#)
- [All Virtual Machines](#)
- [All VMware vCenters](#)

**NOTE**

Out-of-the-box device collections are mostly for use in a lab or in a demo setting. In a production environment, the best practice is to design and configure custom device collections to have granular control and optimal data collection.

**All Devices**

This device collection is an out-of-the-box device collection. Manageable and pingable devices that are detected during a discovery are automatically placed into the All Devices device collection. Inaccessible devices are not included in this device collection.

**IMPORTANT**

This is reserved device collection. To avoid extra SNMP requests being made to pingable-only devices, and sporadic metric family support, do not associate monitoring profiles with this device collection.

**All Routers**

This device collection is an out-of-the-box device collection. Routers that are detected during a discovery are automatically placed into this device collection.

**NOTE**

Routers can appear in both this device collection and in the **All Switches** device collection.

**All Servers**

This device collection is an out-of-the-box device collection. Physical and virtual servers (hosts) that are detected during a discovery are automatically placed into this device collection. Network devices such as routers and switches are not included in this device collection.

**All Switches**

This device collection is an out-of-the-box device collection. Switches that are detected during a discovery are automatically placed into this device collection.

**NOTE**

Switches can appear in both the **All Routers** device collection and in this device collection.

**All Manageable Devices**

This device collection is an out-of-the-box collection. Manageable devices collect advanced performance statistics and are monitored with a protocol such as SNMP. Manageable devices that are detected during a discovery are automatically placed into this device collection.

Pingable devices can only be monitored for availability and do not provide any additional performance metrics. Therefore, pingable devices are not included in this device collection.

**NOTE**

Manageable devices can appear in both the **All Devices** device collection and in this device collection.

**All ESX Hosts**

This device collection is an out-of-the-box device collection. ESX hosts that are detected during discovery are placed into this device collection automatically.

## **All Virtual Machines**

This device collection is an out-of-the-box device collection. VMware virtual machines that are detected during discovery are placed into this device collection automatically.

## **All VMware vCenters**

This device collection is an out-of-the-box device collection. All servers that are running systemEdge with the VCAIM that are detected during discovery are placed into this device collection automatically.

## **View a List of Device Collections**

You can view a list of the device collections and view the monitoring profiles that are applied to each. Administrators can view the device collections for the tenant they are administering. A tenant administrator can view its own list of device collections.

### **Follow these steps:**

1. As a user in the default tenant with the Administrator role, hover over **Administration, Monitored Items Management**, and then click **Monitoring Profiles**.  
The **Monitoring Profiles** page appears. The **Monitoring Profiles** tab is selected under the **Monitoring Configuration** section.
2. Under the **Monitoring Configuration** section, click **Collections**.  
The **Collections** page appears.

## **Manage Groups**

You manage groups by adding, editing, and removing groups.

By default, only users with the Administrator or Tenant Administrator role can manage the groups that other users add. To manage groups that other users have added, you must have the Administer Groups Owned by You and Others role right.

For more information about these role rights, see [Role Rights](#).

### **NOTE**

You can manage groups using NetOps Portal or using the NetOps Portal REST API (the `groups` endpoint). This article describes how to manage them using NetOps Portal.

For more information about how to manage group using the NetOps Portal API, see [Manage Groups Using the Groups Web Service](#).

In this article:

- [Add or Edit a Group](#)
- [Add a Managed Item to a Group](#)
- [View Group Membership](#)
- [Deep Copy a Group](#)
- [Manage Group References](#)
- [Move a Group](#)
- [Remove a Group](#)

The following video examines how to logically organize network items in NetOps Portal by creating and managing groups, and using groups as data set filters for dashboards, views, and context pages:

## **Add or Edit a Group**

You can add a group or edit the properties of an existing group. You cannot modify a group or its contents from a group reference.

For more information about group references, see the ["Manage Group References"](#) section.

### IMPORTANT

To configure views or on-demand reports with the group, after adding the group, wait for the group to be synchronized to the data sources. Newly-added groups are empty until you add managed items to the groups.

### Follow these steps:

1. Log in as an Administrator or as a user with the My Custom Groups functionality enabled (the Administrator has selected the **Enable "My Custom Groups" Functionality** checkbox for your user account).  
For more information about this access permissions setting, see [Manage User Accounts](#).
  2. Do one of the following tasks:
    - Hover over **Administration**, **Group Settings**, and then click **Groups**.  
The **Manage Groups** page appears.
    - Click the name of your user account in the upper-right corner, and then click **My Custom Groups**.  
The groups that display are the groups that the Administrator selected for you (custom groups), based on your responsibilities.
- NOTE**  
Custom groups are read-only groups, and appear only as group references; the group's **Properties** tab shows a path to the original group.  
For more information, see [Manage Subgroups](#).
3. Do one of the following tasks:
    - **Add a group.** Select a location for the new group in the **Groups** tree, and then click **Add Group**. Add groups under the **All Groups** node in the **Groups** tree, or within an existing custom or site group to add a subgroup.

**NOTE**  
You cannot add groups to system groups, which appear locked in the **Groups** tree. However, while the **Custom Collections** group appears locked and cannot be modified, it can contain subgroups.  
For more information, see [Device Collections](#).  
The **Add Group** dialog appears.

    - **Edit a group.** Select the group that you want to edit from the **Groups** tree.
  4. Specify or edit the values for the following fields, and then click **Save**:
    - **Group Name**  
Specifies the name for the group. Use a strategic naming convention. You can use all special characters in the group name except the vertical bar (|), the forward slash (/), and the backslash (\).

**NOTE**  
Groups that are synchronized from DX NetOps Spectrum retain restricted characters except the backslash (\). This character is removed from synchronized group names.

    - **Group Type**  
Defines the type of group to add.

**Options:**

    - **Custom:** Specifies a custom group. Use custom groups to manage items in NetOps Portal.
    - **Site:** Specifies a site group. Site groups contain items and subgroups of items that are grouped by location.

**NOTE**  
If you are deploying business hours, add a site group. Site groups are custom groups that are based on physical locations, such as a city, region, office, or campus. With site groups, you can precisely filter data from business-critical times of day.  
For more information, see [Configure Business Hours Filtering](#).  
For more information, see [Organize Group Items Geographically](#).

**Default:** Custom

    - **Description**

Specifies the description for the group.

– **Include the children of managed items**

Determines whether the group inherits the children of managed items.

**Values:**

- **Selected:** The group inherits the children of managed items that are added to this group.
- **Cleared:** The group does not inherit the children of managed items that are added to this group. For example, if you add a router to the group, the interfaces on that router are not included, and their data is not visible in drilldown views.

**Default:** Selected

– (Site groups) Optionally specify the following parameters:

- **Latitude**
- **Longitude**
- **Elevation**
- **Location**
- **Time zone**
- **Business hours**

(Adding a group) The new group appears in the **Groups** tree. Otherwise, the group is edited.

5. **Next step:** [Add managed items to the custom group.](#)

The group is added or edited. If you added a subgroup within an existing custom or site group, the existing group and its subgroups are added to the selected parent group as group references.

### Add a Managed Item to a Group

Use one of the following methods to add managed items to a custom group:

- [Add managed items to the group \*manually\*.](#)
- [Create group rules that add managed items to the group.](#)

**IMPORTANT**

For groups for a data aggregator data source, to keep reporting times within reasonable limits, limit the group membership to 10,000 managed items, including the children of those managed items.

### Add a Managed Item Manually

Populate custom groups by adding managed items to the groups manually.

**NOTE**

System groups with a lock symbol in the **Groups** tree are read-only. You cannot add items to or delete items from system groups. Custom groups that are added by an Administrator are also locked.

**Follow these steps:**

1. From the **Manage Groups** page, select the group to which you want to add a managed item from the **Groups** tree. Managed items that have already been added to this group appear in the **Group Details** pane. Managed items that you manually add directly to a group appear as **Direct Items** in this pane. Managed items that you add to a group as children of a managed item are **Inherited Items** in this pane.
2. Expand the **Items** section, and then scroll to the managed item type that you want to add to the group. The items that you can add to a custom group depend on the item type (Direct Items, Direct and Inherited Items, Inherited Items, or Excluded Items), the registered data sources, and the items that are discovered. For example, you can add devices, switches, routers, interfaces, device components, virtual interfaces, aggregated components, SSIDs, and tunnels.  
To see more pages of items, click the links below the list. You can also search for an item in the list using the Quick Filter.

3. For example, to add devices, switches, or routers to the group, click **Add Devices**.  
The **Add Items to Group** dialog opens.
4. Select the devices, switches, and routers that you want to add to the group by selecting their checkboxes, and then click **Add**. To select all devices, switches, and routers on a page, select the checkbox in the table header row.  
The devices, switches, and routers are added to the **Selected Devices** list.
5. Click **Save**.

The items are added to the group.

### **View Group Membership**

View a sortable list of all items that have been added to a system group or custom group on the **Manage Groups** page. You can verify group rules, or whether custom scripts have added and populated groups appropriately. You can also view a list of items in a selected group.

Use filters to select the types of items that you want to see, such as all items that were added to the group manually. By default, the list on the Items tab only shows items that are added directly to the group. The items are added either manually or by the application of a rule (Direct Items membership). If you report on a dashboard in the context of a group, all items in the subgroups of the group appear in the dashboard. The items in the subgroups do not appear on the **Group Administration** page.

#### **Follow these steps:**

1. From the **Manage Groups** page, select the group for which you want to view membership details from the **Groups** tree.
2. Scroll to the **Items** section.

#### **TIP**

For optimal performance, collapse the **Items** section until you need to manage items. Likewise, collapse each managed item type until you need to manage a specific item type.

3. To specify the items to display, select a filter from the **Show Items** drop-down list.  
The following membership types are applicable, depending on the type of group that you selected:
  - **Direct Items**  
Includes the items that were added directly to the group, either manually or by the application of a rule. You can add or delete items only when the **Direct Items** membership type is selected. The **Added By** column indicates whether the item was added manually or by a group rule.
  - **Direct and Inherited Items**  
Includes the items in the group, whether they were added directly or inherited as the children of managed items that were added directly to the group. The **Include the children of managed items** checkbox on the **Properties** tab determines whether the group inherits the children of managed items. Excluded items are not inherited.  
For more information about this field, see [Add or Edit a Group](#).
  - **Inherited Items**  
Includes only the children of managed items in the group. When you enable inheritance for this group and add a router, the interfaces that are associated with the router are added to the group.  
You can delete inherited items only by deleting the parent item.
  - **Excluded**  
Group rules do not add items in this list to the group. This list also includes items that you have deleted manually if a group rule originally added that item. To add an item in the **Excluded** membership type to a group, delete the item from the **Excluded** list, and run group rules.
4. Select an item type from the list.  
A list of all items of the selected type that are included in the group appear.

### **Deep Copy a Group**

A group deep copy is a copy of the group and its children (groups, items, group references, and rules).



The following rules apply to deep copying groups:

- If you paste the subgroup to the immediate parent where the pasted group is unique, the group name stays the same.
- If NetOps Portal creates a duplicate name under the immediate parent, it renames the pasted group to `<groupname> Copy`.
- If `<groupname> Copy` is not unique, then NetOps Portal adds a number to the name, for example `<groupname> Copy (1)`.
- If the length of the group name is too large during the renaming, NetOps Portal renames the subgroup without the name `Copy` and without the parenthesis surrounding the number, for example `<groupname> 1`.
- If the length of the subgroup name exceeds the maximum character length (64), an error appears asking you to shorten the original group name.

#### Follow these steps:

1. From the **Manage Groups** page, select the group that you want to deep copy, and then select **Copy Group**.

##### NOTE

If this button is not displayed, click **More Actions**.

The (group copied) message appears at the top of the **Groups** list.

2. Select the group (new destination group) to which you want to copy the group, and then click **Paste Copy**.

##### NOTE

If this button is not displayed, click **More Actions**.

The group is copied to the new location.

### Manage Group References

A group reference is a pointer, or link, to another group. You can manage group references in the following ways:

- [View group references](#)
- [Create a group reference](#)
- [Remove a group references](#)

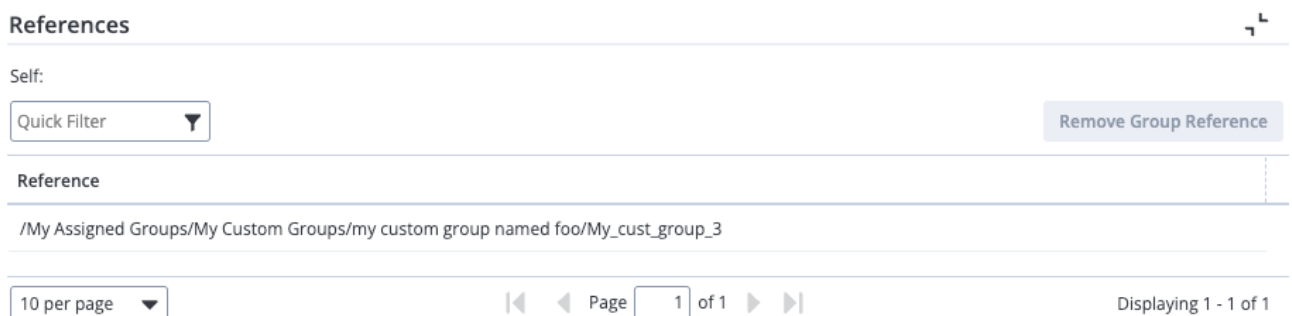
### View Group References

To view group references, from the group that includes the group references that you want to view, scroll to the **References** section.

Groups that include a group reference have a **References** section at the bottom of the **Group Details** page. The following image shows a group with a group reference:

**Figure 12: Group reference**

Group



reference

## Create a Group Reference

Creating a group reference includes the group's subgroups. The original group is displayed in the **References** section. You can view the locations where references of the group have been placed from this section. Any changes that you make to the original group are reflected in the group references.

### Follow these steps:

1. From the **Manage Groups** page, right-click the group to which you want to create a reference, and then select **Copy Group**.

#### NOTE

If this button is not displayed, click **More Actions**.

The (group copied) message appears at the top of the **Groups** list.

2. Select the parent group to where you want to create the group reference, right-click, and then select **Paste Reference**.

#### NOTE

If this button is not displayed, click **More Actions**.

The group and its subgroups display as a group reference in the hierarchy of the selected parent group. The group icons indicate that they are read-only group references.

## Remove a Group Reference

Removing a group that includes the group references removes the group references. Removing a group reference does not affect the original group. Removing a subgroup that is a group reference (a subgroup reference) does not affect the original group or its parent group.

You can remove *copied* group references.

For more information removing subgroups, see [Manage Subgroups](#).

To remove a group reference, from the **References** section, select the group reference that you want to remove, and then click **Remove Group Reference**.

## Move a Group

You can move a group to a new location (destination parent group), which removes it from the previous location. Any existing reports on this group continue to run since the group still exists, but with a different parent. The group's name remains the same.

### Follow these steps:

1. From the **Manage Groups** page, select or right-click the group that you want to move to a new location, and then click or select **Copy Group**.

#### NOTE

If this button is not displayed, click **More Actions**.

The (group copied) message appears at the top of the **Groups** list.

2. Select the new location (destination parent group) to which you want to move the group, and then click **Move Group**.

#### NOTE

If this button is not displayed, click **More Actions**.

The group is moved to the new location and removed from the previous location.

## Remove a Group

The global administrator can remove custom groups, including groups that belong to any tenants. A tenant administrator can also remove custom groups that belong to that tenant definition.

**NOTE**

Removing a group removes its subgroups and its group references. To prevent issues when removing a group that includes group references, remove the group references prior to removing the group.

**Limitations:** You cannot remove system groups. Likewise, you cannot remove the **Default Domain** group.

From the **Manage Groups** page, select or right-click the group that you want to remove from the **Groups** tree, and then click or select **Remove Group**.

## Identify Empty and Unused Groups

As part of a regular maintenance, generate a list of groups that are empty and are unused, and then remove them.

(23.3.5 and higher) Retrieve the list of empty and unused groups by running the `findEmptyGroups.sh` script.

**Prerequisite:** You have a REST API token. If you do not have a REST API token, you can [generate one](#).

The script can generate one of the following types of files:

- [A CSV file](#)
- [An HTML file](#)

### Script usage:

```
./findEmptyGroups.sh [-o filename] [-h filename] [-u username] [-t dir] [-s]
```

- `-o`  
The CSV output filename.  
**Default:** `./emptyGroups.csv`
- `-h`  
The HTML output filename.  
**Default:** none
- `-u`  
Specifies that you want to enter the username for REST calls instead of token authentication. In this case, you are prompted for a password. If you do not specify the `-u` option, you are prompted for a REST API token, which [you can generate](#).
- `-s`  
List only the groups that views, rules, and profiles are not using.
- `-t`  
The location for temporary files.  
**Default:** `/tmp`

## Generate a CSV File of Empty and Unused Groups

### Follow these steps:

1. Run the `findEmptyGroups.sh` script from the `<installation_directory>/PerformanceCenter/Tools/bin/` directory:

```
./findEmptyGroups.sh
```

#### Example:

```
./findEmptyGroups.sh
```

#### – **installation\_directory**

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

**NOTE**

The default CSV output file name is `emptyGroups.csv`. To specify another file name, run the script using the `-o` option, specifying the file name:

```
./findEmptyGroups.sh -o <csv_file_name>
```

**Example:**

```
./findEmptyGroups.sh -o MyemptyGroups.csv
```

The following images shows an example of running the script:

**Figure 13: Image of running the `findEmptyGroups.sh` script**

```
/opt/CA/PerformanceCenter/Tools/bin
[root@ ~]# ./findEmptyGroups.sh
Enter database password:
Enter REST API token:
Mon Dec 18 17:04:54 UTC 2023 ===== Get URLs from DB
Mon Dec 18 17:04:54 UTC 2023 ===== Get Notification Rules from DB
Mon Dec 18 17:04:54 UTC 2023 ===== Get list of empty groups
Mon Dec 18 17:05:14 UTC 2023 ===== Add references column
Mon Dec 18 17:05:14 UTC 2023 ===== Check groups for notifications
Mon Dec 18 17:05:14 UTC 2023 ===== Add Has Notification column
Mon Dec 18 17:05:14 UTC 2023 ===== Check groups for monitoring profiles and threshold profiles
Processing 73 groups
Mon Dec 18 17:05:22 UTC 2023 ===== Add Has Monitoring Profile column
Mon Dec 18 17:05:22 UTC 2023 ===== Add Has Threshold Profile column
Mon Dec 18 17:05:22 UTC 2023 ===== Generate paths for groups
Processing 73 groups
Unable to get path for group 443191. Path will be left blank.
Unable to get path for group 671505. Path will be left blank.
Unable to get path for group 1810103. Path will be left blank.
Unable to get path for group 1810105. Path will be left blank.
Unable to get path for group 1814809. Path will be left blank.
Unable to get path for group 1815332. Path will be left blank.
Unable to get path for group 1815333. Path will be left blank.
Unable to get path for group 1815334. Path will be left blank.
Unable to get path for group 1815343. Path will be left blank.
Unable to get path for group 1815369. Path will be left blank.
Unable to get path for group 1815372. Path will be left blank.
Unable to get path for group 1832863. Path will be left blank.
Mon Dec 18 17:05:26 UTC 2023 ===== Generate URLs for groups
[root@ ~]# cat emptyGroups.csv
itemID,Username,Name,Is Referenced,Is In Rule,Is In Dashboard,Has Notification,Has Monitoring Profile,Has Threshold Profile,path,URL
10929,cd632928,Empty Collection,,,,,,,,/All Groups/Custom Collections/Empty Collection,http://
p/page?pg=99605&GroupTreeID=10929
40881,
40896,
443191
```

2. At the **Enter database password:** prompted, enter the database password.
3. At the **Enter REST API token:** prompt, enter the generated REST API token.

The script generates a CSV file of empty and unused group.

**TIP**

Use the links in the entries in the file to [remove that empty group](#).

## **Generate an HTML File of Empty and Unused Groups**

### **Follow these steps:**

1. Run the `findEmptyGroups.sh` script from the `/opt/CA/PerformanceCenter/Tools/bin/` directory, using the `-h` option, specifying the HTML output file name:

```
./findEmptyGroups.sh -h <html_file_name>
```

**Example:**

```
./findEmptyGroups.sh -h emptyGroups.html
```

2. At the **Enter database password:** prompted, enter the database password.
3. At the **Enter REST API token:** prompt, enter the generated REST API token.

The script generates an HTML file of empty and unused groups.

**TIP**

Use the links in the entries in the file to [remove that empty and unused group](#).

**Next Step**

[Remove the empty and unused groups](#).

## Manage Subgroups

Populate custom groups by adding subgroups that contain managed items.

By default, only users with the Administrator or Tenant Administrator role can administer groups that other users add. To administer groups that other users have added, you require the Administer Groups Owned by You and Others role right. With this role right, you can modify subgroups that other users have added if you have the rights to administer the parent group.

For more information about this role right, see [Role Rights](#).

In this article:

- [Add a Subgroup to a Group](#)
- [Remove a Subgroup from a Group](#)

For more information about the following, see [Manage Groups](#):

- Deep copying subgroups.
- Moving subgroups.

**Add a Subgroup to a Group**

You can create a hierarchical structure by adding groups within custom groups that you previously added. This hierarchical structure is called a subgroup.

You can also add subgroups by copying an existing group into another group.

You can add groups to all groups except system groups, which appear locked in the **Groups** tree. However, while the **Custom Collections** group appears locked and cannot be modified, it can contain subgroups. When you add a subgroup to a group, the existing group and its subgroups are added to the selected parent group as reference groups.

**IMPORTANT**

To keep reporting times within reasonable limits, a maximum of 2000 subgroups is recommended for any parent group.

For more information, see [Manage Groups](#).

**Remove a Subgroup from a Group**

Removing a subgroup that is a group reference (a subgroup reference) does not affect the original group or its parent group. Removing a parent group deletes its subgroups.

For more information, see [Manage Groups](#).

## Manage Group Rules

Keep custom groups up-to-date when systems and networks change using group rules.

Group rules add newly-discovered items, such as devices, that meet rule specifications to groups. Similarly, if those items do not meet rule requirements, or if the items are no longer monitored, group rules delete the item from the group.

#### NOTE

Group rules cannot add items that you have deleted *manually* from a group to the group again.

You can manage group rules by creating, editing, and deleting them from groups. You can use NetOps Portal or use the NetOps Portal API (the `groups` endpoint). This article describes how to manage them using NetOps Portal.

For more information about how to manage group rules using the NetOps Portal API, see [Manage Groups Using the Groups Web Service](#).

In this article:

- [Create a Group Rule](#)
- [Edit a Group Rule](#)
- [Delete a Group Rule](#)
- [Determine Why a Group Rule Added Unexpected Items to a Group](#)

### Create a Group Rule

The following default restrictions apply to group rules:

- Maximum 50 rules per group.
- Maximum 50 conditions per rule.
- Maximum 50 'OR' matches in a condition.
- Maximum 20 'AND' subrules.

#### Follow these steps:

1. Log in as an Administrator or as a user with the My Custom Groups functionality enabled (the Administrator has selected the **Enable "My Custom Groups" Functionality** checkbox for your user account).  
For more information about this access permissions setting, see [Manage User Accounts](#).
2. Do one of the following tasks:
  - Hover over **Administration**, **Group Settings**, and then click **Groups**.  
The **Manage Groups** page appears.
  - Click the name of your user account in the upper-right corner, and then click **My Custom Groups**.  
The **Manage Groups** page shows a tree view of group structure and a tabbed view of group properties.
3. Select the group to which you want to add managed items. Items that have already been added to the group appear in the **Group Details** pane.
4. In the **Group Details** pane, expand the **Rules** section, and then click **New Rule**.  
The **New Rule** dialog opens.
5. Complete the following fields, and then click **Save**:
  - **Add items of type**  
Specifies the type of managed item that you would like to add to the group.  
**Options:** Aggregated Component, Aggregated Interface, Application, BC/MC Interface, CVI, Device, Device Buffer, Device Component, Interface, Interface Buffer, Links, Network, Network Path, SDN Links, Service Chain, SLA Class, SLA Path, SSID, Tunnel, Virtual Interface, VNF, Voice Interface, VoIP Location, Wi-Fi Devices

#### NOTE

The list might include item types that have not been synchronized to your system. Use the field to create rules that are applied later when you discover these items. To view the existing item types on your system, see the **Inventory** menu.

**Default:** Aggregated Interface

- **Rule Name**

The rule name changes based on the selection for **Add items of type**. You can change the name for the rule.

6. In the **Conditions** section, a default 'is a member of' condition exists in the list of conditions. This condition filters added items based on your top-level permission group. This condition includes all items of the type selected for this rule, and includes all items within the nested sub-groups of the selected group. You can adjust the filter to add items based on another group.

**NOTE**

**Best practice:** Filter based on the group with the fewest items.

To adjust the filter to add items based on another group, click the **Select a group** funnel icon, select the group on which to base the filter, and then click **Save**.

7. To add a condition, in the **Conditions** section, complete the following steps:
- Click **Condition**.  
A row of drop-down lists and fields appears.
  - Select a method for identifying managed items. For example, select **Device Type**.  
The remaining lists are updated to match the type of selected item.
  - Select an operator from the second list. For example, select the 'is equal to' operator. 'is a member of' group rule conditions filter and add items to your top-level permission group. This group rule condition includes all items in all nested sub-groups of the group defined in this group rule condition. To return matches on a portion of the string (for example, 'server,WirelessController', 'device,WirelessController', and 'router,WirelessController'), select the 'is like' operator for the condition.  
**Options:** is equal to, is not equal to, is greater than, is greater than or equal to, is like, is not like, is less than, is less than or equal to, is between, is not between, starts with, does not start with, ends with, does not end with, matches regex, does not match regex, is a member of, is not a member of, is in subnet, is not in subnet  
**Default:** is equal to  
**IMPORTANT**  
Use Classless Inter-Domain Routing (CIDR) notation for the IP addresses that you specify for the 'is in subnet' and 'is not in subnet' operators. Use dot-decimal notation for the IP addresses that you specify for the 'is between' and 'is not between' operators.
  - (Optional) Enter a text string to match in the remaining condition field. For multi-character matches, use the wildcard character (\*). For example, to add all routers and servers in the Southwest region, enter a string with the appropriate naming convention, such as 'sw\* '. To have the 'is equal to' operator return only *exact* string matches (for example, 'server,WirelessController'), enter a string field such as 'Device Context Type '. Do not use spaces.  
**IMPORTANT**  
Some methods have a limited list of acceptable condition values. For example, when the method for identifying managed items is device alarm state, the text value must match one of the actual device alarm states.
  - (Optional) To add 'OR' matches, click + at the end of the condition.
  - (Optional) To add 'AND' matches, click **Add Condition**.
8. Confirm that the new rule includes the correct items by clicking **Preview Results**.  
The results are shown in the **Preview** section. You can scroll to each item type to view the specific items that were added.
9. Click one of the following options:
- **Save**  
Saves the rules without running them. The group is populated during the next global synchronization, which occurs approximately every 5 minutes.
  - **Run Rules**  
Populates the group immediately.

The group rule is created and applied to the group.

### **Edit a Group Rule**

You can modify group rules by deleting filters or by adding subrules.

**Follow these steps:**

1. From the **Manage Groups** page, expand the **All Groups** node in the **Groups** tree, and then select the group that contains the rule that you want to modify.
2. In the **Group Details** pane, in the **Rules** section, click the group rule that you want to edit, and then click **Edit Rule**. The **Edit Rule** dialog opens.
3. Edit existing filters, add filters or subrules, or remove filters or subrules as needed, and then click **Save**.
4. Confirm that the modified rule adds the appropriate items to the group by clicking **Preview Results**.
5. Click one of the following options:
  - **Save**  
Saves the rules without running them. The group is populated during the next global synchronization. Global synchronization occurs approximately every 5 minutes.
  - **Run Rules**  
Populates the group immediately.

The group rule is modified.

**Delete a Group Rule**

You can delete the rules that you have created. Deleting a group rule deletes any items in the group to which the rule is applied, but does not delete the items from the inventory. The deleted items are no longer available on the **Items** tab in the **Group Details** pane for the affected group.

**Follow these steps:**

1. From the **Manage Groups** page, expand the **All Groups** node in the **Groups** tree, and then select the group that contains the rule that you want to delete.
2. In the **Group Details** pane, in the **Rules** section, click the group rule that you want to delete, and then click **Remove Rule**. The **Delete Rule** dialog opens.
3. Click **Yes**.

The group rule is deleted and is no longer applied to the group.

**Determine Why a Group Rule Added Unexpected Items to a Group****Symptom:**

You ran a group rule (clicked **Run Rules**), and the group rule added items that you did not expect to be added to the group.

**Solution:**

Prior to running the group rule, click **Preview Results**, and then confirm that the list of items are those that you expect to be listed and added to the group. If the list includes items that you do not expect, then it might be the result of the filtering of an 'is a member of' group rule condition. This condition includes all items of the type selected for this rule, and includes all items within the nested sub-groups of the selected group. Starting from the selected group (the top group), locate the unexpected item. Repeat this step for the sub-groups. After you locate the unexpected item, consider adding another condition that filters out this item, and then re-run the rule.

If you already ran the group rule, add a condition to the group rule that filters out the unexpected item, and then re-run the rule. The group rule automatically removes the items from the group.

**View Groups Change Log**

Changes that are made to groups are logged in the Groups Change Log. You can use this log to troubleshoot unwanted changes. The log shows the username of the user who made the change, the time, and the nature of the change. The



log provides a high-level summary of changes without specific details, such as the old name of a group when the name is changed. To view the Groups Change Log, a user needs the View Groups Change Log role right. By default, the Administrator and Tenant Administrator have this role right.

To view the Groups Change Log, click **Administration, Groups Change Log**. The Owner column specifies the username of the group creator. The Tenant column specifies the tenant under which the group was created. Customize the groups change log as follows:

- To change the Group context, click **[change]** next to 'All Groups'.
- To change the time selection, click **[change]** next to 'Last Hour'.

## Organize Group Items Geographically

After registering a data source to monitor infrastructure usage, status, and performance enterprise-wide, organize infrastructure monitoring and reporting using groups in NetOps Portal. For example, place routers and switches in groups according to their data centers. Group managed items according to their country, region, state, or city.

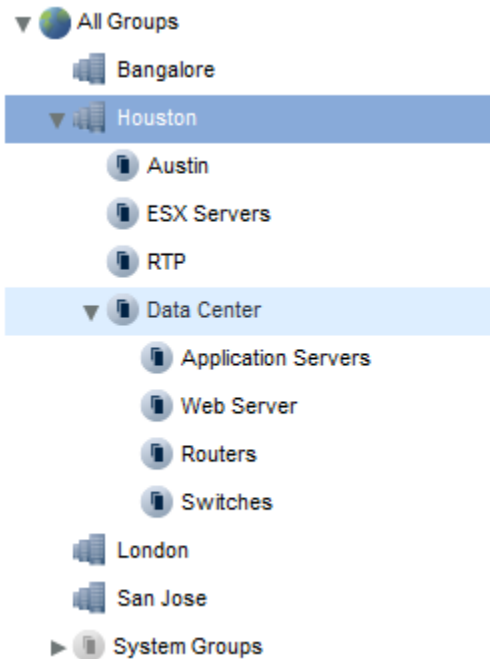
Consider a global enterprise, Mod-Lex Corporation, with the following major branch offices:

- Bangalore
- Houston, TX
- London
- San Jose, CA

Mod-Lex Corporation has a data center at each of these offices. The data centers support branch office operations and smaller sales offices nearby. Mod-Lex Corporation also has small branch offices in Austin, TX and Research Triangle Park, NC. The Austin office does not have a dedicated data center. Its three application servers are connected by a T3 line to the Houston data center. The Research Triangle Park (RTP) office is also connected to the Houston data center. The RTP application servers are virtual machines, running on an ESX server in Houston.

Most IT staff are located at the Mod-Lex Corporation corporate headquarters in London. The staff have NetOps Portal user accounts, most with user-level access. The NetOps Portal server is running in the data center in London. Dedicated IT staff are assigned to each of the major data centers. Staff charged with monitoring Houston are also responsible for Austin and RTP.

The following image shows an example of a grouping strategy that takes geography and logical network structure into account:



To monitor the RTP site effectively, place a copy of the 'Houston\ESX Servers' custom group into the 'RTP' custom group. Metrics from the critical application server are then reflected at the 'RTP'-custom group level. Ensure that a user is monitoring that custom group.

The 'London', 'San Jose', 'Houston', and 'Bangalore' site groups can contain subgroups. Site groups are useful in a geographical-grouping strategy. They contain items and subgroups of items that are grouped by location. Adding site groups to other custom groups in your tree structure allows you to build geographically and logically organized reports. The 'Houston' subgroups break out different types of infrastructure items to be monitored separately. Grant individual users access to data based on entire sites or on subgroups.

For more information about site groups, see [Groups](#).

For example, the Network Managers for Houston, Austin, and RTP can access the entire 'Houston' site group, or they can access the 'Data Center' subgroup. The Austin Network Manager also requires access to the 'Austin' custom group. Add Austin-specific items to the 'Austin' custom group. In this example, the RTP office does not have a dedicated Network Manager.

In this article:

- [Add the Custom Groups](#)
- [Add Managed Items to the Custom Groups](#)
- [Manage Items Using a Group Rule](#)

### **Add the Custom Groups**

Before you add custom groups, consider the types of access permissions that operators require to perform their monitoring duties.

Create the following site groups:

- London
- San Jose
- Houston
- Bangalore

Add the following custom groups:

- Austin
- RTP

To create the grouping structure, copy the 'Austin' and 'RTP' custom groups into the 'Houston' site group as subgroups. Problems that affect users in Austin or Research Triangle Park are visible in NetOps Portal dashboard views for the 'Houston' site group. You can drill down into event or performance information for those specific subgroups. You can quickly pinpoint the source of the issue.

Populate groups by adding subgroups that contain managed items. Copying a group creates a group reference. You cannot modify group references, but you can remove them. Groups that have been copied display an extra tab in the right pane. Any changes that you make to the original group are reflected in all of its references of the group. Removing a group also deletes all of its references.

For more information:

- About how to add custom groups and create group references, see [Manage Groups](#).
- About subgroups, see [Manage Subgroups](#).

### **Add Managed Items to the Custom Groups**

Add managed items to the subgroups in the 'Austin' and 'RTP' custom groups. For the 'Austin' subgroups, add the three application servers. For the RTP custom groups, add the virtual machines that run on the ESX server in Houston.

For more information about how to add managed items to custom groups, see [Manage Groups](#).

### **Manage Items Using a Group Rule**

Keep the 'Austin' and 'RTP' groups up-to-date when systems and networks change by creating group rules. Group rules add newly-discovered items that meet rule specifications to groups. Similarly, if the items do not meet rule requirements, or the items are no longer monitored, the group rules remove the items.

Mod-Lex Corporation consistently uses a geographical prefix to identify the location of devices and servers as they are purchased and brought into service. For example, 'hou-esx-rtr' is the main router that is dedicated to ESX servers in Houston. Use this naming convention to group items automatically using group rules.

Create a group rule to populate the 'Houston\Data Center\Application Servers' parent group. Add the parent group as a condition automatically by selecting it in the group hierarchy. Add a second condition that prevents routers and switches with the 'hou' prefix in their names from being placed in the 'Application Servers' parent group.

For more information about how to create group rules, see [Manage Group Rules](#).

## **Organize Devices and Component Items Using Groups and Group Rules**

Organize your devices and component items into logical groups that the data aggregator can monitor by creating the groups and group rules.

Discover and monitor your end-user switches separately using groups. You can monitor those switches separately because the uplink interfaces on those end-user switches have a significant impact on your network. Separate monitoring lets you report on the inventory that is related to that device collection for better troubleshooting. This scenario shows how to configure device groups in NetOps Portal and the data aggregator using REST web services. This configuration also provides the framework to write proprietary automation scripts that integrate with your third-party or proprietary management system.

**NOTE**

This workflow does not explain the configuration of monitoring profiles in the data aggregator, which you do *after* you have configured the group. Monitoring profiles are associated with device collections, and are used to keep discovered inventory up-to-date.

Use the following process to organize devices and component items using custom groups and group rules :

1. [Create a Custom Group and Group Rules](#)
2. [Verify that the Group Was Added to the Group Tree](#)
3. [Verify that the Device Collection Appears in the Data Aggregator](#)

In this scenario, you issue REST web service calls manually using a REST client editor or an HTTP tool that sends requests and gets responses. You can automate by writing your own application or script that leverages the web services that are described in this article.

**Create a Custom Group and Group Rules**

This scenario involves creating a custom group for only end-user switches. Out-of-the-box device collections do not exist for specific switches in the data aggregator. Instead, the All Switches out-of-the-box device collection is for *all* switches. Since your uplink interfaces have significant performance impact in your network, you want to monitor these specific switches for uplink interfaces separately for better troubleshooting.

You first create a custom monitored group in NetOps Portal. Upon synchronization, the data aggregator creates a corresponding device collection for use in device monitoring.

Add rules to your custom group so that the corresponding the device collection is kept up-to-date with discovered end-user switches. Newly-discovered devices that meet rule specifications are added to device collections. Similarly, if they do not meet rule requirements or they are no longer monitored, devices are not included or removed. For this scenario, we assume that your end-user switches are easily identifiable with the word "EndUser" at the end of their descriptions.

**Follow these steps:**

1. Open a REST client editor that has a connection to the NetOps Portal web server.
2. Set the **Content-type** to application/xml.
3. Provide a valid **Username** and **Password** in the request header for a user account that has Administrator access to NetOps Portal.
4. Enter the **Body** text in a REST client and modify the attributes as needed:
  - **Group Name**  
Specifies a name for the group. For this scenario, type the name **My End-User Switches**.

**NOTE**

Do not use the following special characters in group names: /&\,%,.

- **Group Tree Path**  
Specifies the path where the monitored group is created. Type the path **All Groups/Collections**, which is required for the device collection to show up in the data aggregator Administration pages in NetOps Portal.
- **Rules**  
Specifies the rules that automatically group the devices. In this case, create a rule that groups the end-user switches that have the value "EndUser" at the end of their descriptions.
- **type**  
Indicates the type of group.
- **inherit**  
Indicates whether the group includes child items of group members. In this case, set the "inherit" attribute to **true** so that the device interfaces become group members when the devices are added to the group.

**Example:**

```
<GroupTree path="/All Groups/Collections">
  <Group inherit="true" name="My End-User Switches" desc="" type="user group">
```

```

    <Rules allowDeletes="true" saveRules="true">
      <Rule add="Device" name="Add Devices">
        <Match>
          <Compare readOnly="true" using="MEMBER_OF">
            <Property name="ItemID" type="Device"/>
            <Value reference="/All Groups">1</Value>
          </Compare>
          <Compare readOnly="false" using="LIKE">
            <Property name="DisplayName" type="Device"/>
            <ValueList>
              <Value>EndUser</Value>
            </ValueList>
          </Compare>
        </Match>
      </Rule>
    </Rules>
  </Group>
</GroupTree>

```

5. POST the following URL:

POST `http://<PC_host>:<port>/pc/center/webservice/groups/false/true`

- **PC\_host**  
Specifies the NetOps Portal host name.
- **port**  
Specifies the NetOps Portal required port number.  
**Default:** 8181

For more information about the NetOps Portal server ports that should be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).

The `false` and `true` values refer to the following parameters:

- **uselds**  
Indicates that the `groupId` parameter is used to identify the group. In this example, the XML does not contain a group ID, so the value is 'false'.
- **allowDeletes**  
Enables deletion of the group that you are creating.

The My End-User Switches device collection is created. After synchronization, the data aggregator adds all discovered end-user switches automatically to the My End-User Switches device collection.

### **Verify that the Group Was Added to the Group Tree**

Verify that the group was added to the group tree.

#### **Follow these steps:**

1. Enter the following URL:  
`http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2FCollections`
2. Look for the My End-User Switches group.  
This group change in the web service corresponds to the group change in NetOps Portal.

## Verify that the Device Collection Appears in the Data Aggregator

After automatic synchronization occurs between NetOps Portal and the data aggregator, verify that the new device collection appears in the data aggregator.

### Follow these steps:

1. Enter the following URL:

```
http://<DA_host>:<port>/rest/groups
```

- **DA\_host**

Specifies the data aggregator host name.

- **port**

Specifies the data aggregator required port number.

**Default:** 8581

For more information about the data aggregator server ports that should be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).

2. Look for the My End-user Switches device collection (group).

This change in the web service also appears on the data aggregator user interface administration pages.

The custom group and group rules are configured. You can proceed to create your own automation script with your third-party or proprietary software to populate information in the data aggregator.

## Manage Monitoring Profiles

Monitoring profiles specify which metrics NetOps Portal gathers from which devices and components.

Monitoring profiles define the SNMP monitoring characteristics for devices in associated device collections. They are global across all tenants. They determine the metric families that a device in an associated device collection collects, and how often the data aggregator polls these devices. The metric families in the monitoring profiles determine which components it adds to NetOps Portal for devices in associated device collections, and which metrics it collects for those components.

The relationship of monitoring profiles to device collections governs monitoring. Monitoring is triggered when the following occurs:

- You associate a monitoring profile with a device collection.
- You add a device to a device collection that has an associated monitoring profile.
- You add a metric family to a monitoring profile with an associated device collection.
- You add a vendor certification for an existing metric family that the data aggregator polls in a monitoring profile.
- You change the vendor certification priority for a metric family.
- You change (edit) a device.

For more information about device collections, see [Device Collections](#).

In this article:

- [Out-of-the-box Monitoring Profiles](#)
- [Configure a Monitoring Profile](#)
- [Manage Monitoring Profile Filters](#)
- [Edit or Delete a Monitoring Profile](#)

### Out-of-the-box Monitoring Profiles

NetOps Portal starts monitoring devices in the environment using the out-of-the-box monitoring profiles. These monitoring profiles are listed on the **Monitoring Profiles** page. To view this page, as a user in the default tenant with the Administrator role, hover over **Administration**, **Monitored Items Management**, and then click **Monitoring Profiles**.

The out-of-the-box monitoring profiles provide basic monitoring for common items in NetOps Portal, including reachability, availability, and CPU and memory metrics. To make it easier to customize and manage device polling, you can use these as templates to configure custom monitoring profiles. The out-of-the-box monitoring profiles are locked, as indicated by a lock icon, and are automatically associated with out-of-the-box device collections. For example, the out-of-the-box Router monitoring profile is associated with the out-of-the-box All Routers device collection.

#### NOTE

- The following out-of-the-box monitoring profiles can have a significant performance impact:
  - QoS
  - MPLS
  - Network Interface
 

By default, this monitoring profile is not associated with any devices. The **Interface** metric family is assigned to this monitoring profile, and a high volume of data is associated with this metric family. By default, this monitoring profile does not discover interfaces.
  - Response Path
- A built-in NetOps Portal mechanism uses the **DA Health (Fast)**, **DA Health (Normal)**, and **DA Health (Slow)** monitoring profiles to self-monitor NetOps Portal components. Do not modify these monitoring profiles.

### Configure a Monitoring Profile

Use the following process to configure a custom monitoring profile:

1. [Verify the Prerequisites](#)
2. [Review the Monitoring Profile Configuration Best Practices](#)
3. [Create a Monitoring Profile](#)
4. [Add an Event Rule to the Monitoring Profile](#)
5. [Add a Monitoring Profile Filter to a Metric Family Assigned to the Monitoring Profile](#)
6. [Apply a Monitoring Profile Polling Behavior to Devices](#)

The following video shows how to configure a monitoring profile and assign device collections to the monitoring profile:

### Verify the Prerequisites

Before you configure a monitoring profile, [determine the monitoring requirements](#).

### Review the Monitoring Profile Configuration Best Practices

Consider the following best practices for creating monitoring profiles:

- For monitoring flexibility, configure multiple monitoring profiles. Do not add all metric families that you monitor in your network to a single monitoring profile.
- Some metric families, such as CPU and Memory, apply broadly to all devices. Associate a monitoring profile with these metric families to all-encompassing device collections, such as the All Routers or All Managed Devices device collections.
- Monitoring profiles control the rate at which the data aggregator polls. To poll the same metric family on different devices at different rates, create monitoring profiles with different poll rates.

#### IMPORTANT

Apply fast polling *only* to device collections that include devices that require it. Fast polling increases network traffic and degrades data aggregator performance.

For more information, see [Poll Critical Interfaces Faster than Non-critical Interfaces](#).

- Monitoring profiles control filtering. If you require different filtering criteria for different sets of devices, create a monitoring profile for each set of requirements.
- Apply only those metric families that are applicable to the items in the device collections that you have assigned to the monitoring profile. This minimizes the data aggregator processing footprint during change detection. For example, do not associate a collection of routers to a monitoring profile that contains the VMware metric families.

## Create a Monitoring Profile

To specify which metrics NetOps Portal gathers from which devices and components, create monitoring profiles. For example, create a **Systems** monitoring profile, and then configure the profile to poll resource utilization counters only from SNMP.

Monitoring profiles are available in all tenant work spaces. For example, create a **Service Router Monitoring** monitoring profile and use the profile for all tenants. The device collections that are assigned to monitoring profiles are tenant-scoped.

In an MSP environment, all changes to monitoring profiles apply to all tenants.

### Follow these steps:

1. As a user in the default tenant with the Administrator role, from the **Monitoring Profiles** page, do one of the following:
  - To create a monitoring profile, click **New** at the bottom of the page.
  - To use an existing monitoring profile as a template, select the monitoring profile that you want to use as a template, and then click **Copy**.

#### NOTE

This option does not copy event rules.

The **Create / Edit Monitoring Profile** dialog opens.

2. Complete the following fields, and then click **Save**:

- **Name**

Enter a unique name for the monitoring profile.

**Required:** Yes

- **Description**

Enter a unique description.

- **SNMP Poll Rate**

Defines the polling interval. This rate determines how often NetOps Portal collects data from the devices and collections associated with this monitoring profile for the metrics defined in the monitoring profile.

**Required:** Yes

**Options:** 1 minute, 5 minutes, 10 minutes, 15 minutes, 30 minutes, 60 minutes

**Default:** 5 minutes

#### NOTE

- Polling occurs at the fastest rate for all the monitoring profiles to which the device collection is assigned.
- Changes to the poll rate can take up to two cycles to apply.
- When you use the 60 minutes poll rate for an existing device, a `No Data To Display` message appears in views with a time range of **Last Hour**. If you change the dashboard setting to a prior hour, you can see earlier data. However, the view does not display the latest data until the new poll cycle completes.
- One-minute polling on large device collections or devices with many components can affect NetOps Portal performance.
- To poll critical interfaces at a faster rate, use a separate monitoring profile with a filter.  
For more information, see [Poll Critical Interfaces Faster than Non-critical Interfaces](#).
- **Enable Change Detection**  
Defines whether the monitoring profile detects when the components on a device change.



**Options:**

- **Selected:** The monitoring profile detects changes to the components on a device. NetOps Portal identifies changes to components associated to metric families that it monitors with that profile by cycling through all the devices associated with the monitoring profile within the detection rate, and then updates the components.
- **Cleared:** The monitoring profile *does not* detect changes to the components on a device.

**Default:** Selected

When selected, configure the behavior:

- **Detection Rate**

Specify the rate at which NetOps Portal identifies changes to component items on monitored devices. If a device is associated with multiple monitoring profiles, change detection occurs as the fastest rate that you specify. To determine the configuration changes, NetOps Portal uses the reconciliation algorithm that is defined in the metric family.

**Values:** minutes or hours**Default:** 24 Hours**IMPORTANT**

Consider how frequently the monitored components on the devices are likely to change. Avoid setting change detection rates that are more frequent than necessary. Some metrics, such as availability, never or rarely change. Because change detection uses some system resources, do not enable change detection for such metrics.

- **Automatically Update Metric Families**

Defines whether to have your monitoring profile automatically monitor *newly-detected* components or to manually control when changes to monitoring occur. When this option is selected, the changes to monitoring occur automatically according to the **Detection Rate**. If you choose to manually control when changes to monitoring occur (clear this checkbox), the changes to the device apply when you update the metric family. When the changes are applied to devices, NetOps Portal marks the components that it no longer detects on the devices as "Not Present".

**Options:**

- **Selected:** Have the monitoring profile automatically monitor newly-detected components. The data aggregator starts monitoring new components or stops monitoring retired components. Changes to monitoring occur automatically according to the **Detection Rate**.
- **Cleared:** Manually control when changes to monitoring occur. Clearing this option does not automatically apply the changes to monitoring. The **Events Display** dashboard displays configuration events for the detected changes.

To apply the changes to the device, update the metric family by completing the following steps:

- Watch for configuration events by checking the **Events Display** dashboard.
- Navigate to the data aggregator **Administration** menu, **Monitored Devices**, and then click **Polled Metric Families** tab.
- To help ensure that the data aggregator picks up the latest device reconfiguration, select the metric family that you want to update, and then click **Update Metric Family**. The **Update Metric Family** dialog opens.
- Click **Yes**.

The changes to the device apply when you update the metric family. When the changes are applied to devices, components that are no longer detected on the devices are marked "Not Present".

**Default:** Selected**NOTE**

If you apply an interface filter, the data aggregator monitors only the interfaces that pass the filter conditions after reconfiguration.

- Select the metric families from the **Available Metric Families** tree that you want to assign to the monitoring profile, and then click the right arrow to add it to the **Selected Metric Families** list. The **Selected Metric Families** list shows which metric families NetOps Portal will collect for devices associated with the monitoring profile.

**TIP**

By default, NetOps Portal collects all metrics in each selected metric family. Some metric families have many metrics which might not be relevant to monitor in your environment. You can control what metrics to monitor for each metric family within a monitoring profile, and reduce the data volume required in the database, by configuring metric filters. For example, if you assign the **Interface** metric family to the monitoring profile, the best practice is to then apply a metrics filter to this metric family to control what metrics this metric family monitors.

For more information, see [Configure Metric Filtering](#).

The monitoring profile is created. It is added to the **Monitoring Profiles** list.

**TIP**

**Next Step:** To start monitoring with this monitoring profile, [apply monitoring profile polling behavior to devices](#).

**Add an Event Rule to the Monitoring Profile**

Event rules set conditions to trigger Threshold Violation events. Each rule is based on a single metric family and contains criteria to raise or clear the violation. The event rules apply to each component on all devices in collections that are assigned to the monitoring profile.

You can add event rules to all monitoring profiles except the out-of-the-box monitoring profiles.

**TIP**

For best performance and granular control, configure event rules on threshold profiles instead of monitoring profiles. For more information, see [Configure Threshold Profiles](#).

**Follow these steps:**

1. As a user in the default tenant with the Administrator role, from the **Monitoring Profiles** page, select the monitoring profile to which you want to add an event rule, and then click the **Event Rules** tab.
2. Click **New** to create a rule, or click **Edit** to modify an existing rule.  
The **Create / Edit Event Rule** dialog opens.  
For more information about how to create event rules, see [Configure Threshold Profiles](#).  
NetOps Portal raises and clears events when the devices in the associated device collections satisfy the event rule conditions.

The event rule and event rule conditions are added to the monitoring profile.

**Monitoring Profile Filters**

Monitoring profile filters, also referred to as *component filters*, specify the criteria that govern which components NetOps Portal discovers for metric families that are assigned to the monitoring profile. The data aggregator polls only the component items that match the filter criteria for the associated metric family. Filtering limits SNMP traffic and ensures that the system monitors only relevant components. The filters of each monitoring profile are assessed independently.

**NOTE**

- NetOps Portal monitors device components if any of the monitored components pass the filter criteria of any associated monitoring profile.
- Monitoring profile filters and vendor certification filters reduce SNMP traffic to the same level. Monitoring profile filters apply *after* discovery. NetOps Portal creates the component items but it does not send SNMP requests to monitor the items. NetOps Portal reassesses monitored components according to the change detection interval. Vendor certification filters prevent the creation of component items that do not match the filter criteria. These filters might not recognize changes at the component level.  
Use vendor certification filters only when the monitoring profile filter cannot filter component item.

See the following examples:

- [Example: Multi-Rate Polling.](#)
- [Example: How and When NetOps Portal Combines Filters.](#)

### Example: Multi-Rate Polling

You want to monitor all interfaces with 5-minute polling, and one type of interface with 1-minute polling. One monitoring profile has a poll rate of 5 minutes. This monitoring profile is assigned to a device collection. A second monitoring profile has a poll rate of 1 minute. The interface metric family on that profile has a filter that specifies the type of interface to poll at a 1-minute interval. This monitoring profile is also associated with the same device collection. NetOps Portal polls the specified interface type at a 1-minute rate and all other interfaces at a 5-minute rate.

### Example: How and When NetOps Portal Combines Filters

You want to apply multiple filter profiles to multiple devices. You want to exclude an attribute from one of the filters. However, poll filters are not *exclusion* filters. They are *inclusion* filters. NetOps Portal polls as much as it can, and as fast as it can. So, if two filter criteria conflict, NetOps Portal *includes* interfaces in the polling rather than *exclude* them. The following example explains how filtering works and how you can achieve your goal.

For example, you have two filters with the following configuration:

- **Filter 1:**  
(Description contains "Virtual-Access" or Description contains "Tunnel")
- **Filter 2:**  
Name does not contain "Internal"

This situation means that Filter 1 and Filter 2 are OR'ed. NetOps Portal polls only the components that meet the previously mentioned criteria.

However, if you want to *exclude* any interfaces that contain "Internal" in the Name attribute, add that filter criteria to each monitoring profile as an AND. Filter 1 should appear as follows:

- **Filter 1:**  
(Description contains "Virtual-Access" or Description contains "Tunnel") AND (Name does not contain "Internal")

When you have the same devices that are associated to two monitoring profiles, the two profiles must have filter criteria. In this scenario, effectively, you combine the two filters by stating exclusion criteria in both. Otherwise, the monitoring profile without the filter causes all components on the device to be polled. Therefore, Filter 2 should appear as follows:

- **Filter 2:**  
(Description contains "Virtual-Access" or Description contains "Tunnel") AND (Name does not contain "Internal")

### Manage Monitoring Profile Filters

You can manage monitoring profile filters by:

- Viewing them
- [Adding a Monitoring Profile Filter to a Metric Family Assigned to the Monitoring Profile](#)
- Editing them
- [Deleting \(clearing\) a Monitoring Profile Filter from a Metric Family Assigned to the Monitoring Profile](#)

For more information about monitoring profile filters, see [Monitoring Profile Filters](#).

### Add a Monitoring Profile Filter to a Metric Family Assigned to the Monitoring Profile

The following video shows how to add monitoring profile filters to a metric family:

#### Follow these steps:

1. As a user in the default tenant with the Administrator role, from the **Monitoring Profiles** page, select the monitoring profile to which you want to add a monitoring profile filter.

The **Metric Families** tab is selected by default. This tab shows a list of the metric families that are assigned to the monitoring profile. The list shows the **Name**, **Component Filter**, and **Collected Metrics** columns. If a metric family does not include a monitoring profile filter, an asterisk (\*) appears in the **Component Filter** column for that metric family. In this case, NetOps Portal will discover all components on all devices in device collections associated to that monitoring profile for that metric family.

2. Click the metric family row to which you want to add or edit, and then click **Edit Component Filter**. The **Add / Edit Filter Expression** dialog opens.
3. Complete the following fields:
  - **Attribute**  
Defines the attributes on which to filter.
  - **Operation**  
Defines the operation for the filter condition.
  - **Value**  
Defines the value for the filter condition.
4. Click **Add Condition**.
5. Click **Save**.

The filter applies to the selected metric family on the next poll cycle.

### **Delete a Monitoring Profile Filter from a Metric Family Assigned to the Monitoring Profile**

Clearing a monitoring profile filter from a metric family deletes them from the metric family.

#### **Follow these steps:**

1. As a user in the default tenant with the Administrator role, from the **Monitoring Profiles** page, select the monitoring profile for which you want to delete a monitoring profile filter.  
The **Metric Families** tab is selected by default. This tab shows a list of the metric families that are assigned to the monitoring profile. The list shows the **Name**, **Component Filter**, and **Collected Metrics** columns.
2. Click the metric family row to which you want to clear a monitoring filter, click the metric family row from which you want to clear the monitoring profile filter, and then click **Clear Filter**.

The filter is removed on the next poll cycle.

### **Apply a Monitoring Profile Polling Behavior to Devices**

You can apply a polling behavior of a monitoring profile to devices using the following methods:

- [Assign a device collection to a monitoring profile.](#)

#### **NOTE**

When you assign a device collection to multiple monitoring profiles, the data aggregator uses the fastest specified poll rate.

- [Assign a monitoring profile to a device collection.](#)

#### **NOTE**

Metric families, such as QoS, MPLS, and various Response Path Test metric families, can contribute to significant SNMP requests.

For more information about the implications and restrictions of SNMP requests to your network devices, see the vendor documentation or contact the vendor.

Administrators and Tenant Administrators can modify the associations between monitoring profiles and device collections. Administrators see only the associated device collections for the current tenant.

### **Assign a Device Collection to a Monitoring Profile**

Assign a device collection to a monitoring profile to have NetOps Portal schedule metric family discovery for all devices associated with that profile. For large device collections, this discovery process can take some time.

#### **Follow these steps:**

1. From the **Monitoring Profiles** page, select the monitoring profile that you want to assign a device collection, and then click the **Collections** tab.
2. Click **Manage** at the bottom of the page.  
The **Assign Collections to Monitoring Profiles** dialog opens. The **Available Collections** list shows all the available collections.
3. (If you used an out-of-the-box monitoring profile as a template when you created this monitoring profile) Remove the existing device collections.
4. Select the device collections from the **Available Collections** list that you want to assign to the monitoring profile, and then click the right arrow to add it to the **Assigned Collections** list, and then click **Save**.

#### **IMPORTANT**

To prevent NetOps Portal from making extra SNMP requests to pingable-only devices, which can result in sporadic metric family support, do not assign the out-of-the-box **All Devices** device collection to monitoring profiles.

The device collection is assigned to the monitoring profile. The polling behavior that is defined in the monitoring profile is applied to the assigned device collection. The monitoring profile is applied to the assigned device collections during the next poll cycle.

### **Assign a Monitoring Profile to a Device Collection**

Administrators can view the device collections for the tenant they are administering. A tenant administrator can view its own (tenant) list of device collections.

#### **Follow these steps:**

1. As in the default tenant a user with the Administrator role, under the **Monitoring Configuration** section, click **Collections**.  
The **Collections** page appears. A list of device collections displays on this page.
2. Select the device collection to which you want to assign monitoring profiles, and then, from the **Monitoring Profiles** tab, click **Manage**.  
The **Assign Monitoring Profiles to Collections** dialog opens.
3. Move the monitoring profiles that you want to assign to the collection to the **Assigned Monitoring Profiles** list, and then click **Save**.

The monitoring profile is assigned to the device collection. The monitoring behavior that is defined in the assigned monitoring profiles is applied to the device collection.

### **Edit or Delete a Monitoring Profile**

You can edit and delete only the monitoring profiles that you configure (custom monitoring profiles). To edit or delete a monitoring profile, as a user in the default tenant with the Administrator role, from the **Monitoring Profiles** page, select the monitoring profile that you want to edit or delete, and then click **Edit** or **Delete**.

## **Manage Custom Attributes**

You can add custom attributes to devices or interfaces.

You can manage custom attributes by creating them, by adding them to devices or interfaces, by setting values for them, and by increasing the number of attributes. Populate these items using the `customattributedefinition` and `customattributes` Data Aggregator REST web services. Custom devices and interfaces are then available in reports,

dashboards, and group rules. Custom devices, device components, and interfaces are also available for OpenAPI queries with some delay due to extract, transform, load (ETL) job scheduling.

The custom attributes for devices and interfaces appear in the **Information** section of the **Details** tabs in a context page for an device or interface. For more information, see [Context Pages](#).

### Create or Edit Custom Attributes

By default, you can add up to 5 custom attributes of type `String` and `Integer` to devices or interfaces. To avoid the data aggregator from returning errors, do not add more than the maximum number of custom attributes to devices or interfaces.

#### Follow these steps:

1. Set up a REST client with a connection to the data aggregator server.
2. Specify the following URL:  
`http://DA_host:8581/rest/customattributedefinition`
3. Issue a `POST` request with the XML for setting your custom attributes.  
The following example defines a **CustomerID** attribute for all devices.

```
<CustomAttributeDefinition version='1.0.0'>
  <Label>Customer ID #</Label>
  <Description>Customer account number</Description>
  <Hidden>true</Hidden>
  <Storage>
    <AttributeName>CustomerID</AttributeName>
    <Type>String</Type>
    <ItemType>Device</ItemType>
  </Storage>
</CustomAttributeDefinition>
```

- **<Label>**  
Specify a label that can be displayed in the NetOps Portal user interface for the custom attribute.  
**Limit:** 64 characters
- **<Hidden>**  
Specify whether to hide the property by default from the NetOps Portal user interface. If this tag is set to "true," you can manually display the custom attribute in the user interface. This tag impacts whether custom attributes appear by default in the NetOps Portal user interface only. They are available by default in other areas, such as the OpenAPI QueryBuilder.  
**Default:** `true`  
**NOTE**  
If this tag is set to `true`, custom attributes are hidden from context pages and cannot be manually shown in the user interface. If you want custom attributes to appear on context pages, you must set this tag to "false."
- **<Description>**  
Specify a description for the attribute.
- **<AttributeName>**  
Specify an internal reference name for the custom attribute. Use this name to set a value for the custom attribute in the following procedure.  
**Limit:** 26 characters
- **<Type>**  
Specify whether the custom attribute is a string or an integer.

**Value:** String or Integer

– **<Item Type>**

Specify whether the custom attribute is for a device, component, or interface (Port ).

**Valid entries:** Device , DeviceComponent , and Port

**NOTE**

Custom device component attributes are available only for OpenAPI queries with some delay due to ETL job scheduling. They are *not* available in reports, dashboards, and group rules.

The POST request returns an ID for the definition of the custom attribute.

## Set Values for a Custom Attribute

### Follow these steps:

1. Specify one of the following URLs:
  - To set values for a custom attribute that is for devices, specify `devices` in your URL, along with the device or interface ID:  
`http://da_hostname:8581/rest/devices/customattributes/itemID`
  - To set values for a custom attribute that is for device components, specify `devices/components` in your URL, along with the device or interface ID:  
`http://da_hostname:8581/rest/devices/components/customattributes/itemID`
  - To set values for a custom attribute that is for interfaces, specify `ports` in your URL, along with the device or interface ID:  
`http://da_hostname:8581/rest/ports/customattributes/itemID`
2. Issue a PUT request with the XML for setting your custom attributes.  
 String attributes support up to 255 characters. Integer attributes support up to 64-bit integers.

### Example:

The following example sets the `CustomerID` attribute from the previous example.

```
<DeviceCustomAttributes version='1.0.0'>
  <CustomerID>CustomAttributeValue</CustomerID>
</DeviceCustomAttributes>
```

The following example sets the custom attributes for a device component:

```
<DeviceComponentCustomAttributes version='1.0.0'>
  <AttributeName>CustomAttributeValue</AttributeName>
</DeviceComponentCustomAttributes>
```

The following example sets the custom attributes for an interface:

```
<PortCustomAttributes version='1.0.0'>
  <AttributeName>CustomAttributeValue</AttributeName>
</PortCustomAttributes>
```

## Change the Number of Allowed Custom Attributes

By default, you can add up to five custom attributes of type `String` and `Integer` to devices and interfaces. However, you can increase or decrease these values.

**IMPORTANT**

Consider the performance impact of increasing the maximum number of custom attributes that you can add to devices and interfaces. For example, NetOps Portal allows only a defined number of custom attributes to display. Avoid exceeding a maximum combined total of 30 custom attributes of type `String` and `Integer`.

**Follow these steps:**

1. Create the `<installation_directory>/etc/com.ca.im.item.custattr.CustAttrColumnCache.cfg` file.

- **installation\_directory**

The installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

2. Add the following lines to the file, and then save the file:

**Example:**

```
MaxStringColumnsPerItemType=7
```

```
MaxIntegerColumnsPerItemType=7
```

- **MaxStringColumnsPerItemType**

Defines the maximum number of `String` custom attributes that you can add to devices and interfaces.

**Example:** `MaxStringColumnsPerItemType=7`

**Default:** 5

- **MaxIntegerColumnsPerItemType**

Defines the maximum number of `Integer` custom attributes that you can add to devices and interfaces.

**Example:** `MaxIntegerColumnsPerItemType=7`

**Default:** 5

A message appears logged in the `karaf.log` file, for example:

```
INFO | AttrColumnCache) | 2016-07-29 09:46:34,464 | CustAttrColumnCache |
tem.custattr.CustAttrColumnCache 80 | ository.webservices.impl | | Max Integer custom
attributes per item type:7
INFO | AttrColumnCache) | 2016-07-29 09:46:34,465 | CustAttrColumnCache |
tem.custattr.CustAttrColumnCache 93 | ository.webservices.impl | | Max String custom
attributes per item type:7
```

## Configure Threshold Profiles

Monitor specific devices or components, such as interfaces and IP Service-Level Agreement (SLA) tests, and determine thresholds for given metrics or sets of metrics using threshold profiles.

*Threshold profiles* raise or clear *threshold violation events* when specified *event conditions* occur in associated groups. Threshold profiles require at least one *event rule*. You set event rules to a metric family. The event rule defines the conditions that raise or clear a threshold violation. When metrics violate the threshold, the data aggregator sends a *raised threshold violation event*, and as a result, DX NetOps Spectrum (Spectrum) raises an alarm. When metrics meet the threshold, the data aggregator sends a *clear threshold violation event*, and as a result, Spectrum clears the alarm.

Use the following process to configure a threshold profile:

1. [Create a Threshold Profile](#)
2. [Create an Event Rule and Add Conditions to the Threshold Profile](#)
3. [Assign a Group to the Threshold Profile](#)

The following video examines how to create a threshold profile in NetOps Portal to raise threshold violation or clear events when major changes happen to baseline performance conditions:



**Example:**

You want to monitor the utilization of an interface, and you want the data aggregator to send a raised threshold violation event when the utilization is above 75 percent. When the utilization drops below 75 percent, you want to clear the violation:

- **Create event rules for the threshold profile**

Define the logic in an event rule, which determines when it should raise a threshold violation event in response to a metric exceeding the threshold.

The example in this topic requires at least two event rules. One event rule determines when it raises a threshold violation event, while the other determines when it clears the violation (it raises a clear event).

- **Add conditions to the event rule**

Event conditions appear in the **Event Conditions** view for an event rule when the conditions in the event rule are met. Event rule conditions include the following:

- **Violation:**  
Defines a raised threshold violation event.
- **Clear:**  
Defines a clear threshold violation event.

**Create a Threshold Profile**

Consider the following best practices when creating threshold profiles:

- Increase the granularity and flexibility of thresholding by assigning groups with specific components to the threshold profiles instead of devices.
- Expand threshold monitoring slowly. Start with a small group of components and verify that the monitoring engine does not become degraded.  
For more information, see [Threshold Event Processing Self-Monitoring Metrics](#).
- Thresholds on components with one-minute polling have a high resource cost to the system.
- Threshold evaluations can slow down after a data aggregator restart while NetOps Portal processes cached poll data from the data collectors.

**NOTE**

(If you plan to have the data aggregator evaluate on multiple metric families together by way of linking event rules) In most cases, the data aggregator matches potential raised threshold violation events for each linked event rule at the device level. To prevent noisy events (a flapping condition), do *one* of the following:

- For the threshold profile, (if you select **CPU** or **Memory** as the **Metric Family**) choose the **Aggregate Components by Device** option for **Aggregation**.
- [Select the interfaces that you want to monitor by adding them to the group assigned to the threshold profile.](#)
- [Manage combined metric values from similar components, either on the same device or on different devices, by way of an \*aggregated component\* rather than managing individual metric values.](#)

**Follow these steps:**

1. Hover over **Administration, Monitored Items Management**, and then click **Threshold Profiles**.  
The **Threshold Profiles** page appears.
2. Do one of the following steps:
  - Select an existing threshold profile folder under which you want to create the threshold profile.
  - Create a threshold profile folder for the threshold profile by clicking **New Folder**, and then create the folder.
3. Right-click the threshold profile folder, and then click **New Profile**.  
The (23.3.5 and higher) **New/Edit Threshold Profile** (23.3.4 and lower) **Create / Edit Threshold Profile** page appears.  
The following images show examples of this page based on version:  
(23.3.5 and higher)

New/Edit Threshold Profile

Name\*

Test Profile

Description

Folder\*

Test

Status\*

Enabled

Owner\*

admin

Last Modified By

admin

Creation Time

Feb 5, 2024 8:14:30 PM UTC

Last Modified ...

Feb 12, 2024 8:03:09 PM UTC

When Event Rules Run

When Rules Run

☒ On Schedule

☐ Not During the Specified Times

Run Rules

☒ Only During the Specified Times

☐ Not During the Specified Times

On Days\*

☐ Sun

☒ Mon

☒ Tue

☒ Wed

☒ Thu

☒ Fri

☐ Sat

Start Time/End Time\*

00:30

24:00

Note: Times are in UTC

Event Rules

Quick Filter

NewCopyEditLinkDelete

Rule Name	Metric Family	Linked Rules	Aggregation	D...	Wi...	Se...
CPU	CPU	Memory	No Aggregation	30...	30...	M...
Interface	Interface	Memory	No Aggregation	30...	30...	M...
Memory	Memory		No Aggregation	30...	30...	M...

Event Conditions

Quick Filter

Event Type	Metric	Thr...	Condition Type
Violation	Utilization (%)	> 40...	Fixed Value
Clear	Utilization (%)	< 40...	Fixed Value

(23.3.4 and lower)

Create / Edit Threshold Profile

Name\*

Aggregated Port

Description

Folder\*

Roman

Status\*

Enabled

Owner\*

Last Modified By

Creation Time

May 12, 2023 6:57:48 PM GMT

Last Modified ...

Jul 6, 2023 2:03:24 PM GMT

Event Rules

Quick Filter

NewCopyEditDelete

Rule Name	Metric Family	Aggregation	Durati...	Window	Severity
Aggregated Port	Aggregated Interface	No Aggregation	300 sec	300 sec	Major

Event Conditions

Quick Filter

Event Type	Metric	Thresh...	Condition Type
Violation	Utilization (%)	>= 25.0	Fixed Value
Clear	Utilization (%)	< 25.0	Fixed Value

SaveCancel

4. Complete the following fields:
- Name**  
Defines the name for the threshold profile.  
**Required:** Yes
  - Description**  
Defines the description for the threshold profile.

**Required:** No

– **Folder**

Specifies the folder under which you want to create this threshold profile.

**Required:** Yes

– **Status**

Defines the status of the threshold profile.

**Values:** Enabled or Disabled

**Default:** Enabled

**Required:** Yes

– (Administrators Only) **Owner**

Specifies the owner for the threshold profile. Only the owner or a user with the Administer DA Threshold Profile role right can edit the threshold profile.

**Required:** Yes

– **Last Modified By**

Displays the username who last modified the threshold profile.

**Read-Only:** Yes

– **Creation Time**

Displays the date and time the threshold profile was created.

**Read-Only:** Yes

– **Last Modified Time**

Displays the date and time the threshold profile was modified last.

**Read-Only:** Yes

5. [Create an event rule and add conditions for the threshold profile.](#)

6. Save the threshold profile.

The threshold profile is created. The threshold profile is added to the **Threshold Profiles** pane.

### **Create an Event Rule and Add Conditions for the Threshold Profile**

Threshold profiles require at least one event rule. The data aggregator evaluates event rules on a single metric family. The metrics determine the conditions that raise or clear a threshold violation.

**Prerequisite:** You have created or are in the process of creating a threshold profile.

#### **Follow these steps:**

1. On the (23.3.5 and higher) **New/Edit Threshold Profile** page (23.3.4 and lower) **Create / Edit Threshold Profile** page, in the **Event Rules** pane, do one of the following steps:

- To add an event rule to the threshold profile, click **New**.
- To use an existing event rule as a template for a new event rule, select the event rule that you want to use as a template, and then click **Copy**.

The (23.3.5 and higher) **New/Edit Event Rule** (23.3.4 and lower) **Create / Edit Event Rule** dialog opens.

The following images show examples of this dialog based on version:

(23.3.5 and higher)



## New/Edit Event Rule

Name *	<input type="text"/>	Duration (sec) *	<input type="text" value="300"/>
Description	<input type="text"/>	Window (sec) *	<input type="text" value="300"/>
Metric Family *	<input type="text" value="3G WAN Service"/>	Severity *	<input type="text" value="Major"/>

— A violation occurs when all of these conditions are met: —

Metric: *	Operator: *	Value: *	Condition Type: *
<input type="text" value="3GRoaming Time in Home"/>	<input type="text" value="Above"/>	<input type="text"/>	<input type="text" value="Fixed Value"/>

[+ Add Condition](#)

— A violation is cleared when: —

Metric: *	Operator: *	Value: *	Condition Type: *
<input type="text" value="3GRoaming Time in Home Status (sec)"/>	<input type="text" value="Below"/>	<input type="text"/>	<input type="text" value="Fixed Value"/>

Save

Cancel

(23.3.4 and lower)

## Create / Edit Event Rule



Name *	<input type="text" value="Aggregated Port"/>	Duration (sec) *	<input type="text" value="300"/>
Description	<input type="text"/>	Window (sec) *	<input type="text" value="300"/>
Metric Family *	<input type="text" value="Aggregated Interface"/>	Severity *	<input type="text" value="Major"/>

— A violation occurs when all of these conditions are met: —

Metric: *	Operator: *	Value: *	Condition Type: *
<input type="text" value="Utilization (%)"/>	<input type="text" value="Above Or Equal"/>	<input type="text" value="25.0"/>	<input type="text" value="Fixed Value"/>
<input type="button" value="+ Add Condition"/>			

— A violation is cleared when: —

Metric: *	Operator: *	Value: *	Condition Type: *
<input type="text" value="Utilization (%)"/>	<input type="text" value="Below"/>	<input type="text" value="25.0"/>	<input type="text" value="Fixed Value"/>



2. Complete the following fields, and then click **Save**:

- **Name**  
Defines the name for the event rule.
- **Metric Family**  
Specifies the metric family on which the data aggregator evaluates.
- **Duration (sec)**  
Specifies the total amount of time within a specified range of time, or **Window**, that the threshold must violate what is defined in the event rule condition. The poll cycles that trigger the condition do not need to be consecutive. The duration is cumulative.  
**Default:** 300 (five minutes)
- **Window (sec)**  
Specifies the overall range of time that the data aggregator evaluates the condition and raises a threshold violation event when there are enough values, based on the specified **Duration**. The range of time is a rolling window.  
**Default:** 300 (five minutes)
- **Severity**  
Specifies the severity of the raised event.  
**Options:**
  - **Critical**
  - **Major**
  - **Minor****Default:** Major
- **Aggregation**  
(If you have selected a metric family that supports aggregation, such as a CPU or Memory metric family) Specifies whether the threshold applies to an aggregate value of all components for the device.

**Values:**

- **No Aggregation:** The threshold violation event that the data aggregator raises does not apply to an aggregate value of all components for the device.
- **Aggregate Components by Device:** The threshold violation event that the data aggregator raises applies to an aggregate value of all components for the device.

**Default:** No Aggregation

- In the **A Violation occurs when all of these conditions are met** section, define the condition for the event rule that triggers the data aggregator to send a *raised threshold violation event* by defining the following fields. To add a condition, click **Add Condition**. You can add up to five conditions to an event rule.

- **Metric**

The metric that determines the condition that triggers the event rule to raise a threshold violation event.

- **Operator**

Defines the operator related to the value for the condition at which the threshold violation event is raised.

**Example:** Above

**Options:**

- **Above:** The event rule raises a threshold violation event when the value of the metric exceeds the value.
- **Above Or Equal To:** The event rule raises a threshold violation event when the value of the metric is equal to or exceeds the value.
- **Equal To Or Below:** The event rule raises a threshold violation event when the value of the metric is equal to or lower than the value.
- **Below:** The event rule raises a threshold violation event when the value of the metric is lower than the value.

- **Value**

Defines the threshold at which the event rule raises the threshold violation event.

**Example:** 40

- **Condition Type**

Defines the condition that triggers the event rule to raise a threshold violation event.

**Options:**

- **Fixed Value**
- **Standard Deviations.** See the "[Standard Deviation Event Rule Conditions](#)" section.
- **Percent of Baseline.** See the "[Percent of Baseline Event Rule Conditions](#)" section.

**Default:** Fixed Value

**NOTE**

You can select the **Utilization (%)** metric only if you have selected the CPU or Memory metric family as the family on which this event rule is based. If you select this metric, you must choose **Fixed Value** as the **Condition Type** for the event rule.

- In the **A Violation is cleared when** section, define the condition that triggers the data aggregator to send a *clear threshold violation event*.

The event rule is created.

3. Save the threshold profile.

The event rule and conditions are added to the threshold profile. The event rule is added to the **Event Rules** pane.

### **Standard Deviation Event Conditions**

Event rules that use a standard deviation compare the poll results to the baseline for the device or component. Event rules calculate the baseline and the standard deviation value for the specific hour of the day of the week.

For more information about these calculations, see [Baseline Calculations](#).

Event rules trigger standard deviation rules when the value of the metric differs from the baseline by the specified number of standard deviations. For event rules with conditions that use the **Above** operator, the event rule raises a threshold violation event when the value of the metric exceeds the baseline value plus the number of standard deviations. For event

rules with conditions that use the **Below** operator, the event rule raises a threshold violation event when the value of the metric is lower than the baseline value minus the number of standard deviations.

#### Example:

The baseline is 65% and the standard deviation is 10%. The rule states that an event triggers when CPU utilization is above 2 standard deviations. This event rule condition triggers when the CPU utilization is greater than 85%.

### Percent of Baseline Event Conditions

Event rules with conditions that use the **Percent of Baseline** condition type compare the poll results to the calculated baseline plus or minus a percentage of the calculated baseline for the device or component. Metrics trigger event rules when the qualifying poll data meets the criteria that you have specified for a Percent of Baseline event rule condition. These condition types are useful when there are numerous or minimal variation in the metric values. Consider using this event rule condition when the standard deviation is above 3 or is low, like 0.1 or 0.0.

#### Examples:

The calculated baseline is 60 degrees and the specified **Percent of Baseline** is 50%. The rule states that an event triggers when the temperature rises above 50% of the baseline. This event rule condition triggers when the temperature is higher than 90 degrees.

**Math:**  $60 + (+50\% \times 60) = 60 + 30 \text{ degrees} = 90 \text{ degrees}$

The calculated baseline is 60 degrees and the specified **Percent of Baseline** is -50%. The rule states that an event triggers when the temperature falls below -50% of the baseline. This event rule condition triggers when the temperature is lower than 30 degrees.

**Math:**  $60 + (-50\% \times 60) = 60 - 30 \text{ degrees} = 30 \text{ degrees}$

### Assign a Group to the Threshold Profile

Identify the devices or components that the threshold profile monitors and enable event rules to raise threshold violation events for metrics that are part of a group by assigning a group to the threshold profile. Threshold profiles apply to the components of the devices in those groups that support the selected metric family. Event rules can raise threshold violation events for the metrics that are part of that group.

#### NOTE

The event rules for threshold profiles that you assign to a collection apply only to the devices in that collection, and not to components and interfaces in the collection.

**Best practice:** For components and interfaces, assign groups to threshold profiles for event rules to raise events.

#### Follow these steps:

1. On the **Threshold Profiles** page, from the **Folder View** or **Table View**, select the threshold profile to which you want to assign a group.
2. Click the **Groups** tab in the right-hand pane.  
A list of groups that are assigned to the threshold profile display.
3. Click **Manage**.  
The **Assign Groups to Threshold Profiles** dialog opens.
4. Select the groups from the **Available Groups** tree that you want to assign to the threshold profile, click the right arrow to add them to the **Selected** list, and then click **OK**.

The groups that you added to the **Selected** list are assigned to the threshold profile. The group is added to the **Related Groups Name** table.

## Manage Threshold Profiles and Event Rules

You manage threshold profiles by managing the threshold profiles, the folders, the assigned groups, and the event rules.

In this article:

- [Manage Threshold Profiles](#)
- [Manage Event Rules](#)

### Manage Threshold Profiles

You can manage threshold profiles in the following ways:

- [View a List of Threshold Profiles](#)
- [Create a Threshold Profile](#)
- [Copy a Threshold Profile](#)
- [Edit a Threshold Profile](#)
- [Delete a Threshold Profile](#)

### View a List of Threshold Profiles

Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Threshold Profiles**.  
The **Threshold Profiles** page appears.
2. On the **Threshold Profiles** page, from the **Threshold Profiles** pane, do one of the following:
  - To view the threshold profiles under a specific threshold profile folder, from the **Folder View**, click to expand the threshold profile folder containing the threshold profiles that you want to view.
  - To view all threshold profiles, click the **Table View**.

A list of threshold profiles display.

### Copy a Threshold Profile

Follow these steps:

1. On the **Threshold Profiles** page, from the **Threshold Profiles** pane, do one of the following:
  - From the **Folder View**, click to expand the threshold profile folder containing the threshold profile that you want to copy, click the threshold profile that you want to copy, and then click **Copy Profile**.
  - Click the **Table View**, select the threshold profile that you want to copy, and then click **Copy Profile**.  
The (23.3.5 and higher) **New/Edit Threshold Profile** (23.3.4 and lower) **Create / Edit Threshold Profile** page appears. By default, the Name of the copied threshold profile is appended with "Copy".
2. Edit the name of the threshold profile, including the other required fields.  
For more information about the fields, see [Configure Threshold Profiles](#).
3. Add event rules and conditions, and then save the threshold profile.

### Edit a Threshold Profile

Follow these steps:

1. On the **Threshold Profiles** page, from the **Threshold Profiles** pane, do one of the following:
  - From the **Folder View**, click to expand the threshold profile folder containing the threshold profile that you want to edit, click the threshold profile that you want to edit, and then click **Edit Profile**.
  - Click the **Table View**, select the threshold profile that you want to edit, and then click **Edit Profile**.  
The (23.3.5 and higher) **New/Edit Threshold Profile** (23.3.4 and lower) **Create / Edit Threshold Profile** appears.



2. Edit the threshold profile, and then save your changes.  
For more information about the fields, see [Configure Threshold Profiles](#).

### **Delete a Threshold Profile**

#### **Follow these steps:**

1. On the **Threshold Profiles** page, from the **Threshold Profiles** pane, do one of the following:
  - From the **Folder View**, click to expand the threshold profile folder containing the threshold profile that you want to delete, click the threshold profile that you want to edit, and then click **Delete Profile**.
  - Click the **Table View**, select the threshold profile that you want to delete, and then click **Delete Profile**.
 The **Delete Threshold Profiles** dialog opens.
2. Confirm that you want to delete the threshold profile by clicking **Yes**.

The threshold profile is deleted, and no longer appears in the **Threshold Profiles** pane.

### **Manage Event Rules**

You can manage event rules in the following ways:

- [View a List of Event Rules for a Threshold Profile](#)
- [Add an Event Rule to a Threshold Profile](#)
- [Copy an Event Rule](#)
- [Edit an Event Rule](#)
- [Link an Event Rule to Other Event Rules](#)
- [Delete an Event Rule from a Threshold Profile](#)

### **View a List of Event Rules for a Threshold Profile**

To view a list of event rules for a threshold profile, on the **Threshold Profiles** page, from the **Threshold Profiles** pane, do one of the following:

- From the **Folder View**, click to expand the threshold profile folder containing the threshold profile for which you want to view a list of event rules, and then select the threshold profile.
- Click the **Table View**, and then select the threshold profile for which you want to view a list of event rules.

The event rules display in the **Event Rules** tab in the right-hand pane, which is selected by default.

### **Copy an Event Rule**

#### **Follow these steps:**

1. On the (23.3.5 and higher) **New/Edit Threshold Profile** (23.3.4 and lower) **Create / Edit Threshold Profile** page, in the **Event Rules** pane, select the event rule that you want to copy, and then click **Copy**.  
The (23.3.5 and higher) **New/Edit Event Rule** (23.3.4 and lower) **Create / Edit Event Rule** dialog opens. By default, the Name of the copied event rule is appended with "Copy".
2. Edit the name of the event rule, including the other required fields, and then save your changes.  
For more information about the fields, see [Configure Threshold Profiles](#).
3. Save the threshold profile.

### **Edit an Event Rule**

#### **Follow these steps:**

1. On the **Create / Edit Threshold Profile** page, in the **Event Rules** pane, select the event rule that you want to edit, and then click **Edit**.

The (23.3.5 and higher) **New/Edit Event Rule** (23.3.4 and lower) **Create / Edit Event Rule** dialog opens.

2. Edit the threshold profile.

For more information about the fields in this dialog, see [Configure Threshold Profiles](#).

### Link an Event Rule to Other Event Rules

You can link up to three event rules (a root-linked event rule with one or two dependent-linked event rules) to have the data aggregator evaluate on multiple metric families together:

- When the metrics meet the threshold for *all* the conditions that would trigger the data aggregator to send a *raised threshold violation event* for *all* linked event rules, the data aggregator sends a *raised threshold violation event* for the linked event rules, and as a result, DX NetOps Spectrum (Spectrum) raises an alarm.
- When the metrics meet the threshold for *any* condition that would trigger the data aggregator to send a *clear threshold violation event* for *any* of the linked event rules, the data aggregator sends a *clear threshold violation event* for the linked event rules, and as a result, Spectrum clears the previously-raised alarm.

The data aggregator evaluates the rules at the slowest poll rate of the associated metric families.

**Prerequisites:** The event rules have been created and added to the threshold profile.

**Follow these steps:**

- On the (23.3.5 and higher) **New/Edit Threshold Profile** (23.3.4 and lower) **Create / Edit Threshold Profile** page, in the **Event Rules** pane, select the event rule to which you want to link other event rules, and then click **Link**. This rule will become the root-linked event rule.  
The **Link Event Rules** dialog appears.  
The following image shows an example of this dialog:

### Link Event Rules



You can link up to three rules. When the conditions for all rules is met, NetOps Portal raises an event. NetOps Portal evaluates the rules at the slowest poll rate of the associated metric families.

Linked Event Rule ★

Linked Event Rule ★

Save

Cancel

- From the **Linked Event Rule** drop-downs, select the event rule to link to this event rule, and then click **Save**.

The dependent-linked event rule is linked to the root-linked event rule. On the (23.3.5 and higher) **New/Edit Threshold Profile** (23.3.4 and lower) **Create / Edit Threshold Profile** page, in the **Event Rules** pane, the dependent-linked event rules display in the **Linked Rules** column for the root-linked event rule.

The following image shows an example of a dependent-linked event rule (Memory) linked to a root-linked event rule (CPU) based on version:

(23.3.5 and higher)

New/Edit Threshold Profile

Name \*

Test Profile

Owner \*

admin

Description

Last Modified By

admin

Folder \*

Test

Creation Time

Feb 5, 2024 8:14:30 PM UTC

Status \*

Enabled

Last Modified ...

Feb 12, 2024 8:03:09 PM UTC

When Event Rules Run

When Rules Run

☒ On Schedule

Run Rules

☒ Only During the Specified Times

☐ Not During the Specified Times

On Days \*

☐ Sun

☒ Mon

☒ Tue

☒ Wed

☒ Thu

☒ Fri

☐ Sat

Start Time/End Time \*

00:30

24:00

Note: Times are in UTC

Event Rules

Quick Filter

New

Copy

Edit

Link

Delete

Rule Name	Metric Family	Linked Rules	Aggregation	D...	Wi...	Se...
CPU	CPU	Memory	No Aggregation	30...	30...	M...
Interface	Interface	Memory	No Aggregation	30...	30...	M...
Memory	Memory		No Aggregation	30...	30...	M...

Event Conditions

Quick Filter

Event Type	Metric	Thr...	Condition Type
Violation	Utilization (%)	> 40...	Fixed Value
Clear	Utilization (%)	< 40...	Fixed Value

(23.3.4 and lower)

Create / Edit Threshold Profile

Name \*

CPU-Profile

Owner \*

Description

Last Modified By

Folder \*

Device Level Aggregation

Creation Time

Jul 6, 2023 8:16:58 PM UTC

Status \*

Enabled

Last Modified ...

Jul 6, 2023 8:16:58 PM UTC

Event Rules

Quick Filter

New

Copy

Edit

Link

Delete

Rule Name	Metric Family	Linked Rules	Aggregation	Dur...	Win...	Sev...
CPU	CPU	Memory	Aggregate Compon...	300...	300...	Maj...
Memory	Memory		Aggregate Compon...	300...	300...	Maj...

Event Conditions

Quick Filter

Event Type	Metric	Thres...	Condition Type
Violation	Utilization (%)	> 300.0	Fixed Value
Clear	Utilization (%)	< 300.0	Fixed Value

## Delete an Event Rule from a Threshold Profile

### Follow these steps:

1. On the (23.3.5 and higher) **New/Edit Threshold Profile** (23.3.4 and lower) **Create / Edit Threshold Profile** page, in the **Event Rules** pane, select the event rule that you want to delete from the threshold profile, and then click **Delete**. The **Delete Event Rule** dialog opens.

If the event rule is a dependent-linked event rule, deleting it unlinks it from all root-linked event rules, and the following message appears in the **Delete Event Rule** dialog based on version:

(23.3.5 and higher)

Are you sure you want to delete the event rule "Memory"? This is a dependent-linked event rule. In 24 hours, the data aggregator will delete all events raised or cleared for this event rule.

(23.3.4 and lower)

This is a dependent-linked event rule. In 24 hours, the data aggregator will delete all events raised or cleared for this event rule.

If the event rule is a root-linked event rule, the following message appears in the **Delete Event Rule** dialog based on version:

(23.3.5 and higher)

Are you sure you want to delete the event rule "CPU"? This is a root-linked event rule. In 24 hours, the data aggregator will delete all events raised or cleared for this event rule.

(23.3.4 and lower)

This is a root-linked event rule. In 24 hours, the data aggregator will delete all events raised or cleared for this event rule.

2. Confirm that you want to delete the event rule by clicking **Yes**.

The event rule is deleted, and no longer appears in the **Event Rules** pane.

## Threshold Assessment Logic for Window and Duration

NetOps Portal loads data (points) every minute. The event rules evaluate the threshold against these loaded data points. Threshold processing at the data layer operates in bulk. The monitoring profile poll rate determines the frequency at which the event rules evaluate poll cycle data. For example, if you have set the poll rate to one minute, the event rules evaluate one-minute poll cycle data.

For more information about how to set poll rates, see [Manage Monitoring Profiles](#).

At each poll cycle, the event rule retrieves the items that violate the threshold rules, using the following logic:

```
show me all items that violated x threshold profile for y duration within z window
```

If threshold processing returns the item, the event rule generates a threshold violation event if one is not already active.

The event rule raises a clear event when a metric meets either one of the following conditions:

- The metric does not have enough data points that meet the violation condition for the duration within the window to generate a threshold violation event.
- The metric has enough data points that meet the clear condition for the duration within the window.

Missed polls do not contribute to threshold violation events or event clearing. If you have a threshold profile applied to a mix of items with both one- and five-minute poll rates, with event rules showing 60 seconds (one minute) **Duration** and **Window** values, NetOps Portal evaluates every one-minute poll cycle data point for violations. Every five minutes, NetOps Portal evaluates the five-minute polled item data points along with the one-minute polled item data points. In this scenario, the one minute polled items raise events against each poll cycle, same as the five-minute polled items, while both use the lower one minute **Duration** and **Window** values.

Event rules raise clear events when the metrics no longer have enough violating data points to meet the duration in the window, and have enough data points that meet the clear criteria for the duration within the window. Event rules are always looking at the number of data points in the window.

### **Example Scenarios**

The following example scenarios assume the items targeted poll at the default 5-minute poll rate:

- **Duration:** 300 seconds/five minutes  
**Window:** 900 seconds/15 minutes  
 In this instance, any one poll cycle out of the possible three in the window raises, or continues, the threshold violation event. The poll cycle that triggers the condition must be within the window.
- **Duration:** 600 seconds/10 minutes  
**Window:** 900 seconds/15 minutes  
 In this instance, any two poll cycles out of the possible three in the window raises the threshold violation event. In this example, poll cycle A+B, A+C, and B+C can trigger the condition that raises the threshold violation event. The poll cycles that trigger the condition must be cumulative.
- **Duration:** 900 seconds/15 minutes  
**Window:** 900 seconds/15 minutes  
 In this instance, three poll cycles out of the possible three in the window raises the threshold violation event. Only when poll cycles A, B, and C in a given rolling window violate the threshold does the event rule raise the threshold violation event. The poll cycles that trigger the condition must be consecutive. The duration is cumulative.
- **Duration:** 900 seconds/15 minutes  
**Window:** 1800 seconds/30 minutes  
 In this instance, three poll cycles out of the possible six in the window raises the threshold violation event. The poll cycles can be cumulative or consecutive. As long as any three poll cycle values out of each six evaluated for the rolling window violate the threshold, the event rule raises a threshold violation event.

### **Prevent Noisy Events**

You can prevent noisy events (a flapping condition) using one of the following methods:

- (Recommended) Configure event rules in DX NetOps Spectrum (Spectrum) that recognize DX NetOps Performance Management threshold violation events as a flapping condition, and consolidate them as symptoms/events on a single threshold violation event.  
 For more information about how to configure event rules in Spectrum, see the [DX NetOps Spectrum documentation](#).
- Adjust the event rule to use the window/duration logic. Change the **Duration** from 900 (15 minutes) in a **Window** of 1800 (30 minutes) to a **Duration** of 600 (10 minutes) in a **Window** of 1800 (30 minutes). While this change makes the event rule more sensitive on the front end, it also makes the event rule more sensitive in terms of clearing the threshold violation events.

### **Example: Duration and Window**

A monitored device has a poll cycle of five minutes. An associated threshold profile has an event rule with a duration of 600 (10 minutes) and a window of 3600 (1 hour). The event rule does not raise an event when a metric triggers the event rule conditions for a single poll result because the 5-minute poll does not reach the 10-minute duration. The event rule raises an event only if a metric triggers the event rule conditions for a second poll result within one hour of the first triggering poll.

#### **NOTE**

When a metric breaches a threshold, the event rule creates a threshold violation event. When the event rule clears the violation, it rechecks the threshold with the next poll cycle. If a metric breaches the threshold again, and it has enough data points to meet the duration in the window, the event rule creates a threshold violation event.

## View Event Rules and Threshold Violation Events

Event rules set conditions, or event conditions, to raise threshold violation events, such as when a threshold is violated and when the threshold violation is cleared.

In this article:

- [View Event Rules and Conditions for a Threshold Profile](#)
- [View the Threshold Violation Events for a Threshold Profile](#)
- [View the Details for a Threshold Violation Event](#)

### View Event Rules and Conditions for a Threshold Profile

Follow these steps:

1. Log in as a user with the Create DA Threshold Profile or the Administer DA Threshold Profile role right.
2. Do one of the following steps:
  - (Administrators) Hover over **Administration**, **Monitored Items Management**, and then click **Threshold Profiles**.
  - (Users) Click the name of your user account in the upper-right corner, and then click **Manage Threshold Profiles**. The **Threshold Profiles** page appears.
3. Select the threshold profile for which you want to view event rules and conditions, and then click the **Event Rules** tab in the right-hand pane.

A list of event rules for the threshold profile display. The following image shows this tab:

**Figure 14: The Event Rules tab**

The screenshot shows the DX NetOps interface. The top navigation bar includes Home, Alarms, Performance, Inventory, Reports, System Health, and Administration. The main content area is titled 'Threshold Profiles' and has tabs for Event Rules, Groups, and Events. The 'Event Rules' tab is active, displaying a table of event rules for the 'Aggregated CPU' threshold profile. The table has columns for Rule Name, Metric Family, Linked Rules, Aggregation, and Duration. Below the table, there is a section for 'Event Conditions' with a table showing Event Type, Metric, Threshold, and Condition Type.

Rule Name	Metric Family	Linked Rules	Aggregation	Duration
Aggregated CPU	Aggregated CPU		No Aggregation	300 sec

Event Type	Metric	Threshold	Condition Type
Violation	Utilization (%)	> 40	Fixed Value
Clear	Utilization (%)	< 40	Fixed Value

The **Event Rules** tab shows the rule name, the metric family, the linked rules, aggregation, and duration. Event conditions appear in the **Event Conditions** view for an event rule when the conditions in the event rule are met. Event rule conditions (Event Type) include **Violation** (threshold violation event) and **Clear** (clear event).

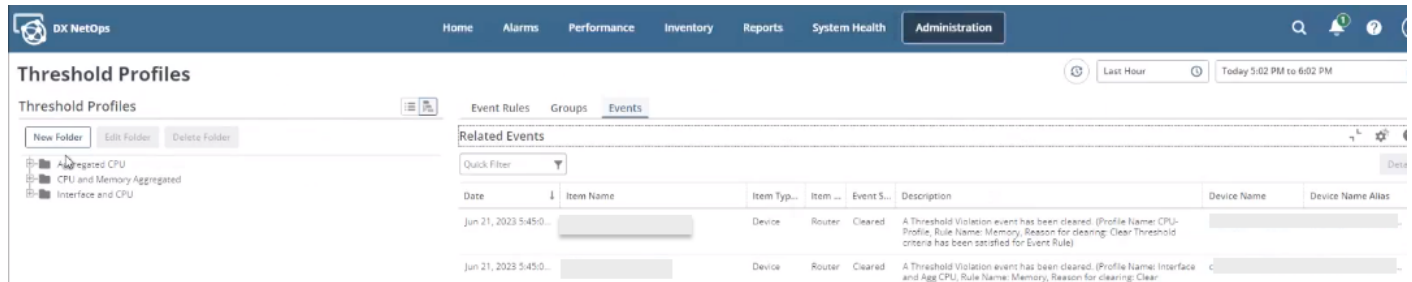
### View the Threshold Violation Events for a Threshold Profile

You can view the threshold violation events that have been generated as a result of a specific event rule on the **Threshold Profiles** page. You can view the threshold violation events that have been generated as a result of all event rules in all threshold profiles on the **Events Display** dashboard. This procedure details how to view the threshold violation events that event rules have raised as a result of a specific event rule.

**Prerequisite:** You are logged in as a user with the Create DA Threshold Profile or the Administer DA Threshold Profile role right.

**Follow these steps:**

1. Do one of the following steps:
  - (Administrators) Hover over **Administration, Monitored Items Management**, and then click **Threshold Profiles**.
  - (Users) Click the name of your user account in the upper-right corner, and then click **Manage Threshold Profiles**.
 The **Threshold Profiles** page appears.
2. Select the threshold profile for which you want to view threshold violation events, and then click the **Events** tab in the right-hand pane.  
 A list of events for the threshold profile, including threshold violation events, display. The following image shows this tab:

**Figure 15: The Events tab**

View the following information from the **Related Events** table on the **Events** tab:

- **Date**  
Specifies the date and time the data aggregator raised or cleared the threshold violation event.
  - **Device Name**  
Specifies the device that is related to the threshold violation event.
  - **Interface Name Alias**  
Specifies the interface name alias that is related to the threshold violation event.
  - **Item Name**  
Specifies the item that is related to the threshold violation event.
  - **Item Type Name**  
Specifies the type of monitored item as a category, such as device, application, and interface.
  - **Item SubType**  
Specifies the subtype of monitored item (subcategory), such as switch, router, Wi-Fi device, manageable, firewall, hub, or printer.
  - **Event SubType**  
Defines the type of threshold violation event that the data aggregator sent.
- Values:**
- **Raised:** For a raised threshold violation event.
  - **Cleared:** For a clear threshold violation event.
- **Description**  
Specifies details for the threshold violation event.
3. (Optional) Select the threshold violation event, and then click **Details**.  
 The **Event Details** dialog opens. For more information, see [the "View the Details for Threshold Violation Events" section](#).
  4. (Optional) Click **Change the selected time window** (the time range) at the top of the page next, and then select one of the default time ranges.

**TIP**

You can also set a different time range by selecting **Custom Time Range**.

## View the Details for a Threshold Violation Event

View the details for a threshold violation event by selecting the threshold violation event from the **Events** tab, and then click **Details**.

The **Event Details** dialog opens. The following image shows this dialog:

**Figure 16: Threshold Violation Event Details**

Event Details

Event ID

21568521

Event Type

Threshold Violation

Event Subtype

Raised

Occurred On

Sep 13, 2023 5:30:00 PM GMT

Description

A Threshold Violation event has been raised on "[c [REDACTED]]". (Profile Name: Interface and CPU, Rule Name: CPU,Interface).

Device Name

[REDACTED]

Quick Filter

Property Name	↑	Event Property Value
Maximum Metric Value(s)		50,138.888888
Metric Name(s)		
Minimum Metric Value(s)		50,94.444445333333
Profile Folder ID		62150
Profile ID		2763784

100 per page

⏪

⏩

Page

1

of 1

⏪

⏩

Displaying 1 - 18 of 18

OK

View the following information:

- **Event ID**  
Specifies the identifier for the threshold violation event.
  - **Event Type**  
For threshold violation events, displays **Threshold Violation** as the event type.
  - **Event Subtype**  
Defines the type of threshold violation event that the data aggregator sent.
- Values:**



- **Raised:** For a raised threshold violation event.
- **Cleared:** For a clear threshold violation event.
- **Occurred On**  
Specifies the date and time the data aggregator raised or cleared the threshold violation event.
- **Description**  
Specifies details for the threshold violation event.
- **Device Name**  
Specifies the device that is related to the threshold violation event.
- **Property Name**  
Specifies the property name that is related to the threshold violation event.
- **Event Property Value**  
Specifies the event property value that is related to the threshold violation event.
- **Interface Name Alias**  
Specifies the interface name alias that is related to the threshold violation event.

## Configure Business Hours Filtering

Business hours represent times of the day and week of increased importance. To filter the data in a dashboard to relevant hours for capacity planning or to report business hours, apply a business hours filter.

As a user with the Administer Business Hours role right (Administrators), organize the devices to which you want to apply business hours into site groups, and then associate the business hours filters with site groups. Site groups are special custom groups that organize managed items, such as devices, that are in the same location or are near each other. Site groups can include time zone and business hours attributes so that you can precisely filter data from business-critical times of day.

Subgroups of managed items do not directly inherit the business hours and time zone filters from the site group to which it is associated. You must associate the business hours definition with each relevant subgroup in the tree.

NetOps Portal renders views that you have applied a business hours filter based on the selected site group. The selected site group's filters apply to all items in that group and in any subgroups. When you change the selected site group to a subgroup, the parent group filters are not applicable. Group references inherit associated business hours and time zone filters from the original site group.

For more information:

- About site groups, including how to create them and add managed items to them, see [Manage Groups](#).
- About group filters, see [Groups](#).
- About subgroups and group references, see [Manage Subgroups](#).

### NOTE

You can also manage (create and edit) business hours definitions by way of REST Web Services.  
For more information, see [Use Web Services to Manage Business Hours](#).

In this article:

- [Create a Business Hours Definition](#)
- [Edit a Business Hours Definition](#)
- [Business Hours Filter Support and Limitations](#)
- [Next Step](#)

### Create a Business Hours Definition

Before you can apply a business hours filter to views, dashboards, or context pages, create a business hours definition. Optionally, associate the business hours definition with site groups. Each definition includes the times of day and days of

the week. Select the hours and days that reflect periods of increased activity. Create definitions for every distinct location in your enterprise.

**Prerequisite:** You have the Administer Business Hours role right.

**Follow these steps:**

1. Hover over **Administration**, and then click **Group Settings, Business Hours**.  
The **Manage Business Hour Definitions** page appears.
2. Click **New**.  
The **Add Business Hours** dialog opens.
3. Complete the following fields:
  - **Name**  
Specify a name for the business hours definition. This name appears in the view subtitle for those views to which you have applied the business hours filter.
  - **Description**  
Specify a description for the business hours definition.
  - **On Days**  
Select the checkboxes for the days of the week that you want included in this business hours definition.
  - **Start Time/End Time**  
Select the start time and the end time for this business hours definition from the drop-down lists.

**NOTE**

You can select only whole-hour increments.

4. (Optional) Associate the business hours definition with a site group that does not already have business hours set:
  - a. Click **Select Site Groups to Associate**.
  - b. Add one or more site groups to the **Selected Sites** list.

**NOTE**

You can select only site groups that have a time zone. Because you can define business hours definitions only in whole-hour increments, only site groups that have time zones with whole-hour offsets are available for selection.

You can specify the business hours and time zone for a site group when you create or edit site groups.

- c. Click **OK**.
5. Click **Save**.

### **Edit a Business Hours Definition**

**Prerequisite:** You have the Administer Business Hours role right.

Follow these steps:

1. Hover over **Administration**, and then click **Group Settings, Business Hours**.  
The **Manage Business Hour Definitions** page appears.
2. Click the business hours definition that you want to edit from the list, and then click **Edit**.  
The **Edit Business Hours** window opens.
3. Modify the name, description, and start and end time, and then click **Save**.

### **Business Hours Filter Support and Limitations**

The following views filter data results based on applied business hours:

- Table view
- Card view
- Dual-severity card view
- Gauge chart
- Pie chart
- Horizontal bar chart
- Composite chart table
- Group scorecard table
- Group scorecard trend

For more information about this view type, see [Group Scorecard Trend Views](#).

Data with the periods outside of the applied business hours are displayed as shaded in the following views:

- Calendar Heat Chart
- Trend Chart (including MultiView Trend and MultiTrend)
- Dynamic Trend
- Trend views in on-demand reports

The following metrics, data, views, dashboards, and charts do not show filter data when a business hours filter is applied:

- Projection metrics
- Daily rollup data and views that shows daily resolution
- Table views in on-demand reports
- Views on dashboards that have a device or interface context
- Time bar chart

### **Next Step**

Now that you have created business hours definitions, you can apply them to views, dashboards, or context pages.

For more information:

- About how to assign a business hours filter and time zone to a view, see [Customize Views](#).
- About how to apply a business hours filter to a context page, see [Manage Context Pages](#).
- About how to apply a business hours filter to a dashboard, see [Manage Dashboards](#).
- About how to apply a business hours filter to items reported on in on-demand reports, see [Manage Scheduled Reports](#).

## **Schedule Maintenance Indicators**

Maintenance indicators represent times when maintenance is occurring. After you schedule maintenance indicators, views indicate maintenance with shading.

Maintenance indicators represent times when maintenance is occurring, and apply to all the devices and components in a site group. With scheduled maintenance indicators applied to a view, the view applies shading to indicate maintenance. Selecting an associated site group in the context displays maintenance indicators in applicable views as you navigate between dashboards. The subtitle of each view indicates whether maintenance indicators apply to the view.

In this article:

### **Maintenance Indicator Inheritance**

Associate maintenance indicators with relevant subgroups, as they do not directly inherit maintenance indicators from site groups. When NetOps Portal renders views, these filters apply to all items based on the selected site group. The filters of the selected site group apply to the items in that group and in any subgroups. When you change the selected site group to a subgroup, the filters of the parent group are no longer applicable.

Reference groups inherit associated maintenance indicators from the original site group.

### **Maintenance Indicators Support and Limitations**

Maintenance indicators can apply to data only in trend chart and calendar heat chart view types.

#### **NOTE**

The data in these views is not limited to the maintenance indicators.

When you have selected an associated site group in the context, maintenance indicators appear as shaded cells in these charts. Shading appears only when you select the site group with maintenance indicators at the page or view level.

The following are known limitations for maintenance indicators:

- The duration must be between 1 and 24 hours.
- Maintenance indicators do not apply to the following:
  - Daily rollup data or to views that show daily resolution.
  - Table views in on-demand reports.
  - Views on context pages.

### **Schedule Maintenance Indicators**

To apply maintenance indicators to views, create maintenance indicators and associate the maintenance indicators with a site group. Each definition includes the times of day and day of the week. Select the hours and day that reflect maintenance periods. Create definitions for every distinct location in your enterprise.

Only users with the Administer Maintenance Indicators role right can add, edit, copy, or delete maintenance indicators.

1. Hover over **Administration, Group Settings**, and then click **Maintenance Indicators**.  
The **Manage Maintenance Indicators** page appears.
2. Click **New**.  
The **Add Maintenance Indicators** dialog opens.
3. Complete the following fields:
  - **Name**  
Specifies the name for the maintenance indicator. This name appears on views displaying maintenance indicators.
  - **Description**  
Specifies the description for the maintenance indicator. This description appears only on the **Manage Maintenance Indicators** page.
  - **Maintenance Date**  
Specifies the day for the maintenance indicator.  
**Default:** today's date
  - **Start Hour**  
Defines the start time for the maintenance indicator. Choose a whole hour increment. The same hour is applied to all selected days.  
**Default:** 00:00  
**Options:** 00:00 to 23:00
  - **End Hour**  
Defines the end time for the maintenance indicator. Choose a whole hour increment. The same hour is applied to all selected days.  
**Default:** 24:00  
**Options:** 01:00 to 24:00
4. Associate the maintenance indicator with a site group by moving it from the **Available Sites** list to the **Selected Sites** list. Associating the maintenance indicator with a site group applies it to that group.

**NOTE**

You can associate maintenance indicators with site groups that already have a time zone. You can associate maintenance indicators only with site groups that have time zones with whole hour offsets.

5. Click **Save**.

The maintenance indicator is scheduled.

## Manage Monitored Devices

You can manage the details for monitored devices (routers, switches, servers, and pingable devices) and can view their associations with device collections, components, monitoring profiles, and metrics.

Monitored devices are those that the data aggregator has discovered, including the SNMP-manageable devices, pingable devices, and DX NetOps Virtual Network Assurance (VNA) and DX NetOps Mediation Manager (DX NetOps MM) devices. During device discovery, the data aggregator adds discovered devices to device collections, which initiates device monitoring and polling. The state of the discovery profile changes to "Running".

NetOps Portal monitors only accessible devices. The data aggregator automatically classifies manageable devices as Router, Switch, and Server primary device types based on the device service information.

You can manage monitored devices in the following ways:

- [View a List of Monitored Devices](#)
- [View the Details for a Monitored Device](#)
- [Edit the Details for a Monitored Device](#)
- [Rediscover a Monitored Device](#)
- [Pause or Resume Polling of a Monitored Device](#)
- [View the Components of a Monitored Device](#)
- [View More Device Details](#)
- [Change the Primary IP Address for a Monitored Device](#)
- [Reconfigure components for configuration updates](#)
- [Manage Device Life Cycles](#)
- [Manage Hostname Changes](#)
- [Override Device Types](#)
- [Device Deduplication](#)
- [Delete Components That Are Not Present on Devices](#)
- [Delete Devices](#)

**NOTE**

Some management tasks require administrator privileges.

### **View a List of Monitored Devices**

You can view the details for monitored devices and can view their associations with device collections, components, monitoring profiles, and metrics from the **Monitored Devices** page. You can also view filter reports. This information helps you see information in context, such as which monitoring profiles that NetOps Portal uses to poll device components.

To view this page, hover over **Administration**, **Monitored Items Management**, and then click **Monitored Devices**.

The **Monitored Devices** pane displays a list of monitored devices. With the **Tree View** tab selected, and you have sorted by **Device by Collection**, the following collections are included:

- **All Clients**  
Displays client devices.

**NOTE**

This collection lists only up to 35k client devices. If there are more than 35K client devices, the list is shown as empty. This is a known issue.

- **All Data Aggregators**

- **All Devices**

Displays all devices.

**NOTE**

- The collection lists all devices *except* client devices.
- This collection list only up to 35k devices. If there are more than 35K devices, the list is shown as empty. This is a known issue.

- **All ESX Hosts**

- **All Imported Devices**

- **All Manageable Devices**

- **All Routers**

- **All Satellite Routers**

- **All Servers**

- **All Switches**

- **All Virtual Machines**

- **All VMware vCenters**

- **Editable groups with rules**

- **Empty Collection**

- **Fast Polling interfaces**

- **Normal Polling interfaces**

- **OI Collection**

- **UTF8 Device**

### View the Details for a Monitored Device

#### Follow these steps:

1. From the **Monitored Devices** page, locate the device that you want to view using one of the following options:
  - With the **Tree View** tab selected, sort by **Device by Collection** or **Device by Monitoring Profile** from the drop-down, and then select a specific device from the corresponding tree view.
  - On the **Search** tab, search by host name, device name, or IP address. You can enter a partial name or IP address to return a list of devices that contain that partial match, but you cannot use wildcards and regular expressions. By default, the **Details** tab displays in the right-side pane.
2. View the following details for the monitored device:
  - **Name**  
Displays the name of the monitored device.
  - **Item ID**  
Displays the item ID for the monitored device.
  - **Host Name**  
Displays the host name for the monitored device.
  - **IP Address**  
Displays the IP address for the monitored device.
  - **IP Domain**  
Displays the IP domain for the monitored device.
  - **DC Host**

Displays the data collector host for the monitored device.

- **Status**  
Displays the status (contact status) of the connection between the data collector and the monitored device. The data aggregator sends the contact status for each device.
- **Life Cycle State**  
Displays the monitoring behavior for the monitored device, such as whether it is active, retired, or in maintenance. For more information about life cycle state, see [Manage Device Life Cycles](#).
- **Creation Time**  
Displays the date and time that the monitored device was created.
- **Device Types**  
Displays the data aggregator's classification of this device, such as pingable, manageable, and other. For more information about the device classifications, see [Override Device Types](#).
- **Polled Items Count**  
Displays the number of devices and components that the associated data collector is polling.

### **Edit the Details for a Monitored Device**

#### **NOTE**

Changes to device attributes can result in changes to the groups and device collections that a device is in. Changes to groups and device collections can potentially add or remove monitoring profiles.

#### **Follow these steps:**

1. From the **Details** tab, edit the details for monitored devices by editing their associations with device collections, components, monitoring profiles, and metrics:
  - **IP address**  
For more information, see [Change the Primary IP Address for a Device](#).
  - **DC Host**  
The data collector host.
  - **SNMP Profiles**  
The SNMP profile that NetOps Portal uses to poll the device.
  - **SNMP version for the device**
2. Save your changes.

The device details are modified.

### **Pause or Resume Polling of a Monitored Device**

From the **Monitored Devices** page, with the device that you want to pause or resume polling selected from the **Monitored Devices** list, from the **Details** tab in the right-side pane, do one of the following to pause or resume polling:

- To pause polling, click **Stop Polling**. At the prompt (the **Stop polling of the device** dialog), click **Yes** to stop polling the device.  
Automatic change detection is disabled. You can manually update metric families or run the rediscovery.

#### **NOTE**

You can stop the polling for any device except DX NetOps MM devices.

- To resume polling, click **Start Polling**.

### **View the Components of a Monitored Device**

#### **Follow these steps:**

1. From the **Monitored Devices** page, with the device that you want to view components selected from the **Monitored Devices** list, click the **Polled Metric Families** tab in the right-side pane.

The table displayed show the total set of metric families that are polled on a device and their poll rates. This set is based on the consolidation of all the monitoring profiles on the device. This table includes the following columns:

- **Metric Family**  
Shows the metric family polled on the device.
- **Vendor Cert**  
Shows the vendor certifications that the device uses for each metric family. When vendor certification priority grouping is occurring on a device, more than one row appears in this column. One row appears for each vendor certification that is used for a metric family.
- **Status**  
Shows whether the device supports the metric family.

**Values:**

- **Supported:** The device supports the metric family.
- **Not Supported:** The device does not support the metric family.

The following video shows how to view the status of newly-discovered devices in NetOps Portal to confirm the metric family, vendor certificate, polling duration, and if the devices are available in the inventory:

- **Last Update Time**  
Displays the date and time the metric family was polled on the device.

The following video shows how to view metric family and vendor certification details for devices in the NetOps Portal inventory to understand and verify monitoring details:

The **Components** table shows the polling status on the components for a metric family component that was previously discovered. One of the following values displays in the **Status** column:

- **Active**  
The component is being polled.

**NOTE**

DX NetOps MM devices show **Not Polled** in the **SNMP Poll Rate** column. NetOps Portal does not poll DX NetOps MM devices.

- **Inactive**  
Polling has stopped on the component because the metric family is no longer monitored for the device.
- **Not Present**  
The component no longer exists on the physical device. Polling is stopped on the component. You can view historical data for reporting purposes. By default, these components are not synchronized with NetOps Portal. To enable this option, select **Synchronize items that are no longer present** in the **Edit Data Source** dialog on the **Manage Data Sources** page in NetOps Portal.

2. To reconfigure components for any configuration updates, see [Device Reconfiguration](#).

### View More Device Details

You can view more device details from the following tabs on right-side pane for a selected device on the **Monitored Devices** page:

- **Monitoring Profiles**  
Select a device collection to view the associated profiles names. Hover over a profile to see the description.
- **Threshold Profiles**  
View the threshold profiles that are applied to the selected device due to the groups to which the device belongs.
- **Metrics**  
View a list of metrics that this device supports. Select a metric family to view the following details:
  - The backing vendor certification.
  - The vendor source (for SNMP vendor certifications, the MIB table source is displayed).
  - Whether the metric is collected.
  - The expression that is used to calculate each metric.
- **Filter Report**



View the interface filter criteria that were used during the component monitoring in the **Interface Filter Criteria** pane. The **Filter Report** tab also shows a report of all of the interfaces that are identified on the device. The tab shows whether the interfaces matched the specified filter criteria.

#### NOTE

If you change the rules on a custom monitoring profile or if you disassociate the monitoring profile from a group, those changes are not reflected in the **Interface Filter Criteria** pane. Filter the interfaces using the changes made to the filter criteria and monitoring profile by [rediscovering the device](#).

- **Events**

View the events that have occurred on the selected device.

## Change the Primary IP Address for a Device

You can change the IP address that DX NetOps Performance Management uses to monitor a device.

The primary IP address for a device is the IP address that DX NetOps Performance Management uses to monitor the device. When a device is first discovered with the IP ranges discovery profile, DX NetOps Performance Management tries to use the IP address that maps to the hostname as the primary IP address.

If the primary IP address on a devices item changes, an event is generated on that device item.

You can change the primary IP address using the following methods:

- [Have DX NetOps Performance Management Change the Primary IP for Devices Automatically](#)
- [Change the Primary IP Address Using NetOps Portal](#)
- [Change the Primary IP Address Through REST](#)

### Have DX NetOps Performance Management Change the Primary IP for Devices Automatically

By default, when the IP address that maps to a hostname on the physical device changes, the primary IP on the device item does not change. To have DX NetOps Performance Management change the primary IP for devices automatically, enable automatic IP change detection. With automatic IP change detection enabled, DX NetOps Performance Management tries to detect the IP address change only after two consecutive polling failures. The data collector uses a reverse hostname lookup to find the new IP address.

Issue a PUT to the following `discoverydefaultconfig` endpoint:

```
http://<da-host>:8581/rest/discoverydefaultconfig/<itemID>
```

- **itemID**

Specifies the NetOps Portal device item ID.

Enter the changed IP change detection in the **Body** tab of the HTTP Request pane, for example:

- To *enable* automatic IP change detection:

```
<DiscoveryDefaultConfig version="1.0.0">
  <DetectIPChange>true</DetectIPChange>
</DiscoveryDefaultConfig>
```

#### IMPORTANT

To prevent automatic IP change detection from causing errors, ensure that your DNS is up to date.

- To *disable* automatic IP change detection:

```
<DiscoveryDefaultConfig version="1.0.0">
  <DetectIPChange>false</DetectIPChange>
</DiscoveryDefaultConfig>
```

## Change the Primary IP Address Using NetOps Portal

Situations can arise when you want to change the primary IP address of a device. For example, you want to change the hostname IP address to the loopback IP address, or the IP address is no longer reachable on the device.

### Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Monitored Devices**.

The **Monitored Devices** page appears.

2. Select **Collections** or **Monitoring Profile**, and then select a specific device.

#### TIP

You can also search by host name, device name, or IP address. You can enter a partial name or IP address to return a list of devices that contain that partial match. Wildcards and regular expressions are not supported.

The **Details** for the device display.

3. Change the primary IP address by taking one of the following steps:

- Edit the **IP Address** field, and then click **Save**.
- Right-click an IP address in the IP Addresses table, select **Set this IP as the device's primary IP**, and then click **Save**.

The primary IP address is changed.

## Change the Primary IP Address Through REST

You can change the primary IP address of a device with a REST client or HTTP tool. If you change the primary IP address of a device to an IP address that another devices uses, the REST call returns an error.

Issue a PUT to the following `devices` endpoint:

```
http://<da-hostname>:8581/rest/devices/deviceitemID
```

Enter the changed primary IP address in the **Body** tab of the HTTP Request pane, for example:

```
<Device version="1.0.0">
  <PrimaryIPAddress><IP></PrimaryIPAddress>
</Device>
```

#### • IP

Specifies the new primary IP address of the device.

### Example:

In this example, the primary IP address on the device changes to 1.2.3.4 :

```
<Device version="1.0.0">
  <PrimaryIPAddress>1.2.3.4</PrimaryIPAddress>
</Device>
```

## Device Reconfiguration

Device reconfiguration includes changes to physical device components and the software configuration, such as monitoring response path tests for protocols. The data aggregator uses the same method to monitor both types of reconfiguration. To keep device components up-to-date, you can monitor and update device reconfiguration changes in the data aggregator.

Additional examples of reconfiguration changes include:

- Adding a board to a device, which adds more ports.
- Adding memory, CPUs, physical interfaces, or any metric family to a discovered device.
- Reconfiguring a virtual switch.
- Changing the configuration of a device to include a discovered device in routing protocols.

When the data aggregator detects a change, it generates reconfiguration events and can update its representation of the metric family to reflect the changes to device components. You can view these reconfiguration events from the dashboards, such as **Operations Displays** and **Events Display**.

Understanding how change detection works in the data aggregator can help you to select the options that are best suited to monitoring device reconfiguration in your environment. For example, you can set the frequency for change detection monitoring.

In this article:

- [Manage Change Detection](#)
- [Update Device Reconfiguration Automatically](#)
- [Update Device Reconfiguration Manually](#)
- [Reconfiguration Detection of Ports](#)

## **Manage Change Detection**

Change detection management planning helps to ensure that the data aggregator detects and monitors device reconfigurations in your environment according to your needs. You can plan ahead for any device reconfiguration when you first set up the data aggregator to discover new devices. You can also edit these options at any time after the devices are discovered.

Refer to the following guidelines when planning change detection:

- Likelihood of change.
- Frequency of change.
- Tolerance for outdated data.

Monitor metric families, such as CPUs, for reconfiguration infrequently. For dynamic metric families, such as virtual systems, choose a more frequent rate.

Manage change detection by [creating or editing a custom monitoring profile](#).

### **TIP**

You can also copy an out-of-the-box monitoring profile, and then edit the copy.

### **TIP**

- If your environment is undergoing major maintenance, clear the **Automatically Update Metric Families** option until the maintenance is complete. For small, regular changes, leave this option selected to ensure that the data aggregator stays up-to-date.
- Monitoring profiles are assigned to device collections. If you want to monitor special devices differently, create a custom device collection and assign a custom monitoring profile with the desired change detection settings. For example, you can monitor critical core routers more frequently than other routers by creating a device collection and assigning a custom monitoring profile that performs change detection hourly. Other routers remain in the 'All Routers' device collection using the out-of-the-box monitoring profile (with no change detection), or a custom monitoring profile that you set to less frequent change detection.

## **Update Device Reconfiguration Automatically**

Reconfiguration changes to a discovered device can affect the metric families that are associated with the device. The device reconfiguration can update automatically in the monitoring profile to which the metric families are assigned, which

applies to any metric families that the monitoring profile includes. This option is set by default when you create a custom monitoring profile, but it can also be edited at any time.

When the metric family is updated, the data aggregator has an accurate representation of the device configuration. Reports that you generate reflect accurate information.

#### Follow these steps:

1. Select the monitoring profile that you want to update automatically, and then click **Edit**.
2. Leave the **Enable Change Detection** checkbox selected.
3. Set the **Detection Rate** to a value that is greater than zero.

#### NOTE

Consider how frequently the metric family is likely to change, and how many devices the monitoring profile is applied to. Avoid setting change detection rates that are more frequent than necessary.

4. Leave the **Automatically Update Metric Families** checkbox selected.
5. Click **Save**.

When you make a configuration change to a device that is associated with this monitoring profile, the device configuration is updated automatically. When a device configuration is updated, the data aggregator does the following steps:

- Generates an event on the monitored device.
- Identifies new components and creates them.
- Identifies components that are no longer present and marks them as Not Present.  
By default, components that are not present are not synchronized with NetOps Portal. To enable this option, select **Synchronize items that are no longer present on the device** on the **Manage Data Sources** page.
- Identifies existing components that have changed from a previous discovery. The **Name** column changes, if applicable. Historical data is accessible and can be reported on.

#### Update Device Reconfiguration Manually

Reconfiguration changes to a discovered device can affect the metric families that are associated with the device. The device reconfiguration can be updated manually when **Automatically Update Metric Families** is not selected in the associated monitoring profile. View the event logs to identify reconfiguration events for which you want to update the metric families.

When the metric family is updated, the data aggregator has an accurate representation of the device configuration. Reports that you generate reflect accurate information.

#### Follow these steps:

1. View the event logs to identify reconfiguration events for which you want to update the metric families.
2. Click **Monitored Devices** from the **Monitored Inventory** menu for a data aggregator data source.  
The Tree View tab appears.
3. Select **Device by Collection** from the drop-down list, and select the monitored device that was updated from the corresponding tree view.  
The Polled Metric Families tab shows the consolidated monitoring profiles that are associated with a device. Devices only have one consolidated monitoring profile. Each consolidated monitoring profile lists every metric family that can be polled on the device, and whether the device supports the metric family.
4. Select the metric family for which you want to update the configuration and click **Update Metric Family**.  
Your device configuration is updated, and the data aggregator does the following steps:
  - Generates an event on the monitored device.
  - Identifies new components and creates them.
  - Identifies components that are no longer present and marks them as "Not Present".
  - Identifies existing components that have changed from a previous discovery. The Name column changes, if applicable.

Historical data is accessible and can be reported on.

### **Reconfiguration Detection of Ports**

Reconfiguration changes to a discovered device may cause ports to be reconfigured. The new port and the old port are identical if either of the following is true for the Interface metric family after the reconfiguration:

- The PortType and Description attributes in the Interface metric family remain the same.
- The PortType and at least two of the following remain the same:
  - Alias
  - Description
  - MacAddress
  - Index

For all other metric families, the new component and the old component are identical if both Name and Index remain the same after the reconfiguration.

## **Manage Device Life Cycles**

You can manage device life cycles by changing the life cycle state (Active, Retired, or Maintenance) and by defining the behavior for devices that are in the "Maintenance" life cycle state.

A monitored device's life cycle state defines the monitoring behavior for the device, such as whether it is active, retired, or in maintenance. Components inherit the life cycle state of the associated device. You can define the usage state of SNMP and ICMP devices by managing the life cycle.

You can manage the device life cycle state using NetOps Portal or using the NetOps Portal API (the `devices` endpoint). This article describes how to manage the life cycle state using NetOps Portal.

For more information about how to manage the life cycle state using the NetOps Portal API, see [Manage Device Life Cycles Using the Devices Web Service](#).

### **NOTE**

**Limitation:** You cannot manage the life cycle state of DX NetOps Mediation Manager (DX NetOps MM) devices in NetOps Portal.

In this article:

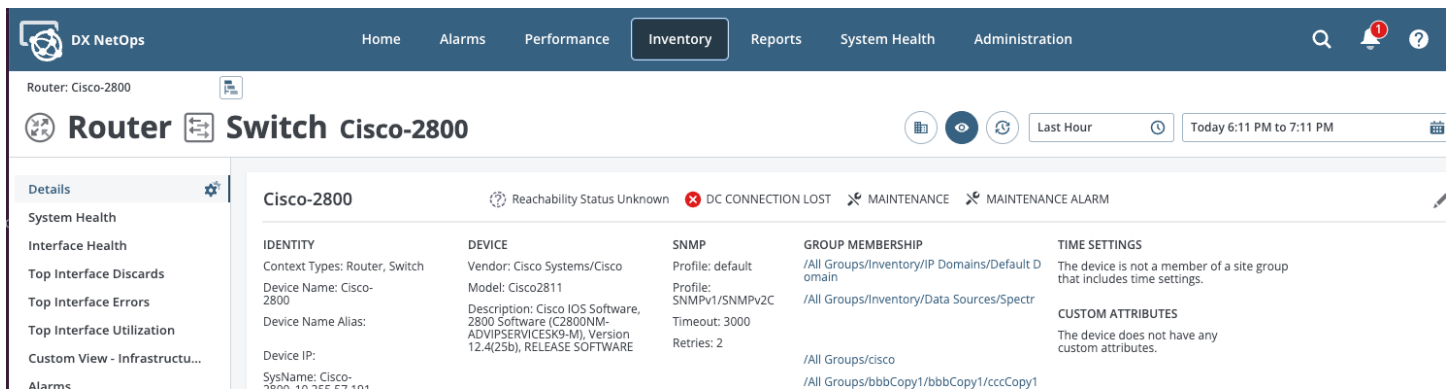
- [View Device Life Cycle State](#)
- [Change the Life Cycle State](#)
- [Specify the Behavior for Devices in Maintenance Life Cycle State](#)
- [Retirement Consolidation](#)

### **View Device Life Cycle State**

Device life cycle state shows when you click a device (the context page for a device), at the top of the page, the third status.

The following example shows an example of a device context page showing the life cycle of the device ("Maintenance"):

### Figure 17: Life Cycle State for a Device



For more information about the other statuses shown on the device context page, see [Reachability Status and Contact Status](#) and [Monitor Device Inventory with Alarm States](#).

## Change the Life Cycle State

To define when a device is active, retired, or in maintenance, change the life cycle state.

## IMPORTANT

- If you have integrated with DX NetOps Spectrum (Spectrum), and you have enabled Spectrum to control the life cycle state of devices in NetOps Portal (the **Synchronize device life cycle state from Spectrum** checkbox is selected for that data source), changes to device states in Spectrum trigger changes to the device life cycle states in NetOps Portal.  
For more information about how to enable Spectrum to control the life cycle state of devices in NetOps Portal, see [Configure a Data Source](#).
- When a change to the life cycle state stops polling a device (the life cycle state changes from "Maintenance" or "Active" to "Retired"), if the data is within the rollup time frame, the polled data is included in rollups. For example, the system stops polling at 9:20 am. The data is included in rollups for the 9:00-10:00 hour, but not for the 10:00-11:00 hour. Baselines use data up to 90 days in the past, even if the device is no longer polled.

**Prerequisites:**

- You have the Administer Life Cycle role right.
- Spectrum is not controlling the life cycle state of devices in NetOps Portal (the **Synchronize device life cycle state from Spectrum** checkbox is cleared for that data source).

**Follow these steps:**

1. Do one of the following steps:
  - Hover over **Administration**, **Monitored Items Management**, and then click **Device Life Cycle**.
  - Hover over **Inventory**, **Items**, and then click **Devices**.The **Devices** page appears.
2. Select the devices that you want to change the life cycle state, and then click **Manage Life Cycle**. The **Manage Device Life Cycle** dialog opens.
3. Select the life cycle state from the **Select Life Cycle State** drop-down, and then click **OK**:
  - **Active**  
Specifies the normal operating status of a device.
  - **Retired**  
Specifies that the device is no longer in use and that no monitoring occurs. This life cycle state disables polling, threshold monitoring, notifications, and change detection. In this life cycle state, NetOps Portal does not update the SNMP profile or change the hostname.

**NOTE**

In this life cycle state, the data collector does not poll virtual devices (Virtual Machine or ESX) discovered with systemEdge for its vCenter statistics.

For more information about the consolidation behavior that occurs in NetOps Portal as a result of a device being put into this life cycle state in Spectrum or NFA, see [the "Retirement Consolidation" section](#).

- **Maintenance**

Specifies that the device is temporarily under maintenance.

**TIP**

You can configure the behavior of this life cycle state.

For more information, see [Configure Maintenance Behavior](#).

**NOTE**

In this life cycle state, NetOps Portal does not update the device during discovery or rediscovery.

NetOps Portal applies the selected life cycle state to the device and logs a life cycle event that marks the change.

### **Specify the Behavior for Devices in Maintenance Life Cycle State**

You can specify the behavior for devices that are in the "Maintenance" life cycle state.

**IMPORTANT**

The changes to the behavior for devices in "Maintenance" life cycle state apply to those devices you put in "Maintenance" life cycle state *after* you save the changes to the behavior. These changes do not apply to devices that are already in the "Maintenance" life cycle state.

#### **Follow these steps:**

1. Hover over **Administration, Configuration Settings**, and then click **Life Cycle Behavior**.  
The **Manage Life Cycle Behavior** page appears.
2. In the **Maintenance State** section, specify the following behavior, and then click **Save**:
  - **Polling**  
Specifies whether DX NetOps Performance Management polls the device.  
**Options:** Disabled or Enabled  
**Default:** Disabled
  - **Threshold Evaluation**  
If polling is enabled, specifies whether DX NetOps Performance Management analyzes event rules for the device.  
**Options:** Disabled or Enabled  
**Default:** Disabled
  - **Event Notification**  
Specifies whether DX NetOps Performance Management sends notifications for the device.  
**Options:** Disabled or Enabled  
**Default:** Disabled

NetOps Portal applies the defined behaviors to devices that you put in "Maintenance" life cycle state.

### **Retirement Consolidation**

The following consolidation behavior occurs in NetOps Portal as a result of a device being retired (put into "Retired" life cycle state) in Spectrum or NFA:

- **Spectrum Retired Devices**  
If you have enabled Spectrum to control the life cycle state of devices in NetOps Portal, when you delete a device in Spectrum, Spectrum changes the life cycle state of that device in NetOps Portal to "Retired".
- **NFA Retired Devices**

You can view replacement devices and devices that are in "Retired" life cycle state coming from NFA in NetOps Portal. Multiple devices with the same IP address display. When a device is in the "Retired" life cycle state, NFA retires the device and creates a new device with an interface. The new and the retired device show in NetOps Portal after synchronization. NetOps Portal can discover the new device using the same IP. You can report on the new and retired device in the same or separate group. To report on the new and old device together, place them in the group. When NFA detects a flow from a different device with the same IP, it changes the life cycle state of the existing device in NetOps Portal to "Retired", and then creates an entry for the new device with a new set of interfaces.

#### NOTE

When you set the life cycle state of a device to "Retired", and then create a new device, as with any new device, you can configure and manage the new device accordingly.

## Manage Hostname Changes

When the hostname of a monitored device changes, NetOps Portal detects the change and updates the hostname for the device item. You can change the hostname change detection rate.

If the device item name uses the hostname, NetOps Portal updates the device name. The default detection rate for hostname changes is 24 hours. However, the change might not appear in NetOps Portal for up to twice the detection period. For example, with the default reevaluation rate of 24 hours, changes to the hostname are updated within 48 hours.

#### NOTE

NetOps Portal looks for hostname changes only on device items with defined DNS hostnames. To add the hostname to a device item, run the discovery profile again.

For more information, see [Run Discovery](#).

### Modify the Hostname Change Detection Rate

The default hostname change detection rate is 24 hours. To reduce DNS traffic or to detect changes faster, modify the reevaluation interval. Each data collector has a separate evaluation interval. To change the interval for a specific device or group of devices, apply the changes to each data collector individually.

#### Follow these steps:

1. Log in to the data collector host that monitors the device.
2. Locate the `<installation_directory>/apache-karaf/etc/com.ca.im.dm.core.collector.cfg` data collector configuration file.
  - **installation\_directory**  
The installation directory for the data collector.  
**Default:** `/opt/IMDataCollector`
3. Add or modify the following line:
 

```
hostname-reevaluation-interval-in-hours=<number>
```

  - **number**  
Specifies the length in hours of the reevaluation interval.
4. Save the file.

The data collector uses the new interval for hostname reevaluation. Changes to the hostname can take up to twice the interval to appear in NetOps Portal.

## Override Device Types

You can override or preserve the device or context types.

Based on the device service information, the data aggregator automatically classifies manageable devices. Classifications include router, switch, server, firewall, load balancer, wireless LAN controller (WLC), and wireless access point (AP).



During synchronization, the data aggregator inspects each device and its classification. It designates a primary device type (Router, Switch, or Server), and forwards any additional types as context types (Firewall, Load Balancer, WLC, and AP). You can view the associated context types, including the primary device types, on the device inventory tables in NetOps Portal.

In this article:

- [Device Type Classifications](#)
- [Override Device Types](#)
- [Example: Map a Device sysObjectID to Another Device Type](#)
- [Preserve Device Types](#)
- [Example: Preserve the Device Types that are already Mapped to the sysObjectIDs](#)

## **Device Type Classifications**

Based on the device service information, the data aggregator automatically classifies devices of type Router, Switch, and Server as manageable. Also, where applicable, the data aggregator associates managed devices with context types. The associated context types, including device types, appear listed in a column within device inventory tables.

The following list describes more device type classifications:

- **Pingable**  
The data aggregator classifies the device type as *pingable* when it *cannot* classify the device type as manageable. For example, the device does not respond to SNMP requests and the data aggregator cannot determine its device type.
- **Manageable**  
The data aggregator classifies the device type as *manageable* when the following criteria are true:
  - It *can* associate the device with a Firewall, Load Balancer, WLC, or AP context type.
  - It *cannot* identify the device type as Router, Switch, or Server.
- **Other**  
The data aggregator classifies the device type as *other* when it *cannot* identify the device type as Router, Switch, Server, or when it *cannot* associate the device with a Firewall, Load Balancer, WLC, or AP context type.

### **TIP**

If the data aggregator does not classify the type of some SNMP-managed devices as expected, you can override the device type and context type.

## **Override Device Types**

If the data aggregator does not classify the device type of some SNMP-manageable devices as expected, you can override device or context type.

### **Scenarios:**

- The data aggregator discovers an existing device as *other* (classification). However, you want the device to be classified as router instead. To override the existing other classification, you can add router to the `DeviceTypes.xml` file.
- The data aggregator discovers an existing device as type Switch. The SNMP agent advertised the wrong type. The device should have been advertised as router and it is also a firewall. To override the incorrectly-discovered types, add the type Router and context type Firewall to the `DeviceTypes.xml` file.

Map the device `sysObjectID` MIB value explicitly to the correct device type in the `DA_installation_directory/data/custom/devicetypes/DeviceTypes.xml` file on the data aggregator host.

In a fault-tolerant environment, a shared directory (for example, `/DASharedRepo`) is defined to help limit data loss. Therefore, in a fault-tolerant environment, map the device in the `DASharedRepo/custom/devicetypes/DeviceTypes.xml` file.

For more information, see [Fault Tolerance](#).

#### NOTE

You cannot add new device types to the `DeviceTypes.xml` file.

The `DeviceTypes.xml` file contains a template to map the `sysObjectID` to appropriate device types. By default, the file does not contain any `sysObjectID`-to-type mapping entry. To classify a device type with a particular `sysObjectID`, modify the template to add the `sysObjectID`-to-type entries into the file. Before you add a `sysObjectID`, uncomment the section where you are adding the `sysObjectID`.

#### NOTE

Updates to the `DeviceTypes.xml` file can take up to one minute to apply.

The data aggregator can classify devices into multiple device types. However, the Device device type is mutually exclusive to other device types. For example, if you add a `sysObjectID` to one or more of the Router, Switch, or Server device types, and you also add that `sysObjectID` to the Device device type, the data aggregator drops the Device device type and does not recognize this device type.

Fully override the device types when the following criteria are true:

- The data aggregator does not find automatically-discovered device types, or the device types it discovers are incorrect.
- You want to override the device types that the data aggregator discovers during device discovery.
- You want to assign the correct device types manually.

### **Example: Map a Device sysObjectID to Another Device Type**

#### IMPORTANT

The `sysObjectID`-to-device type mappings in the following example override any already discovered device types.

#### **Follow these steps:**

1. Open the `DA_installation_directory/data/custom/devicetypes/DeviceTypes.xml` file.  
In a fault-tolerant environment, open the `DASharedRepo/custom/devicetypes/DeviceTypes.xml` file.

#### **– DA\_installation\_directory**

The default installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

2. Add the following information for the XML file, and then save your changes:

```
<DeviceType>
  <Routers>
    <sysObjectID>1.3.6.5.1.34</sysObjectID>
  </Routers>

  <Switches>
    <sysObjectID>1.3.6.5.5.3</sysObjectID>
    <sysObjectID>1.3.6.5.1.34</sysObjectID>
  </Switches>

  <Servers>
    <sysObjectID>1.3.6.5.567.1</sysObjectID>
  </Servers>

  <Devices>
    <sysObjectID>1.3.6.5.49.1</sysObjectID>
```

```
</Devices>
</DeviceType>
```

- Run discovery on the discovery profile that contains the devices.

#### NOTE

The changes that you make to the `DeviceTypes.xml` file take effect on existing devices only *after* you rerun discovery.

The following results occur:

- The data aggregator classifies the devices that have a `sysObjectID` of 1.3.6.5.1.34 as a device type of Router and Switch.
  - The data aggregator classifies the devices that have a `sysObjectID` of 1.3.6.5.5.3 as a device type of Switch.
  - The data aggregator classifies the devices that have a `sysObjectID` of 1.3.6.5.567.1 as a device type of Server.
  - The data aggregator classifies the devices that have a `sysObjectID` of 1.3.6.5.49.1 as a device type of Device.
- Verify your changes by going to the following locations:
    - Hover over **Administration, Monitored Items Management**, and then click **Device Life Cycle**.
    - Hover over **Inventory, Items**, and then click **Devices**.

### Preserve Device Types

Sometimes, a device might have an incomplete classification. For example, a device that the data aggregator identifies as device type Router might also be of device type Firewall. In this case, you want to preserve the identified Router device-type classification. You also want to augment the device-type classification with one or more manually-assigned device types (in this example, Firewall).

#### Scenarios:

- The data aggregator discovers an existing device automatically as a Server. You want to classify the device as a 'Firewall' as well. To preserve the discovered Server device classification, add Firewall to the `DeviceTypes.xml` file, and set the `sysServiceOverride` tag attribute set to `false`.
- The data aggregator discovers an existing device automatically as a Router. You want to classify the device as a Firewall and Switch as well. To preserve the discovered Router device classification, add Firewall and Switch to the `DeviceTypes.xml` file, and set the `sysServiceOverride` tag attribute set to `false`.

You can disable the default override behavior, or preserve the device types, using the `sysServiceOverride` tag attribute. You can assign a `sysObjectID` to one or more device types in the `DeviceTypes.xml` file (for example, Firewall) in addition to the types that the data aggregator automatically discovers (for example, Router).

To preserve the device types, set the `sysServiceOverride` tag attribute to `false` for the `sysObjectIDs` or device types.

Use the `sysServicesOverride` tag attribute when the following criteria are true:

- You want to preserve the type that is discovered during device discovery.
- You accept the automatically discovered device type as correct and do not want to lose it.
- You want to add more types.

### Example: Preserve the Device Types that are already Mapped to the sysObjectIDs

#### Follow these steps:

- Open the `DA_installation_directory/data/custom/devicetypes/DeviceTypes.xml` file.  
In a fault-tolerant environment, open the `DASharedRepo/custom/devicetypes/DeviceTypes.xml` file.
  - **DA\_installation\_directory**  
The default installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`

## 2. Enter the following information:

```
<DeviceType>
  <Routers>
    <sysObjectID sysServicesOverride="false">1.3.6.5.1.34</sysObjectID>
  </Routers>

  <Firewalls>
    <sysObjectID sysServicesOverride="false">1.3.6.1.4.1.8072.3.2.10</sysObjectID>
  </Firewalls>

  <Firewalls sysServicesOverride="false">
    <sysObjectID>1.3.6.1.4.1.9.1.522</sysObjectID>                                <!--
cat6500FirewallSm -->
    <sysObjectID>1.3.6.1.4.1.2620.1.6.123.1.56</sysObjectID>                    <!-- Check Point
21800 -->
    <sysObjectID>1.3.6.1.4.1.2620.1.6.123.1.16</sysObjectID>                    <!-- Smart-1 150
-->
  </Firewalls>

  <WirelessAccessPoints>
</WirelessAccessPoints>

  <WirelessControllers>
    <sysObjectID>1.3.6.1.4.1.9.1.770</sysObjectID>
  </WirelessControllers>

  <LoadBalancers>
    <sysObjectID>1.3.6.1.4.1.6527.1.6.1</sysObjectID>
  </LoadBalancers>

</DeviceType>
```

## 3. Run discovery on the discovery profile that contains the devices.

### NOTE

The changes that you make to the `DeviceTypes.xml` file do not take effect on existing devices until you rerun discovery.

The data aggregator classifies devices as follows:

- It classifies devices that have a `sysObjectID` of 1.3.6.5.1.34 as a device type of Router. The device types that are already mapped to the `sysObjectID` are preserved.
  - It classifies devices that have a `sysObjectID` of 1.3.6.1.4.1.8072.3.2.10 as a context type of Firewall. The device types that are already mapped to the `sysObjectID` are preserved.
  - It classifies devices that have a `sysObjectID` of 1.3.6.1.4.1.9.1.522 as a context type of Firewall. The device types that are already mapped to the `sysObjectID` are preserved.
  - It classifies devices that have a `sysObjectID` of 1.3.6.1.4.1.2620.1.6.123.1.56 as a context type of Firewall. The device types that are already mapped to the `sysObjectID` are preserved.
  - It classifies devices that have a `sysObjectID` of 1.3.6.1.4.1.2620.1.6.123.1.16 as a context type of Firewall. The device types that are already mapped to the `sysObjectID` are preserved.
  - It classifies devices that have a `sysObjectID` of 1.3.6.1.4.1.2620.1.6.123.1.48 as a context type of AP. The device types that are already mapped to the `sysObjectID` are preserved.
  - It classifies devices that have a `sysObjectID` of 1.3.6.1.4.1.9.1.770 as a context type of WLC. The device types that are already mapped to the `sysObjectID` are preserved.
  - It classifies devices that have a `sysObjectID` of 1.3.6.1.4.1.6527.1.6.1 as a context type of Load Balancer. The device types that are already mapped to the `sysObjectID` are preserved.
4. Verify your changes by going to the following locations:
- Hover over **Administration, Monitored Items Management**, and then click **Device Life Cycle**.
  - Hover over **Inventory, Items**, and then click **Devices**.

## Device Deduplication

Without deduplication, NetOps Portal models a single physical device as multiple device items when it discovers them from multiple sources (Simple Network Management Protocol (SNMP), DX NetOps Mediation Manager (DX NetOps MM), DX NetOps Virtual Network Assurance (VNA), or another source). Deduplication models a single device in NetOps Portal.

The following general rules apply to deduplication:

- NetOps Portal deduplicates devices only when they are in the same IP domain.
- NetOps Portal successfully deduplicates devices regardless of discovery order.
- When the deduplication criteria of a device (IP address, hostname, and so on) from one source matches an existing device item from another source, DX NetOps Performance Management does not create a new device item. Instead, the same existing device item is used for both sources.
- If attributes conflict and SNMP is a source, the attributes from SNMP take precedence.

Exceptions to these rules are covered in the following sections.

### SNMP Deduplication

NetOps Portal does not create SNMP devices in the following scenarios:

- The primary IP address matches an existing device. The existing device could be a DX NetOps MM or VNA device with the same primary IP address.

#### NOTE

The primary IP address is the IP address that NetOps Portal uses to monitor a device. When it first discovers a device with the IP ranges discovery profile, NetOps Portal tries to use the IP address that maps to the hostname as the primary IP address.

- The hostname matches an existing device.
- The primary IP address for a new device is in the IP address list of an existing device, and the primary IP address of the existing device is in the IP address list of the new device. If the primary IP address of the new device is listed as a single IP address in the discovery profile, `sysName` is verified. If the `sysName` of the new and existing device match, the new device is not created.
- For devices that support the Device Unique Identifier metric family, the `UniqueID` matches an existing device.

## SNMP and DX NetOps MM Deduplication

DX NetOps MM monitors devices that do not support SNMP or provides metrics that are not accessible through SNMP polling. DX NetOps MM injects data into the data aggregator through one of the data collectors. When the DX NetOps MM device pack provides an IP address or hostname that matches another device in the system, DX NetOps Performance Management deduplicates the devices to a single item.

The following table illustrates that DX NetOps Performance Management deduplicates SNMP and DX NetOps MM devices when the IP address or hostname match:

IP Address Match	Hostname Match	Deduplication
Yes	Yes	✓
Yes	No	✓
No	Yes	✓
No	No	

### Deduplicate Devices by the IP Address Only

To deduplicate devices by the IP address only, edit the discovery profile to exclude hostname.

For more information, see [Discovery Profiles](#).

The following criteria must be met for deduplication to occur:

- Both the data collector and the DX NetOps MM Local Controller are in the same IP domain and poll the same device.
- Either a matching IP address or a matching hostname is present in both SNMP and DX NetOps MM.
- You can deduplicate DX NetOps MM devices for device packs.

#### IMPORTANT

Deduplication with DX NetOps MM devices occurs only for device packs that support deduplication. To determine which device packs support deduplication, see the information file of each device pack.

- For multiple devices packs, the device names on each device pack match.
- DX NetOps MM and SNMP discovery is successful.

DX NetOps Performance Management deduplicates devices only when the data collector and the Local Controller and are in the same IP domain. However, the data collector and the Local Controller can reside on different servers.

### Supported Server Configurations

The following server configurations are supported:

- The data collector server for SNMP polling
- The Local Controller server for DX NetOps MM polling
- The data collector and Local Controller on the same server for SNMP and DX NetOps MM polling

If the combined workload is manageable, you can install the Local Controller on the same server as the data collector for DX NetOps MM polling. However, ensure that your servers have enough capacity to continue operating with the normal DX NetOps MM polling load and the DX NetOps MM requirements.

**Best Practice:** For new installations, install the Local Controller on a data collector in the same IP domain with the devices that you want to monitor. Devices that are discovered through DX NetOps MM and through SNMP are deduplicated. For existing installations, rediscover the devices with SNMP or DX NetOps MM. DX NetOps MM deduplication does not delete any existing duplicated device items. The historical data is still available on the existing device items. Historical data that was captured with the device model from an old IP domain still exists, but is unconnected to the newly reconciled device. Delete or retire the device from the old IP domain to avoid double polling.

from DX NetOps MM. If the device from the old IP domain is not deleted, or the historical data is not aged out, DX NetOps Performance Management continues to have two devices.

NetOps Portal does not reconcile DX NetOps MM components and SNMP components from the same metric family with each other. It delays SNMP component discovery until the next change detection. An error occurs when all the following conditions are true:

- SNMP and DX NetOps MM polled devices in the same IP domain.
- The SNMP and DX NetOps MM devices were polled by the same data collector.
- The SNMP and DX NetOps MM device packs contribute to the same metric family.
- One of the two contributors (DX NetOps MM or SNMP) have late arriving data (greater than 30 minutes after the rollover period ends). Late arriving data is more likely to happen with DX NetOps MM data, especially with a 15-minute polling interval.

When the load on the data collectors are rebalanced, NetOps Portal does not rebalance the DX NetOps MM and VNA devices. Therefore, NetOps Portal does not rebalance the deduplicated SNMP, DX NetOps MM, and VNA devices.

For more information, see [Rebalance the Load on the Data Collectors](#).

## Delete Components That Are Not Present

For best NetOps Portal performance, delete the components that are not present on physical devices. Excessive numbers of components that are marked not present can impact performance.

You can delete the components that are marked not present using the following methods:

- Configure the internal not-present component cleanup process.
- Manually delete the components, using cURL commands.
- Script the deletion of the components.

### Configure the Internal Not-Present Component Cleanup Process

The data aggregator deletes those components that are marked as not present and that have never collected data nightly at 2 am. But there can be many components that are marked not present with data that the data aggregator does not delete by way of this cleanup process.

Configuring the internal not-present component cleanup process defines when it deletes components that are marked not present that have historical data, and have stopped polling (it has no new data). For example, to have the cleanup process delete components marked non-present and then have stopped polling for 30 days, configure it to check for components that have stopped polling for the last 30 days before it deletes the component.

#### Follow these steps:

1. Create the `<installation_directory>/apache-karaf/etc/com.ca.im.dm.core.database.dao.impl.NotPresentItemCleanupDAO.cfg` file.

##### – **installation\_directory**

The default installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

2. Add the following line to the file:

```
ageOutDays=<days>
```

##### **Example:**

```
ageOutDays=30
```

##### – **days**

Defines when (in number of days) the cleanup process deletes components that are marked not present and that have stopped polling before it deletes the component.

The internal not-present component cleanup process is configured.

### **Manually Delete Components that are Marked Not Present**

You can filter for the not-present components associated with a device manually using cURL commands.

Use the following process to delete components that are marked not present:

1. Take note of the IP address for the device with which the components that are marked not present are associated. Use the IP address to determine the device item ID for the device associated to the components that are marked not present.

#### **NOTE**

Filtering by IP address is a two-step process because you cannot do a direct component filter by IP address.

2. Use the device item ID to determine the list of components that are marked not present to delete.
3. Delete the components that are marked not present.

The following procedure uses the cURL command, but you can use any command with which you are familiar.

#### **Follow these steps:**

1. Create the `filterDeviceIP.xml` file, listing the primary IP address for which to filter the device that includes the components that are marked not present:

#### **Example:**

```
<FilterSelect xsi:noNamespaceSchemaLocation="filter.xsd" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance">
  <Filter>
    <And>
      <Device.PrimaryIPAddress type="EQUAL">10.252.1.1</Device.PrimaryIPAddress>
    </And>
  </Filter>
</FilterSelect>
```

2. Return the device item ID for the device with the primary IP address listed in the `filterDeviceIP.xml` file by issuing the following command with the `-u`, `-X`, `-H`, and `-T` options:

```
curl -u <user> -X <filter> -H "<content-type>" -T "<file>" > returnedDeviceID.xml
```

#### **Example:**

```
curl -u admin -X http://myDaHost:8581/rest/devices/
filtered
-H "Content-Type: application/xml" -T "filterDeviceIP.xml" > returnedDeviceID.xml
```

#### **– user**

Indicates the data aggregator REST user to use.

**Example:** admin

#### **– filter**

Creates the filter that you indicate.

**Example:** http://myDaHost:8581/rest/devices/filtered

#### **– content-type**

Indicates the content type of the file that you are posting.

**Example:** Content-Type: application/xml

#### **– file**

Indicates the file that you are posting.

**Example:** filterDeviceIP.xml

The following result is returned as an HTTP response:



```

<?xml version="1.0"?>
  <DeviceList>
    <Device version="1.0.0">
      <ID>107881</ID>
      <PrimaryIPAddress>10.252.1.1</PrimaryIPAddress>
      <supportsOnDemandMFDDiscovery>true</supportsOnDemandMFDDiscovery>
      <SupportedProtocolsList>
        <SupportedProtocols>ICMP</SupportedProtocols>
      </SupportedProtocolsList>
      <DiscProfileID>107503</DiscProfileID>
      <HostName>rtp003723rts.ca.com</HostName>
      <RelatesTo>
        <MonitoredGroupIDList relatesURL="relatesto/monitoredgroups"
rootURL="monitoredgroups">
          <ID>509</ID>
        </MonitoredGroupIDList>
        <GroupIDList relatesURL="relatesto/groups" rootURL="groups">
          <ID>547</ID>
          <ID>530</ID>
          <ID>509</ID>
        </GroupIDList>
      </RelatesTo>
      <IsAlso>
        <IsA name="MetricFamilyDiscoveryHistory" rootURL="devices/
mfdiscoveryhistory"/>
        <IsA name="AccessibleDevice" rootURL="devices/accessible"/>
        <IsA name="Syncable" rootURL="syncable"/>
        <IsA name="IPDomainMember" rootURL="ipdomainmember"/>
      </IsAlso>
      <DataCollectionMgrId version="1.0.0">
        <DcmID>dcname.ca.com:8f53bc55-f442-42fc-9bd5-a907d0261421</DcmID>
      </DataCollectionMgrId>
      <Syncable version="1.0.0">
        <SyncID>-1</SyncID>
      </Syncable>
      <Item version="1.0.0">
        <DisplayName>router.ca.com</DisplayName>
        <CreateTime>Wed Feb 05 10:20:26 EST 2014</CreateTime>
        <Name>router.ca.com</Name>
      </Item>
      <IPDomainMember version="1.0.0">
        <IPDomainID>2</IPDomainID>
      </IPDomainMember>
      <DeviceMonitoringProfile version="1.0.0">
        <ConsolidatedMonitoringProfile>2509</ConsolidatedMonitoringProfile>
      </DeviceMonitoringProfile>

```

```

    </Device>
</DeviceList>

```

In this example, the device with the device item ID of 107881 fits the filter criteria (it has the primary IP address of 10.252.1.1 ), and is returned. Detailed information about the device is also returned in the results.

3. Create the `filterNotPresent.xml` file, listing the device item ID for the device associated to the components that are marked not present for which to filter:

**Example:**

```

<FilterSelect xsi:noNamespaceSchemaLocation="filter.xsd" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance">
  <Filter>
    <And>
      <DeviceComponent.DeviceItemID type="EQUAL">107881</
DeviceComponent.DeviceItemID>
    </And>
  </Filter>
  <Select use="exclude">
    <Item use="exclude">
      <DisplayName use="include"/>
    </Item>
  </Select>
</FilterSelect>

```

4. Return a list of the components that are marked not present associated with the device item ID listed in the `filterNotPresent.xml` file by issuing the following command with the `-u`, `-X`, `-H`, and `-T` options:

```

curl -u <user> -X post <filter> -H "<content-type>" -T "<file>" >
  returnedNotPresentItems.xml

```

**Example:**

```

curl -u admin -X post http://myDaHost:8581/rest/retired/filtered -H "Content-Type:
application/xml" -T "filterNotPresent.xml" > returnedNotPresentItems.xml

```

– **user**

Indicates the data aggregator REST user to use.

**Example:** admin

– **filter**

Creates the filter that you indicate.

**Example:** http://myDaHost:8581/rest/retired/filtered

– **content-type**

Indicates the content type of the file that you are posting.

**Example:** Content-Type: application/xml

– **file**

Indicates the file that you are posting.

**Example:** filterNotPresent.xml

The following result is returned as an HTTP response:

```

<?xml version="1.0"?>
  <RetiredList>
    <Retired version="1.0.0">
      <ID>128452</ID>
      <Item version="1.0.0">
        <DisplayName>GigabitEthernet0/239 - GigabitEthernet0/239</DisplayName>
      </Item>
    </Retired>
  </RetiredList>

```

```

        </Item>
    </Retired>
    <Retired version="1.0.0">
    <ID>128451</ID>
        <Item version="1.0.0">
            <DisplayName>GigabitEthernet0/238 - GigabitEthernet0/238</DisplayName>
        </Item>
    </Retired>
</RetiredList>

```

In this example, the components with the component IDs of 128452 and 128451 fit the filter criteria (they are associated with the device with the device item ID of 107881 and are not present), and are returned.

5. Create the `deleteNotPresentList.xml` file, listing the component IDs of the not-present components to delete.

**Example:**

```

<DeleteList>
    <ID>128452</ID>
    <ID>128451</ID>
</DeleteList>

```

6. Delete the components that are marked not present that are listed in the `deleteNotPresentList.xml` file by issuing the following command with the `-u`, `-X`, `-H`, and `-T` options:

```

curl -u <user> -X post <filter> -H "<content-type>" -T "<file>" >
    deletelistreponse.xml

```

**Example:**

```

curl -u admin -X post http://myDaHost:8581/rest/retired/deletelist -H "Content-Type:
    application/xml" -T "deleteRetiredList.xml" > deletelistreponse.xml

```

– **user**

Indicates the data aggregator REST user to use.

**Example:** admin

– **filter**

Creates the filter that you indicate.

**Example:** http://myDaHost:8581/rest/retired/deletelist

– **content-type**

Indicates the content type of the file that you are posting.

**Example:** Content-Type: application/xml

– **file**

Indicates the file that you are posting.

**Example:** deleteRetiredList.xml

The following result is returned as an HTTP response:

```

<?xml version="1.0"?>
    <DeleteListResult>
        <DeleteResult>
            <ID>128452</ID>
            <Error>SUCCESS</Error>
        </DeleteResult>
        <DeleteResult>
            <ID>128451</ID>
            <Error>SUCCESS</Error>
        </DeleteResult>
    </DeleteListResult>

```

```
</DeleteListResult>
```

The components that are marked not present (the components with the component ID of 128452 and 128451 ) are deleted.

### **Script the Deletion of the Components that are Marked Not Present**

You can delete the components that are marked not present from your network using the `remove_not_present_items.sh` script. This script has two parts: the first part identifies and returns data about the components, which is based on the filter that you set, and the second part issues the delete of the component list.

#### **Follow these steps:**

1. Review the components that are marked not present where the total exceeds 100,000 by completing the following steps:
  - a. From an open command prompt, access the `<installation_directory>/scripts` directory.
    - **installation\_directory**  
The default installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
  - b. Display a list of the components that are marked not present and log the information to a `.csv` file by issuing the following command with the `-h` and `-o` options:
 

```
./remove_not_present_items.sh -h <da_url> -o <ouputfile>
```

**Example:**

```
./remove_not_present_items.sh -h http://my_hostname:8581 -o myOutputFile
```

    - **da\_url**  
The URL for the data aggregator, including protocol, fully qualified host name for the data aggregator, and port.  
**Example:** `http://my_hostname:8581` or `https://my_hostname:8582`

**NOTE**  
If you filter the components by IP domain name or IP domain ID, also specify a specific IP address to return the correct results.

    - **ouputfile**  
The output file name to report the components that are marked not present for all devices. The output is a `.csv` file.

A list of components that are marked not present is displayed.

**NOTE**  
If your filter criteria return too many components, the REST interface cannot return a response. Use other filtering options to narrow the results. More filter criteria are available at the following URL:  
`http://<hostname>:<port>/rest/retired/xsd/filterselect.xsd`
2. Delete the components that are marked not present and log the information to a `.csv` file by issuing the following command with the `-h`, `-o`, and `-c` options:
 

```
./remove_not_present_items.sh -h <da_url> -o <ouputfile> -c <deletion_option>
```

**Example:**

```
./remove_not_present_items.sh -h my_host_name -o mynotpresent.csv -c Yes
```

  - **da\_url**  
The URL for the data aggregator, including protocol, fully qualified host name for the data aggregator, and port.  
**Example:** `my_host_name`
  - **ouputfile**  
The output file name to report the components that are marked not present for all devices. The output is a `.csv` file.  
**Example:** `mynotpresent.csv`
  - **deletion\_option**  
Defines whether to delete the components that are marked not present.

**Options:** Yes or No

The components that are marked not present are deleted.

## Delete Devices

To remove a device item, all associated components, and historic data from the system, delete the device.

Deleting a device:

- Deletes all the associated device components.
- Removes access to historical data on the deleted devices and device components.

### TIP

You can stop the data collector from polling the device and *retain* historical data by changing the device life cycle state to "Retired".

For more information, see [Manage Device Life Cycles](#).

If discovery finds a deleted device, NetOps Portal creates a *new* device item. The new device item has no association with the old device item.

### Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Monitored Devices**. The **Monitored Devices** page opens.
2. (Optional) Search for the device to delete.

### NOTE

Do not use the global Search box at the top of the page.

3. Select the devices to delete, and then click **Delete**.
4. In the confirmation dialog, click **Yes**.  
The devices are deleted and are removed from the **Monitored Devices** inventory. Historical data for the devices is deleted from the database. After the next synchronization cycle, the devices are removed from inventory and groups in NetOps Portal.

### NOTE

If another data source is managing these devices, the devices continue to appear on the **Devices** page and in groups.

## Manage Metric Families

Metric families show their association with device collections, vendor certifications, and monitoring profiles.

You can control how to monitor your devices by understanding the relationships between metric families, device collections, and device types. You can also determine whether you need more metric families to monitor your environment sufficiently.

### Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Monitored Devices**. The **Monitored Devices** page appears.
2. Under **Monitoring Configuration**, click **Metric Families**. The **Metric Families** page appears.
3. To sort the metric family columns, click the heading columns as needed.
4. Select the metric family that you want to manage from the list.
5. Click a tab to get more information:
  - **Metrics** tab

View the metrics that are included in the selected metric family and various properties for each metric:

- **Name**
- **Polled**
- **Min**
- **Max**
- **Percentiles**
- **Projections (Days)**
- **Projections Percentile**
- **Baseline**
- **Rollup Strategy**
- **Standard Deviation**
- **Vendor Certification Priorities** tab  
View a list of device collections that are associated with the selected metric family. Typically, a metric family is associated with a single device collection. When you select a device collection, a prioritized list of MIB sources (vendor certifications) appears. This information shows the order in which the vendor certifications are applied to that device collection for the metric family.
- **Monitoring Profiles** tab  
View a list of the associated monitoring profiles and their poll rates.

## Configure Metric Filtering

Metric filtering reduces the storage footprint of collected performance data by limiting the metrics NetOps Portal collects from a metric family for a particular monitoring profile. For devices in associated collections, NetOps Portal loads only the selected metrics.

Administrators in the default tenant can manage (add or edit) monitoring profile filters. You can apply metric filtering to all monitoring profiles except the out-of-the-box monitoring profiles (they are locked, as indicated by a lock icon). By default, the metric families that you assign to a monitoring profile collect all the metrics for that metric family. In many monitoring environments, users do not view these metrics. As a best practice, remove all metrics, and then add the metrics you require for your environment.

The following video shows how to create a metric filter to define exactly which metrics to monitor for each metric family within a monitoring profile and eliminate collection data for irrelevant metrics:

### NOTE

Metric filtering does not reduce SNMP traffic. NetOps Portal still polls filtered metrics. Metric filtering applies only to metrics that NetOps Portal collects through SNMP.

You can filter metrics while managing dashboards or configuring event rules for thresholds. Consider dashboard and threshold requirements before metric filtering.

The following example helps to describe the behavior for multiple monitoring profiles.

### Example:

- One monitoring profile (A) has no metric filtering.
- Another monitoring profile (B) has some metric filtering to collect only BitsIn, BitsOut.
- Each monitoring profile is associated with a different collection.

**Result:** If a device ends up in both device collections, the configuration from (B) takes priority. The system only collects BitsIn, BitsOut.

### NOTE

If a device is in a device collection assigned to multiple monitoring profiles, NetOps Portal collects and saves all selected metrics from all monitoring profiles for the device at the fastest rate among all monitoring profiles.

Storage savings are not linear.

For more information about how to estimate the storage savings, see [Configure Data Retention Rates](#).

#### TIP

You can view the metrics that NetOps Portal collects for a particular device from the **Metrics** tab for a metric family. You can view where NetOps Portal collects the metric. This tab represents consolidated information across all monitoring profiles assigned to all collections to which the device belongs. If you see a metric that you do not want to collect, view the metric families assigned to each monitoring profile for each collection.

For more information about the **Metrics** tab, see the [Manage Metric Families](#).

#### Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Monitoring Profiles**. The **Monitoring Profiles** page appears. The **Monitoring Profiles** tab is selected under the **Monitoring Configuration** section.
2. Select the custom monitoring profile for which you want configure metric filtering. The **Metric Families** tab is selected by default.
3. Select the metric family that you want to filter, and then click **Edit Collected Metrics**. The **Add / Edit Collected Metrics** dialog opens.
4. Move the metrics that you want to collect to the **Selected** pane, and then click **Save**.

#### TIP

By default, all metrics are selected. To pick a limited set of metrics, select all metrics, and move the metrics to the **Available** pane. Then move the metrics to collect to the **Selected** pane.

NetOps Portal loads only the selected metrics to the data repository for devices in device collections assigned to the monitoring profile.

#### NOTE

If the collection is assigned to another monitoring profile that includes the same metric family, or if the devices are part of another device collection assigned to a monitoring profile with the same metric family, NetOps Portal still collects the metrics.

## Edit a Metric

You can edit metric properties for any metric in a specific metric family using NetOps Portal. When you edit a metric, the metric family is extended automatically with the new values.

For more information about how to extend metrics and metric families through REST services, see [Create or Extend Metric Families](#).

#### NOTE

You can only edit metrics that DX NetOps Performance Management is polling (the value in the **Polled** column on the **Metrics** tab is set to "True").

For more information about how to view this column, see [Manage Metric Families](#).

#### Follow these steps:

1. Hover over **Administration, Data Sources**, and then click the data aggregator data source. The **Monitored Devices** page appears.
2. Under **Monitoring Configuration**, click **Metric Families**. The **Metric Families** page appears. This page displays a list of metric families, including out-of-the-box and custom metric families.
3. Select the metric family from the list related to the metric that you want to edit.
4. From the **Metrics** tab, click the polled metric that you want to edit, and then click **Edit** at the bottom of the page. The **Edit Metric** dialog opens.
5. Complete the following fields, and then click **Save**:

- **95th Percentile:** Enabled or Disabled
- **Percentiles 2 and 3:** 0-99, except 95
- **Projections 1-3:** 0-730
- **Projection Percentile:** 0-99

**WARNING**

Changes to the value for **Projection Percentile** can cause inaccurate projections for up to 90 days. Changes to the value for **Percentile 2** and **Percentile 3** cause a gap in trend views. NetOps Portal recalculates the changes that you make to these percentile values after several days.

For more information, see [Percentiles](#) and [Metric Projection](#).

The selected metric is updated in the parent metric family. The metric family is marked as extended and the Last Modified and Last Modified By fields are updated. The new data is available for reporting within several poll cycles.

## Populate Components List for Response Path Metric Family

The Monitored Devices page can display a Supported status for a response path metric family on the Polled Metric Families tab, but the Components list for that metric family can be empty. To populate the Components list, update the Metric Family.

**Follow these steps:**

1. Verify that the device is configured to run the test that is associated with the metric family.  
For example, if the "Response Path Test DHCP" metric family shows no components, verify that the device is configured to run the DHCP test.
2. Select the row for the metric family, and click Update Metric Family.

## Rediscover Metric Families

Ensure that you have the latest metric family and component support information for devices by rediscovering the metric families for the device.

When you have NetOps Portal rediscover, or update, the metric families for a device, NetOps Portal updates the following:

- The status of each metric family the data aggregator has polled (the **Status** field), such as whether it is newly or no longer supported (the **Status** is "Supported" or "Not Supported"). This includes updating the date and time that the data aggregator polled the metric family on the device (the **Last Update Time** field).
- The components listed in the **Components** view for the metric families that the data aggregator polled. The status for components that NetOps Portal no longer finds during the rediscovery (the **Status** field) changes to "Not Present".

**Follow these steps:**

1. Log in to NetOps Portal as a tenant administrator.
2. Hover over **Administration**, **Monitored Items Management**, and then click **Monitored Devices**.  
The **Monitored Devices** page appears. The **Monitored Devices** list displays a list of the devices that NetOps Portal has discovered.
3. From the **Monitored Devices** list, with the **Tree View** tab selected, sort by **Device by Collection**, expand the **All Manageable Devices** collection, locate and select the device corresponding to the metric families that you want to rediscover.
4. Click the **Polled Metrics Families** tab.  
A list of the metric families that the data aggregator has polled for the device displays. The following image shows an example of this tab for a selected device:

**Figure 18: Discovery History Log**

5. Do *one* of the following:



- Have NetOps Portal rediscover how the metric families are applied to the device by clicking **Update Metric Families**.

The **Update Metric Families** dialog opens.

- Have NetOps Portal rediscover how a metric family is applied to the device *and generate a discovery log* by completing the following:

#### IMPORTANT

Generate a discovery log when directed by Broadcom Support. This information is highly technical and might require assistance from Support to understand what is happening.

- From the list, select the row for a supported metric family (the **Status** is "Supported").
- Click **Update Metric Family with Logging**.

The **Update Metric Family With Logging** dialog opens.

When rediscovery completes, the discovery log appears in the **Discovery History Log** section of the **Details** tab. Only the most recent log is available.

- Have NetOps Portal rediscover how a metric family is applied to the device *without generating a discovery log* by completing the following:

- From the list, select the row for a supported metric family (the **Status** is "Supported").
- Click **Update Metric Family**.

The **Update Metric Family** dialog opens.

- At the prompt, confirm your selection by clicking **Yes**.

NetOps Portal rediscover, or updates, the metric family (or metric families) for the device.

#### NOTE

If you chose the option to rediscover a selected metric family and generate a discovery log, NetOps Portal has rediscovered the metric family for the device, and the discovery log does not display, click **Refresh**.

## Manage Aggregated Components

Aggregated components are a way for you to manage combined metric values from similar components, either on the same device or on different devices, rather than managing individual values.

You can manage aggregated components by creating them, editing existing aggregated components, or by deleting them. Administrators can manage aggregated components using NetOps Portal or by way of the `aggregatedcomponents` endpoint for the Data Aggregator REST web service or use this API in your scripts for managing aggregated components. This topic describes how to manage them using NetOps Portal.

For more information about how to automate aggregated component management by way of the web service, see [Automate Configuring Aggregated Components](#).

You can do the following:

- [View a List of Aggregated Components](#)
- [Create an Aggregated Component](#)
- [View the Context Page Related to an Aggregated Component](#)
- [Edit an Aggregated Component](#)
- [View the List of Aggregated Component Members](#)
- [Override the Speed In or Speed Out Values for a Interface Aggregated Component Member](#)
- [Remove a Member from an Aggregated Component](#)
- [Delete an Aggregated Component](#)

### View a List of Aggregated Components

You can view a list of aggregated components from the **Aggregated Components** page. To view this page, hover over **Administration**, **Monitored Items Management**, and then click **Aggregated Components**.

## Create an Aggregated Component

### Follow these steps:

1. From the **Aggregated Components** page, click **New**. The **Create / Edit Aggregated Component** dialog opens.
2. Complete the following required fields:
  - **Name**  
The name for the aggregated component.
  - **NOTE**  
Single quotes, double quotes, backward slashes (\), forward slashes (/), ampersands (&) are not permitted.
  - **Description**  
The description for the aggregated component.
  - **Aggregated Metric Family**  
Defines the type of (aggregated) metric family for the aggregated component.
  - **Options:**
    - **Aggregated CPU:** An aggregated metric family for devices that contain CPUs, which are non-syncable components.
    - **Aggregated Interface:** An aggregated metric family for interfaces, which are syncable components.
    - **Aggregated Memory:** An aggregated metric family for devices that contain memory, which are non-syncable components.
3. Complete one of the following based on the type of (aggregated) metric family:
  - (For aggregated interface metric families) From the **All Interfaces** table, select one or more interfaces that the data aggregator has discovered and that you want the aggregated component to contain, and then click **Add**.  
The interfaces are added to the **Selected Interfaces** table.
  - (For aggregated CPU or memory metric families) From the **All Devices** table, select one or more devices that the data aggregator has discovered, that contain *two or more* CPUs or memory, and that you want the aggregated component to contain, and then click **Add**.  
The devices containing those components are added to the **Selected Devices** table.
4. Click **Save**.

The aggregated component is created and added to the data aggregator. The new aggregated component is displayed in the **Aggregated Components** list on the **Aggregated Components** page. For aggregated CPU or memory metric families, an aggregated component is created for each device containing *two or more* CPUs or memory, each containing its CPUs or memory. If a device contains only *one* CPU or memory, an aggregated component is not created for that device.

## View the Context Page Related to an Aggregated Component

You can view a list of established aggregated components, and then view the context page related to an aggregated component. To view the list of established aggregated components, hover over **Inventory**, **Items**, and then click **Aggregated Components**.

## Edit an Aggregated Component

### Follow these steps:

1. From the **Aggregated Components** page, select the aggregated component that you want to edit, and then click **Edit**.  
The **Create / Edit Aggregated Component** dialog opens.
2. Complete the following required fields:
  - **Name**  
The name for the aggregated component.

**NOTE**

Single quotes, double quotes, backward slashes (\), forward slashes (/), ampersands (&) are not permitted.

- **Description**

The description for the aggregated component.

3. (If this is an aggregated component for aggregated interface metric families) In the **All Interfaces** and **Selected Interfaces** tables, modify the interfaces that you want the aggregated component to contain.
4. Click **Save**.

Your changes are saved.

### **View the List of Aggregated Component Members**

You can view the list of aggregated component members from the **Components** pane on the **Aggregated Components** page. To view this list, from the **Aggregated Components** page, select the name or the checkbox of the aggregated component for which you want to view members.

The following details appear in the list for interface aggregated component members:

- **Name:** The name of the interface aggregated component member that belongs to a selected device.
- **Device:** The name of the parent device to which the interface aggregated component member belongs.
- **Description:** The description of the interface aggregated component member.
- **Status:** The status of the interface aggregated component member.  
**Values:** Active, Inactive
- **SNMP Poll Rate:** The SNMP poll rate for the interface aggregated component member, in minutes.  
**Example:** 5 minutes
- **Speed In:** The speed in value for the interface aggregated component member, in Megabits per second.  
**Example:** 5 Mbps
- **Speed Out:** The speed out value for the interface aggregated component member, in Megabits per second.  
**Example:** 5 Mbps

The following details appear in the list for CPU aggregated component members:

- **Name:** The name of the CPU aggregated component member that belongs to a selected device.
- **Device:** The name of the device to which the CPU aggregated component member belongs.
- **Description:** The description of the CPU aggregated component member.
- **Status:** The status of the CPU aggregated component member.  
**Values:** Active, Inactive
- **SNMP Poll Rate:** The SNMP poll rate for the CPU aggregated component member, in minutes.  
**Example:** 5 minutes

The following details appear in the list for memory aggregated component members:

- **Name:** The name of the memory aggregated component member that belongs to a selected device.
- **Device:** The name of the device to which the memory aggregated component member belongs.
- **Description:** The description of the memory aggregated component member.
- **Status:** The status of the memory aggregated component member.  
**Values:** Active, Inactive
- **SNMP Poll Rate:** The SNMP poll rate for the memory aggregated component member, in minutes.  
**Example:** 5 minutes

## **Override the Speed In or Speed Out Values for a Interface Aggregated Component Member**

### **Follow these steps:**

1. From the **Aggregated Components** page, select the name or the checkbox of the interface aggregated component for which you want to edit a member.
2. In the **Components** pane, select the interface aggregated component member that you want to edit, and then click **Edit**.  
The **Edit Interface** window opens.
3. Edit the following:
  - **Speed In:** The speed in value for the interface aggregated component member, in Megabits per second.  
**Example:** 5 Mbps
  - **Speed Out:** The speed out value for the interface aggregated component member, in Megabits per second.  
**Example:** 5 Mbps
4. Save your changes.

The speed in or speed out value for the interface aggregated component member is overridden.

## **Remove a Member from an Aggregated Component**

### **Follow these steps:**

1. From the **Aggregated Components** page, select the name or the checkbox of the aggregated component from which you want to remove a member.
2. In the **Components** pane, select the member item that you want to remove, and then click **Remove**.  
The **Remove Aggregated Component Member** window opens.
3. Confirm the removal by clicking **Yes**.

The member is removed from the aggregated component.

## **Delete an Aggregated Component**

### **Follow these steps:**

1. On the **Aggregated Components** page, select the aggregated component that you want to delete, and then click **Delete**.  
The **Delete Aggregated Component** dialog opens.
2. Confirm the deletion by clicking **Yes**.

# **Manage Application Mappings**

You manage application mappings by adding, editing, or deleting them.

Application mappings are a way to group a series of hosts by port (tie a single port to multiple hosts) and define it as an application. They define the application mapping of flow traffic in the Sankey diagram on the **Flow Dashboard**. Each application mapping defines the details—such as the name, IP address, and port—for an application. For example, the application mapping name defines the name that is displayed, as well as the source and destination, for the flow in the Sankey diagram. This dashboard also displays application mapping of flow traffic (by way of NBAR2 or user-defined application mapping).

## **View a List of Application Mappings**

View a list of the application mappings from the **Flow Application Mapping** page.

### **Follow these steps:**

1. Log in as a user with the Administer Advanced Flow Application Mapping administrative role rights.

For more information about this role right, see [Role Rights](#).

2. Hover over **Administration**, **Monitored Items Management**, and then click **Monitored Devices** or **Discovery Profiles**.

The **Monitored Devices** (or **Discovery Profiles**) page appears.

3. Click the **Flow Application Mapping** tab.

### Add an Application Mapping

Define the host and port that make up an application.

#### Follow these steps:

1. From the **Flow Application Mapping** page, click **New**.  
The **Configure Application Mapping** dialog opens.
2. Complete the following fields, and then click **Save**:
  - **Name**  
Defines the name for this application mapping. This name displays in the Sankey diagram for the **Flow Dashboard** for the flow records that match the application's IP/port combinations.  
For more information about the **Flow Dashboard**, see [Flow Dashboards](#).
  - **Description**  
Describes this application mapping.
  - **IP Domain**  
The IP domain for this application mapping.  
**Default:** Default Domain
  - **NBAR2**  
Defines the flow name that you want the Sankey diagram on the **Flow Dashboard** to display for this application mapping.  
**Options:**
    - Cleared: Display the flow name that the router stamped (the NBAR2 application name) for this application mapping.
    - Selected: Display the name defined for this application mapping.**Default:** Cleared
  - **Application IP/Ports**  
Shows the server IP and port combinations that make up this application mapping. Flow records that match these IP/port combinations use the name defined for the application mapping.

The application mapping is added.

## Manage Interfaces

Interfaces represent monitored communications ports, such as Ethernet or serial ports. The articles in this section include the options for managing interfaces.

### Poll Critical Interfaces Faster than Non-critical Interfaces

As an Administrator, frequent data about your most critical systems while maximizing the overall performance of NetOps Portal.

As an Administrator, you can configure the data aggregator to poll critical interfaces as often as necessary. For example, you can configure the data aggregator to poll only critical interfaces at a high rate and poll non-critical interfaces at a normal or slow rate. You can have the data aggregator poll at differing rates by setting a filter on the **Interfaces** metric family that is associated with a monitoring profile. By fast-polling interfaces sparingly, you can reduce unnecessary network traffic and NetOps Portal load while still sufficiently monitoring the health of your network system.

For example, your data center access switch connects many application servers to only two aggregation switches. You decide to configure the data aggregator to poll the interfaces supporting these aggregation switches at a higher rate. These links are critical because they support network traffic to all other connected switches.

### IMPORTANT

Configuring the data aggregator to poll *all* interfaces at a higher rate causes unnecessary network traffic, wastes system resources, and causes network performance issues. Prior to configuring fast polling, consult with your network operations and engineering teams.

**Best practice:** Configure the data aggregator to poll the interfaces that connect each attached server at a normal polling rate. To apply different polling rates, implement two monitoring profiles for interfaces.

### NOTE

NetOps Portal ignores filters that you set on metric families when event rules that are applied to monitoring profile trigger events.

For more information about monitoring profiles, see [Manage Monitoring Profiles](#).

You can minimize unnecessary network traffic that polling *all* interfaces at this fast rate can produce. You can create two monitoring profiles for interfaces—one with normal polling, and the other with fast polling. Use the following process to poll critical interfaces faster than non-critical interfaces:

1. [Create the monitoring profile](#).
2. [Add a monitoring profile filter for the Interfaces metric family](#).
3. [Assign the monitoring profile to a device collection](#).
4. [Verify the results](#).

The following video shows how to discover and monitor interfaces:

### Create the Monitoring Profile

Create a copy of an out-of-the-box monitoring profile, and use it to create a monitoring profile that polls only critical interfaces at a faster polling rate.

For more information, see [Manage Monitoring Profiles](#).

### Add a Monitoring Profile Filter for the Interfaces Metric Family

By default, the out-of-the-box **Network Interface** monitoring profile includes a filter to prevent modeling interfaces that are administratively down. In addition, this monitoring profile does not model interfaces with a type (`ifType`) of 1 (Other) or 24 (Loopback), regardless if those interfaces are administratively up or down. It also does not model IPSLAs with the `rttMonCtrlAdminOwner` MIB object that contains the string "Network Health".

Filtering reduces the number of interfaces that NetOps Portal monitors, which reduces unwanted data collection and network traffic.

In addition to polling only administratively up interfaces, you also want to poll the most critical interfaces more frequently. To isolate and poll only these interfaces faster, you add a second filter condition to the interface filter associated with your custom monitoring profile. This second filter condition isolates the critical interfaces by finding only interfaces that contain a in their description.

### NOTE

The data aggregator applies filtering after discovery. The data aggregator does not poll interface components that do not match the filter criteria. If you add or edit an interface filter *after* you run a discovery, polling on these components stops. NetOps Portal does not display these interface components in dashboards and data views.

### NOTE

Log in as an Administrator to perform this task.

Add a monitoring profile file for the **Interfaces** metric family. Configure the filter conditions with the following options:

- **Attribute:** Description
- **Operation:** Contains
- **Filter Value:** Enter the name of the filter.
- **Case-sensitive:** Yes

Consider the following details about additional attributes you can use for filtering:

- For **Speed In** and **Speed Out**, you can use a decimal in the text field (such as 1.544) and can specify bps, Kbps, Mbps, or Gbps.
- For **Description** and **Alias**, you can use a regular expression for filtering only when you select the **Matches Regex** or the **Does Not Match Regex** operation.
- When you save your changes, the filter criteria display on the **Metric Families** tab. You can then apply this monitoring profile to the appropriate device collection to begin polling your selected interfaces.

### **Considerations for Interface Filters and Multiple Monitoring Profiles**

When multiple monitoring profiles are assigned to a device collection, the filter matching criteria follows the "or" rule. So, the data aggregator monitors all interfaces that satisfy the criteria for any of the monitoring profiles in the group.

Some of the monitoring profiles may have filters and some may not. Plus, these profiles can specify differing poll rates. In this case, the data aggregator monitors the interfaces that match any monitoring profile, but the polling rates can differ. If more than one monitoring profile applies to an interface, the data aggregator polls the interface once, and polls it at the fastest polling rate:

- **Monitoring Profile 1 -- Filter:** Description contains "X," Poll Rate: 1 minute
- **Monitoring Profile 2 -- Filter:** None, Poll Rate: 5 minutes
- **Monitoring Profile 3 -- Filter:** Description contains "Y," Poll Rate: 10 minutes

In this example, interfaces that match Monitoring Profile 1 are polled every minute. All other interfaces are polled every 5 minutes. Interfaces that match Monitoring Profile 3 also match Monitoring Profile 2, which does not include a filter. The fastest poll rate applies, so no interfaces are polled at 10-minute intervals.

In this case, if one monitoring profile has no filter, the result is that many interfaces may be polled more frequently than necessary. Therefore, after you set a filter, remove associations from other monitoring profiles to make sure that only components matching the specified filter are monitored.

### **Assign the Monitoring Profile to a Device Collection**

As the Administrator or a Tenant Administrator, [associate the monitoring profile with a device collection to begin polling](#). The **Switches** device collection is associated with the out-of-the-box **Network Interfaces** monitoring profile. Polling rates are applied to the interfaces in this device collection, as follows:

- **Fast polling:** Interfaces that satisfy the filter criteria of the monitoring profile.
- **Normal polling:** All other interfaces that the monitoring profile discovers.

#### **IMPORTANT**

Custom monitoring profiles are global and visible to tenant administrators. However, you can scope the association of a monitoring profile with a specific device collection to a tenant.

### **Verify the Results**

After you have set up your monitoring profiles, verify that only the critical devices are polled at the higher rate by reviewing the monitored devices and the Filter report. This information helps you to see information in context, such as which monitoring profiles are being used to poll device components. Verifying the results can help you identify any necessary adjustments to help you achieve the polling results that you want.



**NOTE**

Monitored devices are manageable devices and pingable (accessible but not manageable). Inaccessible devices are not monitored devices. You can view components of monitored devices from the **Polled Metric Families** tab.

**Follow these steps:**

1. [Run an on-demand device discovery](#).

**NOTE**

If the discovery profile runs automatically, you can wait for the next scheduled discovery.

2. Click **Monitored Devices** from the **Monitored Inventory** menu for a data aggregator data source.  
A list of devices displays in the **Monitored Devices** page.
3. To locate an aggregation switch device in the corresponding tree view, select *one* of these options from the drop-down list:
  - **Device by Collection**  
The devices appear under the assigned device collection.
  - **Device by Monitoring Profile:**  
The critical interfaces appear under **Devices** under the monitoring profile.

**TIP**

Alternatively, select the **Search** tab to search by host name, device name, or IP address. You can enter a partial name or IP address to return a list of devices that contain that partial match. You cannot use wildcards and regular expressions.

4. Select the device, and then click the **Polled Metric Families** tab.  
This tab shows the consolidated monitoring profiles that are associated with the switch device. Devices have only one consolidated monitoring profile. Each consolidated monitoring profile lists every metric family to poll on the device and whether the device supports the metric family.
5. Select the **Interfaces** metric family.  
The **Components** table for the Interfaces metric family shows one of the following polling statuses for the discovered Interface components:
  - **Active**  
Indicates that the component is being polled.
  - **Inactive**  
Indicates that polling has stopped on the component because the metric family is no longer monitored for the device.
  - **Not Present**  
Indicates that the component no longer exists on the physical device. Polling is stopped on the component. You can view historical data for reporting purposes. By default, the data aggregator does not synchronize retired components with NetOps Portal.
  - **Filtered (interface components only)**  
Indicates that the component does not pass the filter criteria and polling on the component is stopped.

**NOTE**

Dashboards and data views do not display filtered interfaces.

6. (Optional) Select the **Interface** metric family, and then click **Update Metric Family**.  
The **Update Metric Family** dialog opens.
7. Click **Yes**.  
The data aggregator reconfigures components for any configuration updates. For example, if you add a disk drive on a server, you can click **Update Metric Family** to rediscover the configuration update. The configuration update creates a disk component.
8. Click the **Filter Report** tab, and then follow these steps:
  - a. Look at the filters on each of the other Interface monitoring profiles to see if they are monitoring the same device collection that you want to filter.



- b. Remove any relationships between other Interface monitoring profiles and device collections that will block your filter criteria. For example, if your new Interface monitoring profile is associated with the **All Routers** device collection, remove the relationship between *other* Interface monitoring profiles and the **All Routers** device collection.
- c. Run another discovery and review the updated Filter report to verify that the new filter criteria is active. If the Filter report shows that an unwanted monitoring profile was included, repeat the previous steps until you are monitoring only the interfaces that you want.

The **Filter Report** tab shows which interface filter criteria have been used during component monitoring. The tab also shows a report of all of the interfaces that are identified on the device and whether they matched the specified filter criteria.

#### NOTE

If you change the rules on a custom monitoring profile, the **Interface Filter Criteria** pane does not reflect those changes. If you disassociate the monitoring profile from a group, the **Interface Filter Criteria** pane does not reflect those changes. Rediscover the device to filter the interfaces that are based on the changes you made to the filter criteria and monitoring profile.

## Interface Components Naming Convention

The naming convention for interface components that the Interface vendor certification or the High Speed Interface vendor certification backs is based on the following logic:

- If the `ifName` attribute exists and has a value, the interface uses this value for its name.
- If the `ifName` attribute does *not* exist or does *not* have a value, the interface uses the value of the `ifDescr` attribute for its name.

#### NOTE

New certifications for the Interface metric family can provide a different expression for the interface name.

## Override Speed In and Speed Out Values on Interfaces

You can ensure utilization calculations use the appropriate values by overriding the Speed In and Speed Out values in the data aggregator for any interface. For example, you could use the bandwidth command to configure `ifSpeedIn` and `ifSpeedOut` on your router interfaces to affect routing decisions. In this case, ensure that utilization is calculated correctly by providing an override speed with the data aggregator.

The settings that you make on the device can change the value to one that is higher or lower than the actual available data rate. So, the utilization calculations that are made for the interface can appear inaccurate, due to this manipulation of the bandwidth. To ensure that interface utilization is calculated correctly, you want to provide an override speed on the interface within Data Aggregator.

By default, utilization is calculated using the Speed In and Speed Out values that the device, which the interface is a component of, reports. However, you can override these speed values. Reporting on interface utilization can then be more accurate.

#### Follow these steps:

1. Click **Monitored Devices** from the **Monitored Inventory** menu for a data aggregator data source.  
The **Tree View** tab displays.
2. Select **Device by Collection** or **Device by Monitoring Profile** from the drop-down list. Select the device for which you want to override the Speed In and Speed Out values for an interface, and then select the appropriate interface metric family on the **Polled Metric Families** tab.  
The interface components that are monitored on the device appear in the **Interface Components** table.
3. Select the interface component for which you want to override the Speed In and Speed Out values, and then click **Edit**.  
The **Edit Interface** dialog appears. The dialog displays the default discovered Speed In and Speed Out values.

4. Do one of the following:

- **Override** the **Speed In** and **Speed Out** values on the interface by entering values in bits per second, and then clicking **Save**.  
The overridden Speed In and Speed Out values on the interface appear in the **Interface Components** table with asterisks. Going forward, bandwidth utilization charts in NetOps Portal for the interface display utilization use these speed values.
- **Remove** the speed overrides on the interface by clicking **Clear**, and then clicking **Save**.  
Going forward, bandwidth utilization charts in NetOps Portal for the interface display utilization use the speed values that the device reports.

The dialog closes. An event is generated on the interface, indicating that you have *overridden* or *removed* the speed overrides on an interface. You can view this event on the **Events Display** dashboard.

## Configure Counter Behavior

Resolve issue in counter data.

Simple Network Management Protocol (SNMP) uses counters to record data points. The data collector polls the counters, and reports the differences in the counters over time. When the counter hits a preset limit, the counter resets to zero. The data collector assumes that the rollover happens no more than once per poll. If a rollover happens more than once, the resulting data is inaccurate. To resolve the issue of inaccurate data, take one of the following actions:

- Poll more often to ensure that the data is accurate
- Use a 64-bit counter instead of a 32-bit counter
- Show gaps in the data and leave out the inaccurate data

A counter can also go backwards, which indicates that the data is incorrect. The data collector interprets the backward behavior as a large spike in data, and discards the data point.

### Bad Counter Deltas

The delta value for counters is calculated by comparing one poll cycle to the previous poll cycle. When the data collector receives bad delta values, the values are discarded for calculation purposes. Occasionally, the counter value receives a bad poll result, and the value decreases. When the counter value shows a decrease, the data collector skips the bad value.

When `showGapsOnCounterRollover=true`, the data collector discards the previous delta baseline. As a result, the data collector cannot calculate the delta for the next poll cycle.

The following table shows an example of the delta behavior when `showGapsOnCounterRollover=true`:

Poll Cycle	Counter Value	Delta Baseline	Delta Value	Utilization
Poll1	96	95	1	normal
Poll2	97	96	1	normal
Poll3	<b>30 (bad value)</b>	97	<b>null (drop)</b>	<b>gap</b>
Poll4	99	<b>reset</b>	<b>null (drop)</b>	<b>gap</b>
Poll5	100	99	1	normal

### Show Gaps in Data

To show gaps in the data, change the default behavior. Repeat this procedure on each data collector.

#### NOTE

The following global settings apply to all devices that a particular data collector polls.

**Follow these steps:**

1. Create the `<DC_installation_directory>/apache-karaf/etc/com.ca.im.dm.snmp.collector.SnmpCollector.cfg` file on the data collector host.

- ***DC\_installation\_directory***

The installation directory for the data collector.

**Default:** `/opt/IMDataCollector`

2. Add the following line to the file:

```
showGapsOnCounterRollover=true
showGapsOnCounterRollover32=true
```

**Default:**

- `showGapsOnCounterRollover=false`
- `showGapsOnCounterRollover32=false`

3. Save the file.

The new behavior takes effect immediately.

When you hide gaps, the following logic continues to protect against some spikes in the data:

- If a 64-bit counter appears to wrap at the 32-bit mark, the SNMP agent on the device is likely using the 32-bit counter. If the counter goes from an initial value below the maximum value of a 32-bit counter to a lower value, DX NetOps Performance Management computes the delta at the counter wrap as a 32-bit counter.
- If the delta exceeds a certain limit for either a 32-bit or a 64-bit counter, the data point is dropped.

**Limit Delta Values**

For 32-bit or 64-bit counters, the delta value is likely an error when both of the following criteria apply:

- A counter rollover occurs
  - The delta is greater than the maximum allowable delta value
- Default:  $2^{32}$  for 32-bit;  $2^{63}$  for 64-bit

If an error occurs, the value is discarded for calculation purposes, and the report data shows a gap of one poll cycle. You can configure the behavior to increase or decrease the maximum allowable delta value. Repeat this procedure on each data collector.

**Follow these steps:**

1. Open the `<DC_installation_directory>/apache-karaf/etc/com.ca.im.dm.snmp.collector.SnmpCollector.cfg` file.

- ***DC\_installation\_directory***

The installation directory for the data collector.

**Default:** `/opt/IMDataCollector`

2. To set a non-default threshold, add the following line to the file:

```
largeDeltaValueThreshold=integer
largeDeltaValueThreshold32=<integer>
```

- ***Integer***

Specifies the threshold for delta values.

3. **Default:**

- `largeDeltaValueThreshold32=2147483648` (0x80000000, or  $2^{31}$ )
- `largeDeltaValueThreshold=9223372036854775807` (0x7fffffffffffffff, or  $2^{63}-1$ )

When a counter rollover occurs, DX NetOps Performance Management ignores delta values at or above the threshold.

## Counter Rollover Log

The counter rollover log tracks detailed information about each counter rollover for all devices on a data collector. Use this log to verify counter behavior and troubleshoot counter issues.

Locate the log on the data collector that polls the relevant device:

```
<DC_installation_directory>/apache-karaf/data/log/CounterRollover.log
```

- **DC\_installation\_directory**

The installation directory for the data collector.

**Default:** /opt/IMDataCollector

The log contains currently applied configuration values, rollover events, IP addresses, and OIDs.

### Example:

```
2016-03-07 12:54:13,480 | INFO | ecutor-thread-58 | CounterRollover
| .dm.snmp.rdp.impl.SnmpDeltaCache 327 | 186 - com.ca.im.data-collection-
manager.core.interfaces - 2.8.0.SNAPSHOT | | Delta calculated is greater than or
equal to 2147483648; dropping response: previous=1 / current=0 for ip 10.42.96.32, OID
1.3.6.1.4.1.9.9.42.1.3.5.1.39.2222, itemID 906, in poll group 211.
```

## Manage Interface Polling Behavior

If you have the Modify Component Polling role right, you can manage (disable or enable) polling on interfaces using NetOps Portal.

Interface polling allows for more granular polling control than monitoring profile filters alone, since those apply only to common attributes.

Enabled polling generates administrative events.

For more information, see [Event Types](#).

You can also manage polling state on other components using the data aggregator REST web service.

For more information, see [Manage Polling Behavior for Components](#).

By default, polling is enabled for all new components. You can disable polling for all new components that are associated with specific metric families using the data aggregator REST web service.

For more information, see [Manage Default Polling Behavior](#).

### Follow these steps:

1. Hover over **Inventory**, **Items**, and then click **Interfaces**.  
The **Interfaces** page appears.
2. Select the interface for which you want to manage polling behavior, and then click **Update Polling**.  
The **Manage Polling State** dialog opens.
3. Select a polling state for the interface from the **Select Polling State** drop-down, and then click **OK**:
  - **Disable**  
Disallows *and* turns off polling.
  - **Enable**  
Allows polling.

#### IMPORTANT

For polling to occur (to *turn on* polling), set the appropriate monitoring profiles and life cycle states.

For more information:

- About how to configure monitoring profiles, see [Manage Monitoring Profiles](#).
- About how to set life cycle states, see [Manage Device Life Cycles](#).

The polling state is modified.

## Manage Network Flow Processing

When DX NetOps Network Flow Analysis is configured as a data source, as a network administrator, you can manage (enable, disable, edit, merge, and delete) network flow processing for your DX NetOps Network Flow Analysis interfaces.

### NOTE

You can also enable or disable interfaces/devices using DX NetOps Network Flow Analysis.

For more information, see [the DX NetOps Network Flow Analysis documentation](#).

Complete the following procedures on the target routers or interfaces:

### View Network Flow Processing in NetOps Portal

#### Follow these steps:

1. Log in as an Administrator.
2. Hover over **Administration**, **Data Sources**, and then select the Network Flow Analysis data source.
3. Select **Manage Monitoring** in the **Managed Interfaces** section.  
The **Manage Monitoring** page appears showing the list of interfaces.

### NOTE

You can search and sort the list of interfaces. Searching and sorting works on the whole data instead of the current page data. Filter text searches the original values of speed and data columns. It does not search the modified values.

### Enable or Disable an Interface

#### Follow these steps:

1. From the **Manage Monitoring** list, in the right pane, click the **All Interfaces** tab, and then select the interface that you want to enable or disable.  
Select multiple interfaces to enable or disable in bulk.
2. Click **Enable** or **Disable**.
3. Click **Yes** at the prompt.

The interface is enabled or disabled.

### Edit an Interface

#### Follow these steps:

1. In the right pane, click the **Active Interfaces** tab, and then select the interface that you want to edit.  
Select multiple interfaces to perform a bulk edit. However, only some fields are editable in the bulk edit option. To edit all editable fields, select only one interface at a time.
2. Click **Edit Interface**.
3. Set the desired interface parameter (field) values, and then click **Save**.

Your changes are saved.

## **Merge Interfaces**

You can merge interfaces, which merges their data. You can merge only two interfaces at a time.

### **Follow these steps:**

1. In the left pane, select the device (router) related to the interface that you want to merge.
2. In the right pane, click the **Active Interfaces** tab.
3. Select the interfaces that you want to merge, and then click **Merge**.  
The **Merge Interface** dialog appears.
4. (Optional) To merge a different interface than the selection, change the interfaces that you want to merge:
  - Add the interfaces that you want to merge to **Active Interfaces**.
  - Remove the interfaces that you do not want to merge from **Selected Interfaces**.
5. (Optional) To delete the source interface after the merge, click **Delete Source**.
6. (Optional) To merge the interfaces that have overlapping timeframes, click **Merge overlapping interfaces**.
7. Click **OK**.

The interfaces are merged.

## **Delete an Interface**

### **Follow these steps:**

1. In the right pane, click the **Active Interfaces** tab.
2. Select the interface that you want to delete, and then click **Delete**.  
Select multiple interfaces to perform the bulk delete.
3. When prompted, click **Yes**.

The interface is deleted.

## **Create a Custom Virtual Interface**

Create custom virtual interfaces (CVIs) to separate traffic on a particular interface and subnet from other traffic on the interface. You can create a CVI from an existing physical interface. You can also modify an existing CVI.

### **Follow these steps:**

1. In the left pane, select the device (router) related to the interface from which you want to create a CVI.
2. In the right pane, click the **Active Interfaces** tab.
3. Select a *physical* interface, and then click **Add Custom Virtual Interface**.

#### **NOTE**

The **Class** column distinguishes interface type, physical or virtual.

The **Add Custom Virtual Interface** dialog opens.

4. Enter values for the following fields:
  - **Interface Name**  
Replace the default value in the **Interface Name** field with a meaningful name for the interface list.
  - **Description** (*Optional*)  
Enter a text string to help identify the interface.
  - **In Speed** and **Out Speed** (*Optional*)  
Identify the speed of data that is inbound to the parent interface and outbound from the parent interface.
  - **Interface Type** (*Optional*)  
Select the interface type from the list.
  - **Subnet**  
Enter a subnet and mask identification for each subnet filter you want to use for this CVI, then click **Add**. Use the format:

<subnet IP address/subnet mask>

The CVI requires at least one subnet filter.

- Click **Save**.

The CVI is automatically deployed. The **Class** column for the new CVI distinguishes it from a physical interface.

### **Test an SNMP Profile**

You can test the SNMP profile of the device or router using the **Test Profile** option. Selecting an SNMP-profile-assigned router enables this option. The test fails if the SNMP profile is not assigned to the router that you have selected. In such cases, you can assign an SNMP profile to the specified router, and then test the router profile.

For more information about how to assign an SNMP profile for a specified router, see [the "Edit a Device \(Router\)" section](#).

### **Discover SNMP Profiles**

You can discover SNMP profiles for devices (routers). The discovery process is asynchronous. Only the task is scheduled, as it might take a long time for discovery to complete. You can view the current state of discovery from the **Discover Profile State** column in the **Devices** list (hidden by default). Initially, discovery is in "Scheduled" state, and then it moves to "Running".

After discovery is complete, the profile state changes to one of the following states:

- **Profile Found:** Discovery is successful. The **SNMP Profile** column displays the newly found profile for the device (router).
- **Profile Not Found:** Edit the existing SNMP profiles or add a new SNMP profile to poll the device.
- **Failed:** You can view the reason for the failure from the <NFAInstallPath>\REPORTER\Logs\DiscoverProfileLog<timestamp>.log file, and take corrective action.

#### **NOTE**

Discovery does not discover SNMP profiles for retired devices.

To get the latest state of discovery, refresh the devices list.

### **Refresh an SNMP Router**

**Follow these steps:**

1. In the left pane, select the Simple Network Management Protocol (SNMP) router that you want to refresh, and then click **SNMP Refresh**.

#### **NOTE**

You can refresh only one router at a time.

2. When prompted, click **Yes**.

The SNMP router is refreshed.

### **Enable or Disable a Device (Router)**

**Follow these steps:**

1. In the left pane, select the router that you want to enable or disable, and then click **Enable** or **Disable**.  
Select multiple devices to perform bulk action.
2. When prompted, click **Yes**.

The router is enabled or disabled.

### **Edit a Device (Router)**

#### **Follow these steps:**

1. In the left pane, select the router that you want to edit, and then click **Edit**.  
Select multiple routers to perform bulk action. However, only some fields are editable in the bulk edit option. To edit all allowed fields, select only one router.
2. Set the desired router parameter values, including assigning an SNMP profile to the router, and then click **Save**.

Your changes are saved.

### **Delete a Device (Router)**

#### **Follow these steps:**

1. In the left pane, select the router that you want to delete, and then click **Delete**.  
Select multiple devices to perform bulk action.
2. When prompted, click **Yes**.

The router is deleted.

## **Manage Interface Aggregation for Flow**

As a network administrator, you can perform common administrative tasks on all the interfaces by aggregating similar interfaces using NetOps Portal rather than performing them on each interface individually. Interface aggregations show network flow processing by combining traffic from two or more interfaces and reports the traffic together. You can have continuous access to network devices allowing you to manage them from a single interface. Manage interface aggregations for flow by creating, editing, and deleting them.

#### **NOTE**

You can also manage interface aggregations for flow using DX NetOps Network Flow Analysis.  
For more information, see [the DX NetOps Network Flow Analysis documentation](#).

Complete the following procedures on a target aggregated interface:

### **View Network Flow Processing in NetOps Portal**

#### **Follow these steps:**

1. Log in as an Administrator.
2. Hover over **Administration**, **Data Sources**, and then select the Network Flow Analysis data source.
3. Select **Aggregations** in the **Managed Interfaces** section.  
The **Interface Aggregations** page appears showing the list of current interface aggregations.

#### **NOTE**

You can search and sort the list of interface aggregations. Searching and sorting work on the whole data instead of the current page data. Filter text searches the original values of speed and data columns. It does not search for the modified values. Use the **Quick Filter** to find the required interface. You can also use the pagination options to navigate to find the interface for aggregation.

### **Create an Interface Aggregation for Flow**

#### **Follow these steps:**

1. From the **Interface Aggregations** list, click **New**.  
The **New Aggregation** window appears.
2. Complete the following fields:



- **Aggregation Name**  
Specify a name for the interface aggregation.
- **Description**  
Specifies a notation to help identify the interface aggregation.
- **In Speed(bps)**  
Specify the inbound speed of the selected interfaces.
- **Out Speed(bps)**  
Specify the outbound speed of the selected interfaces.
- **Interface Type**  
Select the mode of interface connection from the **Type** list, such as WAN or Ethernet.  
If you do not specify the interface type, the type is set to **Unknown** by default.
- **Active Interfaces**  
Select the interfaces to include in the interface aggregation, and then click **Add**.  
Use the **Quick Filter** to find the required interface. You can also use the pagination options to navigate to find the interface for aggregation.
- **Selected Interfaces**  
Select the interface to exclude from the interface aggregation, and then click **Remove**.  
Use the **Quick Filter** to find the required interface. You can also use the pagination options to navigate to find the interface for aggregation.

3. Click **Save**.

The interface aggregation is deployed within one minute of creation.

### **Edit an Interface Aggregation for Flow**

**Follow these steps:**

1. From the **Interface Aggregations** list, select the interface aggregation that you want to edit, and then click **Edit**. To perform a bulk edit, select multiple interface aggregations.

#### **NOTE**

For bulk edits, only a some fields are editable in the bulk edit option. To edit all editable fields, select only one interface aggregation at a time.

The page changes to the **Edit Aggregation** view, which includes editable options.

2. Make any changes that are needed, such as adding or removing interfaces or changing the aggregation name, description, in/out speed, or type, and then click **Save**.

The changes to the interface aggregation is saved.

### **Delete an Interface Aggregation for Flow**

**Follow these steps:**

1. From the **Interface Aggregations** page, select the interface aggregation that you want to delete, and then click **Delete**. To perform a bulk delete, select multiple interface aggregations.
2. When prompted, click **Yes**.

The interface aggregation is deleted.

## **Configure Round Trip Time (RTT) Tests**

To enable round-trip time tests, configure the tests using the Data Aggregator REST API. DX NetOps Performance Management supports the following test types:

- **DNS**

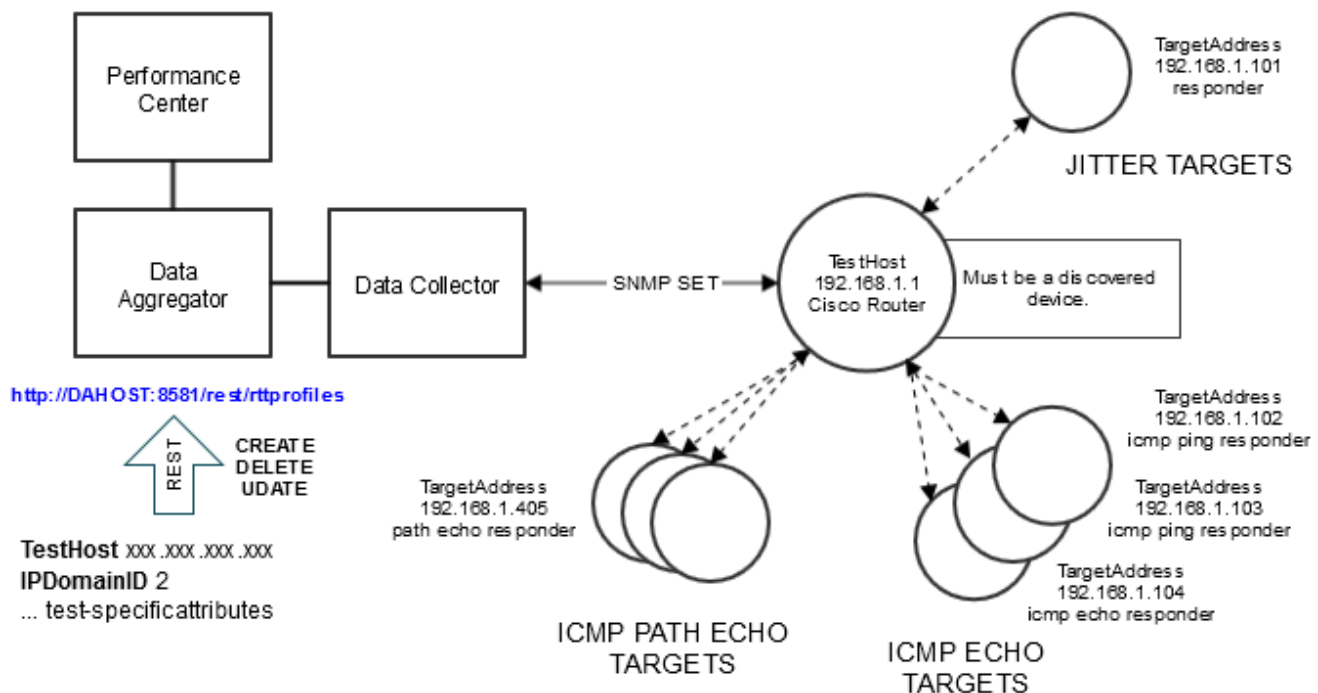
Measures the DNS lookup time.

- **ICMP Echo (Ping)**  
Measures the round-trip delay for the full path.
- **ICMP Path Echo**  
Measures the round-trip delay and hop-by-hop round-trip delay.
- **ICMP Jitter**  
Measures the hop-by-hop jitter, packet loss, and delay measurement statistics in an IP network.
- **HTTP**  
Measures the round-trip time to retrieve a web page.
- **TCP Connect**  
Measures the time to connect to a target device with TCP.
- **UDP Jitter**  
Measures the round-trip delay, one-way delay, one-way jitter, and one-way packet loss. Codec simulation is provided for MOS and ICPIF voice quality scoring capability.

### RTT Test Architecture

The following diagram shows how the various components participate in the RTT test feature:

**Figure 19: RTT Test Architecture**



### Configure the Tests

To configure RTT tests, complete the following steps:

### Gather Required Information

The following information is required to configure a test:

- Required for all RTT tests in your environment:
  - DA host/address
  - IpDomainId

**NOTE**

To determine the IpDomainId, use the Data Aggregator REST API.

- Required for the specific test:
  - TestHost  
The IP address of the device where DX NetOps Performance Management creates the test.
  - Depending on the test type, one of the following targets is required:
    - TargetAddress  
The device that receives the test probes.
    - (DNS) TargetAddressString  
IP address or hostname of the target that receives the probes. Use this string instead of TargetAddress.
    - (HTTP) Url  
The URL that the HTTP probe communicates with.

**Create an SNMP "SET" Profile**

Create an SNMP profile and specify **Yes** for **Use in SNMP SET**.

For more information, see [SNMP Profiles](#).

**Discover Devices Capable of Running RTT Tests**

Apply a monitoring profile with the IPSLA-related metric families that you want to monitor to the target devices. Then, run discovery for those devices.

For more information, see [Discovery](#), [Configure Monitoring Profiles](#), and [IPSLA Polling](#).

**Associate SET Profiles with the Devices**

For each device, associate the SET profile to the device item.

**Through NetOps Portal:**

1. Go to **Administration**, and click the Data Source.
2. Select the monitored device.
3. Select the SNMP SET Profile and the SNMP SET Version.
4. Click **Save**.

You can now use DX NetOps Performance Management to configure RTT tests on the device.

**Through the Data Aggregator API:**

**Documentation:** [http://DA\\_HOST:8581/rest/devices/manageable/documentation](http://DA_HOST:8581/rest/devices/manageable/documentation)

**PUT URL:** [http://DA\\_HOST:8581/rest/devices/manageable/<device item id>](http://DA_HOST:8581/rest/devices/manageable/<device item id>)

**Headers:** [Content-Type=application/xml]

**Body Syntax:**

```
<ManageableDevice version="1.0.0">
  <RWSNMPProfileID>DA_RWSNMP_PROFILE_ID</RWSNMPProfileID>
  <RWSNMPProfileVersion>SNMP_VERSION</RWSNMPProfileVersion>
</ManageableDevice>
```

**TIP**

To find the Item ID of an SNMP profile, go to the following REST URL: [http://DA\\_HOST:8581/rest/profiles/snmpv1](http://DA_HOST:8581/rest/profiles/snmpv1)

**Example:** This example associates the SNMP Profile identified by Item ID 736 with device 1144:

#### PUT URL

```
http://DA_host:8581/rest/devices/manageable/1144
```

#### Body

```
<ManageableDevice version="1.0.0">
  <RWSNMPProfileID>736</RWSNMPProfileID>
  <RWSNMPProfileVersion>SNMPV1</RWSNMPProfileVersion>
</ManageableDevice>
```

#### Create the RTT Profiles

Use a REST client to create RTT profiles. DX NetOps Performance Management uses the RTT profiles to configure the tests on the target devices. Each profile initiates the specified action type to create, update, or delete a test on the device. The profile is not used again after DX NetOps Performance Management performs the specified action. Old RTT profiles are deleted.

For information about how to configure the RTT profiles, see [RTT Configuration Details](#). For configuration examples, see [RTT Configuration Examples](#).

## RTT Configuration Details

Use a REST client to issue REST POST requests to create, delete, or update tests.

#### Endpoints

```
http://DA_host:8581/rest/rttProfiles/TestType
```

The following example shows the basic structure of the POST REST request:

```
<RTT_Profile version="1.0.0">
  <Item version="1.0.0">
    <Name>...</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>...</ActionType>
      <TestHost>...</TestHost>
      <IPDomainID>...</IPDomainID>
      <OptionalAttribute0>...</OptionalAttribute0>
      <OptionalAttribute1>...</OptionalAttribute1>
    </AttrGroup>
  </AttrGroupList>

</RTT_Profile>
```

**RTT\_Profile** specifies the test type and determines which attributes are required and optional. Use one of the following values:

- DNSRoundTripTestProfile
- IcmpEchoRoundTripTestProfile
- IcmpPathEchoRoundTripTestProfile
- ICMPJitterRoundTripTestProfile
- HTTPRoundTripTestProfile
- TCPRoundTripTestProfile
- JitterRoundTripTestProfile

### Action Types

The following table provides details about the operations that you use to configure the tests:

Logical Operation	REST Operation	Action Type	Endpoint	Description
Create	POST	CREATE	/rest/rttpfiles /<TestType>	Creates a test profile.
Force Create	POST	FORCE_CREATE	/rest/rttpfiles /<TestType>	Creates a test profile, but does not verify uniqueness.
Delete	POST	DELETE	/rest/rttpfiles /<TestType>	Deletes a test profile, including all the associated test instances.
Update	POST	UPDATE	/rest/rttpfiles /<TestType>	Deletes a test profile, and creates a new test profile that uses the same underlying MIB object index.
Get all test profiles	GET	GET	/rest/rttpfiles /<TestType>	Fetches all profiles of the specified test type.
Get specific test profile	GET	GET	/rest/rttpfiles /<TestType>/<ID>	Fetches a profile with the specified type and ID.

**TestType** specifies the RTT test type. Use one of the following values:

- dns
- icmpecho
- icmppathecho
- icmpjitter
- http
- tcp
- jitter

### CREATE

This operation creates an instance of a test profile of the specified type with ActionType equal to CREATE. The created profile conceptually represents a background job to perform an SNMP set to create an RTT test on a device. The test is created permanently.

The create operation succeeds only if:

- The TestHost is a discovered device.
- The attributes of the requested test do not match an existing test in DX NetOps Performance Management.
- The attributes are semantically correct. The TargetHost does not reject the attributes because values are incorrect or the combination of attributes is invalid.

The result of a CREATE request is reflected in the read-only <Result> attribute of the created test. Perform a GET to view the result.

## FORCE\_CREATE

This operation is similar to CREATE, but creates the test even if an existing test has the same attribute values.

## DELETE

This operation deletes a test that is identified by ItemIDs or by supplied attributes.

Deletes all tests that match the supplied attributes or the specific tests that matches the ItemIDs.

ItemIDs supersede attributes.

## UPDATE

This operation re-applies attributes to an existing test. The operation deletes the test, and creates a test with new attributes. UPDATE changes only the supplied attributes. UPDATE runs DELETE, then CREATE. To preserve the relationship between the test configuration and the discovered Response Path Test components, the new test reuses the underlying MIB object index.

## Attributes

### NOTE

Depending on the Cisco IOS, some attributes may not apply to your device.

The following attributes apply to all test types:

- **ActionType** The action type of RTT test.
- **TestHost**  
The address of the device on which the test runs.
- **ItemID**  
Specifies the ID of the target test for DELETE and UPDATE actions. To get the ItemID, see [Get the ItemID for a Test](#)
- **IPDomainID** Specifies the IP Domain of the test. IP Domain is required for all actions that do not include ItemID.
  - CREATE always requires IPDomainID.
  - DELETE requires IPDomainID to delete a test without specifying the ItemID.
  - UPDATE never requires IPDomainID.
- (Optional) **Owner**  
Specifies the test owner.
- (Optional) **Tag**  
A short string that identifies the test in logging and notification.
- (Optional) **Threshold**  
Specifies that the test generates a threshold event if test takes longer than specified milliseconds.
- (Optional) **Frequency**  
Duration in seconds between initiating each RTT test.
- (Optional) **Timeout**  
Duration in milliseconds to wait for an RTT operation completion.
- (Optional) **VrfName**

Specifies the VPN name where the RTT operation is used . The agent uses this field to identify the VPN routing table for the operation.

- (Optional) **Persist**

Indicates whether this test configuration should be saved when persisting agent configuration to non-volatile storage. If left unspecified, the default is true.

#### NOTE

Cisco recommends persisting the tests. If the configuration is lost on the device, DX NetOps Performance Management does not re-provision the test.

This table summarizes the attributes that are associated with each IPSLA test. R indicates Required, and O indicates Optional:

Parameter	Description	ICMP Echo (Ping)	ICMP Path Echo	ICMP Jitter	UDP Jitter	UDP Jitter (VoIP)	DNS	TCP	HTTP
SourceAddress	The address of the device on which the test runs.	O	O	O	O	O	O	O	O
SourcePort	Specifies the source address port number. If the port is unspecified, the system selects a port.	N/A	N/A	N/A	O	O	O	O	O
TargetAddress	The destination IP addresses of the RTT test.	R	R	R	R	R	N/A	R	N/A
TargetPort	The destination port to which test probes are sent.	N/A	N/A	N/A	R	R	N/A	R	N/A
RequestSize	The request probe payload size.	(28)	(28)	N/A	(32)	(32)	N/A	N/A	N/A
ResponseSize	The response probe payload size.	N/A	O	N/A	O	O	N/A	N/A	N/A

TypeOfService	The type of service octet in an IP header.	O	O	O	O	O	N/A	O	O
Interval	The inter-packet delay in milliseconds between packets.	N/A	N/A	N/A	O	O	N/A	N/A	N/A
NumPackets	Number of packets to transmit.	N/A	N/A	N/A	O	O	N/A	N/A	N/A
CodecType	The codec type to use with jitter probe.	N/A	N/A	N/A	N/A	O	N/A	N/A	N/A
CodecInterval	The inter-packet delay in milliseconds between packets. Valid only for jitter probe which uses CodecType.	N/A	N/A	N/A	N/A	O	N/A	N/A	N/A
CodecPayload	Number of octets to place into the Data portion of the message. Valid only for jitter probe which uses CodecType.	N/A	N/A	N/A	N/A	O	N/A	N/A	N/A
CodecNumPackets	Number of packets to transmit . Valid only for jitter probe which uses CodecType.	N/A	N/A	N/A	N/A	O	N/A	N/A	N/A
ICPIFAdvFactor	Used while calculating jitter ICPIF values.	N/A	N/A	N/A	N/A	O	N/A	N/A	N/A



TargetAddressString	Specifies the address of the target. This string can be an IP address or a hostname.	N/A	N/A	N/A	N/A	N/A	R	N/A	N/A
NameServer	Specifies the IP address of the name-server.	N/A	N/A	N/A	N/A	N/A	R	N/A	N/A
Url	Specifies the URL that the HTTP probe targets.	N/A	N/A	N/A	N/A	N/A	N/A	N/A	R
Operation	Specifies the HTTP operation that represents the specific type of RTT operation.	N/A	N/A	N/A	N/A	N/A	N/A	N/A	R
HTTPVersion	Specifies the version number of the HTTP server.	N/A	N/A	N/A	N/A	N/A	N/A	N/A	O
String1	Stores the content of a raw HTTP request. If the request cannot fit into String1, use more String attributes. *Required only if <b>Operation</b> is set to <b>httpRaw (2)</b> .	N/A	N/A	N/A	N/A	N/A	N/A	N/A	R*
String2	Continues the raw HTTP request from <b>String1</b> .	N/A	N/A	N/A	N/A	N/A	N/A	N/A	O

String3	Continues the raw HTTP request from <b>String2</b> .	N/A	N/A	N/A	N/A	N/A	N/A	N/A	O
String4	Continues the raw HTTP request from <b>String3</b> .	N/A	N/A	N/A	N/A	N/A	N/A	N/A	O
String5	Continues the raw HTTP request from <b>String4</b> .	N/A	N/A	N/A	N/A	N/A	N/A	N/A	O

**NOTE**

When you specify an attribute value that has an associated enumerated name as defined by the MIB, only the numeric value is recognized.

**Result**

The result attribute is a read-only attribute that indicates the outcome of the request:

Meaning	<Result> value
The operation is in progress.	PENDING
The operation succeeded.	SUCCESS
The TestHost is not discovered.	FAILURE: Unable to identify device at 192.168.96.1
The requested test is already configured.	ALREADY_EXISTS: Existing Count = 1
TestHost is unreachable.	FAILURE: {http://im.ca.com/normalizer}NormalizedJitterInfo.RttType: NO_RESPONSE - {http://im.ca.com/certifications/snmp}CiscoIPSLAJitterMib.rttMonCtrlAdminRttType: SNMP timeout {http://im.ca.com/normalizer}NormalizedJitterInfo.TargetAddress: NO_RESPONSE - {http://im.ca.com/certifications/snmp}CiscoIPSLAJitterMib.rttMonEchoAdminTargetAddress: SNMP timeout {http://im.ca.com/normalizer}NormalizedJitterInfo.Interval: NO_RESPONSE - {http://im.ca.com/certifications/snmp}CiscoIPSLAJitterMib.rttMonEchoAdminInterval: SNMP timeout {http://im.ca.com/normalizer}NormalizedJitterInfo.CodecInterval: NO_RESPONSE - {http://im.ca.com/certifications/snmp}CiscoIPSLAJitterMib.rttMonEchoAdminCodecInterval: SNMP timeout

Invalid credentials, such as SNMP SET profile.	{http://im.ca.com/normalizer}NormalizedICMPInfo.RttType: OPERATION_NOT_ALLOWED - {http://im.ca.com/certifications/snmplib}CiscoRttMonStatsMib.rttMonCtrlAdminRttType: SNMP error 6: No access{http://im.ca.com/normalizer}NormalizedICMPInfo.TargetAddress: ATTRIBUTE_NOT_SET{http://im.ca.com/normalizer}NormalizedICMPInfo.Owner: ATTRIBUTE_NOT_SET{http://im.ca.com/normalizer}NormalizedICMPInfo.Tag: ATTRIBUTE_NOT_SET
The Data Collector associated with the TestHost is down.	<ul style="list-style-type: none"> <li>• SOURCE_NOT_AVAILABLE</li> <li>• SHUTDOWN</li> </ul>
The test was not present on the TestHost because another SNMP client deleted the test.	FAILURE: 1.3.6.1.4.1.9.9.42.1.2.1.1.9.546811228: Commit failed. Error Index: 1

### Get the ItemID for a Test

The UPDATE and DELETE operations use ItemID to identify tests. To get the ItemID, use a filtered REST endpoint:

- *http://DA\_HOST:8581/rest/responsepathtest/filtered*  
Searches all test types.
- *http://DA\_HOST:8581/rest/responsepathicmp/filtered*  
Searches ICMP Echo tests.
- *http://DA\_HOST:8581/rest/responsepathecho/filtered*  
Searches ICMP Path Echo tests.
- *http://DA\_HOST:8581/rest/responsepathjitter/filtered* Searches UDP Jitter tests.
- *http://DA\_HOST:8581/rest/responsepathicmpjitter/filtered*  
Searches ICMP Jitter tests.
- *http://DA\_HOST:8581/rest/responsepathhttp/filtered*  
Searches HTTP tests.
- *http://DA\_HOST:8581/rest/responsepathtcp/filtered*  
Searches TCP Connect tests.
- *http://DA\_HOST:8581/rest/responsepathdns/filtered*  
Searches DNS tests.

### Example Request

This example shows a filtered request for tests with the specified target:

```
<FilterSelect xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance'
  xsi:noNamespaceSchemaLocation='filter.xsd'>
  <Filter>
    <ResponsePathTest.Target type='EQUAL'>10.42.94.250</ResponsePathTest.Target>
  </Filter>
  <Select use="exclude"/>
</FilterSelect>
```

#### TIP

To exclude all attributes except ItemID, the following statement:

```
Select use="exclude"
```

### Example Response

This example shows the response to a filtered request:

```
<ResponsePathIcmpList>
  <ResponsePathIcmp version="1.0.0">
    <ID>955</ID>
  </ResponsePathIcmp>
  <ResponsePathIcmp version="1.0.0">
    <ID>956</ID>
  </ResponsePathIcmp>
</ResponsePathIcmpList>
```

## RTT Configuration Examples

Use the following examples to help configure your tests.

### Manage RTT Tests

The `rtt_rest_configuration_examples.py` script creates an example POST. Use the following example as a template to create, delete, and modify RTT tests:

```
$ ./rtt_rest_configuration_examples.py
Usage: rtt_rest_configuration_examples.py [--quiet] [--dry-run] <DA host> create|
delete ping|jitter [<attribute_name> <attribute_value> ...]
```

#### Options

```
--quiet      Suppress output
--dry-run    Construct REST request, but do not send
```

#### Test Attributes

```
-----
```

```
-----
icmpecho DeviceIp IpDomainId Owner Threshold Tag Frequency Timeout TargetAddress
RttTestName SourceAddress
jitter DeviceIp IpDomainId Owner Threshold Tag Frequency Timeout TargetAddress
RttTestName SourceAddress
```

#### Examples

```
% ./rtt_rest_configuration_examples.py testda1.ca.com create icmpecho DeviceIp
10.165.170.222 IpDomainId 2 TargetAddress 10.132.159.140 SourceAddress 10.117.139.209
Frequency 180
% ./rtt_rest_configuration_examples.py testda1.ca.com delete jitter DeviceIp
10.103.224.132 IpDomainId 2 TargetAddress 10.132.159.140
```

(See script comments for more examples)

## Example RTT Configurations

### ICMPECHO CREATE

This example creates an ICMP Echo test:

```
<IcmpEchoRoundTripTestProfile version="1.0.0">
  <Item version="1.0.0">
    <Name>Example ICMP Echo Test</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>CREATE</ActionType>
      <TestHost>10.250.134.47</TestHost>
      <IPDomainID>2</IPDomainID>
      <TargetAddress>10.0.63.106</TargetAddress>
      <Frequency>62</Frequency>
      <Timeout>5000</Timeout>
      <Threshold>5000</Threshold>
      <Owner>tedison</Owner>
      <Tag>ECHO02</Tag>
      <RequestSize>40</RequestSize>
    </AttrGroup>
  </AttrGroupList>

</IcmpEchoRoundTripTestProfile>
```

### ICMPECHO DELETE

This example deletes three ICMP Echo test with the specified ItemIDs:

```
<IcmpEchoRoundTripTestProfile version="1.0.0">
  <Item version="1.0.0">
    <Name>Delete ICMP Echo Test</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>DELETE</ActionType>
      <ItemId>800</ItemId>
    </AttrGroup>
    <AttrGroup>
      <ActionType>DELETE</ActionType>
      <ItemId>2311</ItemId>
    </AttrGroup>
    <AttrGroup>
      <ActionType>DELETE</ActionType>
      <ItemId>1021</ItemId>
    </AttrGroup>
  </AttrGroupList>
```

```
</IcmpEchoRoundTripTestProfile>
```

## ICMPECHO UPDATE

This example updates the TargetAddress and Owner for the ICMP Echo test with the specified ItemID:

```
<IcmpEchoRoundTripTestProfile version="1.0.0">
  <Item version="1.0.0">
    <Name>Update ICMP Echo Test</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>UPDATE</ActionType>
      <ItemID>921</ItemID>
      <TargetAddress>10.52.217.32</TargetAddress>
      <Owner>admin_andy</Owner>
    </AttrGroup>
  </AttrGroupList>

</IcmpEchoRoundTripTestProfile>
```

## ICMPPATHECHO CREATE

This example creates an ICMP Path Echo test:

```
<IcmpPathEchoRoundTripTestProfile version="1.0.0">
  <Item version="1.0.0">
    <Name>Example ICMP Path Echo Test</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>CREATE</ActionType>
      <TestHost>10.250.134.47</TestHost>
      <IPDomainID>2</IPDomainID>
      <TargetAddress>10.40.29.130</TargetAddress>
      <Owner>att</Owner>
      <Tag>seattle1</Tag>
      <Frequency>180</Frequency>
      <Timeout>7777</Timeout>
      <Threshold>6666</Threshold>
    </AttrGroup>
  </AttrGroupList>

</IcmpPathEchoRoundTripTestProfile>
```

## ICMPPATHECHO DELETE

This example deletes an ICMP Path Echo test with the specified attributes:

```

<IcmpPathEchoRoundTripTestProfile version="1.0.0">
  <Item version="1.0.0">
    <Name>Example ICMP Path Echo Test</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>DELETE</ActionType>
      <TestHost>10.250.134.47</TestHost>
      <IPDomainID>2</IPDomainID>
      <TargetAddress>10.0.63.106</TargetAddress>
    </AttrGroup>
  </AttrGroupList>
</IcmpPathEchoRoundTripTestProfile>

```

## JITTER CREATE

This example creates two Jitter tests:

```

<JitterRoundTripTestProfile version="1.0.0">
  <Item version="1.0.0">
    <Name>Create Jitter Test</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>CREATE</ActionType>
      <TestHost>10.188.72.48</TestHost>
      <IPDomainID>2</IPDomainID>
      <TargetAddress>10.16.75.221</TargetAddress>
      <TargetPort>33333</TargetPort>
      <Owner>admin_dan</Owner>
      <Tag>JITTER022</Tag>
    </AttrGroup>

    <AttrGroup>
      <ActionType>CREATE</ActionType>
      <TestHost>10.84.206.153</TestHost>
      <IPDomainID>2</IPDomainID>
      <TargetAddress>10.42.96.10</TargetAddress>
      <Owner>admin_steve</Owner>
      <Tag>JITTER023</Tag>
      <Frequency>120</Frequency>
      <Timeout>10000</Timeout>
      <CodecType>1</CodecType>
      <CodecInterval>20000</CodecInterval>
      <CodecPayload>172</CodecPayload>
      <CodecNumPackets>100</CodecNumPackets>
    </AttrGroup>
  </AttrGroupList>
</JitterRoundTripTestProfile>

```

```
</AttrGroupList>
```

```
</JitterRoundTripTestProfile>
```

## JITTER DELETE

This example deletes a Jitter test with the specified attributes:

```
<JitterRoundTripTestProfile version="1.0.0">
  <Item version="1.0.0">
    <Name>Example Jitter Test</Name>
  </Item>

  <AttrGroupList>
    <AttrGroup>
      <ActionType>DELETE</ActionType>
      <TestHost>10.250.134.47</TestHost>
      <IPDomainID>2</IPDomainID>
      <TargetAddress>10.237.30.166</TargetAddress>
    </AttrGroup>
  </AttrGroupList>

</JitterRoundTripTestProfile>
```

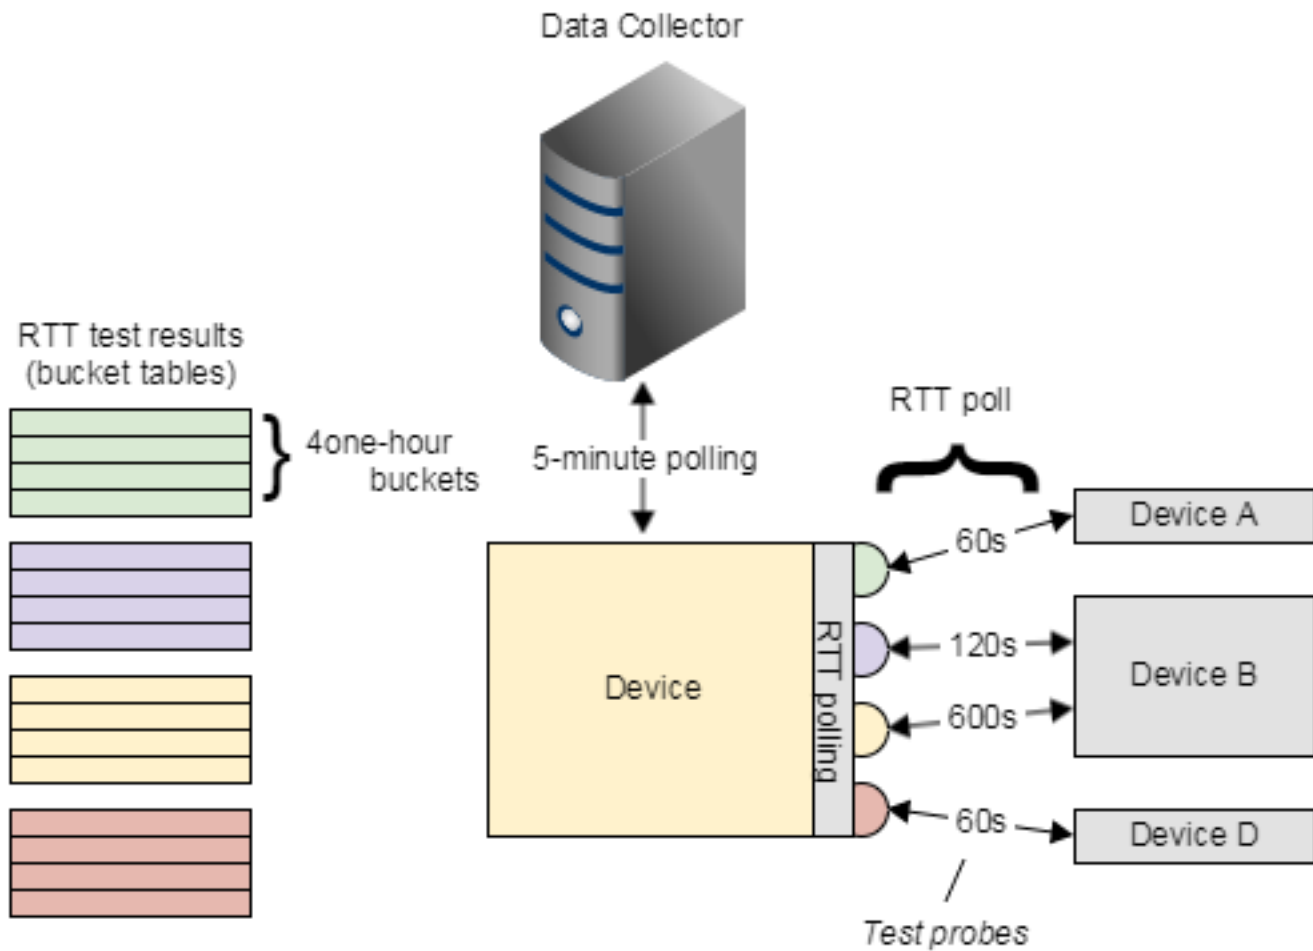
## IPSLA Polling

The data aggregator and data collectors poll IPSLA-enabled devices, which can run RTT tests. An RTT test frequency controls how often an active test runs (for example, every 60 seconds). While a test is active its results are stored in a statistics table. Each row represents one hour of results for a specific test. Test result metrics (for example, `rttMonStatsCollectTimeouts`) are stored as counters and increment after each execution of the test. At the end of an hour a new bucket is created. When the maximum number of buckets is reached, the oldest bucket is dropped. Metrics for new buckets are reset to zero and begin to increment. Disabled tests stop counting. The DX NetOps Performance Management poll rate controls how often the data collectors request these statistics. Each time the data collectors poll the counter, they subtract the last value from the current to produce a delta value for a given metric.

For more information about how to set the poll rate, see [Manage Monitoring Profiles](#).

The following diagram illustrates IPSLA polling:



**Figure 20: IPSLA Polling**

## Set Alias Names Using a Script

You can set the alias names for one or more monitored devices, interfaces, and components simultaneously. The alias appears in the inventory lists for the device, interface, or component.

### NOTE

You can also set the alias name for a monitored device, an interface, or a component using the `devices` endpoint for the NetOps Portal REST web service. Alias names set from the REST web service take precedence over alias names set using the script.

For more information:

- About setting alias names for monitored devices using REST, see [Set Device Alias Names Using the Devices Web Service](#).
- About setting alias names for interfaces using REST, see [Set Interface Alias Names Using the Devices Web Service](#).
- About setting alias names for components using REST, see [Set Component Alias Names Using the Devices Web Service](#).

**Follow these steps:**

1. From an open command prompt, access the `<installation_directory>/PerformanceCenter/Tools/bin` directory.

- **installation\_directory**

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

2. Issue the following command:

```
./update_alias_name.sh
```

3. Export the CSV file by issuing the following command and options:

```
./update_alias_name.sh -h <host_name> -u <username> -p <password> -T <item_type> -o <output_filename>
```

- **host\_name**

Specifies the NetOps Portal host name.

- **username**

Specifies the username of the administrator who sets the alias names.

- **password**

Specifies the password for the NetOps Portal administrator who sets the alias names.

- **item\_type**

Specifies the type of item for which you want to set alias names.

**Values:** device, interface, component

**Default:** device

- **(Optional) output\_filename**

Specifies the name of the CSV file with the total number of monitored devices, interfaces, or components by item ID and device name. Use the `-o` option to specify to override the default file name. If you do not use this option, `DeviceList.csv` is used as the CSV filename.

The CSV file has the following format:

- (If you used the `-T device` option)

```
device Item_ID, <device> Name
```

**Examples:**

- 123,MyDevice123
- 123,MyDevice456
- (If you used the `-T component/interface` option)

```
device Item_ID, <interface/component> Item_ID, <interface/component> Name
```

**Examples:**

- 123,456,MyComponent456
- 123,789,MyInterface789

A list of monitored device, interface, or component item IDs and monitored device, interface, or component names is returned in CSV format.

4. Add the alias names that you want to set for each monitored device, interface, or component to the CSV file using the following format:

- (If you used the `-T device` option when you exported the CSV file)

```
device Item_ID, <device> Alias Name
```

For devices, **URL-encode the Alias Name value**, for example, URL-encode spaces as “%20”. You can use commas in the Alias Name value.

**Examples:**

- 123,MyDevice123Alias
- 123,MyDevice456Alias
- (If you used the `-T component/interface` option when you exported the CSV file)

```
device Item_ID, <interface/component> Item_ID, <interface/component> Alias Name
```

For interfaces and components, XML-encode the `Alias Name` value. XML-encode ampersands (&) as "&amp;," less-than characters (<) as "&lt;," and greater-than characters (>) as "&gt;". You can use commas and spaces in the `Alias Name` value.

**Examples:**

- 123,456,MyComponent456Alias
- 123,789,MyInterface789Alias

If the Item IDs in your CSV file are invalid, the entries are ignored.

5. Import the CSV file by issuing the following command and options:

```
./update_alias_name.sh -h <host_name> -u <username> -p <password> -T <item_type> -i <input_filename>
```

**NOTE**

(Optional) To control the workload when setting the alias names for many monitored devices, interfaces, or components, adjust the batch size and create pauses between batches by issuing the command with the following additional options:

```
./update_alias_name.sh -h <host_name> -u <username> -p <password> -T <item_type> -i  
  <input_filename>  
-b <batch_size>  
-t <time_in_seconds>
```

- **batch\_size**  
Indicates the number of items to process in each batch.  
**Default:** 10000  
**Default with the -i option:** 150
- **time\_in\_seconds**  
Indicates the time, in seconds, to pause between batches.  
**Default:** 1  
**Default with the -i option:** 1  
**Example:**

```
./update_alias_name.sh -h host_name -u username -p password -T device -  
i input_file -b 20 -t 2
```

- **host\_name**  
Specifies the NetOps Portal host name.
- **username**  
Specifies the username of the administrator who sets the alias names.
- **password**  
Specifies the password for the NetOps Portal administrator who sets the alias names.
- **item\_type**  
Specifies the type of item for which you want to set alias names.  
**Values:** device, interface, component  
**Default:** device
- **input\_filename**  
Specifies the name of the CSV file.

The script reads the updated CSV file and sets the alias names for the monitored devices, interfaces, or components. If you do not use the `-i` option with the command, the script locates the item IDs that are required for the specified type, and creates a CSV file with item IDs and item names.

## Using

---

DX NetOps Performance Management monitors the health of your environment, including networks, applications, and devices. NetOps Portal displays relevant data using dashboards and context pages. These dashboards and context pages include views that show different information about your infrastructure.

## Launch NetOps Portal

View performance information and log in to NetOps Portal.

NetOps Portal is the web UI for DX NetOps Performance Management. Use NetOps Portal with any modern web browser to view infrastructure data, configure, collect, and perform administrative tasks.

To view the supported web browsers, see [Installation Requirements and Considerations](#).

The following video shows the NetOps Portal login process and provides more information about the authentication process:

### Follow these steps:

1. Open a web browser.
2. In the address field, enter the following address:  
`http://<PC_host>:<port>`
  - **PC\_host**  
The IP address or hostname of the NetOps Portal host.
  - **port**  
Specifies the NetOps Portal required port number.  
**Default:** 8181
3. Specify your NetOps Portal username and password.
4. (Optional) Select **Remember me on this computer** to save your login session for 15 days. This option prevents your session from being deleted when you log out, time out, or you close the browser.
5. Click **Log In**.  
NetOps Portal opens to your home dashboard.

## Search and Filter in NetOps Portal

Search and filter in NetOps Portal using the search controls that NetOps Portal includes, such as global search and quick filter.

In this article:

- [Global Search](#)  
Use to locate managed items in NetOps Portal.
- [Limit Global Search](#)  
For high scale environments, use to improve performance.
- [Use Wildcards in Filter Expressions](#)
- [Quick Filter](#)  
Use to limit the contents of table views to only those items that contain the search criteria that you enter.

## Global Search

Global searches return lists of the items in the inventory that match the search string. The lists are sorted by item type (such as devices, interfaces, device components).

To perform a global search, click **Search** (the magnifying glass) at the top of any page. The following image shows an example of the global search bar:

**Figure 21: Global\_Search\_bar**



By default, the global search is set to search only device and interface names that contain the search string (Scope: Names).

To widen the search scope and search all managed items (for example, to search for tunnels) or search across other identifying information about an entity (for example, to search for ip-addresses) by selecting **Scope** in the search box, and then selecting **Search All** (Scope: All) from the drop-down. The following image shows the drop-down:

**Figure 22: Global\_Search\_dropdown**



## Limit Global Search

You can limit the search scope in the **Search All** context using filter expressions. In the most basic form, the filter expression uses the following syntax:

```
ViewName:ColumnName=value
```

To search multiple views or columns, use the following syntax:

```
ViewName1:Column1;ColumnN&ViewName2:Column1;ColumnN=value
```

## Filter Expression Examples

The following are examples of filter expressions you can use to limit the search scope:

- [Example: Search Devices Containing the String 'foo' Anywhere in the 'Name' Column](#)
- [Example: Search Interface Names or Descriptions and Devices that Include 'foo'](#)
- [Example: Search All Views](#)

### Example: Search Devices Containing the String 'foo' Anywhere in the 'Name' Column

Use the following search expression:

```
Devices:Name=foo
```

### Example: Search Interface Names or Descriptions and Devices that Include 'foo'

Use the following search expression:

```
Interfaces:Interface Name;Description&Devices:Name=foo
```

### **Example: Search All Views**

Use the following search expression:

```
Name;Description=foo
```

If the view does not include these columns, the expression searches all columns of this view.

### **Example: Search All Columns for a View**

Use the following search expression:

```
Interfaces:*=foo
```

### **Use Wildcards in Filter Expressions**

By default, global search searches for text that is contained in an item string. For example, searching for 'foo' or 'oo' returns 'Foo' as a result. To narrow or broaden your search, for example, to locate column names, add the asterisk (\*) wildcard character to the filter expression.

You can search within interfaces view columns that contain 'Name'--such as "Name Alias", "Interface Name", "Device Name Alias"--using the following filter expression:

```
Interfaces:*Name*=value
```

Using asterisks at the beginning and end of a keyword ('contains' searches), for example '\*Name\*', is equivalent to entering 'Name' as the search string.

#### **NOTE**

'Starts with' searches can return quicker results than 'contains' searches. For example, 'foo\*' (starts with 'foo') executes faster than the '\*foo\*' (contains 'foo') search.

To narrow the search, add multiple search words. For example, searching for devices using the 'server 192.168\*' search string returns all servers on the 192.168.0.0/16 network.

### **Other Wildcard Examples**

- Return all rows with entries that start with 'serv':  
serv\*
- Return all rows with entries that end in 'erver':  
\*erver
- Return all rows with entries that start with 'fo' and end with '200'. For example, the following filter expression returns 'Foo\_5976\_10.92.200.200', but it does not return 'Foo\_5976\_10.92.200.201':  
fo\*200
- Return all rows with entries that:
  - Start with 'fo'
  - Contain at least one character following 'fo'
  - Contain '200'
  - End with any character
 fo\*200\*

## Quick Filter

You can limit the contents of many table views to only those items that contain the search criteria that you enter using the quick filter. Quick filtering a table (for example, an inventory table that shows devices) is similar to performing a global search and looking at the results for that kind of managed item.

The following image shows an example of a quick filter:

**Figure 23: quick\_filter**



To quick filter a table view, type the search string into the quick filter box, and then either click filter (the funnel icon) or press the Return key on the keyboard.

To clear a quick filter, delete the search term, and then click filter (the funnel icon) or press the Return key on the keyboard.

## Customize Your User Settings

You can customize your user account's default dashboard and personal settings, change user passwords, as well as change device and interface display names, and suppressing views.

You can do the following to customize your user account's settings:

- [Set a dashboard as your home page.](#)
- [Customize your user account settings.](#)
- [Change your password.](#)

### Set a Dashboard as your Home Page

To log in to your preferred dashboard, set that dashboard as your home page. By default, the first dashboard in your first menu is your home page.

#### TIP

To return to your home page from any other page, click the logo in the upper left corner.

The following video shows how to set your default dashboard:

#### Follow these steps:

1. Navigate to the dashboard that you want to set as your home page.
2. (Optional) To set a specific context as your default, click the **[change]** link, and then select the group context for the dashboard. The home page saves the context.

#### NOTE

If the selected group is removed from your permission set, your default permission group is used for the dashboard context.

3. Click **More**, and then click **Set as Home Page**.
4. In the confirmation dialog, click **Yes**.

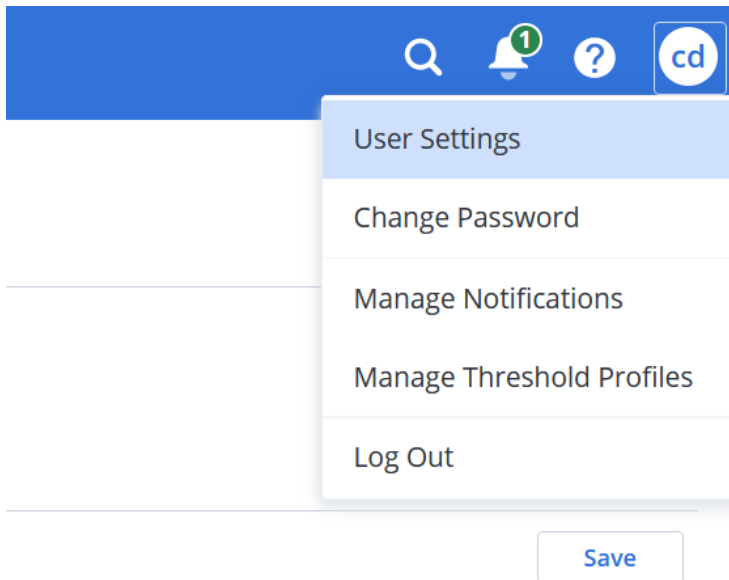
The selected dashboard is now your home page.

### Customize your User Account Settings

User accounts include settings that you can customize, such as the preferred language and the time zone. Your role rights determine the settings that you can customize. If your user account has the required role right, you can change these settings.

**Follow these steps:**

1. Click the name of your user account in the upper-right corner, and then click **User Settings**. The following image shows this option:



The **User Settings** dialog opens.

2. Modify the following user settings in the **Basic Settings** section:

- **Preferred Language**

Specifies the language for NetOps Portal. NetOps Portal displays this language regardless of the language selected for the operating system or for the browser language.

**NOTE**

For a language to display appropriately, the relevant fonts must be installed.

For more information, see [Prepare to Install NetOps Portal](#).

**Options:**

- English (US)
- Japanese
- French (France)

For more information about the languages that DX NetOps Performance Management supports, see [Language Support](#).

- **Email Address**

Defines the email address for the user account.

- **Default Group**

The default context when you log in. The list includes only those groups from your permission groups.

**Default:** /My Assigned Groups/All Groups

3. Modify the following settings in the **Date and Time Display Settings** section:

- **Time Zone**

Defines the time zone.

**NOTE**

Changing the time zone after you have set up email schedules in NetOps Portal can cause incorrect times to appear in lists of email schedules in NetOps Portal.

**Default:** UTC (Coordinated Universal Time)



**Required:** Yes

– **Time Display Format**

**NOTE**

This option is available only if you have chosen English as the language (locale) for NetOps Portal to display dates.

Defines the default time format.

**Options:** 12 Hour Format, 24 Hour Format

**Default:** 12 Hour Format

**Required:** Yes

– **Date Display Format**

Defines the length to display dates in NetOps Portal.

**Options:**

- Short: Displays dates as numbers, for example, 3/24/23.
- Medium: Displays dates using month name abbreviations, for example, Mar 24, 2023.
- Long (long month names): Displays dates using full month name, for example, March 24, 2023.
- Full: Displays dates including the day of the week, for example, Tuesday, March 24, 2023.

**Default:** Medium

**Required:** Yes

– **Date Display Language**

Specifies the language (locale) for NetOps Portal to display dates.

**NOTE**

The date display language is independent from the language that you set for NetOps Portal. For example, if you set the preferred language to French and then set the date display language to German, NetOps Portal displays in French but displays the dates in German.

**Examples:**

- English (United States)
- English (United Kingdom)
- French (France)
- Japanese (Japan)
- German (Germany)

**Default:** Your preferred language

**Required:** Yes

4. Modify the following settings in the **Date and Time Display Settings** section:

– **View Suppression**

You can hide, or suppress, views when a required data source is not registered or when a required technology is not configured. Similar behavior applies to context tabs and custom menus. Menus and tabs that include only suppressed views are hidden. View suppression applies to the default views on out-of-the-box dashboards. View suppression does not hide views from administrators when they use the view categories to edit a dashboard.

When the data source that populates a view is registered, that view is no longer hidden.

**Options:**

- **Display All Views:** Display all views on dashboards (disable view suppression).

**TIP**

Disable view suppression for troubleshooting purposes, or to help you decide whether to configure, or register, another data source.

- **Suppress Views:** Hide, or suppress, all views on dashboards (enable view suppression).

**NOTE**

When views are suppressed and you delete (unregister) a data source, those dashboards and menus containing suppressed views might not appear in NetOps Portal. Dashboards, context tabs, and custom menus must contain at least one unsuppressed view.

For more information about the affects of suppressing views and deleting data sources, see [Configure a Data Source](#).

**Default:** Display All Views

**Required:** Yes

– **Item Name Display Setting**

(Requires the View Item Display Name or Name Alias role right) Determines whether device and interface names appear as the display name or as the alias in dashboards and views.

**Options:**

- **Use Item Display Name:** Users see the device and interface names as the display name in dashboards and views.
- **Use Item Name Alias:** Users see the device and interface names as the alias in dashboards and views.

**Default:** Use Item Display Name

**Required:** Yes

5. Modify the following accessibility settings for NetOps Portal in the **Ease of Access Settings** section:

– **Audio Notifications**

Defines whether this user account receives audio notifications when NetOps Portal adds new alarms and events to the **Alarms** and **Events** views on dashboards that include them, and when **Auto Refresh Page** is turned on (enabled).

**NOTE**

By default, **Auto Refresh Page** is turned off (disabled).

For more information about **Auto Refresh Page**, see [Manage Dashboards](#).

**Options:**

- **Selected (Enabled):** This user account does not receive audio notifications.
- **Cleared (Disabled):** This user account receives audio notifications.

**Default:** This user account does not receive audio notifications.

– **Underline Links**

Specifies whether the links in NetOps Portal are underlined or drawn in blue-colored text.

**Options:**

- **Selected (Enabled):** The links are underlined.
- **Cleared (Disabled):** The links are drawn in blue-colored text.

**Default:** The links are drawn in blue-colored text.

6. Save your changes.

The user account settings are saved.

## **Change your Password**

You can change the password for your user account.

**Prerequisite:** LDAP is *not* configured for your system.

**Follow these steps:**

1. Click the name of your user account in the upper-right corner, and then click **Change Password**. The **Change Password** dialog displays.
2. Enter the old password.
3. Enter the new password.

By default, user passwords must meet the following requirements:

- Be different than the username
- Have a minimum length of 8 characters
- Have a maximum length of 30 characters
- Contain at least three of the following types of characters:
  - Special characters
  - Uppercase letters (A-Z)
  - Lowercase letters (a-z)
  - Numbers (0-9)

**TIP**

You can disable these requirements.

4. Confirm the new password.
5. Click **Save**.

The user password is changed.

## Share Data with Other Users

Share dashboard, context page, and view data with others as reports.

You can share data with others by creating or accessing reusable on-demand report templates, and then downloading (running) a report for sharing with others. From dashboards, views, and context pages, you can also download the data for sharing with others.

The following video shows how to generate digital and printed reports from the content of dashboards and context pages in NetOps Portal:

For more information about how to create or access on-demand report templates, see [On-Demand Reports](#).

Share dashboard, view, and context page data using one of the following options:

- [Download dashboard data as a report in PDF or comma-separated values \(CSV\) format.](#)
- [Set up schedules to have the email server send dashboard data as reports from scheduled reports attached to emails automatically.](#)

**NOTE**

If you have Send Reports to Archive role right and archive reports is enabled, when you use this option, NetOps Portal saves dashboard data as scheduled reports. When the scheduled report runs, NetOps Portal saves the results to disk.

- [Share views with users who do not have access to a dashboard.](#)
- [Download view data as a CSV file.](#)

## Download Dashboard Data as Reports

Download dashboard data as a report in PDF or comma-separated values (CSV) format or preview the dashboard data in PDF or CSV format.

**Follow these steps:**

1. Log in as a user with the Print a Dashboard role right.  
For more information about this role right, see [Role Rights](#).
2. Do one of the following:
  - Select the dashboard for which you want to download data.
  - [Generate an on-demand report.](#)
3. Click the **More options** icon on the toolbar, **Print**, and then click one of the following options:
  - **Print PDF Portrait**

- Specifies that the page layout of the PDF document be portrait.
  - **Print PDF Landscape**  
Specify that the page layout of the PDF document be landscape.
  - **Print CSV Scaled**  
Specifies that the values in the exported dashboards are scaled. Scaled values appear with larger units, for example, 1 KB.
  - **Print CSV Unscaled**  
Specifies that the values in the exported dashboards are not scaled. Unscaled values appear in the raw form for the metric, for example, 1000 bytes.
- The dashboard data is downloaded as a report.

## Manage Scheduled Reports

You can schedule reports, edit scheduled reports, view lists of archived reports, delete archived reports, and send archived reports to others.

You can save dashboard data as reports. These reports are attached to email messages and are saved as scheduled reports. You define the frequency at which NetOps Portal runs the scheduled report and sends them as attachments to emails. You can manage your own scheduled reports. Only users with the Administrator role can manage scheduled reports that belong to other users. Download and save the report as soon as possible.

### NOTE

Scheduled reports include the information only up to the row limit. For example, the limit for the IM Table (Top) view is 2,500 by default. When an IM Table (Top) view contains a mix of core and baseline metrics, the results might be capped at 1,000 rows. This reduced row limit depends on the number of queried items.

In this article:

- [View a List of Scheduled Reports](#)
- [Schedule a Report](#)
- [View or Edit the Settings for a Scheduled Report](#)
- [View a List of Archived Reports for a Selected Scheduled Report](#)
- [View a List of Archived Run Now Reports](#)
- [Send an Archived Report to Others](#)
- [Delete an Archived Report](#)

### View a List of Scheduled Reports

**Prerequisite:** Scheduled reports exist.

Hover over **Administration**, **Configuration Settings**, and then click **All Scheduled Reports**. The **Manage Scheduled Reports** page opens, and a list of scheduled reports displays. The list displays single-page scheduled reports (CSV or PDF) and full-page scheduled reports (CSV or PDF).

### NOTE

(Tenant administrators) Only the scheduled reports that are associated with the tenant are displayed.

### Schedule a Report

**Prerequisite:** The Administrator has configured the email server.

For more information, see [Configure the Email Server](#).

### **Follow these steps:**

1. Log in as a user with one of the following role rights:

- (To run reports and have the email server sends dashboard data as reports from scheduled reports attached to emails) Send Reports by Email
  - (To create a schedule to have the email server sends dashboard data as reports from scheduled reports attached to emails on a recurring basis automatically) Send Reports on a Schedule
  - (To archive, or save, dashboard data as reports from scheduled reports) Send Reports to Archive
- If you have this role right and archive reports is enabled (the **Server Type** field on the **Manage Email Server Settings** page is set to **Use Repository Service Only** or **Use Both SMTP Mail Server and Repository Service**), NetOps Portal saves dashboard data as scheduled reports. When the scheduled report runs, NetOps Portal saves the results to disk.

For more information about how to enable archive reports, see [Configure the Email Server](#).

For more information about these role rights, see [Role Rights](#).

2. Do *one* of the following tasks:
  - Hover over **Reports, On-Demand Report Templates**. On the **Manage On-Demand Report Templates** page, select the on-demand report template that you want to run, and then click **Run**.
  - Navigate to a dashboard.
3. Click the **More options** icon on the toolbar, and then click **Email/Schedule Report**. The **Email Dashboard** dialog opens.
4. Complete the following fields:
  - **Dashboard**  
Identifies the name of the dashboard. This name appears in the filename of the report that is attached to the email, and also appears in the **Dashboard** column in the list of email schedules on the **Manage Scheduled Emails** page.  
**Read-only:** Yes
  - **Send To**  
Specifies the email addresses of the recipients. Use the standard format <name>@<domain> .  
**Required:** Yes  
**NOTE**  
If you have the Send Reports to Archive role right and archive reports is enabled, this field label shows as **Archive Owner** with the name of the user who saved the dashboard data, and is read-only.
  - **Subject**  
Describes the emailed report. Include the dashboard title and any components for which data is included in the report.  
**Required:** Yes
  - **Message**  
Message that accompanies the emailed report.  
**Required:** No
5. In the **Output Options** section, complete the following fields:
  - **Format**  
Define the format for the report.  
**Options:**
    - **CSV:** For single-paged CSV (the **First Page only of multiple page views** option for **Information scope**), only the first page of the report is included as a CSV attachment in the email message that is sent. For full-paged CSV (the **All Pages of multiple page views** option for **Information scope**), a time-sensitive link to the full-page report is included in the email message that is sent.
    - **PDF:** For single-paged PDF (the **First Page only of multiple page views** option for **Information scope**), only the first page of the report is included as a PDF attachment in the email message that is sent. For full-paged PDF (the **All Pages of multiple page views** option for **Information scope**), a time-sensitive link to the full-page report is included in the email message that is sent.
  - Default:** PDF
  - **PDF Orientation**

Specify the page layout of the PDF document.

**Options:** Portrait or Landscape

**Default:** Portrait

– **Information Scope**

Select the pages to include in the report run.

**Options:**

- **First Page only of multiple page views:** Only the first page of the report is included as a single-page report (CSV or PDF) attachment in the email message that is sent.
- **All Pages of multiple page views:** A time-sensitive link to the full-page report (CSV or PDF) is included in the email message that is sent.

**IMPORTANT**

To minimize the performance impact, before you schedule a report using this option, ensure that your servers, especially the data repository, meet the minimum requirements and sizing guidelines.

For more information about these sizing requirements, see [the DX NetOps Sizing Tool](#).

**Default:** First Page only of multiple page views

6. In the **Business Hour Option** section, complete the following fields:

– **Apply Business Hour Filter**

Enable this option. By default, business hours are not enabled.

– **Business Hours Filter**

Select the business hours filter that you want to apply to the data in the report. The list of business hours definitions are those that have been previously defined.

– **Business hour and time zone**

The time zones from which you want to choose in the Time Zone Filter drop-down.

**Options:**

- **Recommended hours/zones:** Lists the time zones that are associated with the business hours you have chosen.
- **System defined hours/zones:** Lists time zones that are associated to site groups.
- **Any hours/zones:** Lists all time zones.

**Default:** Recommended hours/zones

– **Time Zone Filter**

The time zone that you want to assign to the view.

7. In the **Scheduling Options** section, complete the following fields:

– (For NetOps business reports only) **Scheduled Copy** (23.3.4 and higher)

Defines whether NetOps Portal creates a copy of the NetOps business report when you schedule the report so that you can:

- Schedule the NetOps business report with different settings.
- View the settings of the scheduled NetOps business report (by clicking the name of the report (link) in the **Dashboard** column on the **Manage Scheduled Reports** page).
- Edit the settings of the scheduled NetOps business report (for those NetOps business reports that are listed on the **Manage Scheduled Reports** page).

For more information, see [Manage NetOps Business Reports](#).

**Prerequisites:**

- You are in the process of scheduling the NetOps business report.
- The **Frequency** field is set to a value other than Run Now.

**NOTE**

If you are editing an already-scheduled NetOps business report that is listed on the Manage Scheduled Reports page, the **Scheduled Copy** checkbox is disabled.

**Options:**

- **Selected:** NetOps Portal creates a copy of the NetOps business report when you schedule the report so that you can schedule the report with different settings, view the settings, and edit the settings.
- **Cleared:** NetOps Portal does not create a copy of the NetOps business report when you schedule the report. You cannot schedule the report with different settings, view the settings, or edit the settings.

**Default:** Selected

#### – **Frequency**

Defines the frequency at which NetOps Portal saves dashboard data as scheduled reports and the email server sends dashboard data as reports from scheduled reports attached to emails, with a start time, time zone, and time range for the report:

You can define to have NetOps Portal run the report immediately (a Run Now report), or you can define to have NetOps Portal run it on a recurring basis and send dashboard data as reports from scheduled reports attached to emails. For example, you can schedule to have NetOps Portal run interface utilization reports and send dashboard data as reports to the IT department each week for capacity planning.

#### **NOTE**

By default, up to five email tasks can run concurrently.

#### **Options:**

- **Run Now**  
NetOps Portal runs the report and the email server sends the report attached to an email immediately. These reports are one-time-only scheduled reports. If archive reports is enabled, NetOps Portal archives the report, and you can view them. For more information, see the "View a List of Archived Run Now Reports" section.
- **Run Daily**  
The email server runs the report and sends it attached to an email message once per day. If selected, reveals checkboxes where you can select the day of the week when the report is run.  
**Default:** Run the report every weekday (Monday - Friday) at 00:30 hour in the time zone of the logged-in user. The data in the report reflects the previous 24 hours.
- **Run Weekly**  
The email server runs the report and sends it attached to an email message once per week. If selected, you can select the day that the report is run.  
**Default:** Run the report every Sunday at 01:00 hour in the time zone of the logged-in user. The data in the report reflects the previous seven days (Saturday - Sunday).
- **Run Monthly**  
The email server runs the report and sends it attached to an email message once per month. If selected, lets you select the day of the month to run the report.
- **Run Quarterly**  
The email server runs the report and sends it attached to an email message once per quarter. If selected, lets you specify a starting month and day of the month to run the report.
- **Run Yearly**  
The email server runs the report and sends it attached to an email message once per calendar year. If selected, you can specify a starting month. Reports are run on the first day of the specified month.

**Default:** Run Now

**Required** Yes

- **On Days**  
(If you have selected **Run Daily** as the frequency) Determines on which days to run the report.
- **On Day**  
(If you have selected **Run Weekly** as the frequency) Determines on which day to run the report.
- **On Day of Month**  
(If you have selected **Run Weekly** as the frequency) Determines on which day of the month to run the report.
- **Starting Month**  
(If you have selected **Run Quarterly** or **Run Yearly** as the frequency) Determines the first month to run the report.
- **On Day of Month**

(If you have selected **Run Quarterly** as the frequency) Determines the day of the month on which to run the report.

- **Start Time**  
Determines time on which to start running the report.
- **Time Zone**  
Determines the time range on which to run the report.
- **Time Range**  
Determines the time range on which to run the report.
- **Enable**  
Determines if the scheduled report is enabled.

8. Click **OK**.

The report is sent according to the schedule that you selected. For CSV and single-paged PDF, the report is attached and sent in an email message. For full-paged PDF, an email message is sent with a time-sensitive link to the report. Download and save the report as soon as possible.

### **View or Edit the Settings for a Scheduled Report**

**Prerequisite:** You are logged in as a user with the Send Reports by Email role right.

**Follow these steps:**

1. On the **Manage Scheduled Reports** page, select the scheduled report that you want to edit, and then click **Edit**. The **Email Dashboard** dialog opens.
2. View or change the settings for the scheduled report, and then click **OK**.

Your changes to the scheduled report are saved.

### **View a List of Archived Reports for a Selected Scheduled Report**

**Prerequisites:**

- The scheduled report for which you want to view archived reports exists.
- You are logged in as a user with the Send Reports by Archive role right.

On the **Manage Scheduled Reports** page, select the scheduled report for which you want to view a list of archived reports, and then click **Show Archived Reports**. The **Archived Scheduled Reports** dialog appears, showing the list of archived reports for the selected scheduled report.

### **View a List of Archived Run Now Reports**

**Prerequisites:**

- Archive reports is enabled.  
For more information about how to enable archive reports, see [Configure the Email Server](#).
- Archived Run Now reports exist.
- You are logged in as a user with the Send Reports by Archive role right.

On the **Manage Scheduled Reports** page, click **Show Run Now Archived Reports**. The **Archived Run Now Reports** dialog appears, showing the list of archived Run Now reports.

### **Send an Archived Report to Others**

If you are logged in as a user with the Send Reports by Email role right, you can send archived reports to others as attachments in email messages from NetOps Portal.



## Delete an Archived Report

**Prerequisite:** The archived Run Now or scheduled report that you want to delete exists.

### Follow these steps:

1. From the **Archived Run Now Reports** or **Archived Scheduled Reports** dialog, select the archived report that you want to delete, and then click **Delete**.
2. At the prompt, confirm the deletion by clicking **Yes**.

The archived report is deleted.

## Generate a URL for a View

### Generate a URL for a View

You can share a view with users who do not have access to a dashboard by generating a URL for the view. The URL recreates the selected view on demand. The URL lets you add the view to a web page or intranet site to share performance data.

A security token is included with each URL. This token is based on the user who is logged in at the time of URL generation. Any user who accesses the exported view sees the same data as the user who exported the URL.

### Follow these steps:

1. Log in as a user with the Generate URLs from views role right, such as the Administrator role.
2. Open the dashboard that contains the view for which you want to generate a URL.
3. Click the **Edit** (gear) icon on the view, and then select **Generate URL**.  
The **Generate URL** dialog opens. The URL is displayed in the **URL** field. The URL provides access to the selected view.
4. In the **Display Options** section, complete the following fields:
  - **View Container**  
Displays the chart or graph with a surrounding container. The container includes the title of the view and a black outline around the chart or graph.  
**Options:** Enabled or Disabled  
**Default:** Enabled
  - **Copyright**  
Shows the copyright information for the web page in the view.  
**Options:** Enabled or Disabled  
**Default:** Enabled
  - **Drill Down**  
Lets users drill down from the view into the underlying data source for more detailed data.  
**Prerequisites:** The user has access to the data source and has the Drill into Data Sources role right.  
**Options:** Enabled or Disabled  
**Default:** Enabled
  - **Detailed View Logging**  
If you encountered an issue with this view, enable logging. You can use this method to provide the necessary details to Broadcom Support.  
**Options:** Enabled or Disabled  
**Default:** Disabled
5. In the **Time Selection Options** section, complete the following fields:
  - **Start Time**

Specify the start time for the data in the view.

- **End Time**

Specify the end time for the data in the view.

- **Time Range**

6. In the **Token Expiration Options** section, for **Token Expiration**, choose the view expiration, or timeout period. The URL includes an encrypted token that causes the view to expire after the specified timeout period. The token does not enable the user who interacts with the generated view to drill down for more data.

**Options:**

- **Never:** The exported view displays indefinitely.
- **1 Hour:** The exported view displays for one hour.
- **8 Hours:** The exported view displays for eight hours.
- **1 Day:** The exported view displays for one day.
- **1 Week:** The exported view displays for one week.
- **1 Month:** The exported view displays for one month.
- **3 Months:** The exported view displays for three months.
- **1 Year:** The exported view displays for one year.

**Default:** 1 Year

**Next steps:**

- Copy the URL displayed at the top of the page to the clipboard, and paste it to the destination where you want to display the view.
- (Optional) Click **Preview** to preview the view.

## Download View Data

You can download view data as a comma-separated values (CSV) format file.

View content is downloaded to comma-separated values (CSV) files as raw data. The generated files are compatible with spreadsheet applications, such as Microsoft Excel.

**NOTE**

- The generated CSV files include only the information up to the row limit (1,000 for custom views and 5,000 for common views). Items that exceed this limit are not exported from the database.
- The generated CSV files do not include the following leading characters (+, -, @, =).

**Prerequisite:** You are logged in as a user with the Export to CSV role right (23.3.4 and higher) Export/Print Dashboard Views role right.

**Follow these steps:**

1. Navigate to the view that you want to export.
2. Click the gear icon on the view, and then click one of the following options:
  - **Export to CSV Scaled**  
Specifies that the values in the exported view are scaled. Scaled values appear with larger units, for example, 1 KB.
  - **Export to CSV Unscaled**  
Specifies that the values in the exported view are not scaled. Unscaled values appear in the raw form for the metric, for example, 1000 bytes.

The view is downloaded as a file in CSV format.

## Print View Data

You can print view data as a Portable Document Format (PDF) file.

**(23.3.4 and higher)**

**Prerequisite:** You are logged in as a user with the Export/Print Dashboard Views role right.

**Follow these steps:**

1. Navigate to the view that you want to print.
2. Click the gear icon on the view, and then click one of the following options:
  - **Print PDF Portrait**  
Specifies that the view is printed as a PDF file in portrait orientation, where the height of the document is greater than the width.
  - **Print PDF Landscape**  
Specifies that the view is printed as a PDF file in landscape orientation, where the width of the document is greater than the height.

The view is downloaded as a file in PDF format.

## Inventory Pages and Views

You can view lists of managed items from all data sources from the Inventory pages and views.

The **Inventory** page displays categories of items that are available to NetOps Portal from registered data sources, such as devices, device components, interface addresses, aggregated components, interfaces, network paths, virtual interfaces, tunnels, Wi-Fi devices, and client devices. The categories are limited to items types that are members of the groups in your user account permission set. To view this page, click the **Inventory** menu.

You can do the following from inventory pages and views:

- **Drill down to the context page for an item.** Click the item in the list.
- **Sort by a column.** Click the column heading.  
For more information about how to manage a view, see [Customize Views](#).
- **Add or remove columns.** Hover over a column heading, and then click the gear icon. Expand **Columns**, and then select or clear columns.  
For more information about how to manage a view, see [Customize Views](#).
- **Run and On-Demand report.** Select an item, and then click **On Demand**.  
For more information, see [On-Demand Reports](#).
- **Retire a device, or mark it for maintenance.** Select the device, and then click **Manage Life Cycle**.  
For more information, see [Manage Device Life Cycles](#).
- **Enable or disable polling on specific interfaces.** Select the interface, and then click **Select Polling State**.  
For more information, see [Manage Interface Polling Behavior](#).

The dashboards and context pages that show relevant inventory list views show the same information as inventory pages. Inventory lists provide minimal information to identify each item, such as device hostnames or IP addresses.

You can view a list of inventory for the following items:

- [Device Inventory](#)
- [Devices with Alarm States](#)
- [Device Component Inventory](#)
- [Interface Address Inventory](#)
- [Aggregated Components Inventory](#)
- [Interfaces Inventory](#)
- [Network Path Inventory](#)
- [Virtual Interfaces Inventory](#)
- [Tunnel Inventory](#)
- [Wi-Fi Device Inventory](#)
- [Client Device Inventory](#)

### **Device Inventory**

You can view a list of devices—such as monitored devices, routers, and switches—from the **Devices** page. To view this page, hover over **Inventory**, **Items**, and then click **Devices**.

The following columns show/display on the **Devices** page by default:

- **Name**  
Defines the name of the device.
- **Type**  
Defines the device type (Router, Switch, or Server).
- **Domain**  
Defines the IP domain for the device.
- **Address**  
Defines the IP address for the device.
- **Description**  
Defines the description for the device.
- **Contact Status**  
Defines the status of the monitored device, such as Down, DC Connection Lost, or Not Monitored. The data aggregator sends the contact status for each device.  
**TIP**  
To view the date and time that the device contact status changed, show/display the **Contact Status Change** column in the table.  
For more information about the contact status, see [Reachability Status and Contact Status](#).
- **Current Alarm State**  
Defines the current alarms state for the device. For example, Normal, Minor, Major, and Critical.  
For more information current alarm state, see [Monitor Device Inventory with Alarm States](#).
- **Life Cycle State**  
Defines the device life cycle state for the device. For example, Active, Retired, or Maintenance.  
For more information about the life cycle state, see [Manage Device Life Cycles](#).
- **Context Types**  
Defines the context types for the device. For example, Switch, Router, Device, Wireless Sensor.

**TIP**

You can show/display (or hide) these columns to any view that shows devices in a table.  
For more information, see [Customize Views](#).

### **Devices with Alarm States**

If you have integrated with DX NetOps Spectrum (Spectrum), you can [view the alarm states for devices](#).

## **Device Component Inventory**

You can view the components for devices, such as environmental sensors, from the **Device Components** page. To view this page, hover over **Inventory**, **Items**, and then click **Device Components**.

## **Interface Address Inventory**

You can view the interface addresses, such as ethernet interfaces, from the **Addresses** page. To view this page, hover over **Inventory**, **Items**, and then click **Interface Addresses**.

## **Aggregated Components Inventory**

You can view a list of the established aggregated components, such as aggregated CPUs and aggregated interfaces, and then view the context page related to an aggregated component. To access the aggregated components inventory, hover over **Inventory**, **Items**, and then click **Aggregated Components**.

For more information, see [Manage Aggregated Components](#).

You can view details for the aggregated component from the following context tabs:

- **Details:** Displays details for the aggregated component, including a description of the aggregated component.
- **Interface Health:** This context tab includes the following views:
  - **Utilization/Discard Trend with Events (Out)**
  - **Utilization/Discard Trend with Events (In)**
  - **Discard Out Trend**
  - **Discard In Trend**
  - **Errors Out Trend**
  - **Errors In Trend**
- **CPU Health:** This context tab includes the **CPU Utilization** view.
- **Memory Health:** This context tab includes the following views:
  - **Memory Utilization**
  - **IM Trend Chart (Interface - Component) - Aggregated Interface**
- **Interface Performance:** This context tab includes the following views:
  - **Utilization/Discard Out Trend with Events**
  - **Utilization/Discard In Trend with Events**
  - **Interface Utilization Out Trend/Baseline Details with Events**
  - **Interface Utilization In Trend/Baseline Details with Events**
- **Interface Calendar Heat Charts:** This context tab includes the following views:
  - **Calendar Heat Chart - Average Percent Utilization In**
  - **Calendar Heat Chart - Average Percent Utilization Out**
  - **Calendar Heat Chart - Percent Interface Discards Out**
  - **Calendar Heat Chart - Percent Interface Discards Out**
- **Custom View - Infrastructure Management**
- **Alarms:** Displays the alarms for the aggregated component.

## **Interfaces Inventory**

If DX NetOps Network Flow Analysis (NFA) is registered as a data source, you can view a list of the interfaces on the **Interfaces** page. To access the context page related to an interface, click the name of the interface from the list. The **Details** tab is selected by default.

For more information, see [Network Flow Analysis Views in NetOps Portal](#).

## Network Path Inventory

If DX NetOps Virtual Network Assurance (VNA) is registered as a data source and the AppNeta plug-in is configured, you can view a list of the network paths in NetOps Portal, and then view the context page related to a network path. Network paths are the paths/routes that a network packet takes from a particular source to a destination/target. AppNeta monitors the network using these network paths.

To access the list of network paths, hover over **Inventory**, **Items**, and then click **Network Paths**. To access the context page related to a network path, click the name of the network path from the list. The **Details** tab is selected by default.

For more information:

- About this parameter in the AppNeta plug-in, see [the DX NetOps Virtual Network Assurance documentation](#).
- About AppNeta, see [the AppNeta documentation](#).

You can view details for the network path from the following context tabs:

- **Details:** Displays details for the network path, including a description of the network path, the network type, and the monitoring point to which this network path is associated. This context tab includes the **Alarms** and **Event List** views. From the **Details** context tab, if the API token from AppNeta (the `AUTH_TOKEN` parameter that is configured in the AppNeta plug-in) is associated to an organization in AppNeta to which you have access, you can access the network path in AppNeta by clicking the network path name (the link).
- **Utilization and Performance:** Displays metrics for the network path. This context tab includes the following views:
  - **Utilization Out**
  - **Utilization In**
  - **Utilization Out Over Threshold**
  - **Utilization In Over Threshold**
  - **Utilization Out Over Time**
  - **Utilization In Over Time - In**
  - **Latency Over Threshold**
  - **Round Trip Time Over Threshold**
  - **Latency Over Time**
  - **Round Trip Time Over Time - Round Trip Average**
- **Data:** Displays data metrics for the network path, such as data loss and data jitter. This context tab includes the following views:
  - **Data Jitter Out**
  - **Data Jitter In**
  - **Average Data Jitter Out Over Threshold**
  - **Average Data Jitter In Over Threshold**
  - **Average Data Jitter Out Over Time**
  - **Average Data Jitter In Over Time**
  - **Data Loss Out**
  - **Data Loss In**
  - **Average Data Loss Out Over Threshold**
  - **Average Data Loss In Over Threshold**
  - **Average Data Loss Out Over Time - Packet Loss Average**
  - **Average Data Loss In Over Time - Packet Loss Average**
- **Voice:** Displays voice metrics for the network path, such as Mean Opinion Score (MOS), voice loss, and voice jitter. This context tab includes the following views:

- MOS Out
- MOS In
- MOS Out Over Threshold
- MOS In Over Threshold
- MOS Out Over Time
- MOS In Over Time
- Voice Jitter Out
- Voice Jitter In
- Average Voice Jitter Out Over Threshold
- Average Voice Jitter In Over Threshold
- Average Voice Jitter Out Over Time
- Average Voice Jitter In Over Time
- Voice Loss Out
- Voice Loss In
- Average Voice Loss Out Over Threshold
- Average Voice Loss In Over Threshold
- Average Voice Loss Out Over Time - Packet Loss Average
- Average Voice Loss In Over Time - Packet Loss Average
- **Custom View - Network Path**
- **Alarms:** Displays the alarms for the network path.

### **Virtual Interfaces Inventory**

You can view a list of virtual interfaces, or vSwitches, and then view the context page related to a virtual interface. The context page provides performance information for the virtual interface. To access the virtual interfaces inventory, hover over **Inventory**, **Items**, and then click **Virtual Interfaces**.

For more information, see [Monitor SDN/NFV vSwitch Performance](#).

You can view details for the virtual interface from the following context tabs:

- **Details:** Displays details for the virtual interface, including a description of the virtual interface, the IP addresses, and the MAC address. This context tab includes the **Event List** view.
- **Health:** This context tab includes the following views:
  - **Virtual Interface Utilization/Discard In Trend with Events**
  - **Virtual Interface Utilization/Discard Out Trend with Events**
  - **Virtual Interface Errors Trend Out**
  - **Virtual Interface Errors Trend In**

### **Tunnel Inventory**

If you have integrated with DX NetOps Virtual Network Assurance (VNA) and you have configured a plug-in that collects tunnels, you can access tunnels from the **Tunnels** page. To access this page, hover over **Inventory**, **Items**, and then click **Tunnels**. Tunnels represent the connection between two devices and has polled statistics (jitter, latency, and packet loss).

For more information, see [Modern Network Monitoring](#).

### **Wi-Fi Device Inventory**

If you have integrated with VNA and you have configured a plug-in that monitors Wi-Fi devices, such as the Cisco Meraki and Aruba Central plug-ins, you can access these devices from the **Wi-Fi Device** page. To access this page, hover over **Inventory**, **Items**, and then click **Wi-Fi Devices**.

For more information, see [Monitor Wi-Fi Device Inventory](#).

### **Client Device Inventory**

If you have integrated with VNA and you have configured a plug-in that monitors client devices, such as the Cisco Meraki and Aruba Central plug-ins, you can access these devices from the **Clients** page. To access this page, hover over **Inventory**, **Items**, and then click **Clients**.

For more information, see [Monitor Client Device Inventory](#).

## **Dashboards**

Dashboards provide high-level information about individual managed items, such as average performance of monitored items in a group. Most are composed of views of summary data, such as hourly rollups or averages from a group of items. You can use them to view the polled data as meaningful information and to generate reports. They often provide a drill-down path to more detailed, related pages from a selected context (context pages).

For more information about context pages, see [Context Pages](#).

## **Manage Dashboards**

Custom dashboards are pages with custom set of views.

To add a page with a custom set of views, create a custom dashboard. To modify the views or layout, edit an existing dashboard. To use an existing dashboard as a template, copy the dashboard.

You can manage dashboards in the following ways:

- [View a Dashboard](#)
- [Create a Dashboard](#)
- [Edit a Dashboard](#)
- [Copy a Dashboard](#)
- [Turn on \(enable\) Page Auto Refresh](#)
- [Customize a Dashboard](#)
- [Change the Context of a Dashboard](#)
- [Filter Data Based on a Specific Time Range](#)
- [Apply a Business Hours Filter to a Dashboard](#)
- [Drill Down to a Context Page from a Dashboard](#)

### **View a Dashboard**

You can view and filter the performance data on a dashboard with group contexts and time ranges.

Use the group context to filter the data that appears in views on the dashboard. When you select a group for a standard dashboard, you apply a filter to all views on the page. When you select a group context, items from all subgroups that are available to you appear in views on the dashboard.

### **Create a Dashboard**

To add a dashboard to the dashboard menu, create a dashboard.

#### **Follow these steps:**

1. Log in as a user with the Create a Dashboard role right.
2. From any dashboard, click the **More options** icon on the toolbar, and then click **Add Dashboard**.  
The **Add Dashboard** page appears.



3. Complete the following fields:
  - **Dashboard Menu**  
Specify the menu where this dashboard will display.
  - **Menu Item**  
Specify the name of the dashboard in the menu.
  - **Dashboard Title**  
Specify the title that appears at the top of the dashboard page.
4. In the **Layout** section, select a layout template for the dashboard. Each layout treats the page as a table with rows and columns for views. The layout buttons indicate the number of views in each column and row on the page. Select a layout before you add views.

**NOTE**

Some views, such as scorecards, MultiView, and MultiTrend views, include a lot of detail and require more screen space. These views do not render well in layouts with more than one column.

5. In the **Views** section, click and drag views to the page layout. The maximum number of views per dashboard is 25. You can limit the list of views by completing the following steps:
  - a. Click **Select Context**.  
The **Select View Context** dialog appears.
  - b. Select a group, item, or device context.  
The views that you add to the layout are pinned to the selected context.  
To customize the view settings, click the **Edit** (pencil) icon for the view.  
For more information about how to configure views, see [Customize Views](#).
6. Click **Save**.

The dashboard is saved, and is added to the selected menu. NetOps Portal refreshes the dashboard to reflect your changes, and persists these changes across login sessions.

**Edit a Dashboard**

Add or remove views, or rearrange views, by editing the dashboard. You can save the changes to your own user account.

**Follow these steps:**

1. Log in as a user with the Administer Shared Dashboards role right or Create a Dashboard role right.
2. From the dashboard that you want to edit, click the **More options** icon on the toolbar, and then click **Edit Dashboard**.  
The **Edit Dashboard** page appears.
3. Edit the dashboard as required, and then click **Save**.

The dashboard reloads and shows edited settings.

**Copy a Dashboard**

Copy a dashboard to use an existing dashboard as a template for a new dashboard. You can also copy dashboards to other menus, for example, to share a dashboard on the **My Dashboard** menu with others.

**Follow these steps:**

1. Do one of the following:
  - (To copy a dashboard) Log in as a user with the Create a Dashboard role right.
  - (To copy a dashboard on the **My Dashboards** menu) Log in as a user with the Create a Dashboard, Edit Context Pages, Edit Shared Views, Generate URLs from views, and Save Changes to Shared Views role rights.
2. From the dashboard that you want to copy, click the **More options** icon on the toolbar, and then click **Copy Dashboard**.  
The **Copy Dashboard** page appears.
3. Complete the following fields, and then click **Save**:

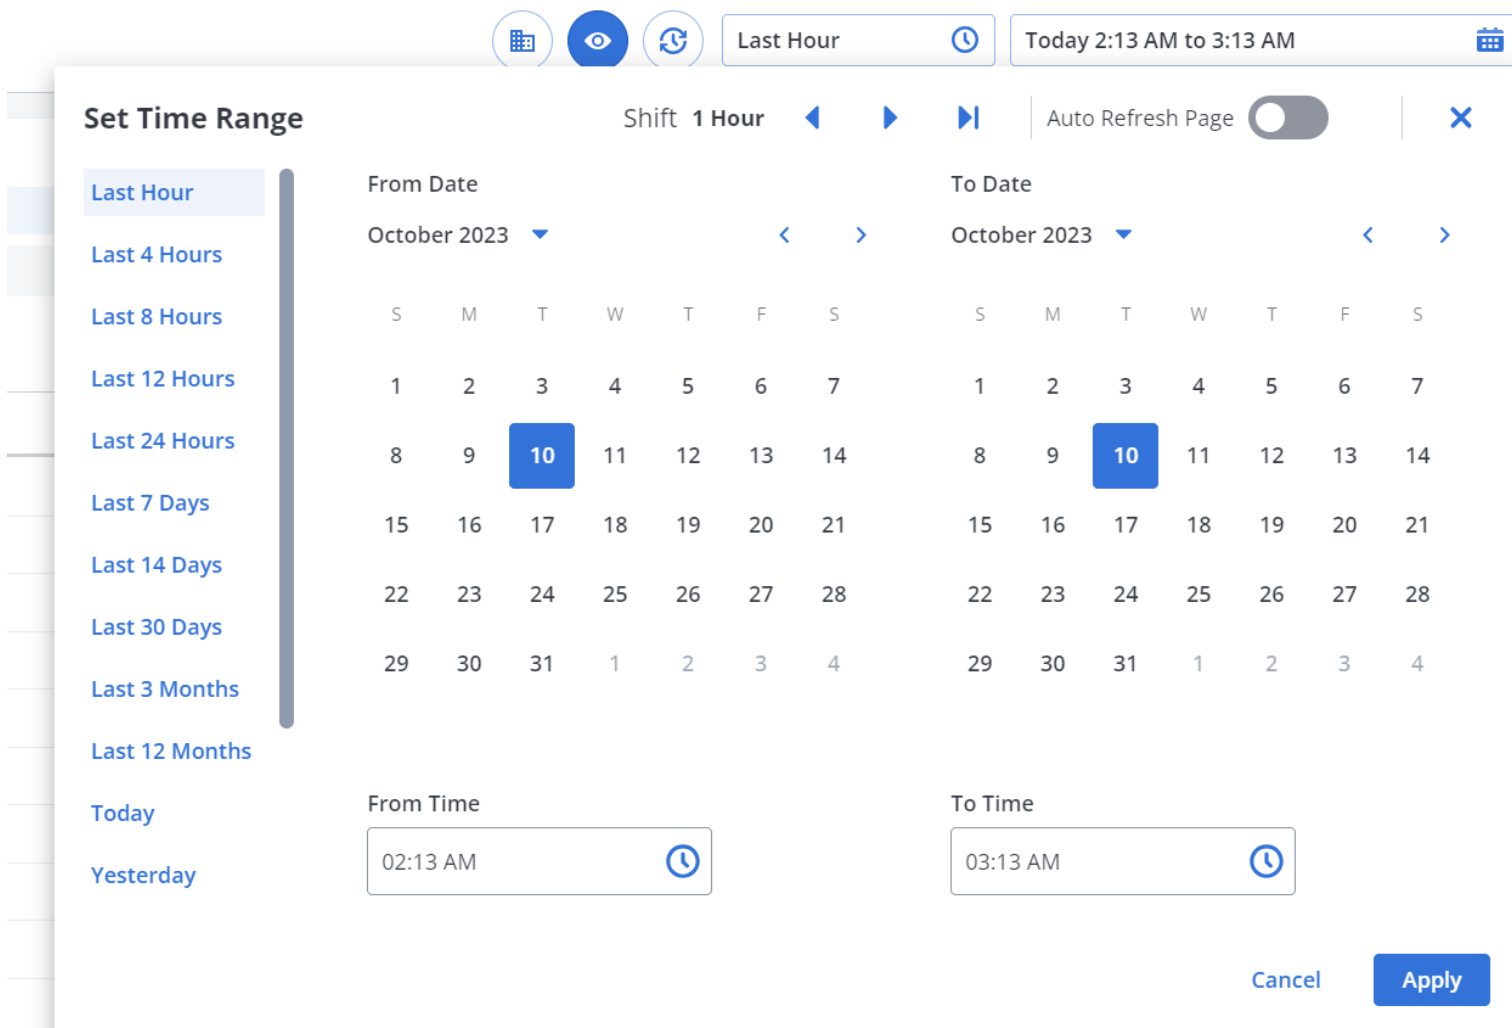
- **Dashboard Menu**  
Select the dashboard menu where you want the copied dashboard to appear.
- **Menu Item**  
Specify the name for the dashboard.
- **Dashboard Title**  
Specify the title that appears at the top of the dashboard page.

A copy of the dashboard is created. The dashboard opens.

### Turn on (Enable) Page Auto Refresh

You can have NetOps Portal refresh and display the most recent data in the **Alarms** and **Events** views on dashboards by turning on **Auto Refresh Page** or by manually refreshing the dashboard. By default, **Auto Refresh Page** is turned off (disabled). You can turn it on and off by clicking the **Auto Refresh Page** option in the upper-right corner of NetOps Portal.

The following image shows this option:



### NOTE

By default, NetOps Portal refreshes and displays the most recent data every 60 seconds, however, you can configure this (**Dashboard Auto Refresh Rate**) setting.

For more information, see [Configure the Dashboard Settings](#).

### **Customize a Dashboard**

Custom dashboards are dashboards that you have customized. You can customize dashboards with different set of views, context, and layouts. You can use custom dashboards to troubleshoot issues or long term to monitor categories of items.

You can use the custom dashboards that the Administrator has created and the out-of-the-box dashboards and the views that are associated with all registered data sources. Log in to access the dashboards that are assigned to your user account.

#### **Examples:**

- A regional manager uses a dashboard that pins views to each site group in that region.
- A systems administrator uses a dashboard to monitor all servers.
- A network engineer uses a dashboard with views pinned to problematic routers or critical interfaces.

The following video shows how to customized dashboards in NetOps Portal with views and context pages to display meaningful information and visualizations for specific network monitoring requirements:

### **Change the Context of a Dashboard**

The dashboard context filters the data that appears in views on the dashboard. When you select a group for a standard dashboard, you apply a filter to all views on the page. When you select a group context, items from all subgroups that are available to you appear in views on the dashboard. However, the same subgroups might not appear on the **Items** tab in Group administration.

For more information about how to manage groups, see [Groups](#).

You can also view dashboards in multiple windows, and apply a different data context to each dashboard.

#### **Follow these steps:**

1. From the dashboard that you want to change context, click the **Change Filter** link that is above the title of the dashboard.  
The **Change Group Filter** dialog box opens.
2. Select a group from the group hierarchy, and then click **OK**.

The views on the page with dynamic context are refreshed to reflect the new data context. The filter refines the content in each dashboard in NetOps Portal to only the data in the selected group. The filter remains in place from one dashboard to another until you change the filter or until you log out.

### **Filter Data Based on a Specific Time Range**

Filter data based on specific time periods to help with troubleshooting performance issues. For example, you can show data from the last seven days by changing the time range. In this case, the time range helps you to determine whether an issue is occurring regularly.

You can select a precise time range for the performance data shown in the current dashboard. Use the time period selectors to select the day, the start time, and the end time. When you change the time range for a dashboard, the change applies to all dashboards in that window. To compare content in different time ranges, open the dashboard in multiple browser windows or modify specific views to show a fixed custom time range.

#### **Follow these steps:**

1. From the dashboard that you want to filter data, click **Change the selected time range** (the down arrow) in the upper-right corner of a dashboard.  
The following image shows the options that appear:

The screenshot displays a time range selection dialog box. At the top, a status bar shows 'Today 2:18 PM TO 3:18 PM' and 'Last Hour' with a refresh icon and a close button. Below this, the 'Select a Time Range' section offers various preset options: 'Last Hour' (selected), 'Last 4 Hours', 'Last 8 Hours', 'Last 24 Hours', 'Last 7 Days', 'Last 14 Days', 'Last 30 Days', 'Last 3 Months', 'Last 12 Months', 'Today', 'Current Week', 'Current Month', 'Yesterday', 'Previous Week', and 'Previous Month'. To the right of these options is a 'Live Update Auto Refresh' toggle switch, which is currently turned off. Below the preset options is the 'Select a Custom Time Range' section, which includes two calendar pickers for 'From' and 'To' dates. The 'From' calendar is set to September 2022, and the 'To' calendar is set to July 2022. Both calendars show the 16th of the month selected. Below the calendars are input fields for the time range, showing 'From: 2:18 PM' and 'To: 3:18 PM'. At the bottom right are 'Cancel' and 'Apply' buttons.

2. Select the default time range by which to filter data in the dashboard from the list or specify a custom time range:

- **Select a Time Range**

Defines the time range by which to filter data.

**Options:** Last Hour, Last 4 Hours, Last 8 Hours, Last 12 Hours, Last 24 Hours, Last 7 Days, Last 14 Days, Last 30 Days, Last 3 Months, Last 12 Months, Today, Current Week, Current Month, Yesterday, Previous Week, Previous Month

**Default:** Last Hour

- **Live Update Auto Refresh**

Defines whether NetOps Portal refreshes the views on dashboards automatically using the latest time range. When enabled, NetOps Portal performs a live update, and displays the most recent data in the view.

**Default:** Disabled

- **Select a Custom Time Range**

To select a *custom* time range by which to filter data, choose the time range, including the day and hour.

3. Click **Apply**.  
The selected time range is applied to the dashboard.

### **Apply a Business Hours Filter to a Dashboard**

To show data in a dashboard for particular business hours, you can apply a business hours filter to it as part of the user-session page context. This sets the reporting profile that persists during your user session. The filter filters the data that appears in the views on the dashboard. User-session level business hours filters remain applied when switching between context pages, navigating between dashboards, and changing time ranges.

Most views on dashboards show filtered data. For a complete list of the views that support business hours filtering, see [Configure Business Hours Filtering](#).

You can also set the reporting profile by specifying a site group that is associated with a business hours definition to a page context.

#### **NOTE**

The applied business hours on dashboards (session business hours filter) supersede the page context for the site group if the site group is associated with a business hour filter.

**Best Practice:** When using a dashboard with applied business hours, ensure that you have applied the business hours filter in the most meaningful manner by selecting a site group as the page context before applying the business hours filter.

When applying business hours filters to a dashboard as part of your user session, the business hours override other business hour settings in the following priority order:

1. Business hours and time zone assigned to a view.
2. Site group that is associated with a business hours filter applied to a view.
3. Business hour filters applied to a dashboard or a context page (session-level business hours).
4. Site group that is associated with a business hours filter specified at the context level.

#### **Follow these steps:**

1. On the dashboard to which you want to apply a business hours filter, click the **Apply Business Hours** (clock) switch icon in the upper-right corner of the dashboard.

#### **TIP**

This icon indicates whether business hours are or are not being applied to the dashboard.

The **Apply Business Hours** dialog opens. The **Context Business Hour** shows the context page-related business hours filter that you have applied. The **Session Business Hour** shows the user-session level business hours filter that you have applied to the dashboard.

2. Complete the following fields, and then click **Apply**:
  - **Apply Business Hour Filter**  
Enable this option. By default, business hours are not enabled.
  - **Business Hours Filter**  
Select the business hours filter that you want to apply to the dashboard. Or, to disable the business hours filter that are associated with a site group that is applied to the context page, select **No business hour selected**.
  - **Business hour and time zone**  
The time zones from which you want to choose in the **Time Zone Filter** drop-down.

#### **Options:**

- **Recommended hours/zones:** Lists the time zones that are associated with the business hours you have chosen.
- **System defined hours/zones:** Lists the time zones that are associated to site groups.
- **Any hours/zones:** Lists all time zones.

**Default:** Recommended hours/zones

– **Time Zone Filter**

If you selected a business hours filter, select the time zone that you want to apply to the dashboard.

– **Apply Changes**

Select the user-session level to which to apply this business hours filter.

**NOTE**

The option you choose changes the priority order in which business hours override other business hours.

**Options:**

- **My Current Session:** Applies this business hours filter to *all* dashboards during your user session.
- **Current Session for Page only:** Applies the business hours filter to *this* dashboard only during your user session. This option does *not* override the priority of business hours filter explicitly applied to a view.

**NOTE**

This option supersedes the **My Current Session** option if associated to a business hours filter.

- **My Current Session override views:** Applies the business hours filter to *all* context pages during your user session and overrides the properties for applied business hours filters to views.

**NOTE**

This option requires the Administer Shared Dashboards and Edit Shared Views role rights.

- **Current Session for Page override views:** This applied business hours filter applies to *all* dashboards during your user session and overrides the options for the business hours filter applied to the views on the dashboards.

**NOTE**

This option requires the Administer Shared Dashboards and Edit Shared Views role rights.

**Default:** My Current Session

The business hours are applied to the dashboard.

### **Drill Down to a Context Page from a Dashboard**

**Follow these steps:**

1. Log in as a user with the Drill into Views role right.
2. On the dashboard from which you want to drill down to a context page, take *one* of the following steps:
  - Right-click the hyperlink for an item, and then select a context page.
  - Click the hyperlink for an item.

## **Organize Dashboards in Menus**

Dashboards are organized in customizable menus.

The available dashboards and menus display when you hover over the **Alarms** and **Performance** tabs. Initially, before you add custom menus, only the predefined menus are included in the list on the **Performance** tab. Your user account role determines the menus that you can access. The custom menus that you define are shared within a single tenant. The out-of-the-box menus are shared among all tenants.

Users with the required administrative role rights can manage menus using the following options:

- [Add a Menu](#)
- [Add or Remove a Dashboard to or from a Menu](#)
- [Associate Menus with User Roles](#)
- [Reorganize Menus for User Roles](#)
- [Delete a Menu](#)

## Add a Menu

You can organize dashboards and make them available to users with specific roles by adding them to custom menus. Administrators and designers can add custom menus, and can select dashboards for each menu. The administrator must add the menu to a role before the custom menus are available to users with that role.

### Follow these steps:

1. Hover over **Administration, User Settings**, and then click **Menus**.  
The **Manage Menus** page opens.
2. Click **New**.  
The **Add Menu** dialog appears.
3. Specify a name and description for the menu.
4. Move the dashboards that you want to include in the menu from the **Available Dashboards** column to the **Selected Dashboards** column by selecting the dashboard, and then clicking the **Select item** arrow.
5. Click **Save**.
6. **Next step:** [Associate this menu to user roles](#).

The menu is added to the **Performance** tab.

## Add or Remove a Dashboard to or from a Menu

With the Generate URLs from views role right, you can add, move, or copy dashboards to a menu.

Dashboards in the **My Dashboards** menu are not visible to other users. With the Generate URLs from views, Create a Dashboard, Edit Context Pages, Edit Shared Views, and the Save Changes to Shared Views role rights, you can copy and customize a dashboard from another menu to the **My Dashboards** menu.

### TIP

With the Proxy Users role right, you can edit the **My Dashboards** menu for another user by proxying that user account.

## Delete a Menu

You can delete menus that are no longer in use, which removes them from the **Performance** tab. Dashboards that are assigned to the menu are not affected, and remain available to other menus. Deleting a user account that is associated with a custom menu does not delete that menu.

## Configure the Dashboard Settings

Configure the rate at which NetOps Portal refreshes dashboards that contain normal and rapidly-pollled data when you have Live Update Auto Refresh turned on (enabled).

### Follow these steps:

1. Hover over **Administration, Configuration Settings**, and then click **Dashboard Settings**.  
The **Dashboard Settings** page appears.
2. Change the values for the following settings, and then save your changes:
  - **Dashboard Auto Refresh Rate (seconds)**  
The rate at which NetOps Portal refreshes the views—including the **Alarms** and **Events** view—on dashboards containing normal data.  
**Default:** 60 (seconds)  
**Values:** 15 - 300 seconds (5 minutes)  
For more information about the **Alarms** view, see [Alarms View](#).

**IMPORTANT**

Set this rate to a rate that is equal to or higher than the monitoring profile's poll rate. For example, if the poll rate is set to 5 minutes, set this rate to 5 minutes.

For more information about the poll rate, see [Manage Monitoring Profiles](#).

- **High Resolution Dashboard Auto Refresh Rate (seconds)**

The rate at which NetOps Portal refreshes the views on dashboards containing rapidly-pollled data.

**Default:** 10 (seconds)

**Values:** 5 - 300 seconds (5 minutes)

The dashboard settings are configured.

## Out-of-the-Box Dashboards

The out-of-the-box dashboards display data immediately *after* you run discovery and begin collecting metrics from your environment.

The out-of-the-box dashboards are examples of ones that you can build to monitor your environment, and provide views for different users across your organization. Use them as a starting point to create your own dashboard.

**NOTE**

Many of these dashboards include views that include data from DX NetOps Virtual Network Assurance (VNA), CA Application Delivery Analysis, and DX NetOps Network Flow Analysis (NFA). These views appear only if these data sources are installed in your monitoring system and are registered as a data source in NetOps Portal (you have integrated with that product). For dashboards without a listed data source, data comes from the data aggregator into NetOps Portal.

You can access the dashboards from the following **Performance** menus:

- [Infrastructure Health](#)
- [Application Health](#)
- [Capacity Planning](#)
- [Management](#)
- [Operations Displays](#)
- [ACI Reporting](#)
- [SD-WAN Reporting](#)
- [SDN/NFV Reporting](#)
- [NSX-T Reporting](#)
- [Cisco DNA Center Reporting](#)
- [Flow](#)
- [Alarms](#)
- [Analytics](#)

To access these menus, hover over **Performance**.

### **Infrastructure Health**

The **Infrastructure Health** menu includes the following summary dashboards of overall system and device health and performance, events, and thresholds:

- **Infrastructure Overview**  
View key health metrics across multiple technologies in the environment.  
**User:** Operations, Network Engineer  
**Data Source:** Data aggregator, CA Application Delivery Analysis, NFA
- **Server Device Health**



View consumption and status-related metrics specific to servers.

**User:** Operations, System Administrator

- **Server Performance**

View server performance metrics.

**User:** Operations, System Administrator

**Data Source:** CA Application Delivery Analysis

For more information about this dashboard, see [Server Performance Dashboard \(ADA\)](#).

- **Network Device Health**

View performance and consumption metrics for routers, switches, and interfaces.

**User:** Operations, Network Planner

**Data Source:** Data aggregator, NFA

- **Network Performance**

View network, site, and component performance and health metrics for network devices.

**User:** Operations, Capacity Planner

**Data Source:** Data aggregator, CA Application Delivery Analysis, NFA

For more information about this dashboard, see [Network Performance Dashboard \(ADA\)](#).

- **Wi-Fi Health**

View the list of inventory, such as wireless access points (AP) and wireless LAN controllers (WLC), and performance metrics, such as most and least utilized APs and WLCs by CPU and memory usage.

**User:** Operations, System Administrator, and Network Engineer

**Data Source:** Data aggregator

For more information about this dashboard, see [Monitor Wi-Fi Device Inventory](#).

- **Network Interface Performance**

## **Application Health**

The **Application Health** menu includes the following dashboards that summarize performance from an application perspective:

- **Application Performance**

View the worst performing applications based on a particular metric.

For more information about this dashboard, see [Application Performance Dashboard \(ADA\)](#).

- **Unified Communications - Worst Performance**

For more information about this dashboard, see [Worst Performance Dashboard](#).

- **Unified Communications - Volume and Utilization**

For more information about this dashboard, see [Top Volume and Utilization Dashboard](#).

- **Unified Communications - Performance Overview**

For more information about this dashboard, see [Performance Overview Dashboard](#).

- **APM - Applications Summary**

- **APM - Business Services Summary**

## **Capacity Planning**

The **Capacity Planning** menu includes the following dashboards of projections, thresholds, and recent changes to systems or devices for capacity planning:

- **Interface Capacity Watch Lists**

View interfaces with the top utilization and baseline deviation.

**User:** Capacity Planner

- **Router/Switch Capacity Watch Lists**

View device level flow and consumption metrics.

**User:** Capacity Planner, Network Engineer

**Data Source:** Data aggregator, NFA

- **Server Capacity Watch Lists**

View server consumption metrics and baseline deviation.

**User:** Capacity Planner, System Administrator

## **Management**

The **Management** menu includes the following high-level summary, overview, and comparison dashboards for management:

- **Management Overview**

View availability, reachability, and flow.

**User:** Operations

**Data Source:** Data aggregator, NFA

- **Network Overview**

View high-level key health metrics across multiple technologies in the environment.

**User:** Operations, Network Engineer

**Data Source:** Data aggregator, NFA

- **Server Overview**

View a multi-technology focus on server consumption metrics.

**User:** Operations

**Data Source:** Data Aggregator, CA Application Delivery Analysis

## **Operations Displays**

The **Operations Displays** menu includes the following high-level overview dashboards for display in the Operations Center:

- **Interfaces Display**

View interface utilization and traffic health.

**User:** Operations, Network Engineer

**Data Source:** Data aggregator, NFA

- **CPU/Memory Display**

View top-utilized CPU and memory for all devices in the monitoring environment. Includes the **MultiView CPU and Memory Utilization** view.

**User:** Operations, Capacity Planner

- **Interface Trends**

View interface utilization, errors, discards trend graphs.

**User:** Capacity Planner, Network Engineer

## **ACI Reporting**

The **ACI Reporting** menu includes the following Cisco ACI-related dashboards:

- **ACI Console**
- **ACI Health**
- **ACI Switches Overview**

For more information, see [Monitor Cisco ACI](#).

## **SD-WAN Reporting**

The **SD-WAN Reporting** menu includes the following Software Defined Wide Area Networking (SD-WAN)-related dashboards:

- **SD-WAN Tunnel Statistics**
- **SD-WAN Application Statistics**

For more information, see [Monitor SD-WAN Devices](#).

### **SDN/NFV Reporting**

The **SDN/NFV Reporting** menu includes the following dashboards for Software Defined Networking (SDN) and Network Functions Virtualization (NFV):

- **SDN/NFV Virtual Inventory Overview**
- **SDN/NFV Virtual Compute Usage Overview**
- **SDN/NFV Virtual Storage Usage Overview**
- **SDN/NFV Physical Server Usage Overview**
- **SDN/NFV vSwitch Performance Overview**

For more information, see [Monitor SDN/NFV Virtual Inventory](#).

### **NSX-T Reporting**

The **NSX-T Reporting** menu includes the following VMware NSX-T-related dashboards:

- **NSX-T Console**
- **NSX-T Health**

For more information, see [Monitor VMware NSX-T](#).

### **Cisco DNA Center Reporting**

The **Cisco DNA Center Reporting** menu includes the following Cisco DNA Center-related dashboards:

- **Cisco DNA Center Console**
- **Cisco DNA Center Summary**

For more information, see [Monitor Cisco DNA Center](#).

### **Flow**

The **Flow** menu includes the following NetOps Flow-related dashboards:

- **Flow Statistics**
- **Flow Dashboard**

For more information, see [Flow Dashboards](#).

### **Alarms**

The **Alarms** menu includes the **Alarm Console OOTB** alarm-related dashboard.

For more information, see [Alarms View](#).

### **Analytics**

#### **NOTE**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We fully intend to

make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per capability basis.

The **Analytics** menu includes the **Interface Volatility** analytics-related dashboard that you can use to identify volatility in network performance.

#### NOTE

The **Analytics** menu and the **Interface Volatility** dashboard are not visible by default. An Administrator must add the **Analytics** menu to a user account role to which your user account is assigned.

For more information about this dashboard, see [Interface Volatility](#).

## Technology-Specific Dashboards

The following dashboards are associated with specific technologies. The views on each dashboard are customized to fit the needs of each technology:

- **CBQoS Device Component Views**  
View the overall status of a device component and determine if the component is causing an issue.
- **CBQoS Overview Dashboard and Device Views**  
Use the CBQoS Overview dashboard and device views to analyze the performance of your class-based quality of service policies. CBQoS policies help you manage bandwidth. For example, policies can determine the handling of VoIP and video conferencing data to minimize packet loss.  
You can gauge the effects of policies on packet discards from some of the views. You can discover disparities in the discard rates of prepolicy and postpolicy data for a device. A view such as the **Top CBQoS Policing Packets Transmitted** view shows the number of packets that violated or exceeded the rate limits set by Class-Based Policing on the device. Use this information to pinpoint the problem.
- **DS1 Interface Statistics Device Component Views**  
View the overall status of a device component and determine if the component is causing an issue.
- **Firewall Statistics Device Component Views**  
View the overall status of a device component and determine if the component is causing an issue.
- **IP Statistics Device Component Views**  
View the overall status of a device component and determine if the component is causing an issue.
- **Mobile Wireless Device Component Views**  
View the overall status of a device component and determine if the component is causing an issue.
- **Modem Statistics Device Component Views**  
View the overall status of a device component and determine if the component is causing an issue.
- **MPLS Device Component Views**  
View the overall status of a device component and determine if the component is causing an issue.
- **MPLS Overview Dashboard and Device Views**  
View the status of the Multi-Protocol Label Switching (MPLS) environment. Use the information to gain a broad sense of the overall MPLS performance for the enterprise or a specific device.
- **QoS Statistics Device Components**  
View the overall status of a device component and determine if the component is causing an issue.
- **Response Path Tests Overview Dashboard and Device Views**  
This dashboard and device views help you see the overall status of response tests for devices that support service level agreements (SLA) using various protocols. Use this information to gain a broad sense of overall SLA performance for the enterprise or a specific device.  
For example, as a network infrastructure manager, you can identify which protocols are experiencing the highest latency and slowest response times from the dashboard. You can then look at the device-level views to determine which devices are experiencing the problems. This information can help determine if the cause is increased activity, a device that is down, or some other cause.
- **Response Path Device Component Views**

For any supported protocol, view the overall status of a device component, and determine if the response path tests are causing an issue.

- **Sonet/SDH Device Component Views**

View the overall status of a device component and determine whether the component is causing an issue.

## Vendor-Specific Dashboards

The following dashboards are associated with specific vendors. The views on each dashboard are customized to fit the needs of each vendor:

- **Cisco UCS Overview Dashboard and Device Views**

The Cisco Unified Computing System (UCS) platform integrates networking, virtualization, storage access, and management of the UCS components chassis, blade server, and Fabric Interconnect (a FCoE switch).

This dashboard provides performance and status information about Cisco UCS blade servers, chassis, and fabric interconnects. From the **Overview** page, you can drill down to the **Blade**, **Chassis**, and **Fabric Interconnect** overview pages. From these overview pages, you can drill down to component-level views to view trends.

The views exploit the capabilities of the Cisco UCS manager that is integrated within Cisco UCS. You can easily identify and replace a malfunctioning component, such as a memory module or a fan, before it becomes a severe problem.

**Users:**

- Engineers can use this information to troubleshoot problems with the blade chassis physical components, including power supplies, fans, CPUs, and memory modules.

- **Cisco UCS Blade Overview Dashboard and Device Views**

Provides performance and status information about Cisco UCS blade servers. From the **Overview** page, you can drill down to component-level views that allow you to view trends.

For example, a single blade can have up to 24 memory modules. Using this view, an overheated memory module can easily be identified. You can also spot a malfunctioning fan at a single glance.

- **Cisco UCS Chassis Overview Dashboard and Device Views**

Provides performance and status information about Cisco UCS chassis. From the **Overview** page, you can drill down to component-level views that allow you to view trends.

For example, a single chassis can have up to eight blades. Using this view, an overheated power supply can easily be identified. You can also spot a malfunctioning fan at a single glance.

- **Cisco UCS Fabric Interconnect Overview Dashboard and Device Views**

Provides performance and status information about Cisco UCS fabric interconnects. From the **Overview** page, you can drill down to component-level views that allow you to view trends.

Use this view to identify CPU consumption of a heavily loaded fabric interconnect. You can also spot a malfunctioning fan or power supply at a single glance.

- **Citrix VDI Overview Dashboard and Device Views**

Provides performance and status information on the Citrix VDI Management Servers and Desktops. Citrix Virtual Desktop Infrastructure (VDI) is the market leading technology for providing virtualized desktops using Hypervisor technologies. DX NetOps Performance Management supports VMware vSphere Hypervisors. From the **Overview** page, you can drill down to device-level views to isolate a problem.

**Users:**

- Engineers can use this information to troubleshoot problems.
- Capacity planners can use this information to see which services are overloaded and plan for more VDI servers.

**NOTE**

The CPU and Memory reports available in both the **Citrix VDI Overview** dashboard and the **Citrix VDI Server Overview** dashboard can display metrics that are not explicit to the Citrix VDI metric family. To avoid this problem, edit the view and save its context to a particular group to ensure the dashboard only reports on Citrix VDI controllers.

Desktop information provided includes:

- Number of desktops
- Desktops that are showing exceptional utilization
- Content of desktop groups
- Available VMs in the VDI catalog

Engineers can use this information to get a quick overview of the infrastructure and its utilization.

Infrastructure servers (called desktop controllers) information provided includes:

- State
- CPU
- Memory
- Storage
- Network utilization

Engineers can use this information to identify bottlenecks on the controllers before they become problems.

- **Citrix VDI Desktop Overview Dashboard and Device Views**

Provides performance and status information on Citrix Virtual Desktop Infrastructure (VDI) Desktops. Citrix VDI provides virtualized desktops using Hypervisor technologies. DX NetOps Performance Management supports VMware vSphere Hypervisors. From the **Overview** page, you can drill down to device-level views to isolate a problem.

Desktop information provided includes:

- The number of desktops
- The desktops that are showing exceptional utilization
- The content of desktop groups
- Available VMs in the VDI catalog

**Users:**

- Engineers can use this information to get a quick overview of the infrastructure and its utilization.

- **Citrix VDI Server Page**

This page contains the following views:

- Citrix VDI Controller - Registered Desktops (Table)
- Citrix VDI Controller - Services Health (Table)
- Citrix VDI Controller - State (Time Chart)
- Citrix VDI Controller - Configuration (Table)
- Top Citrix VDI Desktop - CPU/Latency/Logon Time (Table)
- Citrix VDI Desktop Group - Used Desktops/Group Enabled (Table)
- Citrix VDI Desktop Group - Desktops Available/Connected (Table)
- Citrix VDI Catalog - Used VMs (Table)
- Citrix VDI Catalog - Assigned VMs (Table)

On this page, the metrics are in the scope of its controller, except for desktop group and catalog metrics.

- **IBM LPAR Overview Dashboard and Device Views**

View the status of overall LPAR performance for the enterprise or a specific device.

- **IBM LPAR Device Component Views**

View the overall status of a device component and determine if the component is causing an issue.

- **Microsoft Hyper-V Overview Dashboard and Device Views**

View the status of overall Microsoft Hyper-V performance for the enterprise or a specific device.

- **Microsoft Hyper-V Device Component Views**

View the overall status of a device component and determine if the component is causing an issue.

- **Microsoft Exchange Server and Active Directory Overview Dashboard and Device Views**

View the overall status of the monitored devices in your Exchange 2007, Exchange 2010, and Active Directory 2008 environments. This information includes directory domain services as well as mailbox and hub-transport mail services.

For those services you can identify the key performance metrics, which includes mail traffic and directory traffic metrics. You can then access device-level views to isolate a problem.

**Users:**

- Engineers can use this information to troubleshoot problems.
- Capacity planners can use this information to see which services are overloaded and plan for more domain and mail servers.

- **Microsoft Exchange Server and Active Directory Device Component Views**

View the overall status of a specific device component and determine whether the component is causing an issue.

- **Microsoft Cluster Service Overview Dashboard and Device Views**

View the overall status of the monitored devices that are part of your Microsoft Cluster Service (MSCS).

You can identify which devices are the top offenders for:

- Resource failures
- Crypto and registry checkpoint restoration and saves
- CPU and memory usage
- Node traffic
- Time spent in various states.

You can then isolate a problem by accessing device-level views.

**Users:**

- Engineers can use this information to troubleshoot problems.
- Capacity planners can use this information to see which MSCS devices are nearing their utilization capacity and plan for more resources.

- **Microsoft Cluster Service Device Component Views**

Displays the overall status of a device component and determine if the component is causing an issue.

- **Solaris Zones Overview Dashboard and Device Views**

Displays the status of overall Solaris Zones performance for the enterprise or a specific device.

- **Solaris Zones Device Component Views**

Displays the overall status of a device component and determine if the component is causing an issue.

- **VMware Overview Dashboard and Device Views**

Displays the overall status of monitored devices in your VMware environment, such as virtual centers, virtual machines, ESX hosts, and other supported devices. Use this information to identify which devices are the top offenders for attributes such as CPU and memory usage, bytes in/out, capacity, and the percentage of time that was spent in various states. You can then access device-level views to isolate a problem.

**Users:**

- Engineers can use this information to troubleshoot problems.
- Capacity planners can use this information to see which virtual devices are maximizing their CPU shares and plan for more resources.

**To access this dashboard:**

- Hover over **Dashboards**, and then click **VMware Overview**.
- Drill down from the dashboard view to a specific device, and then select a VMware-related tab for more detail about that device.

**To access the device pages directly:**

- Hover over **Inventory**, **Items**, **Devices**, and then select a supported device.
- Select either the **VMware**, **VMware ESX Host**, or **VMware Virtual Machine** tab.

- **VMware Device Component Views**

View the overall status of a device component and determine whether the component is causing an issue.

- **IWF Statistics Device Component Views**

View the overall status of a device component and determine if the component is causing an issue.



## Context Pages

Context pages within NetOps Portal provide focused performance and status data that is scoped to a specific managed item, such as a single router or server. They resemble dashboards with a fixed context. When additional data is available from the data source and you have the Drill into Views role right, you can access context pages as drill-down links from dashboards.

The following video shows how to drill down from a dashboard in NetOps Portal to context pages, as well as how to configure context pages to display meaningful views, including information about device components and interfaces:

For more information:

- About dashboards, see [Dashboards](#).
- About role rights, see [Role Rights](#).
- About how to view status data for a managed device, see [Reachability Status and Contact Status](#).

## Manage Context Pages

You can customize the context pages in NetOps Portal, including adding context tabs to context page.

In this article:

- [View a Context Page](#)
- [Set the Context for the Views on Context Pages](#)
- [View a Context Page](#)
- [Filter Data from a Narrow Context](#)
- [Pick a Context](#)
- [Manage the Context Tabs on a Context Page](#)
- [Add or Edit a Context Tab](#)
- [Edit Item Properties](#)
- [Apply a Business Hours Filter to a Context Page](#)
- [Rearrange Context Tabs](#)

The following video shows how to customize context pages:

### View a Context Page

You can access context pages from inventory lists or as drill-down links from dashboards. Unlike standard dashboards, item context pages are clustered in sets of tabbed pages (accessible from the context tabs).

From an inventory list or dashboard, do one of the following tasks:

- Right-click a hyperlink on an item, and then select a context tab.  
For more information, see [Manage Dashboards](#).
- Click a hyperlink on an item to open the default context tab.

### Set the Context for the Views on Context Pages

For all context pages, selecting a row in the top table view sets the corresponding context for the other views in the page. You can change the measurement (Rate/Bytes/Utilization) and direction (In/Out) in the settings for each of the views.

#### **NOTE**

You can send or schedule these views in a report by email. By default, the first row is selected in the Top table and the data for the following views are of the first-row selection.

For more information, see [Share Data with Other Users](#).



## **View a Context Page**

You can access context pages from inventory lists or as drill-down links from dashboards. Do one of the following tasks:

- Right-click a hyperlink on an item, and then select a context tab.  
For more information, see [Manage Dashboards](#).
- Click a hyperlink on an item to open the default context tab.

## **Filter Data from a Narrow Context**

The views on context pages show filtered data from a narrow context, such as a view of data from a single managed item. Determine the source of a performance problem by drilling down into specific data using the links.

In some table views, right-click the name of an item (the link). For example, right-click the link that corresponds to an item name in the **Inventory** section. In the resulting menu, select a related context page, containing more granular data.

## **Pick a Context**

You can pick a context using the context picker that is available next to the name of the managed item. Context pickers are also available for the children of the managed item.

To pick a different context, click **Change**, select the context, and then click **OK**.

## **Manage the Context Tabs on a Context Page**

**Prerequisite:** You are an Administrator or Designer or you have the Edit Context Pages and Drill into Views role rights.

Editing a context page entails managing (adding or editing) the context tabs. You can edit the predefined context tabs and can change the views that are displayed on those context tabs. You can add context tabs and rearrange them in an item context to change their order.

Where applicable, the data aggregator associates managed devices with context types (Firewall, Load Balancer, Wireless Controller, and Wireless Access Point (AP)). The associated context types, including primary device types, appear listed in the **Type** column within device inventory tables.

When you add or edit tabs for a device context page, you can associate them with a primary device type or context type. Tabs that are associated with a primary device type or context type appear on the device context pages that are associated with the specified types.

### **Examples:**

- If you add or edit a tab and associate it with the server type, it appears for all server types.
- If you add or edit a tab for a server type and associate it with firewall, it appears only for all server, firewall types.
- If you add or edit a tab for a router type and associate it with AP, it appears only for all router, AP types.
- If you add or edit a tab for a server type and associate it with wireless controller, it appears only for all server, wireless controller types.

### **TIP**

Associating a tab with a primary device type or context type impacts all manageable devices that are associated with the primary device type and context type. For change control, minimize edits to the tab after you associate it with a primary device or context type.

## **Add or Edit a Context Tab**

Revert changes to all tabs in the current context by selecting **Restore Tabs to Defaults** from the **Edit Tab** menu.

Modifications apply to the current tenant.

**Follow these steps:**

1. Navigate to the item context to which you want to add or edit a context tab. For example, click the link for a router on a dashboard to open the **Router** context pages.  
The tab that is selected includes an **Edit** icon so that you can access the **Edit** menu.
2. Select the tab that you want to modify.  
The **Edit** icon appears.
3. Click the **Edit** icon, and then select **Add Tab** or **Edit Tab**.  
The **Add Context Tab** or **Edit Context Tab** appears.
4. Complete the following fields:
  - **Default Tab Templates**  
Select one of the default tab templates from **Default Tab Templates** menu. Each template populates the page with the default views for that type of page.  
**Required:** No
  - **Tab Title**  
Enter a tab title. The tab title determines the name that appears at the top of the tabbed context page.  
**Required:** Yes
  - **Context Types**  
To associate the tab with a primary device type or context type, select a context page type.  
**Required:** No
5. In the **Layout** section, do the following:
  - a. Select a layout template for the page from the layout buttons.
  - b. (Optional) Remove unwanted views.
  - c. Click one of the following:
    - **Clear Layout:** Changes the positioning of all views on the page.
    - **[X]:** Removes an individual view from the page.

The views that you can add to the page are shown in categorized lists. The lists are filtered by the selected group or item context.

All registered data sources are represented. However, the available views are limited to those views that are applicable to the context.

**NOTE**

The item context for the page is preselected for the present context.

6. Click to expand the categories of views.
7. Select a view, drag it to the **Layout** pane, and drop it where you want it to appear.
8. Click **Save**.

The context page refreshes to reflect your changes. The changes persist across login sessions, but they are only applied to the current tenant.

**Edit Item Properties**

You can edit the item properties (fields) that are listed for an item in the **Information** section of the **Details** tabs in a context page.

**Prerequisite:** You have the following role rights:

- Modify Device Alias
- Modify Interface Alias
- Modify Device IP Address
- Modify Interface Speed Overrides

**Follow these steps:**

1. On the context page to which you want to edit item properties, select the **Details** context tab if it is not already selected.
2. Click the **Edit properties for this <device/interface>** icon in the **Information** section.  
The **Edit Properties** dialog opens.
3. Edit the available details as desired.

**NOTE**

If **Alias**, **Speed In Override (bps)**, or **Speed Out Override (bps)** are not set, they appear blank.

If you edit an device/interface that does not exist in the data aggregator, the following fields are unavailable:

- **Speed In Override (bps)**
- **Speed Out Override (bps)**

4. Click **Save**.

If you cleared the **Speed In Override (bps)** or **Speed Out Override (bps)** fields, the change takes a few minutes to appear.

**Apply a Business Hours Filter to a Context Page**

To show data in a context page for particular business hours, you can apply a business hours filter to it. This sets the reporting profile that persists during your user session. The filter filters the data that appears in the views on the context page. You cannot modify the business hours within a view's configuration.

You can also set the reporting profile that persists during your user session by selecting a site group that is associated with a business hours definition in context.

**NOTE**

When navigating to a dashboard page from a context page, the applied business hours on context pages (session business hours filter) supersede the page context for the site group if the site group is associated with a business hour filter.

Applied business hours filters to context pages remain applied when switching between context pages, navigating between dashboards, and changing time ranges. Moving to a top-level dashboard or changing the context page removes the business hours filter.

**Follow these steps:**

1. On the context page to which you want to apply a business hours filter, click the **Apply Business Hours** (clock) switch icon in the upper-right corner of the page.

**TIP**

This icon indicates when business hours are and are not being applied to a view on a dashboard.

The Apply Business Hours dialog opens.

If you have applied a business hours filter with a time zone to the context page, **Context Business Hour** shows the applied business hours filter and time zone. Otherwise, **No business hours selected** displays. If you have applied a business hours filter with a time zone as part of your user session, **Session Business Hour** shows the applied business hours filter and time zone.

2. Complete the following fields, and then click **Apply**:

- **Apply Business Hour Filter**

Enable this option. By default, business hours are not enabled.

- **Context Filter Type**

Defines the profile of the context filter type for the business hours filter.

**Options:**

- **Site associated business hours:** Lists the site groups that are associated with a business hours filter and selected on the context page in the Business Hours Filter drop-down.
- **System defined business hours:** Lists all available business hours filters in the **Business Hours Filter** drop-down.

**Default:** System defined business hours

#### – **Business Hours Filter**

Select the business hours filter that you want to apply to the context page. Or, to disable the business hours filter that is being applied during the user session, select **No business hour selected**. The list of business hours definitions are those that have been previously defined.

#### – **Business hour and time zone**

The time zones from which you want to choose in the **Time Zone Filter** drop-down.

##### **Options:**

- **Recommended hours/zones:** Lists the time zones that are associated with the business hours you have chosen.
- **System defined hours/zones:** Lists time zones that are associated to site groups.
- **Any hours/zones:** Lists all time zones.

#### – **Time Zone Filter**

If you selected a business hours filter, select the time zone that you want to apply to the context page.

#### – **Apply Changes**

Select the user-session level to which to apply this business hours filter.

##### **Options:**

- **My Current Session:** Applies this business hours filter to all context pages during your user session.
- **Current Session for Page only:** Applies the business hours filter to this context page only during your user session.

#### **NOTE**

This option supersedes the **My Current Session** option if associated to a business hours filter.

**Default:** My Current Session

The business hours for the specified site are applied to the context page.

## **Rearrange Context Tabs**

Each item context page consists of set of tabbed pages (accessible from the context tabs). In addition to modifying individual tabbed pages, you can rearrange the context tabs in an item context. Modifications are only saved to the current tenant.

**Prerequisites:** You are an Administrator or Designer or you have the Edit Context Pages and Drill into Views role rights.

### **Follow these steps:**

1. Click the **Edit** icon for the context tab that you want to rearrange, and select **Reorder Tabs**. The **Reorder Tabs** dialog opens. A list of **Current Context Page Tabs** appears. The list reflects the current ordering of tabs, from left to right.
2. Select a tab to move, and then drag it to another location in the list.
3. Click **Save**.

The context page refreshes to reflect your changes. The tabs are displayed in a new order from left to right.

If too many tabs are available for the context to display without horizontal scrolling, an arrow appears on the right. Click the arrow to see additional tabs.

## **Views**

Views report collected data in a chart and/or a table format. You can add views to dashboards and context pages.

Views that show data for a group contain collated and aggregated data from data sources. Views that show data for individual items provide a drill-down path to the context page for the item.

In this article:

- [View Contexts](#)
- [Out-of-the-box and Custom Views](#)
- [Add a Custom View to a Dashboard or Context Page](#)
- [View Types](#)

By default, the views on a dashboard show data from the same group and for the same time frame. You can show a specific group by changing the data context for the view.

### **View Contexts**

View contexts act as filters that determine the nature of the data that is displayed in a view. The context of a view can be one of the following:

- **Dynamic**  
Indicates that the context of the view changes with the context of the dashboard.
- **Fixed**  
Indicates that the view uses a specified group, device, or component as a context for the data. For views with fixed context, changing the dashboard context does not affect the view.

Views on a context page default to the context of an item. However, from a context page, you can view the context of another item by dynamically changing the filter.

The context level of a view determines the data that appears in the view:

- **Group**  
Group context level views show data for devices and component items that belong to the selected group. Views that show a single value aggregate the results to the group level.
- **Device**  
Device context level views show data for a specific device. Views that show a single value aggregate the results to the device level.
- **Component / Interface**  
Component or interface context level views show data only for the individual selected device component or interface. Because this context shows only a single item, aggregation does not occur.

### **Out-of-the-box and Custom Views**

**Out-of-the-box views** show data for a preselected metric or a limited set of metrics. These views include only limited customization options. Each registered data source includes associated out-of-the-box views.

**Custom views** provide flexible configuration options using the information from the data aggregator. With custom views, you can select from many collected metric in multiple graphical and tabular formats.

### **Add a Custom View to a Dashboard or Context Page**

You add custom views to dashboards and context pages while adding or creating them.

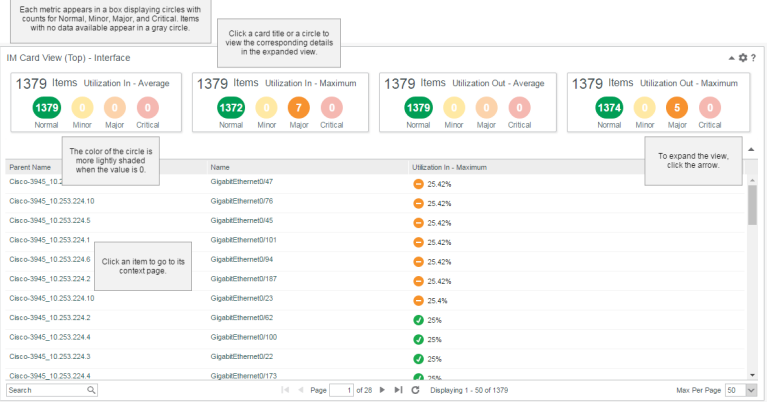
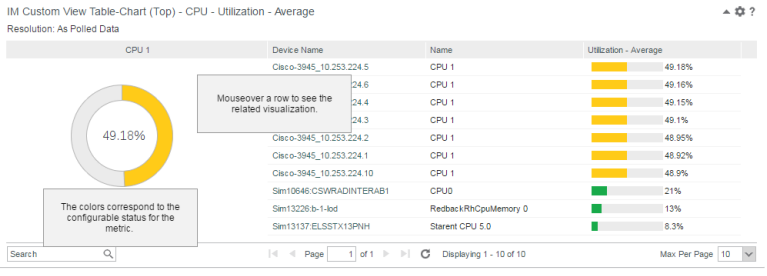
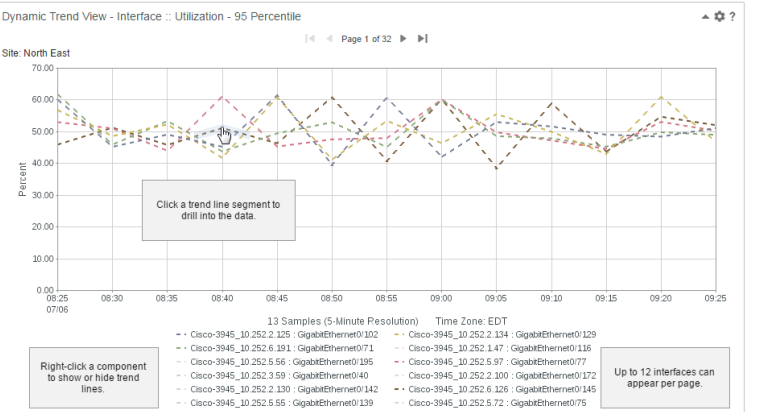
For more information:

- About how to add custom views to dashboards, see [Manage Dashboards](#).
- About how to add custom views to context pages, see [Manage Context Pages](#).

## View Types

The following list shows the view types that you can add to dashboards or context pages:

Preview	View Type	
<p><b>Figure 24: Alarms View Example</b></p>	<p><b>Alarms</b></p> <p>If you have DX NetOps Performance Management, the Alarms view shows Spectrum alarm information.</p>	
<p><b>Figure 25: Bar Chart Example</b></p>	<p><b>Bar Chart</b></p> <p>Display bar charts, or other visualizations of data. The chart shows the expected value of a metric, and the difference from the expected value differs from the others.</p>	
	<p><b>Browser</b></p> <p>An approach to pulling data from a source can show the content of the data. The data can have aspects of your data, and you can update internal data.</p>	
<p><b>Figure 26: Calendar Heat Chart Example</b></p>	<p><b>Calendar Heat Chart</b></p> <p>Provide a month-long view of data on custom thresholds. The chart shows one month with a color-coded view of the data.</p>	

Preview	View Type	
<p><b>Figure 27: Card View Example</b></p>  <p>Each metric appears in a box displaying circles with counts for Normal, Minor, Major, and Critical. Items with no data available appear in a gray circle.</p> <p>Click a card title or a circle to view the corresponding details in the expanded view.</p> <p>IM Card View (Top) - Interface</p> <p>1379 Items Utilization In - Average</p> <p>1379 Items Utilization In - Maximum</p> <p>1379 Items Utilization Out - Average</p> <p>1374 Items Utilization Out - Maximum</p> <p>The color of the circle is more lightly shaded when the value is 0.</p> <p>Click an item to go to its context page.</p> <p>To expand the view, click the arrow.</p> <p>Search</p> <p>Page 1 of 28</p> <p>Displaying 1 - 50 of 1379</p> <p>Max Per Page 50</p>	<p><b>Card</b></p> <p>Display cumulative counts and ranges of each specific metric.</p>	
<p><b>Figure 28: Radial Bar/Table Example</b></p>  <p>IM Custom View Table-Chart (Top) - CPU - Utilization - Average</p> <p>Resolution: As Poled Data</p> <p>CPU 1</p> <p>Device Name</p> <p>Name</p> <p>Utilization - Average</p> <p>49.18%</p> <p>Mouseover a row to see the related visualization.</p> <p>The colors correspond to the configurable status for the metric.</p> <p>Search</p> <p>Page 1 of 10</p> <p>Displaying 1 - 10 of 10</p> <p>Max Per Page 10</p>	<p><b>Chart/Table</b></p> <p>Display a combination of chart and table view.</p>	
<p><b>Figure 29: Dynamic Trend View Example</b></p>  <p>Dynamic Trend View - Interface :: Utilization - 95 Percentile</p> <p>Site: North East</p> <p>Click a trend line segment to drill into the data.</p> <p>Right-click a component to show or hide trend lines.</p> <p>Up to 12 interfaces can appear per page.</p> <p>13 Samples (5-Minute Resolution) Time Zone: EDT</p> <ul style="list-style-type: none"> <li>- Cisco-3945_10.252.2.125 : GigabitEthernet0/102</li> <li>- Cisco-3945_10.252.2.134 : GigabitEthernet0/129</li> <li>- Cisco-3945_10.252.6.191 : GigabitEthernet0/71</li> <li>- Cisco-3945_10.252.2.147 : GigabitEthernet0/136</li> <li>- Cisco-3945_10.252.5.56 : GigabitEthernet0/195</li> <li>- Cisco-3945_10.252.5.97 : GigabitEthernet0/77</li> <li>- Cisco-3945_10.252.3.59 : GigabitEthernet0/40</li> <li>- Cisco-3945_10.252.2.100 : GigabitEthernet0/172</li> <li>- Cisco-3945_10.252.2.130 : GigabitEthernet0/142</li> <li>- Cisco-3945_10.252.6.126 : GigabitEthernet0/145</li> <li>- Cisco-3945_10.252.5.55 : GigabitEthernet0/139</li> <li>- Cisco-3945_10.252.5.72 : GigabitEthernet0/75</li> </ul>	<p><b>Dynamic Trend</b></p> <p>Display combined data trends over time. Trend lines are useful to compare performance of managed items, in order to determine how data is changing.</p>	

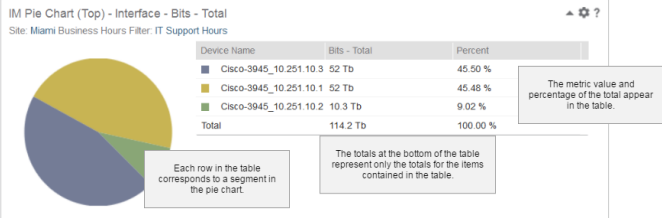
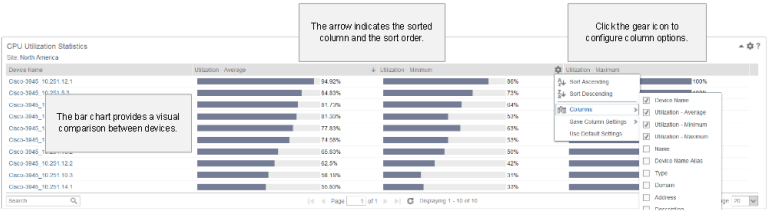
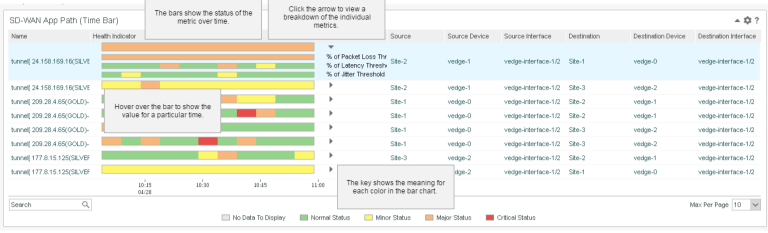
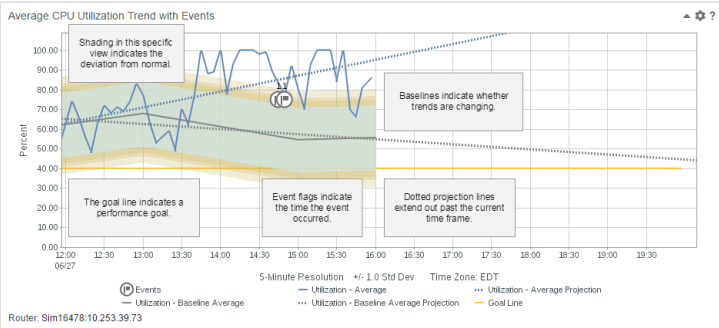
Preview	View Type																																																										
<div><div>Figure 30: Dynamic Table View Example</div><div><div>Dynamic Table Time Series View - Memory</div><div><div>User Group: Group1_2_3</div><div>Time Range: Previous Month</div></div><div><div>View Type: Date Table for Month by Day per Device with Multiple Metrics</div><div>Resolution: Daily Roll-up (Default)Time Interval: 1 DayMetric Calculate Level: by Device[Time Zone: EDT]</div></div><div><div>Quick Filter</div><div><table><tr><th>End Time</th><th>Device Name</th><th>Memory Utilization - Average</th><th>Memory Utilization - Maximum</th><th>Memory Utilization - Minimum</th></tr><tr><td>August 1, 2022 8:00:00 PM</td><td>Cisco [REDACTED].201</td><td>50.38%</td><td>90%</td><td>10%</td></tr><tr><td>August 1, 2022 8:00:00 PM</td><td>Cisco [REDACTED].202</td><td>47.68%</td><td>90%</td><td>10%</td></tr><tr><td>August 1, 2022 8:00:00 PM</td><td>Cisco [REDACTED].203</td><td>50.44%</td><td>90%</td><td>10%</td></tr><tr><td>August 2, 2022 8:00:00 PM</td><td>Cisco [REDACTED].201</td><td>49.46%</td><td>90%</td><td>10%</td></tr><tr><td>August 2, 2022 8:00:00 PM</td><td>Cisco [REDACTED].202</td><td>50.47%</td><td>90%</td><td>10%</td></tr><tr><td>August 2, 2022 8:00:00 PM</td><td>Cisco [REDACTED].203</td><td>49.72%</td><td>90%</td><td>10%</td></tr><tr><td>August 3, 2022 8:00:00 PM</td><td>Cisco [REDACTED].201</td><td>50.68%</td><td>88.24%</td><td>10%</td></tr><tr><td>August 3, 2022 8:00:00 PM</td><td>Cisco [REDACTED].202</td><td>50.52%</td><td>90%</td><td>10%</td></tr><tr><td>August 3, 2022 8:00:00 PM</td><td>Cisco [REDACTED].203</td><td>46.02%</td><td>90%</td><td>10%</td></tr><tr><td>August 4, 2022 8:00:00 PM</td><td>Cisco [REDACTED].201</td><td>51.35%</td><td>90%</td><td>10%</td></tr></table></div><div><div>10 per page</div><div>Page 1 of 10</div><div>Displaying 1 - 10 of 96</div></div></div></div></div> <tr><td></td><td>Dynamic Table</td><td>Compare data from h in a tabular format.</td></tr> <tr><td><div><div>Figure 31: Gauge View Example</div><div><div>IM Gauge Chart (Top) - Interface - Utilization - Average</div><div>Resolution: Hourly Roll-up</div><div><div>Severity thresholds are based on percentage values</div><div><div>The gray region shows the minimum and maximum values.</div><div>The gauge needle shows the average value for the time range.</div></div><div>Min: 25%49.62%Max: 50.03%</div></div></div></div></td></tr>	End Time	Device Name	Memory Utilization - Average	Memory Utilization - Maximum	Memory Utilization - Minimum	August 1, 2022 8:00:00 PM	Cisco [REDACTED].201	50.38%	90%	10%	August 1, 2022 8:00:00 PM	Cisco [REDACTED].202	47.68%	90%	10%	August 1, 2022 8:00:00 PM	Cisco [REDACTED].203	50.44%	90%	10%	August 2, 2022 8:00:00 PM	Cisco [REDACTED].201	49.46%	90%	10%	August 2, 2022 8:00:00 PM	Cisco [REDACTED].202	50.47%	90%	10%	August 2, 2022 8:00:00 PM	Cisco [REDACTED].203	49.72%	90%	10%	August 3, 2022 8:00:00 PM	Cisco [REDACTED].201	50.68%	88.24%	10%	August 3, 2022 8:00:00 PM	Cisco [REDACTED].202	50.52%	90%	10%	August 3, 2022 8:00:00 PM	Cisco [REDACTED].203	46.02%	90%	10%	August 4, 2022 8:00:00 PM	Cisco [REDACTED].201	51.35%	90%	10%		Dynamic Table	Compare data from h in a tabular format.	<div><div>Figure 31: Gauge View Example</div><div><div>IM Gauge Chart (Top) - Interface - Utilization - Average</div><div>Resolution: Hourly Roll-up</div><div><div>Severity thresholds are based on percentage values</div><div><div>The gray region shows the minimum and maximum values.</div><div>The gauge needle shows the average value for the time range.</div></div><div>Min: 25%49.62%Max: 50.03%</div></div></div></div>
End Time	Device Name	Memory Utilization - Average	Memory Utilization - Maximum	Memory Utilization - Minimum																																																							
August 1, 2022 8:00:00 PM	Cisco [REDACTED].201	50.38%	90%	10%																																																							
August 1, 2022 8:00:00 PM	Cisco [REDACTED].202	47.68%	90%	10%																																																							
August 1, 2022 8:00:00 PM	Cisco [REDACTED].203	50.44%	90%	10%																																																							
August 2, 2022 8:00:00 PM	Cisco [REDACTED].201	49.46%	90%	10%																																																							
August 2, 2022 8:00:00 PM	Cisco [REDACTED].202	50.47%	90%	10%																																																							
August 2, 2022 8:00:00 PM	Cisco [REDACTED].203	49.72%	90%	10%																																																							
August 3, 2022 8:00:00 PM	Cisco [REDACTED].201	50.68%	88.24%	10%																																																							
August 3, 2022 8:00:00 PM	Cisco [REDACTED].202	50.52%	90%	10%																																																							
August 3, 2022 8:00:00 PM	Cisco [REDACTED].203	46.02%	90%	10%																																																							
August 4, 2022 8:00:00 PM	Cisco [REDACTED].201	51.35%	90%	10%																																																							
	Dynamic Table	Compare data from h in a tabular format.																																																									
<div><div>Figure 31: Gauge View Example</div><div><div>IM Gauge Chart (Top) - Interface - Utilization - Average</div><div>Resolution: Hourly Roll-up</div><div><div>Severity thresholds are based on percentage values</div><div><div>The gray region shows the minimum and maximum values.</div><div>The gauge needle shows the average value for the time range.</div></div><div>Min: 25%49.62%Max: 50.03%</div></div></div></div>																																																											

Preview	View Type																																																						
<div><div>Figure 32: Group Scorecard Trend View Example</div><div><div>Group: My Assigned Groups &gt; All Groups &gt; &gt; Reporting Groups &gt; Small Group 2 [image]</div><div>Previous Month: Aug 1, 2018 12:00 AM - Aug 31, 2018 11:59 [image]</div></div><div><div>IM Group Scorecard Trend - Interface - Utilization - Average</div><div><div>User Group: Small Group 2</div><div>Timeframe: Previous MonthMetric Calculate Level: by GroupResolution: Daily Roll-up</div><div>▲ Critical Status (80.0)▲ Major Status (80.0)▲ Minor Status (40.0)▲ Normal Status</div><div><div>The header shows the timeframe and calculation level.</div><div>Each column represents a time interval or aggregate value</div></div><div><table><tr><td>Group/Sub-Group</td><td>Jan 31, 2018</td><td>Feb 28, 2018</td><td>Mar 31, 2018</td><td>Apr 30, 2018</td><td>May 31, 2018</td><td>Jun 30, 2018</td><td>Jul 31, 2018</td><td>Aug 31, 2018</td><td>Overall Average...</td><td>Projection: Sep 30...</td><td>Projection: Oct 31...</td><td>Projection: Nov 30...</td></tr><tr><td>Small Group 2</td><td>38.21%</td><td>38.24%</td><td>37.99%</td><td>38.23%</td><td>38.23%</td><td>38.23%</td><td>38.23%</td><td>38.21%</td><td>38.21%</td><td>38.21%</td><td>38.21%</td><td>38.21%</td></tr><tr><td>Small Group</td><td>26.45%</td><td>26.47%</td><td>26.45%</td><td>26.46%</td><td>26.45%</td><td>26.47%</td><td>26.46%</td><td>26.46%</td><td>26.46%</td><td>26.47%</td><td>26.47%</td><td>26.47%</td></tr><tr><td>Small Group</td><td>49.98%</td><td>50%</td><td>49.97%</td><td>50%</td><td>50%</td><td>50%</td><td>50%</td><td>50%</td><td>50%</td><td>50%</td><td>50%</td><td>50%</td></tr></table></div><div><div>The first line shows the aggregate for all items across the entire group. The following lines show the subgroups.</div><div>For percentile metrics, the first line would show the maximum value for all items across the entire group.</div><div>The colored icons indicate the status.</div></div></div></div></div> <tr><td></td><td>Group Scorecard Trend</td><td>Display performance group. These views p how key metrics perfo</td></tr>	Group/Sub-Group	Jan 31, 2018	Feb 28, 2018	Mar 31, 2018	Apr 30, 2018	May 31, 2018	Jun 30, 2018	Jul 31, 2018	Aug 31, 2018	Overall Average...	Projection: Sep 30...	Projection: Oct 31...	Projection: Nov 30...	Small Group 2	38.21%	38.24%	37.99%	38.23%	38.23%	38.23%	38.23%	38.21%	38.21%	38.21%	38.21%	38.21%	Small Group	26.45%	26.47%	26.45%	26.46%	26.45%	26.47%	26.46%	26.46%	26.46%	26.47%	26.47%	26.47%	Small Group	49.98%	50%	49.97%	50%	50%	50%	50%	50%	50%	50%	50%	50%		Group Scorecard Trend	Display performance group. These views p how key metrics perfo
Group/Sub-Group	Jan 31, 2018	Feb 28, 2018	Mar 31, 2018	Apr 30, 2018	May 31, 2018	Jun 30, 2018	Jul 31, 2018	Aug 31, 2018	Overall Average...	Projection: Sep 30...	Projection: Oct 31...	Projection: Nov 30...																																											
Small Group 2	38.21%	38.24%	37.99%	38.23%	38.23%	38.23%	38.23%	38.21%	38.21%	38.21%	38.21%	38.21%																																											
Small Group	26.45%	26.47%	26.45%	26.46%	26.45%	26.47%	26.46%	26.46%	26.46%	26.47%	26.47%	26.47%																																											
Small Group	49.98%	50%	49.97%	50%	50%	50%	50%	50%	50%	50%	50%	50%																																											
	Group Scorecard Trend	Display performance group. These views p how key metrics perfo																																																					



Preview	View Type	
<p><b>Figure 33: Group Scorecard Table View Example</b></p> <p>The screenshot displays the 'Group Scorecard Table View'. It features a left-hand navigation pane with a tree structure of groups. The main area is a table with columns: Group/Sub-Group Name, Device Name, Health Indicator, Percent Discards - Average, Percent Errors - Average, and Utilization - Average. The table lists various devices and their associated metrics. Callouts provide additional context: 'You can sort the top level of the hierarchy. You can also sort the results within the hierarchy.', 'The health indicator column shows an overall score based on the severity and weight of each metric for each item.', 'Multiple metrics appear in the same multi-metric scorecard.', 'The first line shows the aggregate across the group. The following lines show the metric aggregated to the selected metric calculable level.', 'When the calculable level is Component Hierarchy, the view shows sub-groups and the items in those groups.', 'Each metric has its own threshold for the colored score.', 'The highest severity for any of the reported metrics determines the color of the Unhealthy Score circle.', 'A relative weighting of metrics prioritizes metrics that are more indicative of overall health.', and 'No Data To Display' for some metrics.</p>	<p><b>Group Scorecard Table</b></p> <p>Display multiple metrics component for a selected indicators of performance of user-defined threshold</p>	
<p><b>Figure 34: Inventory Hierarchy View Example</b></p> <p>The screenshot displays the 'Inventory Hierarchy View'. It features a left-hand navigation pane with a tree structure of devices. The main area is a table with columns: Name Alias, Type, Domain, Address, Description, Current, Life Cycle, and Context Type. The table lists various devices and their associated metrics. Callouts provide additional context: 'On Demand', 'Quick Filter', 'Page 1 of 4', and 'Displaying 1 - 100 of 315'.</p>	<p><b>Inventory Hierarchy</b></p> <p>Display the group tree pinned to the current</p>	
<p><b>Figure 35: Map View Example</b></p> <p>The screenshot displays the 'Map View'. It features a map of the world with various locations marked. Callouts provide additional context: 'Geo Map', 'Automatic Group: Sites', 'Click on a site to see its tunnels', 'Click the magnifying glasses to zoom in or out.', 'Click a site to view the connections to and from the site.', 'Click the link in the hover details to drill in further.', 'Click the arrows to view the details for each site metric.', 'Click the arrows to view the details for each tunnel metric.', 'Double-click anywhere on the map to focus on a particular location.', 'Hover on a site to view details for the site.', 'Hover on a connection to view details for the tunnel.', 'Jitter - Average: 0.00', 'Latency - Average: 0.30', 'Packet Loss Percentage - Average: 13.83', '1.1.1.5-gold-1.1.1.7-gold', '1.1.1.5-gold-1.1.1.7-silver', 'Jitter - Average: 0.00', 'Latency - Average: 0.04', 'Packet Loss Percentage - Average: 13.97', 'Locations: Framingham Data Center, Portsmouth Data Center, ITC Development Center', and 'Tunnels: 1.1.1.5-gold-1.1.1.7-gold, 1.1.1.5-gold-1.1.1.7-silver'.</p>	<p><b>Map</b></p> <p>For SD-WAN tunnels DX NetOps Virtual Network</p>	

Preview	View Type	
<p><b>Figure 36: MultiTrend Example</b></p>	<p><b>MultiTrend</b></p> <p>Display combined trend chart.</p>	
<p><b>Figure 37: MultiView Example</b></p>	<p><b>MultiView</b></p> <p>Display combined static data</p>	
<p><b>Figure 38: On-Demand Report Example</b></p>	<p><b>On-Demand Report</b></p> <p>Display static data set in chart and table.</p>	

Preview	View Type	
<div>Figure 39: Pie Chart View Example</div> <div></div>	<div>Pie Chart</div> <div>Display relative values for small groups of items as a percentage of the total.</div>	
<div>Figure 40: Table View Example</div> <div></div>	<div>Table</div> <div>Display vertically and horizontally many items, or many items so you can view the details.</div>	
<div>Figure 41: Time Bar Chart View Example</div> <div></div>	<div>Time Bar Chart</div> <div>Displays a bar to show the status of a metric over time. These view slices represent the status of the metric over time.</div>	
<div>Figure 42: Trend View Example</div> <div></div>	<div>Trend</div> <div>Displays the value of a metric over time and help identify when trends are changing. Provide one or more trends.</div>	

## Customize Views

NetOps Portal uses the view options that you customize to determine how it filters the data it reports in views.

You can collapse views, modify view settings, and access context-sensitive help using the icons in the upper right of a view.

For some view types, you can customize the table or chart format. You can also select the data that is reported in custom views.

**NOTE**

The customizations that you make to a view *during* a DX NetOps Performance Management upgrade do not appear in views. To use the new customizations, reset the view to the default options.

You can customize views in the following ways:

- [Edit a View](#)
- [Include Metrics in a View](#)
- [Set a Custom Time Range](#)
- [Assign a Business Hours Filter and a Time Zone to a View](#)
- [Change the Data Context for a View](#)
- [Set the Scope for View Settings](#)
- [Restore the View Settings to Default Values](#)
- (23.3.4 and higher) [Set a Goal or Threshold Line Trend Line](#)

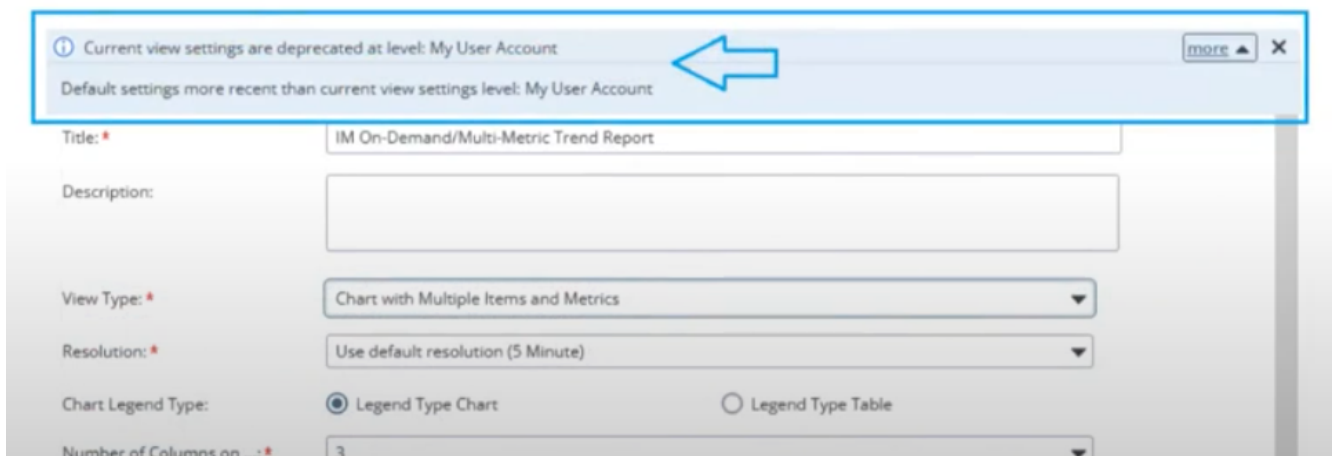
**Edit a View**

To edit a view, click the View Settings (gear) icon on the view, and then select **Edit** from the menu. The view settings dialog opens.

**IMPORTANT**

If you have previously customized the view, a notification similar to the following appears at the top of the view settings dialog when editing the view:

**Figure 43: Notification\_view**



The notification specifies that the settings have been configured, and the level at which the settings have been configured. These configurations are retained.

For more information, see [Set the Scope for View Settings](#).

**Include Metrics in a View**

Metrics fall into the counter profile type or gauge profile type. NetOps Portal reports metrics in several categories, such as bytes, and in percentages, such as utilization. NetOps Portal sums up metrics with a counter profile type over time, and rolled up as such. NetOps Portal averages metrics with a gauge profile type over time.

To include a metric in a view, from the view settings dialog, select a metric that you want to include in the view from the **Metric Value** list.

**NOTE**

The following options impact the available metrics in the **Metric Value** list:

- **Metric Family**  
The selected **Metric Family** populates the **Metric Value** list with the metrics associated with the **Metric Family**. A metric family is a normalized categorization of metrics associated with device collections, vendor certifications, and monitoring profiles. The metric family contains information about the supported units and rollup strategy for when an associated metric is populated.
- **Metric Filtering**  
Select to limit the **Metric Value** list to only the metrics that are appropriate to the metric profile type (counter or gauge). Clear this option to make available the metrics for the selected **Metric Family**. If you select a metric that is inconsistent with the metric profile type, NetOps Portal flags the metric with a warning.
- **Baseline Metrics**  
Select to add baselines to the **Metric Value** list. Baseline data helps to characterize past performance for selected monitored parameters, assess present performance, and estimate future performance.  
For more information, see [Baseline Calculations](#).

**Set a Custom Time Range**

You can filter data based on specific time periods by selecting a custom time range for the view. The custom time range that is set for a view overrides the time range that is set for the dashboard.

You can specify a custom time range for all views except the calendar heat chart and on-demand report views.

For more information about these view types, see [Views](#).

**Follow these steps:**

1. From the view settings dialog, in the **Customize Time Range** section, complete the following:
  - **Custom Time Range**  
Specifies whether to define a custom time range. Select **Enabled**.  
**Options:** Enabled, Disabled
  - **Time Range**  
When the custom time range is enabled, defines the custom time range.  
**Options:** Last Hour, Last 4 Hours, Last 8 Hours, Last 12 Hours, Last 24 Hours, Last 7 Days, Last 14 Days, Last 30 Days, Last 3 Months, Last 12 Months, Yesterday, Previous Week, Previous Month, Today, Current Week, Current Month  
**Default:** Last Hour
2. Click **Save**.

The custom time range is indicated in the subtitle of the view.

**Assign a Business Hours Filter and a Time Zone to a View**

You can show data in a view for particular business hours by assigning a business hours filter and time zone to the view. You can apply business hours filters to out-of-the-box views that report data from the data aggregator and to custom views. The view subtitle indicates if you have assigned a time zone and applied a business hours filter.

You can also apply a business hours filter to a view by associating the business hours definition with a site group and then apply this site group to the view. Views show data only if the group context includes groups with the assigned time zone. NetOps Portal applies the filter to all the devices and components in the site group.

Trend views (including MultiTrend views and MultiViews), dynamic trend views, calendar heat chart views, and trend views in on-demand reports display data with the periods *outside* of the applied business hours as shaded. The data in these views is not limited to the business hours.

Data are displayed as shaded only in the following circumstances:

- The data is on a page that has a specified site group with associated business hours.
- The data is on a page that has a specified business hours set at user-session level.
- The data is in a view that has a specified site group with associated business hours.
- The data is in a view that has an assigned a business hours filter and time zone.

When applying a business hours filter to a page as part of your user session or view, the business hours override other business hour settings in the following priority order:

1. Business hours and time zone explicitly assigned to a view.
2. Site group that is associated with a business hours filter explicitly applied to a view.
3. Business hours filter applied to a dashboard or to a context page (user-session level business hours).
4. Site group that is associated with a business hours filter specified at the context level.

#### NOTE

When assigning a business hours filter and a time zone to a view, if the time frame is *outside* the business hours, the view displays the following message:

No data is within business hours for this time frame

For more information:

- About business hours filters, see [Configure Business Hours Filtering](#).
- About how trend views show business hours, see [Trend Views](#).
- About how dynamic trend views show business hours, see [Dynamic Trend Views](#).
- About how dynamic table views show business hours, see [Dynamic Table Views](#).
- About how trend views in on-demand reports show business hours, see [On-Demand Reports](#).
- About how calendar heat chart views show business hours, see [Calendar Heat Chart Views](#).

#### Follow these steps:

1. From the view settings dialog, in the **Apply Business Hours** section, complete the following fields:  
In this section, the **Current Context Business Hour** shows the applied business hours filter and time zone to the context page or session. The **Current View Business Hour** shows the business hours filter that you have applied to the view.
  - **Apply Business Hour Filter**  
By default, the view inherits the business hours filter applied to context pages (user-session level business hours). Enable this option to override this filter and assign a business hours filter and time zone to the view.
  - **Business Hours Filter**  
Select the business hours filter that you want to apply to the view. Or, to disable the business hours filter that is being applied to the context page, select **No business hour selected**. The list of business hours definitions are those that have been previously defined.
  - **Business hour and time zone**  
The time zones from which you want to choose in the **Time Zone Filter** drop-down.  
**Options:**
    - **Recommended hours/zones:** Lists the time zones that are associated with the business hours that you have chosen.
    - **System defined hours/zones:** Lists time zones that are associated to site groups.
    - **Any hours/zones:** Lists all time zones.**Default: Recommended hours/zones**
  - **Time Zone Filter**  
The time zone that you want to assign to the view.
2. (Optional) In the **Context Settings** section, to show a particular site group in the view, select **Fixed**, and then select a site group.

**NOTE**

Ensure that the view shows data by selecting a site group that has the selected time zone and business hours.

3. Click **Save**.

The view filters data to the selected business hours according to the view type.

### **Change the Data Context for a View**

By default, the views on a dashboard show data from the same group and for the same time range. You can change the data context for a particular view to show a specific group. View contexts act as filters that determine the nature of the data that views display. Views show data only from a selected item or group. The view subtitle indicates the selected context, and includes a lock icon.

You can compare data that was collected at different times by selecting a custom time range for a view.

From the view settings dialog, in the **Context Settings** section, complete the following, and then click **Save**:

- **Context**  
Defines the context for the view. Select the context according to the view type. Select **Fixed**.  
**Options:**
  - **Dynamic:** Indicates that the context of the view changes with the context of the dashboard page.  
For more information, see [Dashboards](#).
  - **Fixed:** Indicates that the view shows data from a specified item, group, device, or component as the context for the data.
- Default:** Dynamic
- **Context Type**  
Defines the context type for the view.  
**Options:** Summary
- **Current Item**  
Defines the group for the view context.

### **Set the Scope for View Settings**

To determine which users see view changes, set the scope when changing the view settings.

From the view settings dialog, select one of the following options from the **Apply Changes** drop-down, and then click **Save**:

**NOTE**

The options are scoped according to your user account permissions.

- **My Current Session**  
Save the changes to the current session. Changes do not persist after the logged in session ends. Values set at this level supersede settings at the **My User Account** and **For All Tenant Users** levels during the logged in session. This option is available only after a view has been configured and has been added to a dashboard page.  
For more information about how to add a view to a dashboard, see [Manage Dashboards](#).
- **My User Account (Default)**  
Save the changes to the current user account rendering the view. Changes persist after the logged in session ends and the user logs back in. Values set at this level supersede settings at the **For All Tenant Users** level for the user account.
- **For All Tenant Users**  
Save the changes to all user accounts associated with the tenant rendering the view. Values set at this level supersede the out-of-the-box default settings for all user accounts associated with tenant.

## **Restore the View Settings to Default Values**

You can return view settings to the default values.

From the view settings dialog, click **Use Defaults**. The view reverts to the default options. These changes apply only to the selected scope when you save the view.

Expect the following behavior when restoring defaults for each scope:

- **My Current Session**  
The settings set at the **My Current Session** level are cleared. Views use the values set at the **My User Account** level, the values set at the **For All Tenant Users** level, or the out-of-the-box default settings.
- **My User Account**  
The settings set at the **My User Account** level are cleared. Views use the values set at the **For All Tenant Users** level or the out-of-the-box default settings.
- **For All Tenant Users**  
The settings set at the **For All Tenant Users** level are cleared. Views use the out-of-the-box default settings.

## **Set a Goal or Threshold Line Trend Line**

### **(23.3.4 and higher)**

You can set to show a device/interface threshold (line) in dynamic trend views, MultiViews, and MultiTrend views, and trend views.

For more information:

- See [Dynamic Trend Views](#).
- See [Trend Views](#).

## **Set the Resolution for Reported Data**

In this article:

- [Determine the Resolution for a View](#)
- [Set the Resolution for Reported Data](#)
- [View Report Resolution Settings](#)

### **Determine the Resolution for a View**

Trend charts and related views show data as a time series. Time series data shows values for multiple time points across the time range of the report. The resolution setting for a trend view determines the amount of time that each data point in the chart represents. For trend views, the *chart* subtitles indicate the data resolution.

Table views also include data resolution settings. For table views, the resolution represents the data rollup level to show: as polled, hourly rollup, or daily rollup. For table views, the *view* subtitle indicates the data resolution.

The following examples describe some table view subtitles:

- **Resolution: Hourly Roll-Up (Overridden)**  
Hourly resolution is requested for the selected time range for the current user.
- **Resolution: Hourly Roll-Up (Overridden) Resolution Returned: Daily Roll-Up**  
Hourly resolution is requested for the selected time range for the current user. The data source does not have hourly data for the entire time range, so the view provides daily resolution.

### **Follow these steps:**

1. Click the View Settings (gear) icon on the view, and then select **Edit** from the menu.  
The view options appear in the dialog that opens.



2. For **Resolution**, determine the appropriate resolution using the time range by selecting **Use default resolution**.

### **Set the Resolution for Reported Data**

You can set the resolution for reported data in some view types, such as table views and trend views.

#### **NOTE**

You can only set the number of days to a value that is *lower* than the corresponding data aggregator data retention value.

#### **Follow these steps:**

1. Hover over **Administration, Configuration Settings**, and then click **Report Resolution**.  
The **Manage Report Resolution Settings** page appears.  
The following image shows an example of this page:
2. Complete the following fields under the **Number of days for user without Run Dashboards At Higher Resolution role right** section:
  - **As Polled**  
The data aggregator as polled data retention value is 34.  
**Maximum time range:** Less than, or equal to, 24 hours  
**Default:** 1
  - **Hourly**  
The data aggregator hourly data retention value is 64.  
**Maximum time range:** Greater than 24 hours to less than 14 days  
**Default:** 31
  - **Daily**  
The data aggregator daily data retention value is 150.  
**Maximum time range:** 14 days to less than one year
  - **Roll-up**  
The data aggregator roll-up data retention value is 720.  
**Maximum time range:** One year or longer
3. Complete the following fields under the **Number of days for user with Run Dashboards At Higher Resolution role right** section:
  - **As Polled**  
The data aggregator as polled data retention value is 34.  
**Maximum time range:** Less than, or equal to, 24 hours  
**Default:** 31
  - **Hourly**  
The data aggregator hourly data retention value is 64.  
**Maximum time range:** Less than 30 days, and more than 24 hours

#### **NOTE**

This value must be less than or equal to the data aggregator data retention value.

**Default:** 90

4. Click **Save**.

The report settings are saved.

### **View Report Resolution Settings**

You view the resolution settings on the **Manage Report Resolution Settings** page. These settings determine the default data resolution that NetOps Portal applies to trend views and table views. When a view uses a default resolution, these settings determine the resolution NetOps Portal applies to the time range that is in force for the view.

The view shows the resolution only when the following conditions are true:

- The user can view report data at that resolution for that time range.
- The data resolution is within the data retention policy of the data source.

If the data is unavailable or the time range is too great, NetOps Portal overrides the specified resolution automatically. NetOps Portal returns the closest permissible resolution for the time range. For trend views, the *chart* subtitles indicates the data resolution. For tables views, the *view* subtitle indicates the data resolution.

The time range for a view determines the default data resolution. The following values show the maximum time range for each data resolution:

#### Number of days for user with Run Dashboards At Higher Resolution role right

- **As Polled**  
The data aggregator as polled data retention value is 34.  
**Maximum time range:** Less than, or equal to, 24 hours  
**Default:** 1
- **Hourly**  
The data aggregator hourly data retention value is 64.  
**Maximum time range:** Greater than 24 hours to less than 14 days  
**Default:** 31
- **Daily**  
The data aggregator daily data retention value is 150.  
**Maximum time range:** 14 days to less than one year
- **Roll-up**  
The data aggregator roll-up data retention value is 720.  
**Maximum time range:** One year or longer

The currently selected time range determines the possible resolution options. The following values show the maximum time range for each data resolution:

#### Number of days for user with Run Dashboards At Higher Resolution role right

- **As Polled**  
The data aggregator as polled data retention value is 34.  
**Maximum time range:** Less than, or equal to, 24 hours  
**Default:** 31
- **Hourly**  
The data aggregator hourly data retention value is 64.  
**Maximum time range:** Less than 30 days, and more than 24 hours  
**NOTE**  
This value must be less than or equal to the data aggregator data retention value.  
**Default:** 90
- **Daily**  
**Maximum time range:** More than, or equal to, 30 days

Users *without* the Run Dashboards At Higher Resolution role right use the following maximum time ranges:

- **As Polled**  
**Maximum time range:** Less than, or equal to, 31 days
- **Hourly**  
**Maximum time range:** More than 31 days
- **Daily**  
**Maximum time range:** More than three months

## Alarms View

Quickly focus on resolving the most impactful issues using the **Alarms** view.

If you have configured DX NetOps Spectrum (Spectrum) as a data source, you can view and manage Spectrum alarms from the **Alarms** view in NetOps Portal. To view this view, hover over **Alarms**, and then click **Alarm Console**.

For more information about how to register Spectrum as a data source, see [Integrate with DX NetOps Spectrum for Fault Management](#).

The **Alarms** view provides prioritized list of Spectrum alarms. This list provides visibility into other potentially related issues on the same device or connected to a device. The view provides standard view configuration options, such as context settings.

For more information about how to change the data context of a view, see [Customize Views](#).

In this article:

- [Role Rights](#)
- [Context Settings](#)
- [View Alarms Data](#)
- [Display the Most Recent Data in the View](#)
- [Alarms View Elements](#)
- [Sort Multiple Columns](#)
- [Manage Alarm Filters](#)
- [Filter the Alarms View](#)
- [Manage Alarms](#)
- [Configure an Alarms View](#)

### Role Rights for the Alarms View

The following role rights are associated with the **Alarms** view:

- **Allow Alarm Modification Actions**  
Allows users to acknowledge, clear, or assign a troubleshooter to alarms.
- **Allow Alarm Triage Actions**  
Allows users to perform alarm triage actions.
- **Allow Alarm Filter Creation**  
Allows users to create and manage alarm filters.
- **Allow Alarm Filter Management**  
Allows users to manage and assign alarm filters to other users.

For more information, see [Data Source Role Rights](#).

### Context Settings for the Alarms View

In Spectrum, you can view the alarms asserted on a global collection by selecting it. Group filters in NetOps Portal behave similarly. You can view the alarms asserted on a group, which can correspond to your global collections, by selecting the group.

When you select a custom group for reporting, NetOps Portal scopes the **Alarms** view to the elements within that custom group. NetOps Portal groups can include Spectrum landscapes and global collections. Landscapes and global collections are available under **All Groups**, **Inventory**, **Data Sources**, and **your Spectrum data source**. If the group context includes landscapes or global collections, an alarms view can show alarms for devices that are not synced to NetOps Portal.

For more information, see [Manage Groups](#).

## **View Alarms Data**

The following views display Spectrum alarms data:

- **Alarms**

Displays a list of active alarms and their details based on selected filters.

**NOTE**

If you experiencing issues, such as Spectrum alarms are not displaying in the view, review the symptoms and solutions in [Spectrum Alarms Are Missing from the Alarms View](#).

- **Alarm Details**

Displays the details of the selected alarm. You can add or remove the attributes to display in this pane. By default, this pane displays the following attributes:

- Severity
- Date/Time
- Item Name
- IP Address
- Model Type
- Acknowledged
- Contact Person
- Troubleshooter
- Trouble Ticket ID
- Number of Occurrences
- Impact

- **Impact: Management Lost**

Displays the devices that this alarm impacts.

- **Impact: Symptoms**

Displays the alarms leading up to this alarm.

- **Neighbor Topology**

Displays the neighboring models that are connected to the device of the selected alarm.

- **Interfaces**

Displays the interfaces associated with the selected alarm.

- **Events**

Displays the events associated with the selected alarm. If you have the Export to CSV role right, you can export the events list to CSV using the Gear icon that is in this pane. You can also filter and search the events list from this pane.

**NOTE**

When you search or filter the **Alarms** view, the returned results show only the alarms up to the maximum alarms retrieved, which is specified in the view settings.

- **Log Events**

Displays the **Event Timeline** and **Event History** log events views.

**NOTE**

To prevent the following message from appearing when viewing log events for the router (clicking the **Log Events** tab in context of a router), configure the OI Connector:

No log information available for the device. DX AIOps must be configured here.

For more information about how to install and configure Log Analytics for Insights, see [Install and Configure Log Analytics for Insights](#).

## **Display the Most Recent Data in the View**

NetOps Portal displays the most recent data when you refresh the view manually or when **Auto Refresh Page** is turned on (enabled). By default, NetOps Portal does not refresh alarms automatically (**Auto Refresh Page** is turned off).

For more information about how to turn on **Auto Refresh Page**, see [Manage Dashboards](#).

### NOTE

**Best Practice:** When viewing the **Alarms** and **Events** views, turn on **Auto Refresh Page**.

## Alarms View Elements

The following example shows the important elements of the **Alarms** view:

**Figure 44: Alarms\_Views**

The screenshot displays the DX NetOps Alarms view. At the top, there is a table of alarms with columns: Severity, Date/Time, Item Name, Model Type, IP Address, Alarm Title, Impact, Number of Occurrences, Acknowledged, Troubleshooter, and Trouble Ticket ID. A row is highlighted in blue, and a callout box points to it with the text: "Click on a row in the table to see more details about an alarm." Another callout box points to the device name in the table with the text: "Click on the device name in the table to go to the context page of that device." Below the table, there are buttons for "Acknowledge", "Unacknowledge", "Clear", "Troubleshooter", "Poll", "Ping", "Traceroute", "On Demand", "Manage Life Cycle", and "Create Ticket". A callout box points to these buttons with the text: "Manage and troubleshoot alarms with the available buttons." Below the buttons, there is a section for "Alarm Details" for the selected alarm. It includes a sidebar with tabs for "Impact: Management Lost", "Impact: Symptoms", "Neighbor Topology", "Interfaces", and "Events". The main area shows the alarm details, including Severity (Critical), Date/Time (August 13, 2018 12:12:51 PM EDT), Item Name (cis7505-96.1122.ca.com), IP Address (136.42.96.11), Model Type (Cisco7505), Acknowledged (No), Contact Person (comm), Troubleshooter (None), Trouble Ticket ID (None), Number of Occurrences (1), and Impact (621). A callout box points to the "Events" tab with the text: "Select a tab to view the related pane." The right side of the details section shows the alarm description: "Device cis7505-96.1122.ca.com of type Rtr\_Cisco has stopped responding to polls and/or external requests. An alarm will be generated." It also lists symptoms, probable causes, and actions.

## Sort Multiple Columns

You can sort by up to three columns in the **Alarms** view.

### NOTE

When you sort the **Alarms** view, the returned results show only the alarms up to the maximum alarms retrieved, which is specified in the view settings.

### Follow these steps:

1. In the column header that you want to sort first, click the Gear icon.
2. Select **Sort First**, and then click one of the following options:
  - **Sort Ascending**  
Sort the column in ascending order.
  - **Sort Descending**  
Sort the column in descending order.
  - **Remove Sort**  
Remove the sort order from the column.

### NOTE

You can remove the sort order by selecting CTRL+clicking a column heading using your keyboard.

3. In the column header to sort second, click the Gear icon.
4. Select **Sort Second**, and then click a sort option.

**NOTE**

You can add it as a secondary sort by SHIFT+clicking a column heading using your keyboard.

5. In the column header to sort third, click the Gear icon.
6. Select **Sort Third**, and then click a sort option.

**Manage Alarm Filters**

With the Allow Alarm Filter Management role right, you can manage alarm filters (create them, edit them, and assign them to other users):

- [Create an Alarm Filter](#)
- [Assign Filters to Specific Users or Roles](#)
- [Manage Saved Filters](#)

**Create an Alarm Filter**

With the Allow Alarm Filter Creation role right, you can create alarm filters. You can add up to five custom filter attributes of each type (String, Boolean, and Integer) to the **Alarm Console** dashboard. NetOps Portal does not display attributes that you add over this limit in the **Alarms** view or in alarm filters.

**TIP**

Users with the Administer Users and the Allow Access to REST Services rights can manage alarm filter attributes using the NetOps Portal web services.

For more information, see [Alarm Attributes Web Service](#).

**Follow these steps:**

1. To the right of the filter drop-down, click the plus icon.  
The **Create Filter** dialog opens.
2. Complete the fields, specify the conditions for the alarm filter, and then click **Save**:
  - **For My Use**  
The alarm filter is available only for your user account.
  - **For All Users With My Role**  
The alarm filter is available only for users with your role.
  - **For All Tenant Users**  
The alarm filter is available for all the tenant users.

**NOTE**

If you add a role, users with that new role do not have access to the shared filter. You can assign the filter to the new role.

- **For Specific Users or Roles**  
The alarm filter is available for selected users or roles.

The alarm filter appears in the filter drop-down.

**Assign an Alarm Filter to a Specific User or Role**

With the Allow Alarm Filter Management role right, you can assign shared alarm filters to specific users or roles.

**Follow these steps:**

1. From the filter drop-down, select the filter that you want to assign.
2. Click **Users** or **Roles**.
3. Select the user or role to which to assign the alarm filter, and then click **Save**.

The alarm filter is assigned to a the user or role.

## **Manage Saved Alarm Filters**

### **Follow these steps:**

1. Click the filter drop-down, and then select **Manage Saved Filters**.  
The **Manage Saved Filters** dialog opens.
2. Manage (add, edit, copy, or delete) alarm filters as necessary, and then save your changes.

## **Filter the Alarms View**

### **NOTE**

When you filter the **Alarms** view, the returned results show only the alarms up to the maximum alarms retrieved, which is specified in the view settings.

You can filter alarms by the following attributes:

- **Acknowledged**  
Indicates whether the alarm is acknowledged.  
**Attribute ID:** 0x11f4d  
**Type:** BOOLEAN
- **Cause Code**  
Identifies the cause of the alarm.  
**Attribute ID:** 0x11f50  
**Type:** HEX
- **Clearable**  
Indicates whether the alarm is clearable.  
**Attribute ID:** 0x11f9b  
**Type:** BOOLEAN
- **Contact Person**  
The person to contact when device problems occur.  
**Attribute ID:** 0x23000c  
**Type:** STRING
- **IP Address**  
The IP address for the item.  
**Attribute ID:** 0x12d7f  
**Type:** ADDRESS\_RANGE
- **Location**  
The location of the device.  
**Attribute ID:** 0x23000d  
**Type:** STRING
- **Manufacturer**  
The manufacturer of the device.  
**Attribute ID:** 0x10032  
**Type:** STRING
- **Model Class**  
The model class of the item.  
**Attribute ID:** 0x11ee8  
**Type:** ENUM
- **Name**  
The device or model name.  
**Attribute ID:** 0x1006e

- Type:** STRING
- **Number of Occurrences**  
The number of times that the alarm occurred.  
**Attribute ID:** 0x11fc5  
**Type:** INTEGER
- **Severity**  
The severity of the alarm.  
**Attribute ID:** 0x11f56  
**Type:** ENUM
- **Show Symptoms**  
Indicates whether the alarm is the result of symptoms.  
**Attribute ID:** 0x12a07  
**Type:** BOOLEAN
- **Symptom Count**  
The number of symptoms for the alarm.  
**Attribute ID:** 0x12a06  
**Type:** INTEGER
- **System Name**  
The system name of the device.  
**Attribute ID:** 0x10b5b  
**Type:** STRING
- **Troubleshooter**  
The troubleshooter assigned to the alarm.  
**Attribute ID:** 0x11fc6  
**Type:** ENUM
- **Trouble Ticket ID**  
The Trouble Ticket ID assigned to the alarm.  
**Attribute ID:** 0x12022  
**Type:** STRING

## Manage Alarms

You can manage alarms on the **Alarms** view using the following options:

- **Acknowledge**  
Acknowledge the selected alarms. "Yes" appears in the **Acknowledge** column in the top **Alarms** pane.
- **Unacknowledge**  
Unacknowledge the selected alarms. Click to remove "Yes" from the **Acknowledged** column in the top **Alarms** pane.
- **Clear**  
Remove the selected alarms from the top **Alarms** pane. This option is available only for clearable alarms. To determine whether an alarm is clearable, show the **Clearable** column.
- **Troubleshooter**  
Manage the troubleshooter assignment for the selected alarms.
- **Poll**  
Poll the devices for the selected alarms. The poll is initiated from the SpectroSERVER.
- **Ping**  
Send an ICMP ping to the devices for the selected alarms. The ICMP ping is initiated from the SpectroSERVER.
- **Traceroute**



Determine the route (path) to the device of the selected alarm over a maximum of 30 hops. The round-trip time of packets that are received from each hop is also measured. Traceroute is initiated from the SpectroSERVER. This option is unavailable when you select multiple alarms.

- **On Demand**

Launch On-Demand reports for the selected alarms. Alarms and items that are not synchronized to NetOps Portal are filtered out. For more information, see [On-Demand Reports](#). This option is available when you have the Create On-Demand Report Templates role right and the Run On-Demand Report Templates role right.

- **Create Ticket**

Create a service desk ticket for a single alarm that is selected in the table. To launch the trouble ticket system from an alarms view, a trouble ticket management system at your organization must be set up with Spectrum.

For more information, see the "Alarms Tab Preferences" in the [DX NetOps Spectrum documentation](#).

This option is available when you have either the Allow Alarm Modification Actions role right or the Allow Alarm Triage Actions role right.

- **Manage Life Cycle**

Change the life cycle state of a device to "Active" or "Maintenance". This option is available when you have the Administer Life Cycle role right. Select the **Synchronize device life cycle state from Spectrum** checkbox for the Spectrum data source.

For more information, see [Manage Device Life Cycles](#).

## **Configure an Alarms View**

To display the **Alarms** view, go to the **Alarm Console** dashboard or add the **Alarms** view that is available from the **Alarms and Events** section. The **Alarms** view is also available from device, router, server, switch, and interface context pages. For example, the context pages for Spectrum devices include an **Alarms** tab.

For more information about tabs on context pages, see [Context Pages](#).

Configure the following **Alarms** view settings:

- **Max Alarms to Retrieve**

Specify the maximum number of alarms to retrieve from Spectrum starting with the most recent.

**Default:** 20,000

**Maximum:** 20,000

- **Ping Count**

Specify how often to send an ICMP ping to each device.

**Default:** 3

**Values:** 1 through 5

- **Grid Height**

Specify the height of the **Alarms** view grid.

**Default:** 300

- **Sort First**

Select a column on which to sort first.

**Default:** Date/Time

- **Sort Second**

Select a column on which to sort second.

**Default:** None

- **Sort Third**

Select a column on which to sort third.

**Default:** None

- **Sort Direction**

Select whether to sort the selected column in ascending or descending order.

**Default:** Descending

- **Customize Panels**

Select the panes to enable or disable.

- **Attributes to Show**

Add or remove the attributes to display in the **Alarm Details** pane.

- **Filter Selection**

Select whether to pin an alarm filter to the view. If enabled, select the alarm filter to pin to the view.

**Default:** Disabled

## Browser Views

You can show the content of a web page or application directly in NetOps Portal and have aspects of your web browser view adjust dynamically by configuring the web browser views.

### Configure a Web Browser View

Follow these steps:

1. On the **Add Dashboard** or **Edit Dashboard** page, edit the browser view by clicking the **Edit** (pencil) icon. The **Browser View** dialog opens.

2. Complete the following fields:

- **Title**

Defines the title for the view.

**Default:** Browser View

- **URL**

Specifies the URL for the view.

**IMPORTANT**

Enture a URL for a web page that supports being an embedded iFrame.

3. (Optional) To have aspects of your browser view adjust dynamically, to add context-sensitive NetOps Portal information to the URL, such as time range or group, and to provide specific information to OpenAPI applications, complete the following fields:
  - a. In the **URL Parameter** field, select one of the following URL parameters, and then click **Append to URL**:  
The parameters appear in the drop-down with two columns. The **Property** column shows the parameter. The **Value** column shows the current values for the logged in user. The values in the **Value** column are based on the current context and time on NetOps Portal. They resolve in the following manner:
    - **{Culture}**  
Resolves to the culture of the logged in user.  
**Example:** en-US
    - **{ItemDesc}**  
The description of the item or group of the current context. Resolves to the context that is specified at the top of the page or within the Dashboard Builder.  
**Example:** Includes every group and item ty...
    - **{ItemId}**  
The NetOps Portal item ID of the current item or group.  
**Example:** 1
    - **{ItemIdDA}**  
The data aggregator item ID of the current item or group.  
Appears only when you have a data aggregator data source, and resolves to the ID of the interface for the selected context page.
    - **{ItemName}**  
The name of the current item or group.  
**Example:** All Groups
    - **{ItemSubType}**

The subtype of the current item or group.

**Example:** tenant

- **{ItemSubTypeName}**  
The subtype name of the current item or group. The subtype name is usually the same as the subtype, but not always.  
**Example:** tenant
- **{ItemType}**  
The type of the current item or group.
- **{ItemTypeLabel}**  
The singular form of the item type label.
- **{ItemTypeLabels}**  
The plural form of the item type label.
- **{ItemTypeName}**  
The item type name.
- **{Locale}**  
The locale of the logged in user.  
The locale is similar to culture, but has a slightly different format.
- **{ModelHandleSpectrum}**  
The model handle from DX NetOps Spectrum (Spectrum) for the device or interface.  
This parameter is unavailable for groups. This parameter appears only when you have registered a Spectrum data source.
- **{PageSize}**  
The number of items that are specified to display in the view.  
**Example:** 10
- **{Resolution}**  
The data resolution in seconds.  
**Example:** 300
- **{ResolutionLabel}**  
The data resolution in a descriptive text.  
**Example:** 5-Minute Resolution
- **{ServerName}**  
The NetOps Portal server name or IP address.
- **{ServerNameDA}**  
The data aggregator server name or IP address.  
This parameter appears only when you have a data aggregator data source.
- **{ServerPort}**  
The NetOps Portal server port.  
**Example:** 8181
- **{ServerPortDA}**  
The data aggregator server port. This parameter appears only when you have a data aggregator data source.  
**Example:** 8581
- **{TimeEndUTC}**  
The number of seconds of the *end* of the time range (based on UNIX time) that is selected in the time picker for the context page. This parameter resolves to the end time that is selected in the time picker.  
**Example:** 1687294311
- **{TimeEndUTCExpanded}**  
The expanded form of the end of the time range.
- **{TimeSpan}**  
The description of the time range that is specified at the top of the page.
- **{TimeStartUTC}**

The number of seconds of the *start* of the time range (based on UNIX time) that is selected in the time picker for the context page. This parameter resolves to the start time that is selected in the time picker.

**Example:** 1687290711

- **{TimeStartUTCExpanded}**

The expanded form of the start of the time range.

**Example:** 2023-06-20\_2051

- **{UserEmailAddress}**

The email address of the logged in user (if specified).

- **{UserId}**

The NetOps Portal ID of the current user.

- **{UserName}**

The current user name.

- **{UserRoleId}**

The NetOps Portal role ID of the current user.

- **{UserRoleName}**

The role name of the current user.

**Example:** Administrator

- **{UserSsoToken}**

The single sign-on token for the current user.

You user use this token to log in to other data sources that honor single sign-on (for example, DX NetOps Network Flow Analysis (NFA)).

- **{UserTimeZone}**

The time zone of the current user.

**Example:** GMT

The URL parameter is added to the end of the URL in the **URL** field.

- To help you determine the meaning of each parameter in the URL, in the **URL** field, add an identifier *before* each parameter.

Enter a question mark ( ? ) *before* the first parameter. Enter an ampersand ( & ) *before* each subsequent parameter.

**Example:**

The following example displays the content of a web page or application directly in NetOps Portal at the following URL, appended with the NetOps Portal item ID of the current item or group, the start time that is selected in the time picker, and the end time that is selected in the time picker:

```
http://<PC_host>:<port>/pc/apps/user/<appsubdirectory>/<MyPageFile.html>?
id={ItemId}&startTime={TimeStartUTC}&endTime={TimeEndUTC}
```

- Complete the following fields, and then click **Save**:

- **Height**

Defines the height for the view, in pixels.

- **Apply Changes**

To determine which users see view changes, [Set the scope for the view settings](#).

The web page or application that the URL references appears in the web browser view.

## Bar Chart Views

Bar charts, or horizontal bar charts, show a comparative visualization of data. These charts are useful to compare items in a small group when the expected value of the metric is similar for all items. You can use this visualization to identify which item differs from the others. Bar charts can be useful for capacity planning.

Business hours filtering can apply to the data in bar chart views. The applied business hours filter appears in the subtitle.

For more information, see [Configure Business Hours Filtering](#).

For more information about how to configure a view, see [Customize Views](#).

In this article:

- [Side-by-Side Bar Charts](#)
- [Stacked Bar Charts](#)
- [Configure a Horizontal Bar Chart View](#)
- [Bar Chart Tables](#)


## Side-by-Side Bar Charts

Side-by-side bar charts compare pairs of related metrics for interfaces.

For the out-of-the-box views, select the **Metric Value** to display on the view: **Percentage**, **Rate**, or **Volume**.

The following example shows the important elements of the **Side-by-Side Bar Chart Interface Utilization**:

**Figure 45: Side-by-Side Bar Chart Elements**




## Stacked Bar Charts

Stacked bar charts show a horizontal bar with segments representing each selected metric.

The following example shows the important elements of the **IM Horizontal Bar Chart (Top)** with the stacked-chart type:

**Figure 46: Stacked Bar Chart Elements**



## Configure a Horizontal Bar Chart View

The **IM Horizontal Bar Chart** is a custom view. From this view, you can select the metrics to display. The chart can display multiple metrics for each device or component in the view context.

### Follow these steps:

1. On the **Add Dashboard** or **Edit Dashboard** page, edit the browser view by clicking the **Edit** (pencil) icon. The chart dialog opens.
2. Complete the following fields:
  - **Title**  
Defines the title for the view.  
**Default:** Browser View
  - **Metric Family**  
Defines the metric family.
  - **Metric Fields**  
Defines the metrics to show in the chart. Each selected metric appears as different colored bar in the chart.  

**TIP**  
 The custom bar chart is best used to compare a few similar metrics for a small group of items. If you select too many items or too many metrics, the view is difficult to read.  
  
 Select metrics with consistent unit types (bits, percentage, and so on) for more valuable comparisons.  
  
 Also, select metrics with a consistent rollup or aggregation strategy. For example, compare two sets of averages for two different metrics. This comparison is more valuable than comparing the total for one metric and the average for another metric.
  - **Metric Sort**  
Defines the sort order for the view. The metric that you select from this list represents the first bar on the bar chart.
  - **Sort Direction**  
Defines the direction to sort the metrics.  
**Options:** Descending, Ascending  
**Default:** Descending
  - **Metric Calculate Level**  
Determines the level of aggregation each row in the view represents.  
**Options:** by Device, by Component  
**Default:** by Component
  - **Chart Type**  
Defines the chart type.  
**Options:**
    - **Stacked Chart:** For each device or item, this chart type shows a horizontal bar with segments representing each selected metric.
    - **Side-by-Side Chart:** For each device or item, this chart type shows a horizontal bar representing each selected metric.
  - **Max Rows**  
Defines the maximum number of rows to display for the bar chart view.  
**Options:** 20, 15, 10, 5  
**Default:** 10
  - **Metric Filtering**  
For more information, see [Customize Views](#).
  - **Baseline Metrics**  
For more information, see [Customize Views](#).

3. To define the fields in the **Customize Time Range**, **Apply Business Hours**, and **Context Settings** sections, see [Customize Views](#).
4. Click **Save**.

The horizontal bar chart view is configured.

### **Bar Chart Tables**

Some out-of-the box table views show a column with a bar graph. These tables display a bar chart column for metrics with percentage values.

For more information about these views, see [Table Views](#).

## **Calendar Heat Chart Views**

Calendar heat chart views provide a month-long overview with hourly intervals that are color coded based on custom thresholds. The calendar shows one month with a colored block to show the status for each hour.

### **NOTE**

If the hourly data retention is less than 1 month, the calendar heat chart still shows hourly data.

As a network engineer, you can identify patterns in percentage-based metrics such as utilization, latency, and loss metrics, using calendar heat chart views. Identifying patterns can be a critical step in locating the source of performance issues that might appear to be intermittent. Usage patterns also aid in capacity planning.

You can limit the data that appears in the view by applying one of the following pattern-matching filters:

- **Busy Hour**  
View the hour each day with the highest value for the selected metric.
- **Business Week Pattern**  
View patterns where the same hours have similar values on three or more days in the business week.
- **Calendar Week Pattern**  
View patterns where the same hours have similar values on four or more days in the calendar week.
- **Repeating Hours by Business Day**  
View patterns where the same hour on the same day of the business week has a similar value for three or more times in the month.
- **Repeating Hours by Calendar Day**  
View patterns where the same hour on the same day of the calendar week has a similar value for three or more times in the month.

Maintenance indicators apply to all the devices and components in a site group. When you have selected the associated site group in the context, maintenance indicators appear as shaded cells in the calendar heat chart.

For more information, see [Schedule Maintenance Indicators](#).

Applying a business hours filter to a calendar heat chart view displays the same data in the view, but the data with the periods *outside* of the applied business hours are shaded.

For more information:

- About business hours filtering, including how to define business hours definitions, see [Configure Business Hours Filtering](#).
- About how to apply a business hours filter to a dashboard, see [Dashboards](#).
- About how to apply a business hours filter to a context page, see [Context Pages](#).

The following example shows the important elements of a calendar heat chart:





**TIP**

You can configure calendar heat chart views to reflect reverse value severity. For metrics where low values are bad and high values are good, set **Green** highest and **Red** lowest. To omit a severity level, set the threshold value to zero. For example, setting **Orange** to zero removes the major severity when profiling metric thresholds.

## Card Views

Card views display cumulative counts of groups, devices, or components for the threshold ranges of each specified metric.

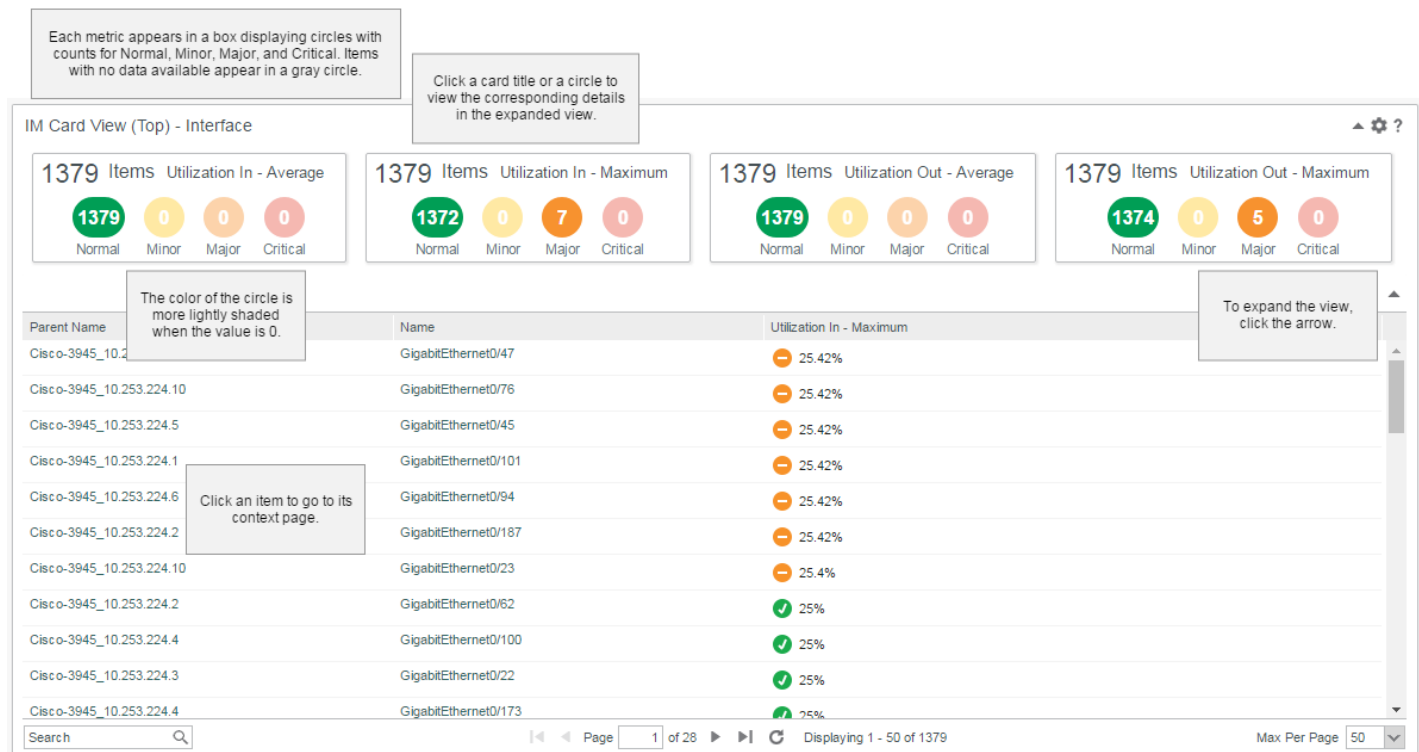
Card view counts appear within green, yellow, orange, and red circles (indicators). The color of each circle is based on a set of user-defined thresholds. You can specify separate thresholds for each metric.

Card views include standard view configuration options, such as time filters and context settings.

For standard view configuration details, see [Customize Views](#).

The following image shows the important elements of a card view:

**Figure 48: Card\_Views**



### Configure a Card View

#### Follow these steps:

1. On the **Add Dashboard** or **Edit Dashboard** page, edit the browser view by clicking the **Edit** (pencil) icon. The card view dialog opens.
2. Complete the following fields:
  - **Title**  
Defines the title for the card view.

**Default:** Card View

– **Items Limit**

Specifies the limit for the number of groups, devices, or components (counter) shown in the indicator for each card. If the number of groups, devices, or components for a threshold, such as utilization, is equal to or greater than the set limit, then at the top of the view, a warning message displays with a list of cards with truncated results.

**Default:** 5000

**Example:**

If you set the **Items Limit** to 7, and there are more than eight groups, devices, or components in the green threshold, the green indicator shows truncated results (7) and at the top of the view, with the following message:

Showing first 7 results for Utilization In - Maximum, Utilization Out - Average

– **Metric Family**

Defines the metric family.

– **Metric Fields**

Defines the metrics to show in the card view, such as current utilization. Each selected metric appears as a box in the view.

– **Metric Calculate Level**

Determines the level of aggregation for the cumulative counts.

**Options:** by Device, by Component, by Groups

**Default:** by Component

– **Resolution**

Determines the appropriate resolution.

**Options:**

- **Use default resolution:** Determines the appropriate resolution using the time range.
- **1 Minute:** The resolution is set to one minute.
- **5 Minute:** The resolution is set to five minutes.
- **10 Minute:** The resolution is set to ten minutes.
- **15 Minute:** The resolution is set to 15 minutes.
- **30 Minute:** The resolution is set to 30 minutes.
- **1 Hour:** The resolution is set to one hour.
- **1 Day:** The resolution is set to one day.

**Default:** Use default resolution

– **Metric Filtering**

For more information, see [Customize Views](#).

– **Baseline Metrics**

For more information, see [Customize Views](#).

3. For each metric, specify the threshold values for the status indicators (**Minor Status**, **Major Status**, and **Critical Status**). If the value for the metric is at or above the specified value, the icon indicates the status by color.

**Options:**

- **Green:** Indicates the state is normal.
- **yellow:** Indicates the state is minor (Minor Status)
- **orange:** Indicates the state is major (Major Status)
- **red:** Indicates the state is critical (Critical Status)

**TIP**

You can remove the status indicators by setting the threshold values for all status levels to zero.

You can specify thresholds with up to two decimal places with the lowest value being 0.01. Values with more than two decimal places are rounded.

**TIP**

For metrics where low values are bad and high values are good, set **Minor Status** highest and **Critical Status** lowest. To omit a severity level, set the threshold value to zero.

4. To define the fields in the **Customize Time Range**, **Apply Business Hours**, and **Context Settings** sections, see [Customize Views](#).
5. Click **Save**.

The card view is configured.

## Dynamic Trend Views

Dynamic trend views combine data from multiple managed items, or groups of items, in a single view.

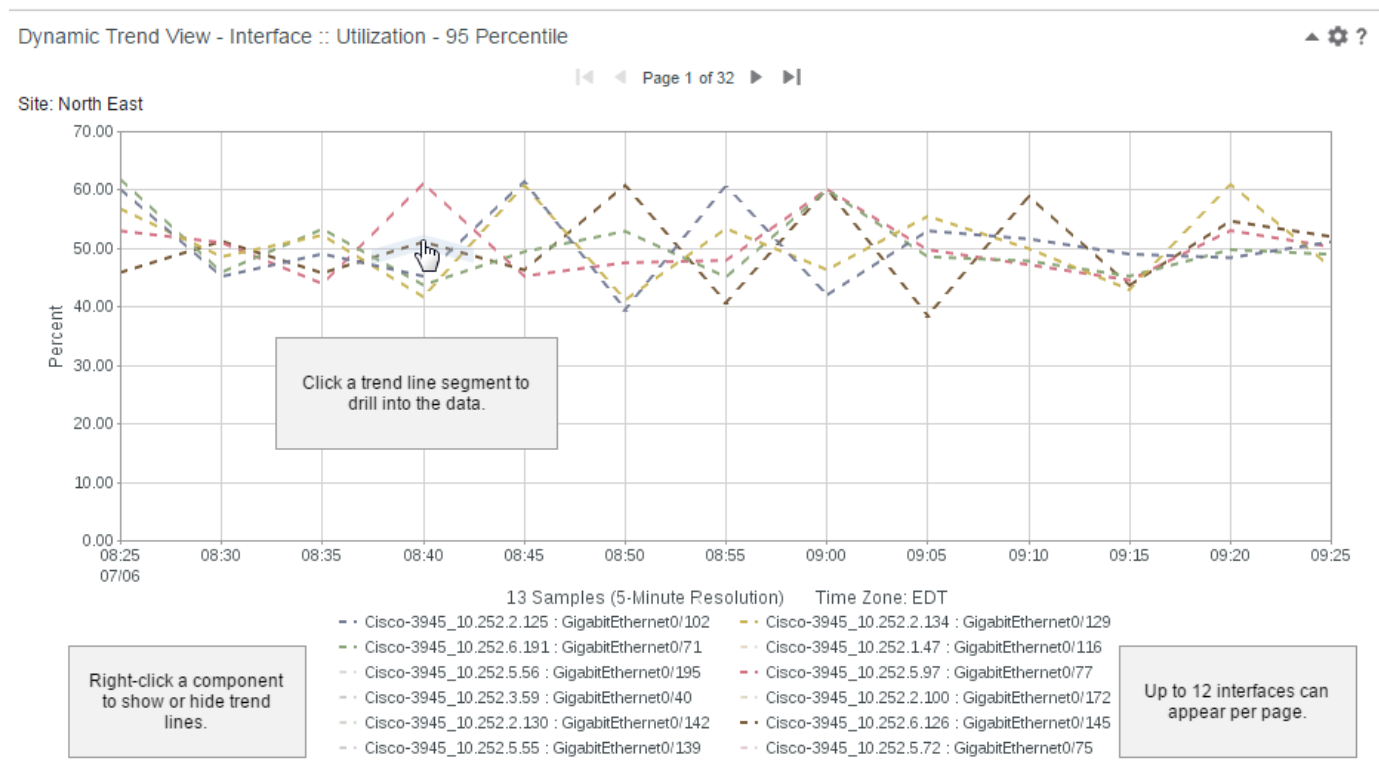
You can compare the data from hundreds of managed items, or groups of managed items, in one or more charts in dynamic trend views. You can select display units and determine how data is plotted using the view settings.

**NOTE**


Dynamic trend views support only the data aggregator data source.

The following image shows the important elements of a dynamic trend view:

**Figure 49: DynamicTrendView**



The following image shows the important elements of a dynamic trend MultiView:

**Figure 50: DynamicTrendMultiView**

The trend charts in dynamic trend views are for selected metrics. These views are "dynamic" because they display the following:

- Individual charts for selected metrics for each selected item
- A chart for all selected items, for a single metric
- Trend lines for selected metrics that reflect aggregated data from all the selected items

#### NOTE

Percentile metrics appear as dashed lines.

Use dynamic views during advanced troubleshooting. Use these views to build a complex comparison of data from multiple groups, or a temporary grouping of items. Quickly select items from across a wide range of possibilities, without creating a group in advance.

For example, you can use dynamic views to analyze data from any of the following items:

- Multiple interfaces in a group
- The interfaces in your inventory with the worst performance metrics
- The interfaces on a single device

Trend data in dynamic views represents averages of metric values across all items that are included.

#### TIP

Dynamic trend views require more page space than other types of views. To add a dynamic trend view to a dashboard, add it to a single-column area in the layout. Do not add a dynamic trend view to a small section on a dashboard layout.

Dynamic trend views also require more space for the PDF output.

Maintenance indicators apply to all the devices and components in a site group. When the associated site group is selected in the context, maintenance indicators appear as shading in dynamic trend views.

For more information, see [Schedule Maintenance Indicators](#).

### Configure a Dynamic Trend View

To configure a dynamic trend chart, customize or edit a dynamic trend view. Dynamic trend views do not contain data until you edit them to select settings.

For standard view configuration details, see [Customize Views](#).

You can add dynamic trend views to dashboards or context pages.

For more information:

- About how to add custom views to dashboards, see [Manage Dashboards](#).
- About how to add custom views to context pages, see [Manage Context Pages](#).

#### Follow these steps:

1. From the settings dialog for the dynamic trend view, in the **Dynamic Trend View - Settings** section, complete the following fields:
  - **View Type**  
Defines the type of view.  
**Options:** MultiView (Draws individual chart for each item), MultiTrend (Draws all items on one chart - Top N items as a time), or Composite Trend (Draws single trend line for all items)  
**Default:** MultiTrend (Draws all items on one chart - Top N items as a time)
  - **Metric Family**  
(If you are configuring a **MultiView (Draws individual chart for each item)** or a **MultiView (Draws all items on one chart - Top N items at a time)** dynamic trend view) Defines the metric family for the view.  
For more information about the options for customizing MultiView trend views, see [Trend Views](#).
  - **Resolution**  
Defines the resolution frequency.  
**Default:** Use default resolution (5 minute)
  - **Sort Direction**  
Determines which interfaces are reflected in the first pages of trend charts.  
**Options:**
    - **Descending:** highest values first
    - **Ascending:** lowest values first**Default:** Descending
  - **Chart Legend Type**

Defines how NetOps Portal displays the information about the trend chart below the chart, either in a table or chart legend.

**Options:**

- **Legend Type Chart:** The trend chart information is displayed in a chart legend.
- **Legend Type Table:** The trend chart information is displayed in a table legend.

**Default:** Legend Type Chart

– **Chart Type**

Defines the type of chart, either lines or area.

**Options:**

- **Trend Chart**  
Show trend lines for each metric.
- **Stacked Chart**  
Show lines for each metric stacked with area fills under each.

**Default:** Trend Chart

– **Chart Display Order**

(If you are configuring a **MultiView (Draws individual chart for each item)** or a **MultiView (Draws all items on one chart - Top N items at a time)** dynamic trend view) Defines the order in which to display interfaces.

**Options:**

- **Display by Metric:** Sort the interface charts by metric value, with the highest (most severe) values shown first, top to bottom.
- **Display by Name:** Sort the interface charts in alphanumeric order, by interface name.

**Default:** Display by Metric

– **Metric Calculate Level**

(If you are configuring a **MultiView (Draws individual chart for each item)** or a **MultiView (Draws all items on one chart - Top N items at a time)** dynamic trend view) Defines whether to show metric values for devices or for the device components that are plotted on the charts in the view.

**Options:**

- **by Component:** Show metric values only for the device components.
- **by Device:** Show metric values only for devices.

**Default:** by Component

– **Standardized Axis**

Defines the range of the Y-axis of each chart in the view.

**Options:**

- **Calculated:** The first Y-axis adjusts dynamically, based on the range of metric values that are included. The calculated setting is applied to all charts in the view.
- **Fixed at 0 to 100:** The Y-axis maintains a static range, 0 through 100.
- **Scale per Chart:** The Y-axis adjusts dynamically, based on the range of metric values of the chart. The Y-axis scaling is on a per-chart basis.

**Default:** Calculated

– **Number of Charts on Page**

(If you are configuring a **MultiView (Draws individual chart for each item)** or a **MultiView (Draws all items on one chart - Top N items at a time)** dynamic trend view) Defines the number of charts for each page.

**Default:** 15

**Limit:** 36

– **Maximum Number of Charts**

(If you are configuring a **MultiView (Draws individual chart for each item)** or a **MultiView (Draws all items on one chart - Top N items at a time)** dynamic trend view) Defines the maximum number of charts.

**Default:** 60

**Limit:** 1200

– **Metric Filtering**

Specifies whether to filter metrics.

**Options:**

- **Enabled:** Metric are filtered.
- **Disabled:** Metric are not filtered.

**Default:** Enabled

– **Baseline Metrics**

Determines whether NetOps Portal changes utilization trends using the baseline trend line. The baseline data that NetOps Portal plots in many views shows statistical deviations from normal performance for a given statistic.

**Options:** Selected (enabled), Cleared (disabled)

**Default:** Cleared (disabled)

For more information, see [Baseline Calculations](#).

2. (23.3.4 and higher) In the **Set Trend Goal/Threshold Line** section, complete the following fields:

– **Goal/Threshold Line**

Specifies whether to show or hide a goal/threshold line trend line in the view. Goal line trend lines serve as visual indicators of fixed metric values.

**Options:**

- **Disabled:** Do not show a goal or threshold line trend line in the view.
- **Enabled:** Show a goal or threshold line trend line in the view.

**Default:** Disabled

– **Goal Line**

(If you have decided to show a goal/threshold line trend line) Specifies the type or trend line to show in the view, and where NetOps Portal shows it.

**Options:**

- **Hide:** Do not show a goal/threshold line trend line in the view.
- **Show Goal Line:** Show a goal line trend line in the view and in the chart legend (golden color).
- **Show Threshold Line (Above):** Show a threshold line trend line in the view with shading above the line with pink offset of critical color. Show the line in trend charts and in the chart legend.
- **Show Threshold Line (Below):** Show a threshold line trend line in the view with shading below the line with pink offset of critical color. Show the line in trend charts and in the chart legend.
- **Show Threshold Line (Above and Below):** Show a threshold line trend line in the view with shading above the line with pink offset of critical color and shading below the line with green offset of normal color. Show the line in trend charts and in the chart legend.

**Default:** Hide

3. (23.3.3 and lower) Complete the following fields:

– **Goal Line**

Determines if NetOps Portal shows a line that plots the value that you selected as the goal for the metric in the view.

**Options:**

- **Hide:** NetOps Portal does not show a goal line.
- **Show:** NetOps Portal shows a goal line.

**Default:** Hide

4. In the **Customize Time Range** section, [filter data based on specific time periods by selecting a custom time range for the view](#).

5. In the **Apply Business Hours** section, [assign a business hours filter and a time zone to the view](#).

Applying a business hours filter to a dynamic trend view displays the same data in the view, but the data with the periods *outside* of the applied business hours are shaded.

6. (If you are adding this view to a dashboard) In the **Items or Groups to Include** section, complete the following field:
  - **Context**  
Defines the context of the view.
  - Options:**
    - **Dynamic:** Indicates that the context of the view changes with the context of the dashboard. With this option, the page context is honored for site groups that are associated with a business hours definition.
    - **Fixed:** Indicates that the view uses a specified group, device, or component as a context for the data. With this option, the page context is honored only if the view is configured with a single site group that is associated with a business hours definition.
  - Default:** Fixed
7. If the view supports business hours, select whether to **Compress Non Business Hours**.  
**Default:** Cleared (non-business hours are not compressed in this view)  
For more information about compressed non-business hours in trend views, see [Trend Views](#).
8. Click **Save**.

The dynamic trend view is configured.

## Dynamic Table Views

Dynamic table views combine data from multiple managed items, or groups of items, in a single tabular view.

Compare data from hundreds of managed items, or groups of managed items, in a tabular format using dynamic table views. These views are often used during advanced troubleshooting to build a complex comparison of data from multiple groups, devices, or components, such as interfaces, in a time-series tabular format. You can also configure a temporary grouping of items. You can quickly select items from across a wide range of possibilities, without creating a group in advance.

For example, you can use dynamic table views to analyze data from any of the following items:

- Multiple interfaces in a group
- The interfaces in your inventory with the worst performance metrics
- The interfaces on a single device

The tabular data in these views represent metric values across the items that are included as a time-series report, sorted by date and time.

You can add dynamic table views to dashboards or context pages. You can select display units and determine how data is plotted using the view settings.

Dynamic table views are termed "dynamic" because they have the following options:

- Report aggregated metric values.
- Report metric values aggregated by group/device/component.
- Report metrics aggregated values by time intervals.
- Report metrics values aggregated by group/device/component by time intervals.
- Report metrics values aggregated device by time intervals for week as hourly resolution.
- Report metrics values aggregated by device by time intervals for month as daily resolution.

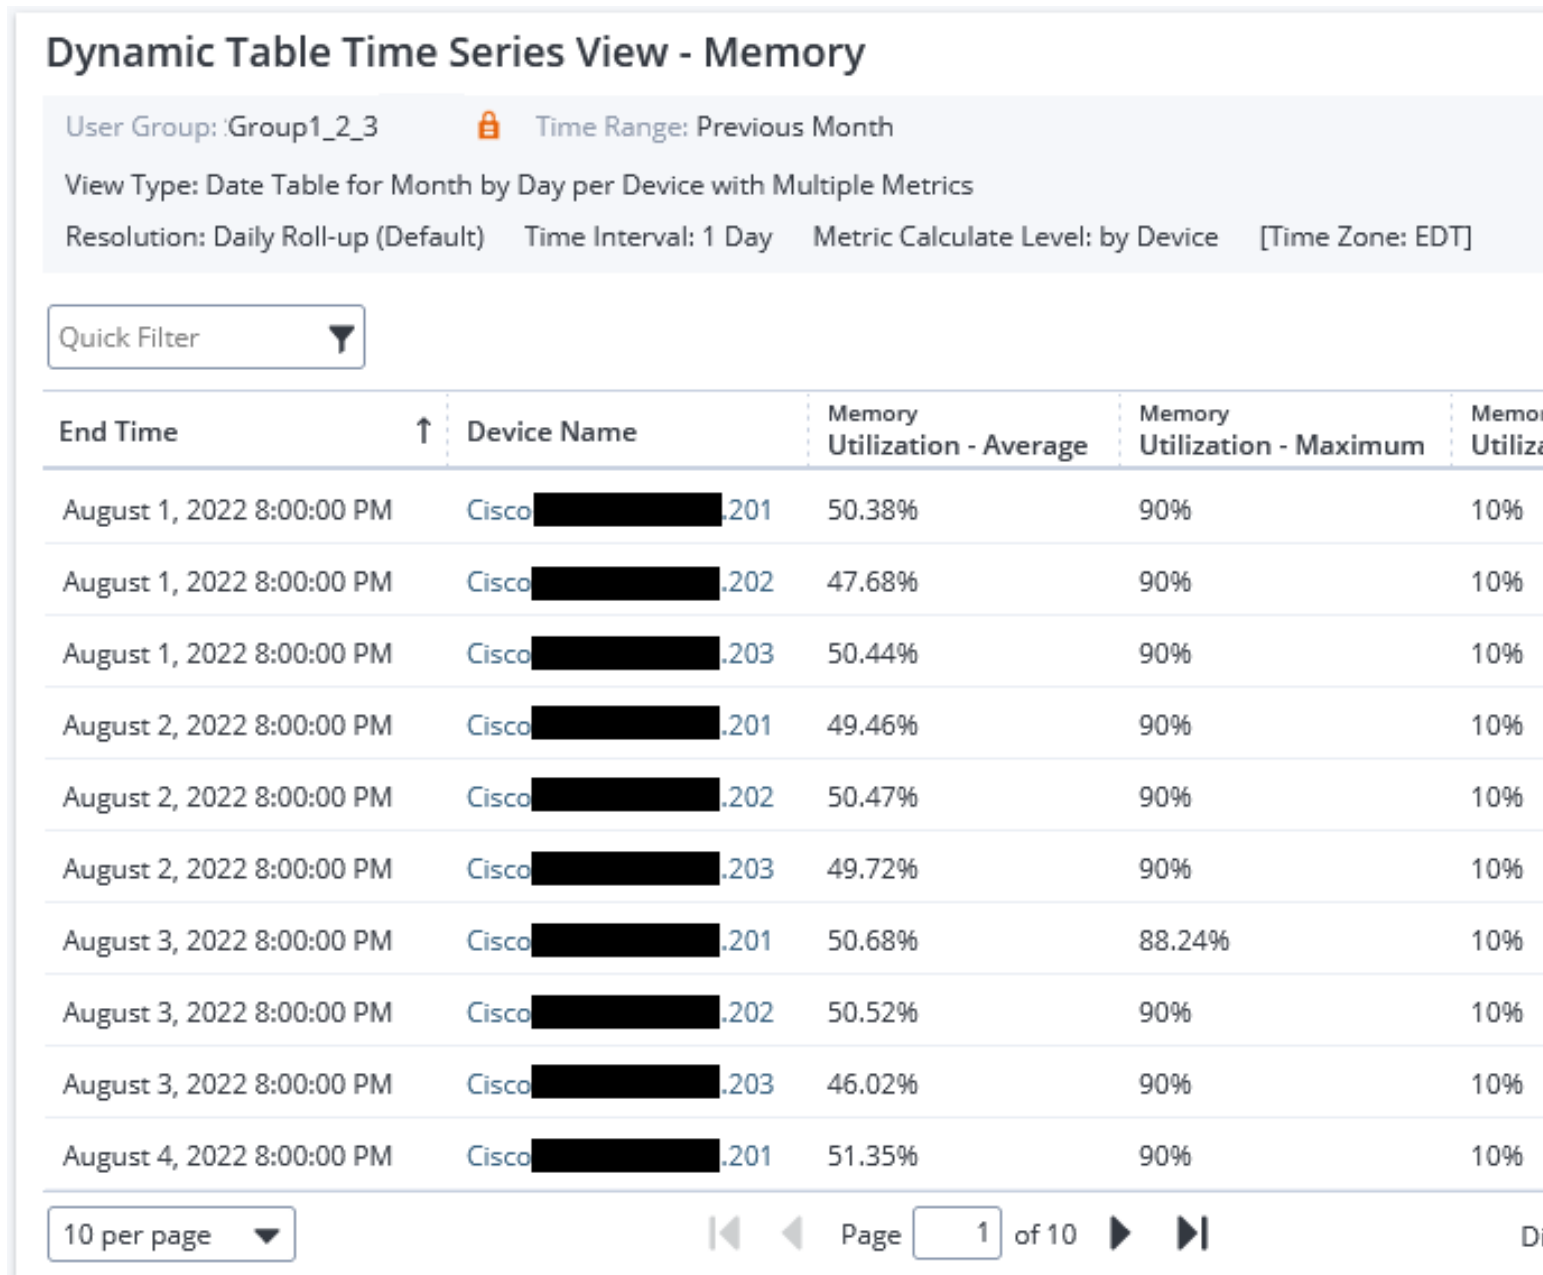
### NOTE

This view supports core, baseline, percentile, and projection metrics.

The following image shows an example of a dynamic table view:



Figure 51: Dynamic Table View Example



#### Configure a Dynamic Table View

Show tabular data for a selected metric by configuring a dynamic table view. These views contain data only *after* you edit them to select settings.

For standard view configuration details, see [Customize Views](#).

#### Follow these steps:

1. From the settings dialog for the dynamic table view, in the **Dynamic Table View - Settings** section, complete the following fields:
  - **View Type**

Defines the type of view.

**Options:**

- **Table with Multiple Metrics**  
The table aggregates metric values.
- **Table per Item with Multiple Metrics**  
The table aggregates metric values by group/device/component.
- **Date Table with Multiple Metrics**  
The table aggregates metric values by time intervals.
- **Date Table per Item with Multiple Metrics**  
The table aggregates group/device/component metric values by time intervals.
- **Date Table for Week by Hour per Device with Multiple Metrics**  
The table aggregates device metric values by time interval for week as hourly resolution.
- **Date Table for Month by Day per Device with Multiple Metrics**  
The table aggregates device metric values by time interval for month as daily resolution.

**Default:** Date Table per Item with Multiple Metrics

– **Metric Family**

(If you are configuring a custom view) Defines the metric family for the view.

– **Metric Value**

(If you are configuring a custom view) Defines the metric family for the view.

– **Resolution**

Defines the resolution frequency.

**Default:** Use default resolution (5 minute)

– **Include Inventory Columns**

Specifies whether to include inventory columns in the dynamic table view.

**Default:** Selected

– **Result Limit (Max Rows)**

Defines the maximum number of rows in the dynamic table view.

**Limit:** 5000

**Default:** 60

– **Metric Filtering**

Specifies whether to filter metrics.

**Default:** Selected

– **Baseline Metrics**

Specifies whether utilization tables are changing using the baseline table line. The baseline data that is plotted in many views shows statistical deviations from normal performance for a given statistic.

**Default:** Cleared

For more information, see [Baseline Calculations](#).

2. In the **Customize Time Range** section, [filter data based on specific time periods by selecting a custom time range for the view](#).
3. In the **Apply Business Hours** section, [assign a business hours filter and a time zone to the view](#).  
Assigning a business hours filter to a dynamic table view displays the same data in the view, but the data with the periods *outside* or *inside* of the applied business hours are depicted by the **Business Hour Status** column in the table.
4. In the **Items or Groups to Include** section, complete the following field:
  - **Context**  
Defines the context of the view.

**Options:**

- **Dynamic:** Indicates that the context of the view changes with the context of the dashboard. With this option, the page context is honored for site groups that are associated with a business hours definition.
- **Fixed:** Indicates that the view uses a specified group, device, or component as a context for the data. With this option, the page context is honored only if the view is configured with a single site group that is associated with a business hours definition.

**Default:** Fixed

5. If the view supports business hours, select whether to **Compress Non Business Hours**. Compressed non-business hours in table views do not have rows that contain time periods outside the applied business hours.

**Default:** Cleared

For more information about compressed non-business hours in table views, see [Trend Views](#).

The dynamic table view is configured.

## Gauge Views

Gauge views highlight where the value of a metric differs from the expected value. The out of the box gauge views display percentage metrics. However, custom gauge views do allow for configuring metric values that are not of type percentage.

You can apply business hours filtering to the data in gauge views. The applied business hours filter appears in the subtitle. For more information about how to apply business hours filters to view, see [Customize Views](#).

### Gauge View Elements


The gauge shows the minimum, maximum, and average values in the time range. The gauge needle shows the average value. The gray inner arc shows the minimum and maximum range of the metric value.

The outer arc uses the following colors to show customizable threshold ranges:

- **Green:**The normal status range
- **Yellow:**The moderate status range
- **Red:**The critical status range

The gauge chart legend shows the resolution, which specifies the granularity of the reporting data: as polled, hourly rollup, or daily rollup. For custom views, the minimum and maximize thresholds are based on this resolution. For out of the box views, the minimum and maximum thresholds are based on the as polled resolution.

The following example shows the important elements of a gauge view:

**Figure 52: Gauge View Elements****Configure Gauge Views**

All gauge views let you customize the threshold values for metric status. The IM Gauge Chart custom view provides more configuration options.

For more information about how to configure standard views, see [Customize Views](#).

**Follow these steps:**

1. Select the **Metric Needle** from the list of metrics for the selected metric family. This option determines the selected metric for the view.
2. Do one of the following tasks:
  - If you are configuring a custom view, select whether **Gauge End Points** are calculated or fixed. This option determines the range of the gauge.
    - If fixed, define the minimum and maximum values.
    - If calculated, the view determines the gauge end points based on the minimum and maximum values of the data from the appropriate item level.
  - If you are configuring an out of the box view, the gauge end points are fixed. Define the minimum and maximum values.
3. If you are configuring a custom view, select whether **Determine Threshold** is by percentage value or by numeric value:
  - **By Percentage Value**  
The status start values are a percentage of the range between the gauge endpoints.
  - **By Numeric Value**  
The status start values are fixed numeric values.
4. Define the thresholds for the gauge:

- **Moderate Status Start** The gauge shows yellow for this zone.
- **Critical Status Start**  
The gauge shows red for this zone.

**TIP**

For metrics where low values are bad and high values are good, such as availability, set Moderate Status Start higher than Critical Status Start.

- If you are configuring a custom view, select **Scaled** or **Unscaled** to specify whether the values in the view are scaled. Scaled values appear with larger units, for example, 1 KB. Unscaled values appear in the raw form for the metric, for example, 1000 bytes.

**Gauge/Table Views**

Gauge/table views highlight where the value of one metric differs from the expected value when compared to other metrics. The gauge shows the minimum, maximum, and average values for one item in the time range. The table displays the values for all items. You can select an item to show on the gauge.

For more information about these views, see [Table Views](#).

**Group Scorecard Trend Views**

Group scorecard trend views display performance metrics by subgroup, device, or component for the selected group. Group scorecard trend views provide line-of-business owners a group-level summary of how key metrics perform over time. Performance is based on a set of user-defined thresholds. These views incorporate red, orange, yellow, and green icons as visual indicators of performance levels.

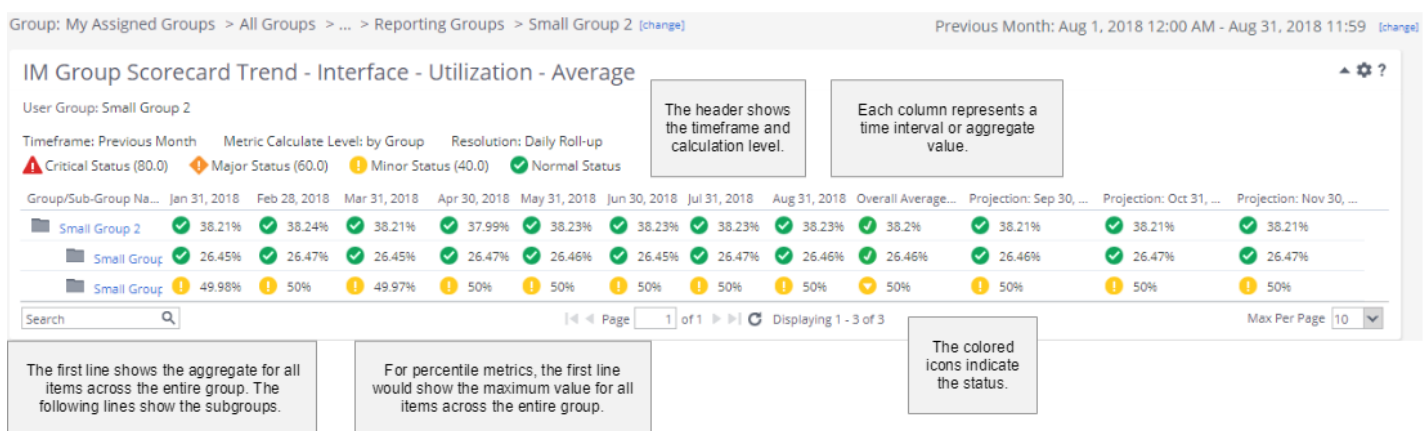
Group scorecard trend views show the average, minimum, and maximum utilization, or the selected metric at different time intervals for a group. Each interval length is equal to the time range selected for the dashboard. Out-of-the-box views show either the utilization average or the 95 percentile. The IM Group Scorecard Trend view shows any selected metric.

You can select counter metrics (for example, Bits Out - Total). However, group scorecard trend view thresholds are intended for gauge metrics (for example, Bits Out - Average Rate).

Group scorecard table views are also available, which let you view multiple metrics in the same scorecard. For more information, see [Group Scorecard Table Views](#).

The following images show the important elements of a scorecard view.

The Metric Calculate Level in the following image is set to Group:

**Figure 53: Group Scorecard Trend**

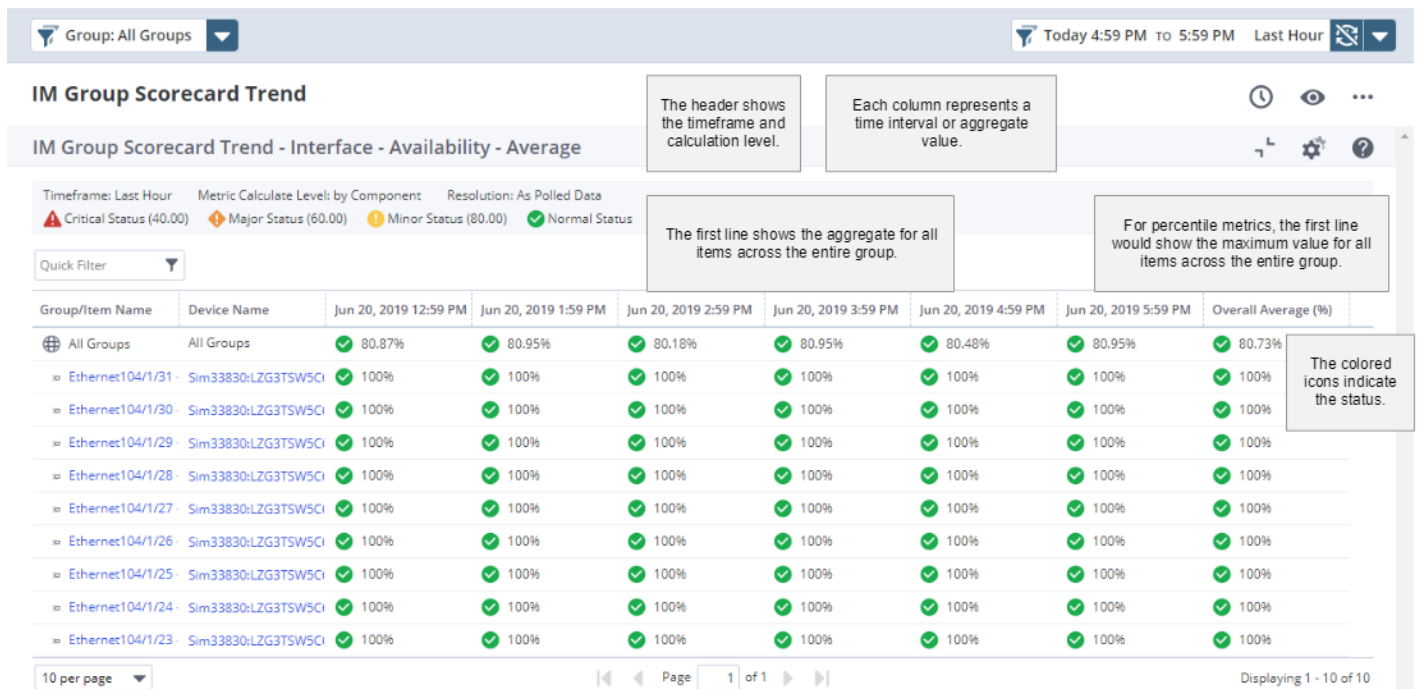
The Metric Calculate Level in the following image is set to Device:

Figure 54: Group Scorecard Trend Device



The Metric Calculate Level in the following image is set to Component:

Figure 55: Group Scorecard Trend Component



**NOTE**

If a time period renders nonnumeric value for the reporting metric, the scorecard view and header contains a gray icon. This icon indicates that the data for that timeframe includes an invalid data point.

Group scorecard trend views calculate and display projected values as follows:

- For monthly intervals, the first projection is the last day of the following month.
- For weekly intervals, the first projection is midnight on the following Sunday.
- For daily intervals, the first projection is the following midnight.
- For hourly intervals, the first projection is the top of the following hour.

In this article:

- [IM Group Scorecard Trend View](#)
- [Interface 95th Maximum Percentile Trend Scorecard View](#)
- [Configure Group Scorecard Trend Views](#)
- [Scorecard Projection](#)
- [HEADING TITLE](#)
- [HEADING TITLE](#)

**IM Group Scorecard Trend View**

The IM Group Scorecard Trend view is a fully customizable scorecard. This scorecard displays historical time intervals for any metric, and up to three projected values. The projected values are calculated from the trend of historical data in the time range of the view. The items in a group determine the metrics that are available in this view.

**Interface 95th Maximum Percentile Trend Scorecard View**

Most group scorecard trend views always attempt to use daily data unless your data retention policy prevents it. However, the Interface 95th Maximum Percentile Trend Scorecard view uses polled data for accuracy.

When you select larger time ranges (for example, Previous Month), you can easily surpass the configured data retention rate, which limits the number of columns available for comparison.

**Configure Group Scorecard Trend Views**

Group scorecard trend views include standard view configuration options, such as time filters and context settings. For standard view configuration details, see [Customize Views](#).

**TIP**

On dashboards, scorecards render best in single-column layouts. Scorecards with fewer columns also render well in two-column layouts. Do not use scorecards in three-column layouts.

The following configuration options apply only to scorecard views.

**Follow these steps:**

1. With the view open in edit mode, (Trend and Custom Scorecard Views Only) select the **Number of Time Intervals**. Each time interval is one column on the view.

**Default:** 6

**NOTE**

This setting directly relates to the **Time Range** setting. The time range for the view determines the length of each time interval. For example, if you select Last 30 days for the Time Range, each interval equals 30 days.

2. (Custom Scorecard View Only) Configure scorecard projection:
  - Select the **Number of Projections** to display.

Each projection is a column on the view.

- For **Projection Calculation Method**, select the method in which to calculate projections.

**Values:** Approximation or Detailed Data

For more information about these methods, see [the "Scorecard Projection" section](#).

3. Select the **Metric Calculate Level**.

This option determines what level of aggregation each row in the scorecard view represents: a subgroup, a device, or a component.

**Values:** by Group, by Device, by Component

**Default:** by Group

4. Specify the thresholds for the status indicators:

- **Critical Status**

**Default:** 80

- **Major Status**

**Default:** 60

- **Minor Status**

**Default:** 40

If the value for the metric is at or above the specified value, the status icon in the view indicates the status by color. Red indicates critical, orange indicates major, yellow indicates minor, and green indicates normal. To remove the status indicators, set the threshold values for all status levels to zero.

**TIP**

For metrics where low values are bad and high values are good, set **Minor Status** highest and **Critical Status** lowest.

5. Set the **Results Limit**.

This limit determines the number of items included in the data results shown on each page of the scorecard.

**Default:** 10

## **Scorecard Projection**

Scorecard projections use a customizable set of data points to predict future values for metrics. DX NetOps Performance Management calculates projected values when it renders the view. These values are based on the historical time frame of the view. You can add these values to the IM Group Scorecard Trend View.

**TIP**

Do not use scorecard projections for error metrics. Each error is a discrete event that is not affected by historical errors.

The scorecard view includes the following methods to calculate projections:

- **Approximation**

This method uses the average from each time frame in the view to calculate the projection values. DX NetOps Performance Management calculates a least squares regression on the averages, then uses the line equation to project future values. This calculation method is faster than the Detailed Data method.

- **Detailed Data**

This method uses the polled data for the entire time frame of the view. DX NetOps Performance Management calculates a least squares regression for the entire set of data points. This calculation is more statistically accurate than the **Approximation** method, and provides extra columns in the view.

**NOTE**

Detailed data scorecard projections are supported only for gauge metrics (for example, Bits Out - Average Rate). Detailed data scorecard projections are *not* supported for counter metrics (for example, Bits Out - Total). Projection values are calculated on the As Polled (rate) data to ensure precision.

The following columns are hidden by default:

- **Slope**



- Indicates the slope of the line equation.
- **Intercept**  
Indicates the intercept of the line equation.
- **Degrees**  
Degrees of freedom, which indicates the sample size.
- **Linear Fit**  
Indicates the confidence level of the projected values as related to the sample data.
- **Days to Threshold**  
Indicates the projected number of days before the specified critical threshold is reached.

## Group Scorecard Table Views

Group scorecard table views let you view multiple metrics in the same scorecard. Group scorecard table views display the metrics by subgroup, device, or component for the selected group. These views incorporate red, orange, yellow, and green icons as visual indicators of performance levels. The color of the status icons is based on a set of user-defined thresholds. You can specify separate thresholds for each metric.

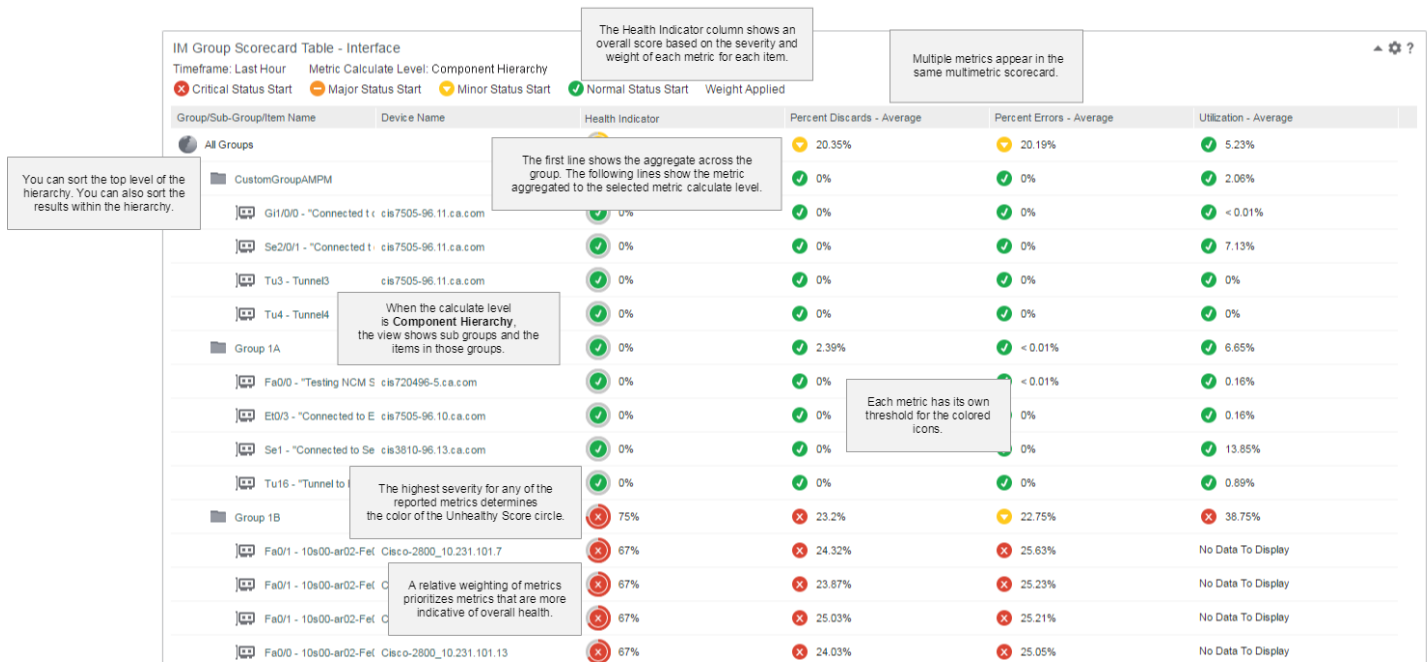
You can select counter metrics (for example, Bits Out - Total). However, group scorecard table view thresholds are intended for gauge metrics (for example, Bits Out - Average Rate).

Group scorecard trend views are also available, which provide line-of-business owners a group-level summary of how key metrics perform over time. For more information, see [Group Scorecard Trend Views](#).

The following images show the important elements of a group scorecard table view. The Health Indicator column is shown.

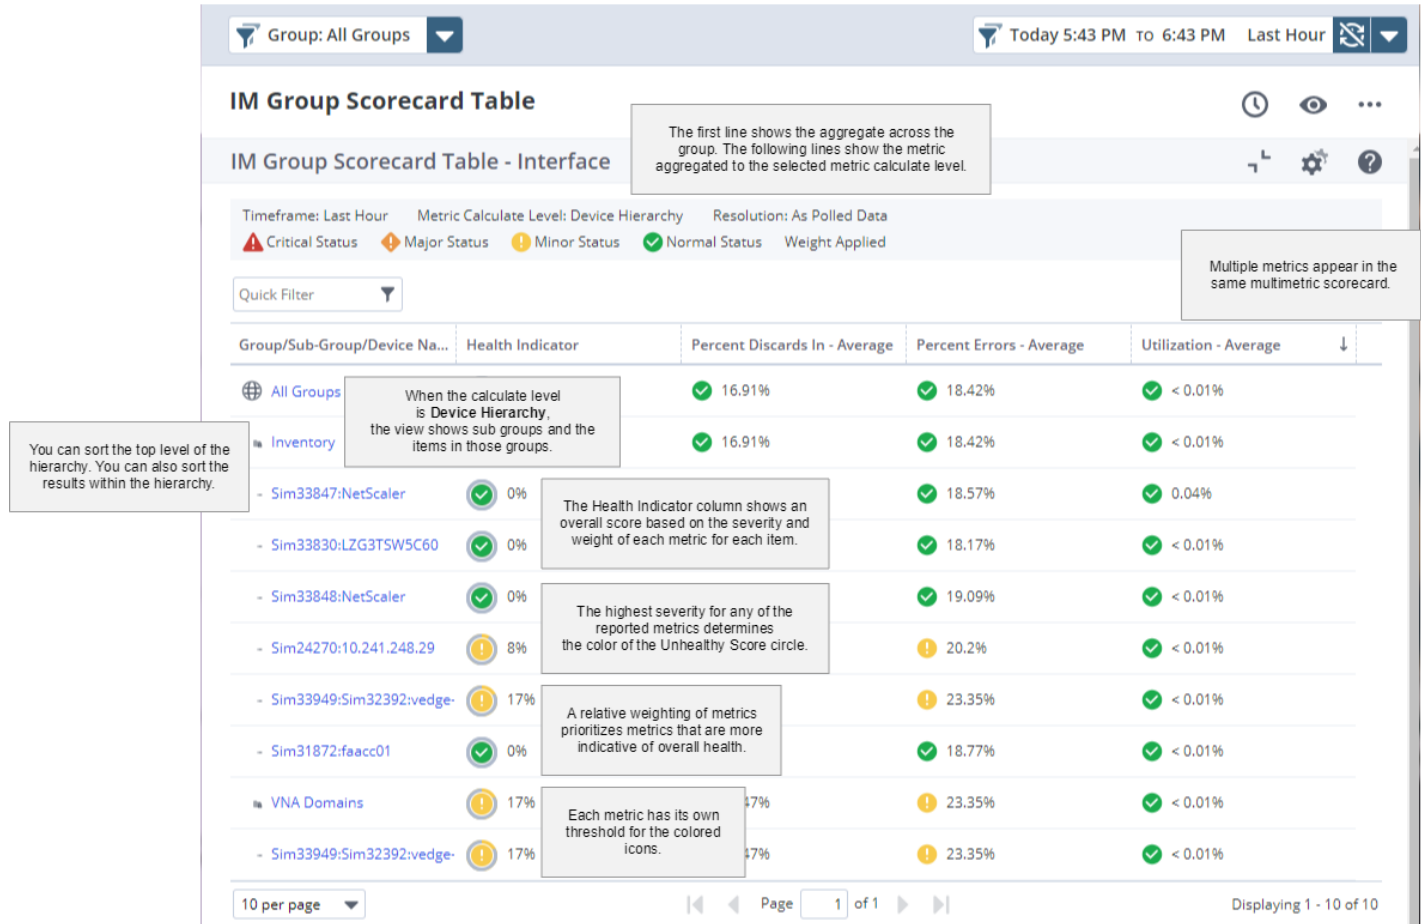
The Metric Calculate Level in the following image is set to Component Hierarchy.

**Figure 56: Group Scorecard Table Component Hierarchy**

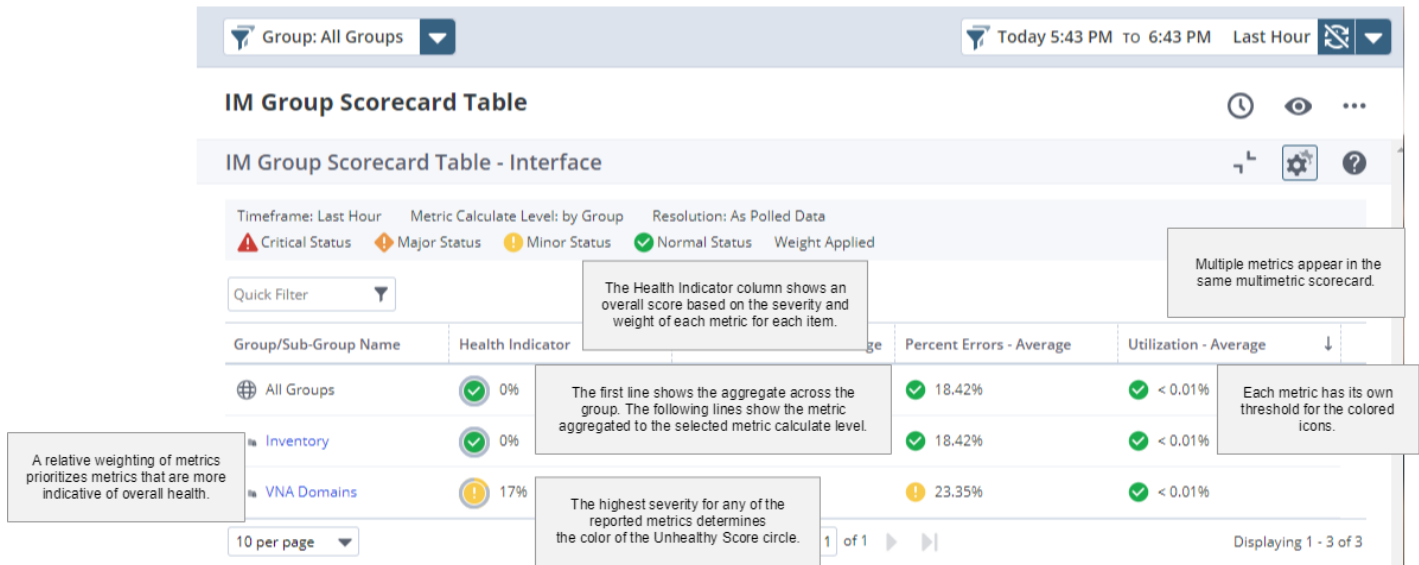


The Metric Calculate Level in the following image is set to Device Hierarchy:

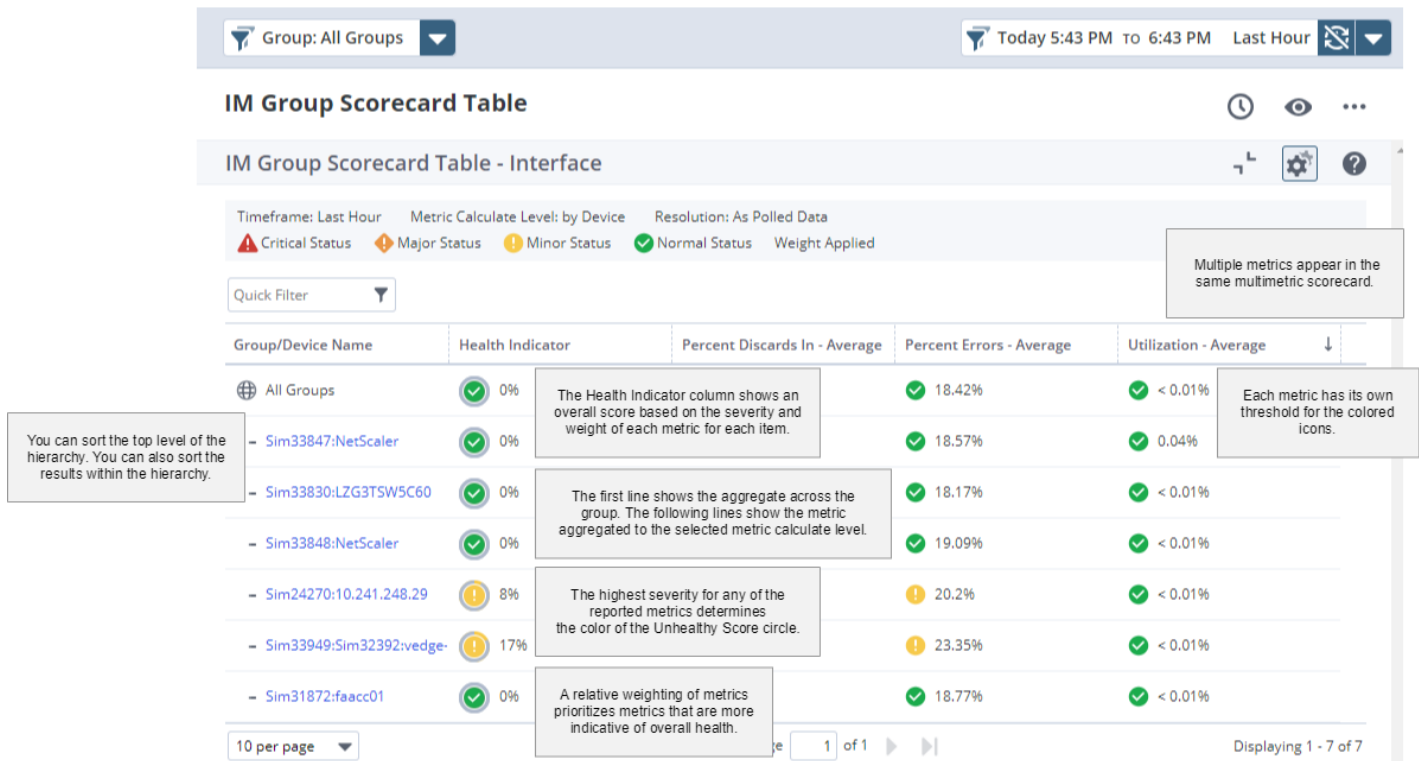
Figure 57: Group Scorecard Table Device Hierarchy



The Metric Calculate Level in the following image is set to Group:

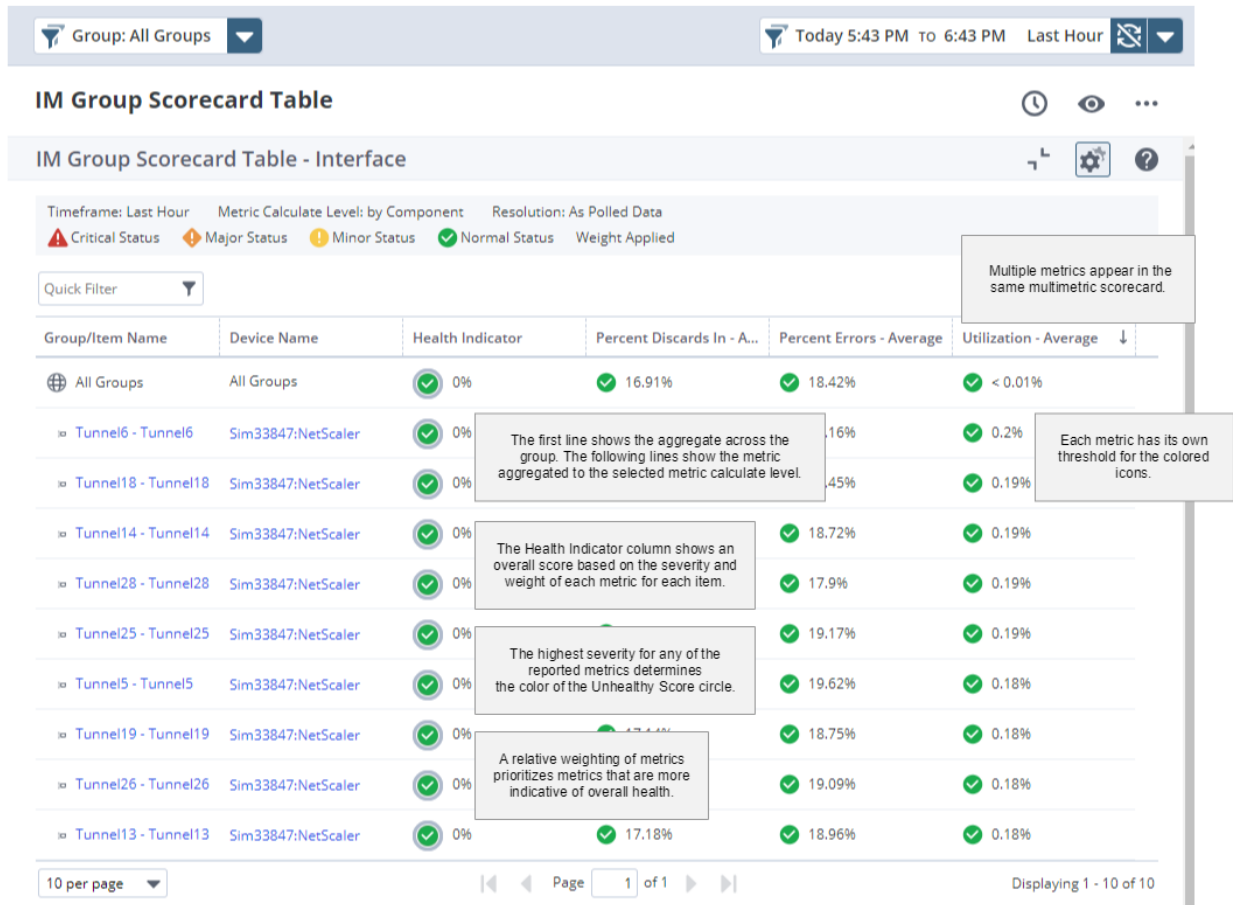
**Figure 58: Group Scorecard Table Group**

The Metric Calculate Level in the following image is set to Device:

**Figure 59: Group Scorecard Table Device**

The Metric Calculate Level in the following image is set to Component:

Figure 60: Group Scorecard Table Component



### Configure a Group Scorecard Table View

To display a group scorecard table view, add the IM Group Scorecard Table view. Group scorecard table views include standard view configuration options, such as time filters and context settings. For standard view configuration details, see [Customize Views](#).

#### TIP

On dashboards, group scorecard tables render best in single-column layouts. Group scorecard tables with fewer columns also render well in two-column layouts. Do not use group scorecard tables in three-column layouts.

#### Follow these steps:

1. **Show or Hide** the Health Indicator column.
2. Select the **Metric Calculate Level**.  
This option determines what level of aggregation each row in the scorecard view represents: a group, a device, a component, a device hierarchy, or a component hierarchy.  
**Component Hierarchy** shows aggregation to the current group and to each group that is a child of that group.
3. Set the **Results Limit**. This limit determines the number of items that are included in the data results shown on each page of the scorecard.

**NOTE**

If the **Results Limit** is less than the child item count, only the data results within the limit are aggregated. For accuracy, the **Results Limit** should exceed the child item count. However, your child item count might exceed the supported maximum limit to meet performance expectations. In this scenario, each parent group should contain a smaller set of child items.

4. For each metric, specify the thresholds for the status indicators.

If the value for the metric is at or above the specified value, the icon indicates the status by color.

Red indicates critical, orange indicates major, yellow indicates minor, and green indicates normal.

To remove the status indicators, set the threshold values for all status levels to zero.

You can specify thresholds with up to two decimal places with the lowest value being 0.01. Values with more than two decimal places are rounded.

**TIP**

For metrics where low values are bad and high values are good, set Minor highest and Critical lowest. To omit a severity level, set the threshold value to zero.

5. If the Health Indicator column is shown, for each metric, specify a weight:

- **0**

This weight shows the metric in the table, but does not factor the metric into the Health Indicator.

- **1**

This weight applies a standard weight to the metric when calculating the Health Indicator.

- **2**

This weight doubles the metric value when calculating the Health Indicator.

- **3**

This weight triples the metric value when calculating the Health Indicator.

## Inventory Hierarchy Views

You can view the group tree, devices, and device components, or interfaces.

The **Inventory Hierarchy** view shows the group tree, devices, and device components, or interfaces. The group tree is pinned to the current group context for the dashboard or context page.

The following image shows an example of this view:

**Inventory Hierarchy**

Europe

- Austria
- Belgium
- Bulgaria
- Croatia
- Cyprus
- Czech Republic
- Denmark
- Estonia
- Finland
- France
- Greece
- Greenland
- Hungary
- Iceland
- Ireland
- Italy
- Latvia
- Lithuania
- Luxembourg
- Malta
- Netherlands
- Norway
- Poland
- Portugal
- Romania
- Slovakia
- Slovenia
- Spain
- Sweden
- Switzerland
- United Kingdom

**Devices**

Quick Filter

On Demand

<input type="checkbox"/>	Name Alias	Type	Domain	Address	Description	Current...	Life Cycl...	Context Typ...
<input type="checkbox"/>	<a href="#">cis7606-96.37.ca.com.37.ca.c...</a>	Router	Default Do...	10.35.58.23	Cisco IOS Software, c7600s72033_rp Software (...)	✓ Norma	✓ Active	Router
<input type="checkbox"/>	<a href="#">cis7606-CPU TEST2.ca.com</a>	Router	Default Do...	10.35.57.13	Cisco IOS Software, c7600s72033_rp Software (...)	✓ Active	Active	Router
<input type="checkbox"/>	<a href="#">cis7606-CPU TEST2.ca.com</a>	Router	Default Do...	10.35.57.141	Cisco IOS Software, c7600s72033_rp Software (...)	✓ Active	Active	Router
<input type="checkbox"/>	<a href="#">cis7606-CPU TEST2.ca.com</a>	Router	Default Do...	10.35.57.14	Cisco IOS Software, c7600s72033_rp Software (...)	✓ Active	Active	Router
<input type="checkbox"/>	<a href="#">cis7606-CPU TEST2.ca.com</a>	Router	Default Do...	10.35.57.14	Cisco IOS Software, c7600s72033_rp Software (...)	✓ Norma	✓ Active	Router
<input type="checkbox"/>	<a href="#">Cisco-2800_10.35.57.191</a>	Router	Default Do...	10.35.57.191	Cisco IOS Software, 2800 Software (C2800NM-...)	⚠ Minor	✓ Active	Router
<input checked="" type="checkbox"/>	<a href="#">Cisco-2800_10.35.57.32</a>	Router	Default Do...	10.35.57.32	Cisco IOS Software, 2800 Software (C2800NM-...)	✓ Active	Active	Router, Swit...
<input type="checkbox"/>	<a href="#">Cisco-2800_10.35.57.32</a>	Router	Default Do...	10.35.57.32	Cisco IOS Software, 2800 Software (C2800NM-...)	⚠ Major	✓ Active	Router

Page 1 of 4

Displaying 1 - 100 of 315

**Interfaces (Cisco-2800\_10.35.57.32)**

Quick Filter

On Demand

<input type="checkbox"/>	Interface Name Alias	Description	Device Name Alias	If...	In...	Iftype	Domain	Speed I...	Speed O...	Fl...	Po...	IP Address...	Life Cy...
<input type="checkbox"/>	<a href="#">Fa0/0</a>	FastEthernet0/0	Cisco-2800_10.35...	1...	3	ethernetCs...	Default D...	100.00 ...	100.00 ...	n/a	E...		✓ Active
<input type="checkbox"/>	<a href="#">Fa0/0.1</a>	FastEthernet0/0.1	Cisco-2800_10.35...	R...	8	l2vlan (135)	Default D...	100.00 ...	100.00 ...	n/a	E...	170.12.1...	✓ Active
<input type="checkbox"/>	<a href="#">Fa0/0.729</a>	FastEthernet0/0.729	Cisco-2800_10.35...	R...	9	l2vlan (135)	Default D...	100.00 ...	100.00 ...	n/a	E...	10.11.18...	✓ Active
<input type="checkbox"/>	<a href="#">Fa0/1</a>	FastEthernet0/1	Cisco-2800_10.35...	1...	4	ethernetCs...	Default D...	100.00 ...	100.00 ...	n/a	E...		✓ Active
<input type="checkbox"/>	<a href="#">Mu1</a>	Multilink1	Cisco-2800_10.35...	M...	10	pppMultili...	Default D...	4.61 M...	4.61 Mb...	n/a	E...	170.12.2...	✓ Active
<input type="checkbox"/>	<a href="#">Se0/0/0</a>	Serial0/0/0	Cisco-2800_10.35...	1...	1	ppp (23)	Default D...	1.54 M...	1.54 Mb...	n/a	E...		✓ Active
<input type="checkbox"/>	<a href="#">Se0/1/0</a>	Serial0/1/0	Cisco-2800_10.35...	1...	2	ppp (23)	Default D...	1.54 M...	1.54 Mb...	n/a	E...		✓ Active
<input type="checkbox"/>	<a href="#">Se0/2/0</a>	Serial0/2/0	Cisco-2800_10.35...	1...	11	ppp (23)	Default D...	1.54 M...	1.54 Mb...	n/a	E...		✓ Active

Page 1 of 1

Displaying 1 - 8 of 8

You can do the following from this view:

- **Filter the devices, components, and interfaces to the members of that group.** Select a group.
- **Filter the interfaces or components in the bottom pane.** Select a device in the top pane.

#### NOTE

The filter only shows results for a single device. If you select multiple devices, the bottom pane is not filtered.

### Configure an Inventory Hierarchy View

For standard view configuration details, see [Customize Views](#).

The **Inventory Hierarchy** view can display devices, components, and interfaces. To configure which item types appear, configure the following properties:

- Select whether to show devices, components, or both.
- If the view shows components, select the **Component Context** type. This property determines the type of components that appear in the view.

## Map Views

Map views are available for SD-WAN tunnels and application/SLA paths only. Map views show the location of each DX NetOps Virtual Network Assurance-based site for a selected group. Sites with no data appear as gray icons. Map views show only a single level below the selected group.

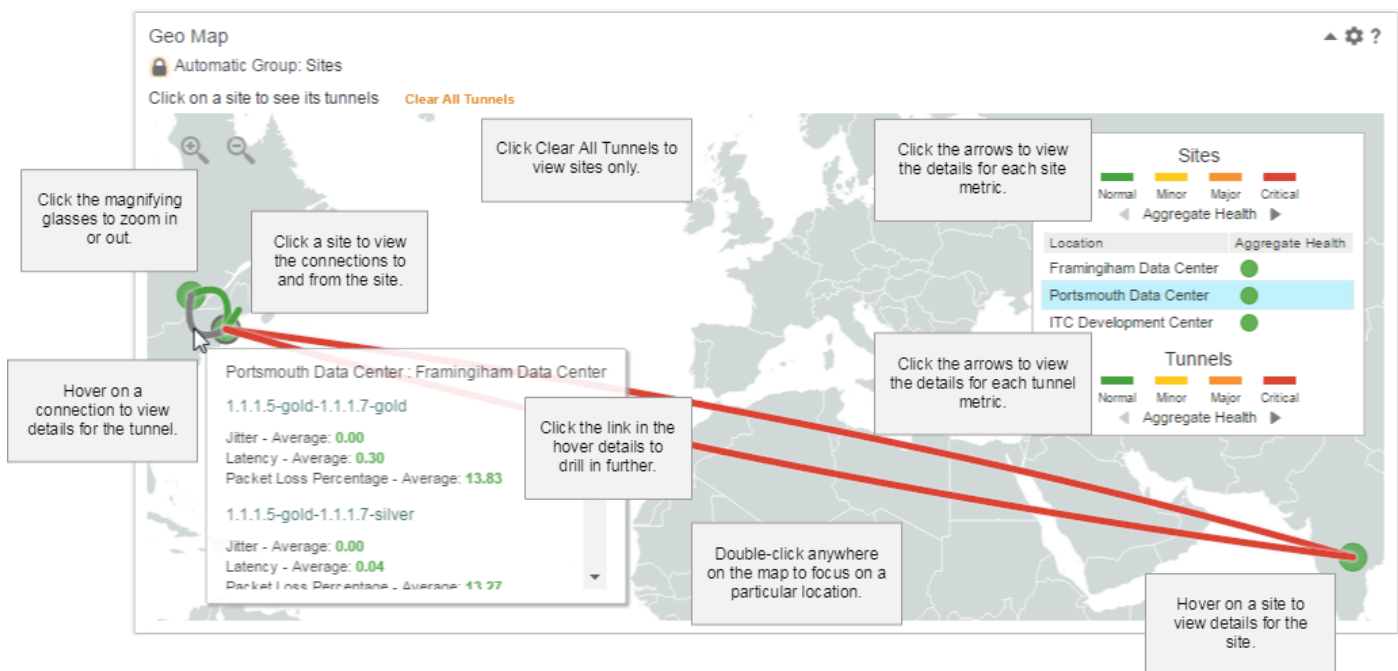
**Example:** A USA site group contains site groups for states. The site groups for states contain site groups for cities. If you select USA, the map view shows the site groups for states, but not for cities. If you select a state, the map view shows the site groups for cities within the state.

When a site is selected, the connections to and from other sites appear. The connection lines are color-coded based on health metrics. Site router details appear when you hover over a site. Tunnel or application/SLA path details appear when you hover over a connection.

For more information about how to monitor SD-WAN device inventory, see [Monitor SD-WAN Devices](#).

The following example shows the important elements of a map view for tunnels:

**Figure 61: Map\_View**



### Configure a Map View

Map views include standard view configuration options, such as time filters and context settings. For standard view configuration details, see [Customize Views](#).

#### NOTE

Map views are available only for SD-WAN tunnels and application/SLA paths. Other managed item types are unsupported.

#### Follow these steps:

1. Specify the thresholds for each site metric (CPU utilization, memory utilization, and virtual interface utilization). If the value for the site is at or above the specified value, the icon indicates the status by color. Green indicates normal, yellow indicates minor, orange indicates major, and red indicates critical.

You can specify thresholds with up to two decimal places with the lowest value being 0.01. Values with more than two decimal places are rounded.

**TIP**

For metrics where low values are bad and high values are good, set Minor highest and Critical lowest. To omit a severity level, set the threshold value to zero.

- Specify the thresholds for each tunnel or application/SLA path metric.

## Pie Chart Views

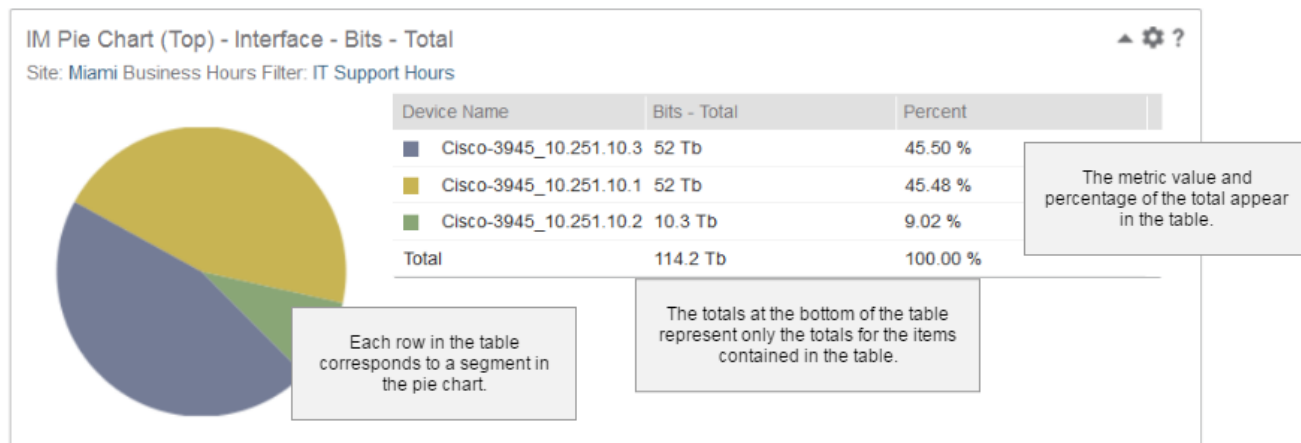
Pie charts show the relative values of a metric. Use pie charts to view metric values that represent parts of a whole. Pie charts are best used for small groups of items. This view includes a table that shows the metric value and the percentage of the total value for each item.

You can apply a business hours filter to the data in pie chart views. The applied business hours filter appears in the subtitle.

For more information about how to configure business hours definitions, see [Configure Business Hours Filtering](#).

The following example shows the important elements of a pie chart view:

**Figure 62: Pie Chart Elements**



To set up a pie chart, configure the following properties:

- Metric Label**  
 The selected attribute (Device Names, Item name, or Description) appears as the first column of the table. When the selected attribute is Device Name, the metrics are aggregated at the device level.
- Sort Direction**  
 The pie chart includes the items for the selected metric in descending or ascending order.
- Max Rows** The maximum number of table rows determines the maximum segments that appear in the pie chart.

For more information about how to configure a standard view, see [Customize Views](#).

### Pie/Table Views

For more information about pie/table views, see [Table Views](#).

## Table Views

Show vertically and horizontally dense data, include multiple metrics, many items, or many properties using a table view.



You can sort the data according to any metrics, and can view the best and worst results for each metric.

In this article:

- [Bar Chart Table Views](#)
- [Deviations from Normal Table Views](#)
- [Chart/Table Views](#)
- [Configure a Chart/Table View](#)
- [Configure a Table View](#)
- [Apply Business Hours Filtering to a Table View](#)
- [Manage a Table View](#)


For more information about how to manage table views, see [Customize Views](#).

### Bar Chart Table Views

Some out-of-the-box views show tables with a bar chart column. These views show metrics with percentage values. Use these views to identify problematic devices quickly.

The following example shows the important elements of a bar chart table:

**Figure 63: Table Bar Chart Elements**



### Deviations from Normal Table Views

The **Top Deviations from Normal** views compare actual values from the selected time frame to a calculated baseline value. Use these views to identify places where performance has changed. These views display the items that deviated the most from that normal value. The way the baseline is calculated varies by data source.

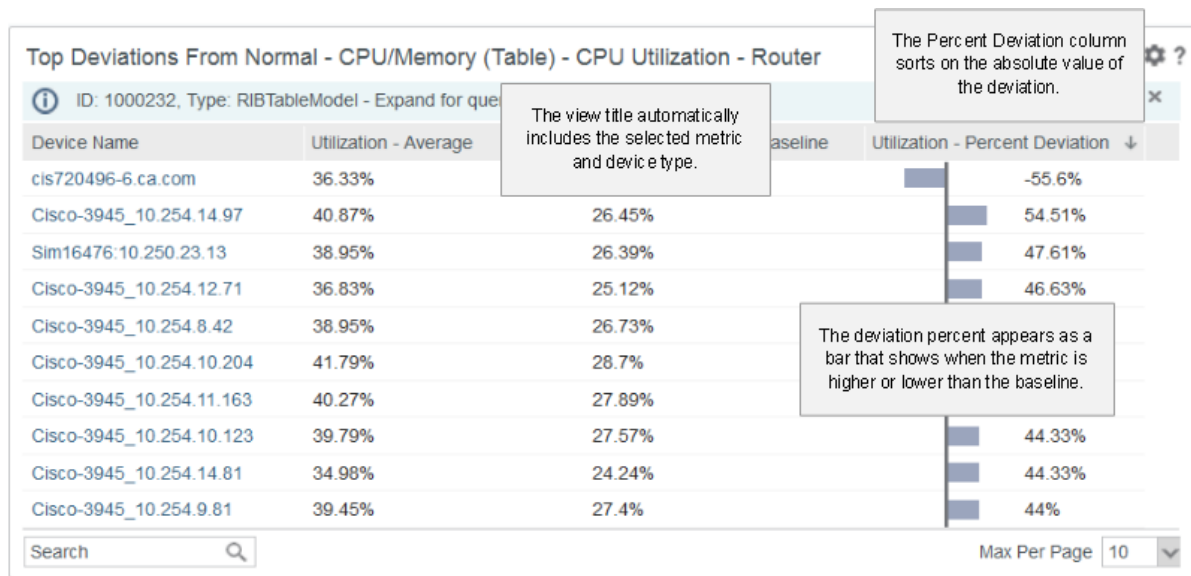
The table shows the following information:

- The average for the selected metric
- The average baseline for the selected metric
- The percent deviation with a relative visualization sorted by absolute value

Edit the view to specify the following information:

- The type of device: router, server, or switch
- The selected metric

The following example shows the important elements of a deviations from normal table:

**Figure 64: Deviations from Normal Table Elements**

### Chart/Table Views

Chart/table views combine aspects of other visualizations with table views. The following views support only percentage-type metrics:

- [Gauge/Table Views](#)
- [Trend/Table Views](#)
- [Radial Bar/Table Views](#)
- [Pie/Table Views](#)

#### NOTE

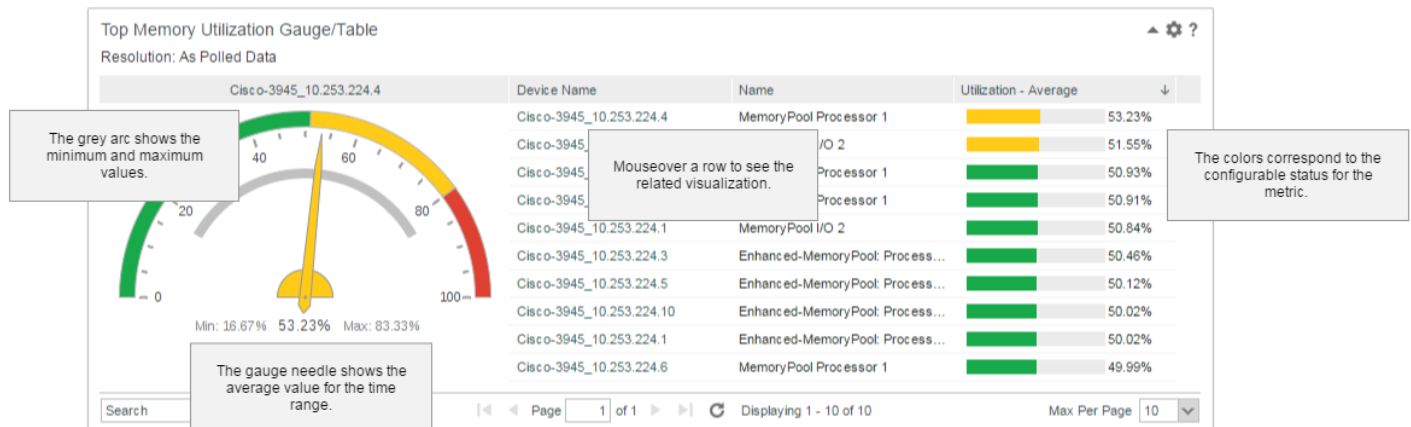
- Printing a combi-view displays only the table portion of the view in the PDF.
- Some settings are available for only certain table view types.  
For a list of all settings, see [Customize Views](#).

### Gauge/Table Views

Gauge/table views highlight where the value of one metric differs from the expected value when compared to other metrics. These views are good for metrics where the value is a percentage. The gauge shows the minimum, maximum, and average values for one item in the time range. In these tables, you can see the values for all items and select the item to show on the gauge.

The gauge chart legend shows the resolution, which specifies the granularity of the reporting data: as polled, hourly rollup, or daily rollup. For customized views, the minimum and maximize thresholds are based on this resolution. For out of the box views, the minimum and maximum thresholds are based on the as polled resolution.

The following example shows the important elements of a gauge/table view:

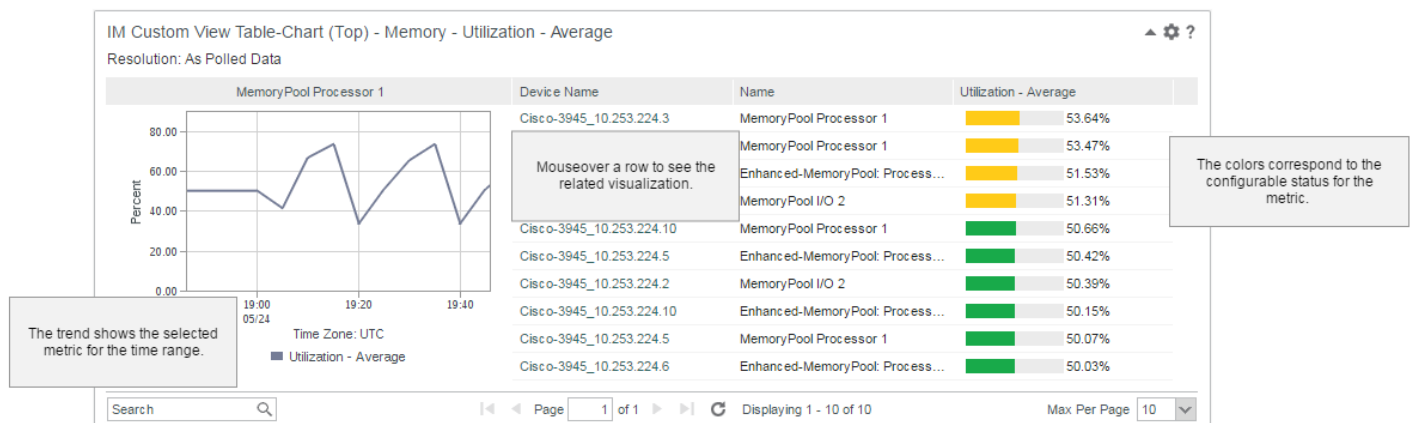
**Figure 65: Gauge Table Elements**

### Trend/Table Views

Trend/table views provide tables with data. The trend shows the change in a metric over a time range. In these tables, you can see the values for all items and select the item to show on the trend. Percentage-based metrics are represented with a bar. As you hover a metric, a related trend chart displays. The trend chart shows the metric value over the selected time range for each item. The table shows the values across the entire time range for each item being charted. Some out-of-the box trend/table views show a column with a line graph.

Most trend/table views have a trend chart that shows the value of a single metric. Some trend/table views have a trend chart that shows the value of multiple metrics. Each metric is represented with its own trend line.

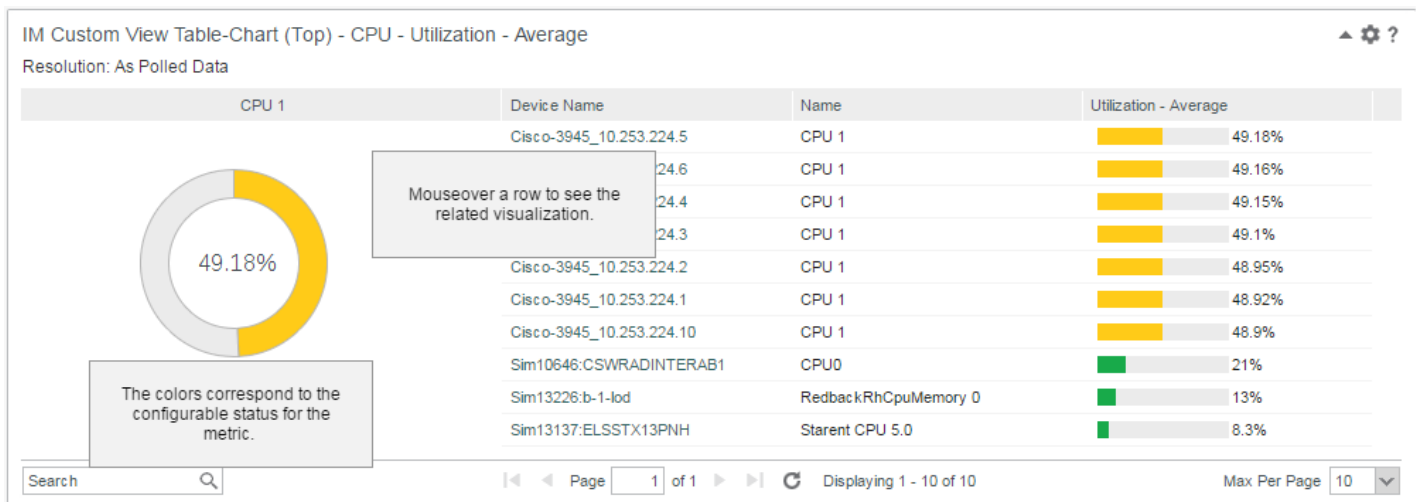
The following example shows the important elements of a trend/table view:

**Figure 66: Trend Table Elements**

### Radial Bar/Table Views

Radial bar/table views show percentage values within a circular bar.

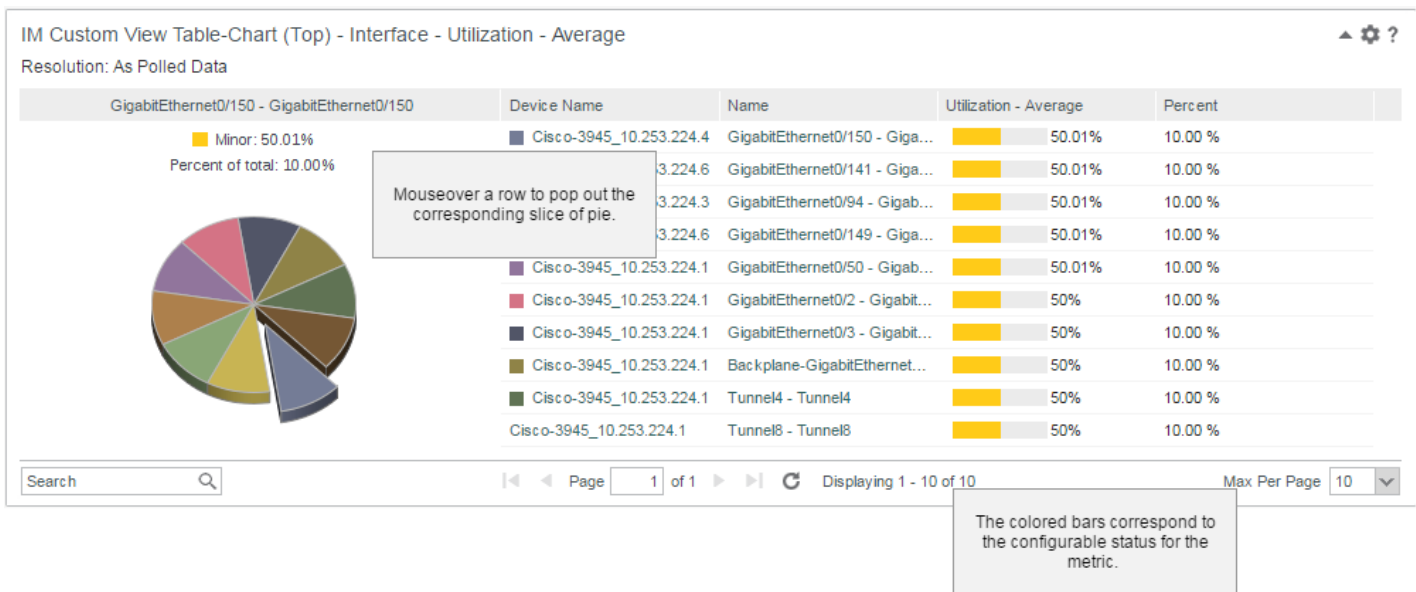
The following example shows the important elements of a radial bar/table view:

**Figure 67: Radial\_Bar**

### Pie/Table Views

Pie/table views show the relative values of a metric. When you hover over the corresponding item in the table, the corresponding slice of pie pops out. Use pie charts to view metric values that represent parts of a whole. Pie charts are best used for small groups of items.

The following example shows the important elements of a pie/table view:

**Figure 68: Pie\_Table\_View**

### Configure a Chart/Table View

The **IM Custom Chart-Table** view is a custom view that reports data from the data aggregator data source.

**Follow these steps:**

1. Select a **Chart Type**.
2. Select the **Metric Name** from the list of metrics for the selected metric family. This option determines the selected metric for the view.
3. If you are configuring a custom view, select the **Metric Calculate Level**.  
This option determines what level of aggregation each row in the view represents: a device, or a component.
4. Change definitions for the status levels:
  - **Minor Status Start**  
Metrics before this value have a normal status indicated with green.  
Metrics at or after this value have a minor status indicated with yellow.
  - **Major Status Start** Metrics at or after this value have a major status indicated with orange.
  - **Critical Status Start**  
Metrics at or after this value have a critical status indicated with red.

**NOTE**

Some out-of-the-box views show **Moderate Status Start**. For these views, moderate status is yellow and orange is unavailable.

**TIP**

You can configure chart/table views to reflect reverse value severity. For metrics where low values are bad and high values are good, set Green highest and Red lowest. To omit a severity level, set the threshold value to zero. For example, setting Orange to zero removes the major severity when profiling metric thresholds.

5. Select the **Max Rows**. This limit defines the total number of rows that the table can display. If the number of results exceeds the max rows, the table shows the top metrics according to the sorted column.

The chart/table view is configured.

**Configure a Table View**

The IM Table (Top) view is a custom view for data from the data aggregator.

**Follow these steps:**

1. Select the **Metric Fields** from the selected metric family to display in the table. Each selected field is a column in the table.

**NOTE**

If you select Device Name without Name, the metrics are aggregated to the device level. If you select Name or Description with or without Device Name, the metrics are aggregated to interface or component level.

2. Select the **Metric Sort**, which defines the column to use to sort the table by default.
3. Select the default **Sort Direction** for the **Metric Sort** column.
4. [Select a requested resolution](#).
5. Select the **Max Rows**. This limit defines the total number of rows that the table can display. If the number of results exceeds the max rows, the table shows the top metrics according to the sorted column.

**Apply Business Hours Filtering to a Table View**

You can apply business hours filtering to the data in table views. The applied filter appears in the subtitle.  
For more information, see [Configure Business Hours Filtering](#).

**Manage a Table View**

From a table view, you can do the following:

- Drill down to detailed data for individual items by clicking the item name.
- Sort by a particular column or reverse the sort order by clicking the column heading. An arrow icon indicates the sorted column.

The following image shows an example of the **Address** column sorted:

**Figure 69: Example of the Address column sorted**

## Devices

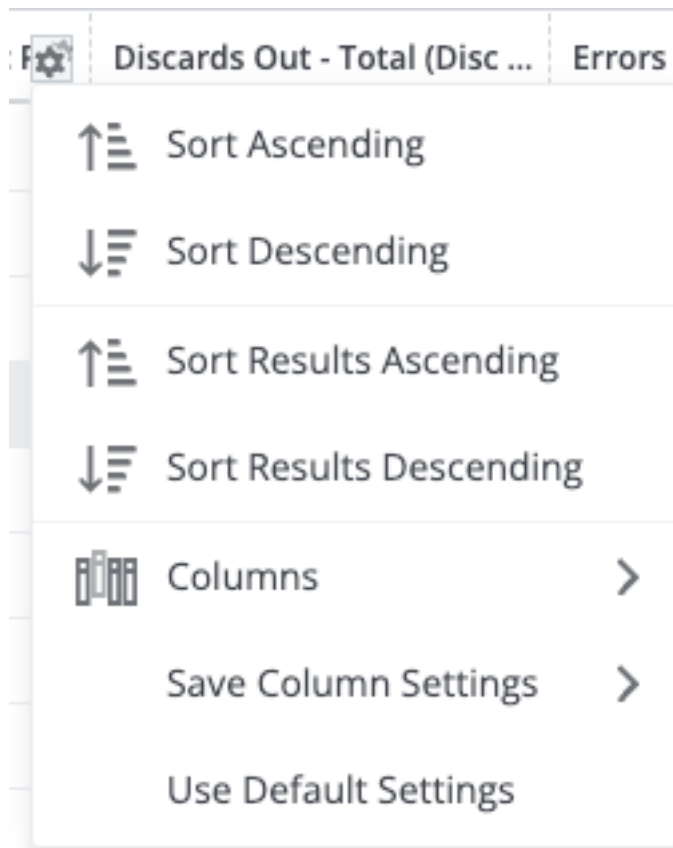
### Devices

Quick Filter 				
<input type="checkbox"/>	Name	Type	Domain	Address 

When you sort by a different column, NetOps Portal queries the data source to get new results for the table that match the top values for the new sorted column. The sort order applies only to the currently rendered view. The sort order does not persist and cannot be saved on the view.

- To move columns, click and drag the column header to the desired location.
- To resize columns, hover over the separator line between two columns in the grid header, and click and drag the column to be wider or narrower.
- To manage columns, mouse over a column heading, and then click the gear icon.

The following image shows an example of these column options:

**Figure 70: Gear Icon for Columns**

Column options:

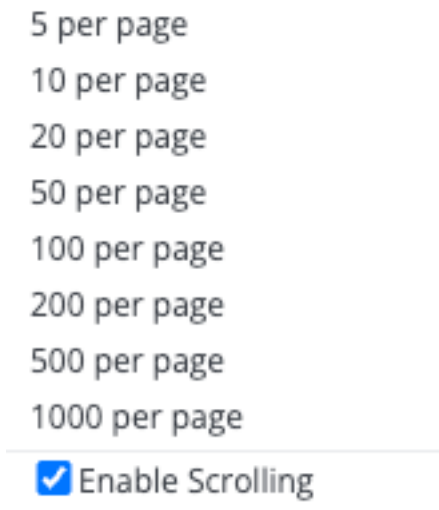
- **Sort Ascending/Sort Descending**  
Sort the entire data set from the data source.
- **Sort Results Ascending/Sort Results Descending**  
If available for the view, sort the already-retrieved and rendered result set.
- **Columns**  
To show/display or hide columns in the table, select **Columns**, and then select the columns to show/display and/or clear the checkbox for the columns to hide. This list includes the selected metrics from the view configuration and more details about the items, such as **Name Alias** and **Life Cycle State**.
- **Save Column Settings**  
To save the view column settings, select Save Column Settings, and then select the level at which to save the column settings changes, such as column selection and the order of columns. Saving the view column settings does not preserve the column sort order.  
**Options:**
  - **For All Tenant Users:** Save the view column settings at the tenant level.
  - **My User Account:** Save the view column settings at the user level.
  - **My Current Session:** Save the view column settings at the session level.
- **Use Default Settings**  
Revert the changes applied to the view column settings.

#### NOTE

By default, when you display more than 20 rows per page (the **Rows per page** drop-down in the lower-left corner), the **Enable Scrolling** setting is selected and the Search bar is locked in place. To change this setting, clear the **Enable Scrolling** setting.

The following image shows an example this setting:

**Figure 71: Enable Scrolling**

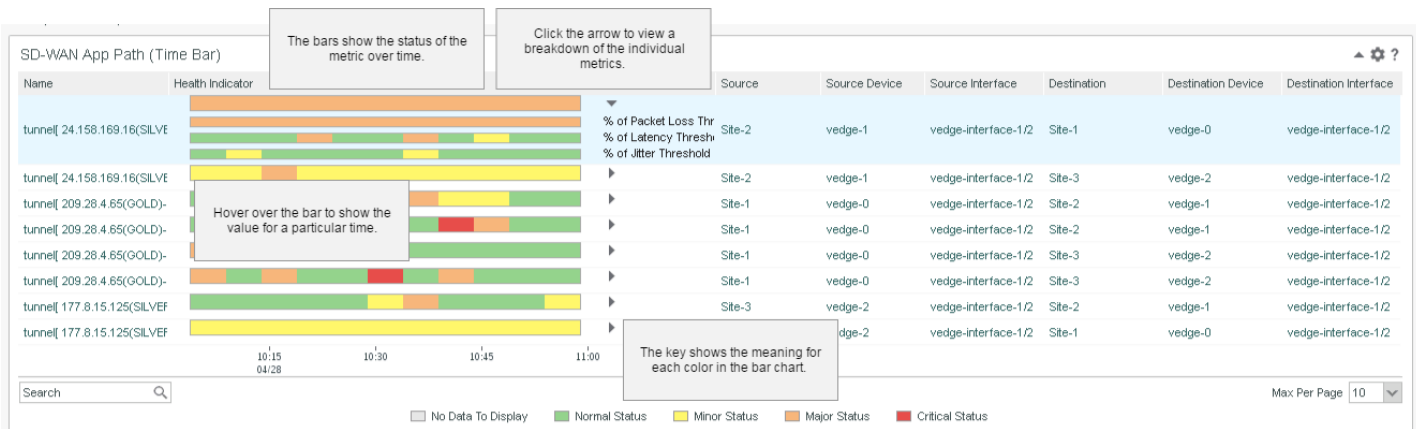


## Time Bar Chart Views

Time bar chart views use a bar to show the status of percentage metrics by color over the selected time range. These views show threshold values within a time series. The time slices represent the statuses as bars of colors. Hover over a section of the time range to view a tooltip with the corresponding value.

The following example shows the important elements of a time chart:

**Figure 72: Time Chart Elements**



## Sort a Time Bar Chart

Time bar views report only percentage metrics. Percentage values range inclusively from 0% to 100%. Values that are outside that range are marked as invalid, shown as a gray bar, and excluded from the sort order.

Use the Sort Direction in the view settings to sort the highest and lowest items based on the maximum rows across the entire time range. This also sorts the colored bars within the time bar itself.

You can also click the gear icon in the table heading to sort as follows:

- **Sort Ascending/Sort Descending**



Sort the already retrieved and rendered result set.

### **Configure a Time Bar Chart**

The IM Time Bar Chart is available on dashboards, device context pages, and group context pages. This view is not available on interface or component context pages. For standard view configuration details, see [Customize Views](#).

#### **TIP**

On dashboards and in context pages, time bar charts render best in single-column or two-column layouts.

To configure the behavior of a custom time bar chart, specify the following properties:

- **Metric Value**  
Select a specific metric to include in the view. To view a time bar chart aggregating the associated metrics, select the following option:
  - **Health Indicator**  
You can click to expand and view a time bar chart for each metric.
- **Metric Calculate Level**  
This option determines what level of aggregation each row in the view represents: a device, or a component.
- Define the thresholds for the bar:
  - **Critical Status Start**  
The bar shows red for metrics with this value for this time.
  - **Major Status Start**  
The bar shows orange for metrics with this value for this time.
  - **Minor Status Start**  
The bar shows yellow for metrics with this value for this time.

### **Time Bar/Table Views**

For more information about these views, see [Table Views](#).

## **Trend Views**

Trend views show the value of a metric across time. Trends show spikes or dips in activity and help identify whether a metric is increasing or decreasing.

Trend views provide one or more of the following chart types:

- **Trend Chart**  
View a traditional trend line for each device, component, or metric.
- **Stacked Chart**  
View a stacked line on top of each other for each device, component, or metric.

#### **NOTE**

- Trend views displaying data up to the present moment might contain dips or spikes for the present moment as data is processed.
- Maintenance indicators apply to all the devices and components in a site group. When the associated site group is selected in the context, maintenance indicators appear as shading in trend views.  
For more information, see [Schedule Maintenance Indicators](#).

In this article:

- [Trend View Options](#)
- [Trend View Element Examples](#)
- [Compress Non Business Hours in a Trend View](#)
- [Configure a Trend View](#)
- [Hide Gaps in a Trend View](#)
- [Configure a MultiTrend View](#)
- [Configure a MultiView](#)
- [Applied Business Hours Filters](#)

For information about trend/table views, see [Table Views](#).

### **Trend View Options**

With trend views, you can quickly change the trend lines that are displayed on the graph. The following options also apply to MultiTrend views:

- To remove a metric from the chart temporarily, right-click the metric in the chart legend, and then click **Hide**.
- To add a metric back into the chart, right-click the metric in the chart legend, and then click **Show**.
- To hide all other metrics, right-click a metric in the chart legend, and then click **Focus**.
- To narrow the view on a precise time frame, click and drag across the time frame in the chart legend. Select a time frame of at least 30 minutes. For MultiTrend views, drill into an individual trend view, then select a time frame.

#### **NOTE**

To zoom, non-business hours compression must be disabled (the **Enabled** checkbox for the **Compress Non Business Hours** field is cleared). By default, it is disabled.

For more information about this field, see [the "Configure a Trend View" section](#).

- To apply the precise time frame to the entire dashboard, click **Apply to Dashboard** at the bottom of the chart.
- To revert the view to the default settings, click **Undo** in the lower-left corner.

### **Trend View Element Examples**

Trend views can include the following elements:

#### **NOTE**

A specific trend view might or might not include the following elements. A specific trend view might include elements in addition to following elements.

- **Baselines**

If available, you can use baseline trend lines to determine whether utilization trends are changing. The baseline data that is plotted in many views shows statistical deviations from normal performance for a given statistic.

For more information, see [Baseline Calculations](#).

#### **NOTE**

Baseline metrics are processed up to hourly and daily increments. When the resolution is less than an hour, data points are interpolated between known hourly data points.

For more information about the resolution for trend views, see [Set the Resolution for Reported Data](#).

- **Events**

If available, events appear as flags indicating the times that the events occurred.

#### **NOTE**

Events are based on as-pollled data. If a rollup is applied to the trend line, a threshold event might appear outside the trend line.

Trend charts with events are disabled when the time range is greater than three (3) months.

- **Trend Projections**

If available, projections on charts appear as dashed trend lines that predict future performance of metrics. Each projection is based on trends in recent values in the time frame of the chart. Projections help determine whether current performance is within normal bounds.

By default, each projection includes a future time frame equivalent to one time beyond the current time frame. For example, if the time frame is the Last Hour, the projection shows values for the next hour.


For views that support projections, you can show or hide the projection data by editing the view settings.

The following image shows the important elements of a trend view:

#### NOTE

Percentile metrics appear as dashed lines.

**Figure 73: TrendView**




#### Compress Non Business Hours in a Trend View

If the resolution is less than one day and business hours filtering is assigned to the view, trend views apply shading. The shading differentiates between the selected business hours and non business hours.

For more information:


- About how to assign business hours filters to a view, see [Customize Views](#).
- About how to manage the resolution for trend views, see [Set the Resolution for Reported Data](#).

The following image shows a trend view with shaded non business hours:

**Figure 74: Non\_Business\_Hours**

You can also configure a trend view to compress non business hours. When non business hours are compressed (the **Enabled** checkbox for the **Compress Non Business Hours** field is selected), the shaded regions are compressed into a single region equal to the resolution period. The resolution is listed at the bottom of the chart. When non business hours compression is enabled, you cannot zoom.

The following image shows a trend view with compressed non business hours:

**Figure 75: Compressed\_Non\_Business\_Hours**

## Configure a Trend View

To configure a trend chart, customize or edit a trend view. For standard view configuration details, see [Customize Views](#).

### Follow these steps:

1. Do one of the following:
  - To edit an existing trend view, click the **View Settings** (gear) icon on the view, and then select **Edit** from the menu.
  - To configure a new trend view, click the **Edit** icon (the Pencil icon).

The view settings dialog opens.
2. To control any of the following trend lines or data points, select **Show** or **Hide** from the drop-down list:
  - **Projection**  
A line that plots projected average values for twice the current time period.  
**Default:** Hide
  - **95th Percentile**  
A line that plots the values in the top 5 percent of measurements that were taken during the reporting interval.
  - **Events**  
Data points that show when events were reported during the reporting interval.

**NOTE**  
Not all trend line options are available for all trend chart types.
3. (23.3.4 and higher) In the **Set Trend Goal/Threshold Line** section, complete the following fields:
  - **Goal/Threshold Line**  
Specifies whether to show or hide a goal/threshold line trend line in the view. Goal line trend lines serve as visual indicators of fixed metric values.  
**Options:**
    - **Disabled:** You do not want to show a goal or threshold line trend line in the view.
    - **Enabled:** You want to show a goal or threshold line trend line in the view.**Default:** Disabled
  - **Goal Line**  
(If you have decided to show a goal/threshold line trend line) Specifies the type or trend line to show in the view, and where NetOps Portal shows it.  
**Options:**
    - **Hide:** A goal/threshold line trend line is hidden from the view.
    - **Show Goal Line:** A goal line trend line is shown in the view and in the chart legend (golden color).
    - **Show Threshold Line (Above):** A threshold line trend line is shown in the view with shading above the line with pink offset of critical color. The line is shown in trend charts and in the chart legend.
    - **Show Threshold Line (Below):** A threshold line trend line is shown in the view with shading below the line with pink offset of critical color. The line is shown in trend charts and in the chart legend.
    - **Show Threshold Line (Above and Below):** A threshold line trend line is shown in the view with shading above the line with pink offset of critical color and shading below the line with green offset of normal color. The line is shown in trend charts and in the chart legend.**Default:** Hide
  - **Goal Line Value**  
Specifies a value for the goal line trend line. Goal line trend lines serve as visual indicators of fixed metric values. For example, select a value within a critical threshold range to determine whether performance is approaching a degraded level quickly.  
**Default:** 40
  - **Value Format**  
Specifies whether NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.  
**Options:**

- **Scaled:** NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.
- **Unscaled:** NetOps Portal does not scale the goal line or threshold line value. Instead, it uses the raw value on the trend line on the chart Y-axis.

**Default:** Scaled

– **Goal Line Label**

Specifies the label for the goal line trend line.

**Default:** Goal Line

4. (23.3.3 and lower) Complete the following fields:

– **Goal Line**

A line that plots the value that you selected as the goal for the metric in this view. This option is available only for some of the trend views.

**Default:** Hide

– (If you selected to show the goal line trend line) **Goal Line Value**

Specifies a value for the goal line trend line. Goal line trend lines serve as visual indicators of fixed metric values. For example, select a value within a critical threshold range to determine whether performance is approaching a degraded level quickly.

**Default:** 40

– (If you selected to show the goal line trend line) **Goal Line Label**

Specifies the label for the goal line trend line.

**Default:** Goal Line

5. Select whether to **Compress Non Business Hours** by selecting or clearing the **Enabled** checkbox.

**Options:**

- **Selected:** Non-business hours are compressed.
- **Cleared:** Non-business hours are not compressed.

**Default:** Cleared (Disabled)

For more information about non business hours compression, see [the "Compress Non Business Hours in a Trend View" section](#).

6. Click **Save**.

The trend view is configured.

### **Hide Gaps in a Trend View**

By default, DX NetOps Performance Management shows gaps in trend chart types when DX NetOps Performance Management does not have data for the item for a timeframe (data does not exist for that data point). You can configure to hide these gaps in trend views, and to show a straight line from the last data point to the next data point.

#### **NOTE**

When configured to hide the gaps, the connected line *does not* reflect the real behavior for the device during the period where data does not exist for that data point. Gaps are visible in exported data.

The following circumstances cause gaps in trend views:

- The device does not respond to a poll request.
  - The device is down.
  - A network problem causes the poll response to be lost.
  - The data collector is down.
  - For as-pollled data, a counter rollover or bad poll occurred.
- For more information about counters, see [Configure Counter Behavior](#).

**Follow these steps:**

1. Log in to the data aggregator host.
2. Create the `<installation_directory>/IMDataCollector/apache-karaf/etc/com.ca.im.dm.ribsource.RIBSourceWsImpl.cfg` file.
  - ***installation\_directory***  
The installation directory for the data aggregator.  
**Default:** `/opt`
3. Edit the file and insert the following property configuration to that file:  
`gapShowingOn=false`

The trend views now hide gaps. To apply the new behavior, refresh the view.

**Configure a MultiTrend View**

With MultiTrend views, you can combine trend data from multiple components in a single-line trend chart. You can compare the data from up to 12 components in a single MultiTrend view. You can use MultiTrend views to view data from the following interfaces:


- Multiple interfaces in a group
- Interfaces in your inventory with the worst performance metrics
- Interfaces on a single device

**TIP**

MultiTrend views require more page space than other types of views. To add a MultiTrend view to a dashboard, add it to a single-column area in the layout. Do not add a MultiTrend view to a small section on a dashboard layout.

MultiTrend views also require more space for the PDF output.

The following image shows the important elements of a MultiTrend view:

**Figure 76: MultiTrendView****Follow these steps:**

- From the view settings dialog, specify the following fields:
  - Sort Direction**  
The sort direction determines which interfaces are reflected in the first pages of trend charts.  
**Options:**
    - Descending:** highest values first
    - Ascending:** lowest values first**Default:** Descending
  - Resolution**  
Defines the resolution frequency.  
**Default:** Use default resolution (5 minute)
  - Chart Type**  
Defines the type of chart, either lines or area.  
**Options:**
    - Trend Chart**  
Show trend lines for each metric.
    - Stacked Chart**  
Show lines for each metric stacked with area fills under each.**Default:** Trend Chart
  - Chart Display Order**  
Defines how the interface charts are sorted.  
**Options:**
    - Display by Metric**



Sort the interface charts by metric value, with the highest (most severe) values shown first, top to bottom.

- **Display by Name**

Sort the interface charts in alphanumeric order, by interface name.

**Default: Display by Metric**

- **Metric Calculate Level**

Defines whether to show metric values for devices or for the device components that are plotted on the charts in the view.

**Options:**

- **by Component:** Show metric values only for the device components.
- **by Device:** Show metric values only for devices.

**Default:** by Component

- **Standardized Axis**

Defines the options for the Y-axis for the trend chart.

**Options:**

- **Calculated**

Let the Y-axis adjust dynamically, based on the range of metric values that are included. The calculated setting is applied to all charts in the view.

- **Fixed at 0 to 100**

Maintain a static range, 0 through 100, for the Y-axis.

- **Scale per Chart**

Let the Y-axis adjust dynamically, based on the range of metric values for each chart in the view. The Y-axis scaling of one chart in a view does not affect the scaling of any other chart.

**Default: Calculated**

- **Number of Charts on Page**

(If you are configuring a **MultiView (Draws individual chart for each item)** trend view) Defines the number of charts for each page.

**Default:** 15

**Limit:** 36

- **Maximum Number of Charts**

(If you are configuring a **MultiView (Draws individual chart for each item)** trend view) Defines the maximum number of charts.

**Default:** 60

**Limit:** 1200

2. (23.3.4 and higher) In the **Set Trend Goal/Threshold Line** section, complete the following fields:

- **Goal/Threshold Line**

Specifies whether to show or hide a goal/threshold line trend line in the view. Goal line trend lines serve as visual indicators of fixed metric values.

**Options:**

- **Disabled:** You do not want to show a goal or threshold line trend line in the view.
- **Enabled:** You want to show a goal or threshold line trend line in the view.

**Default:** Disabled

- **Goal Line**

(If you have decided to show a goal/threshold line trend line) Specifies the type or trend line to show in the view, and where NetOps Portal shows it.

**Options:**

- **Hide:** A goal/threshold line trend line is hidden from the view.
- **Show Goal Line:** A goal line trend line is shown in the view and in the chart legend (golden color).
- **Show Threshold Line (Above):** A threshold line trend line is shown in the view with shading above the line with pink offset of critical color. The line is shown in trend charts and in the chart legend.
- **Show Threshold Line (Below):** A threshold line trend line is shown in the view with shading below the line with pink offset of critical color. The line is shown in trend charts and in the chart legend.
- **Show Threshold Line (Above and Below):** A threshold line trend line is shown in the view with shading above the line with pink offset of critical color and shading below the line with green offset of normal color. The line is shown in trend charts and in the chart legend.

**Default:** Hide

– **Goal Line Value**

Specifies a value for the goal line trend line. Goal line trend lines serve as visual indicators of fixed metric values. For example, select a value within a critical threshold range to determine whether performance is approaching a degraded level quickly.

**Default:** 40

– **Value Format**

Specifies whether NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.

**Options:**

- **Scaled:** NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.
- **Unscaled:** NetOps Portal does not scale the goal line or threshold line value. Instead, it uses the raw value on the trend line on the chart Y-axis.

**Default:** Scaled

– **Goal Line Label**

Specifies the label for the goal line trend line.

**Default:** Goal Line

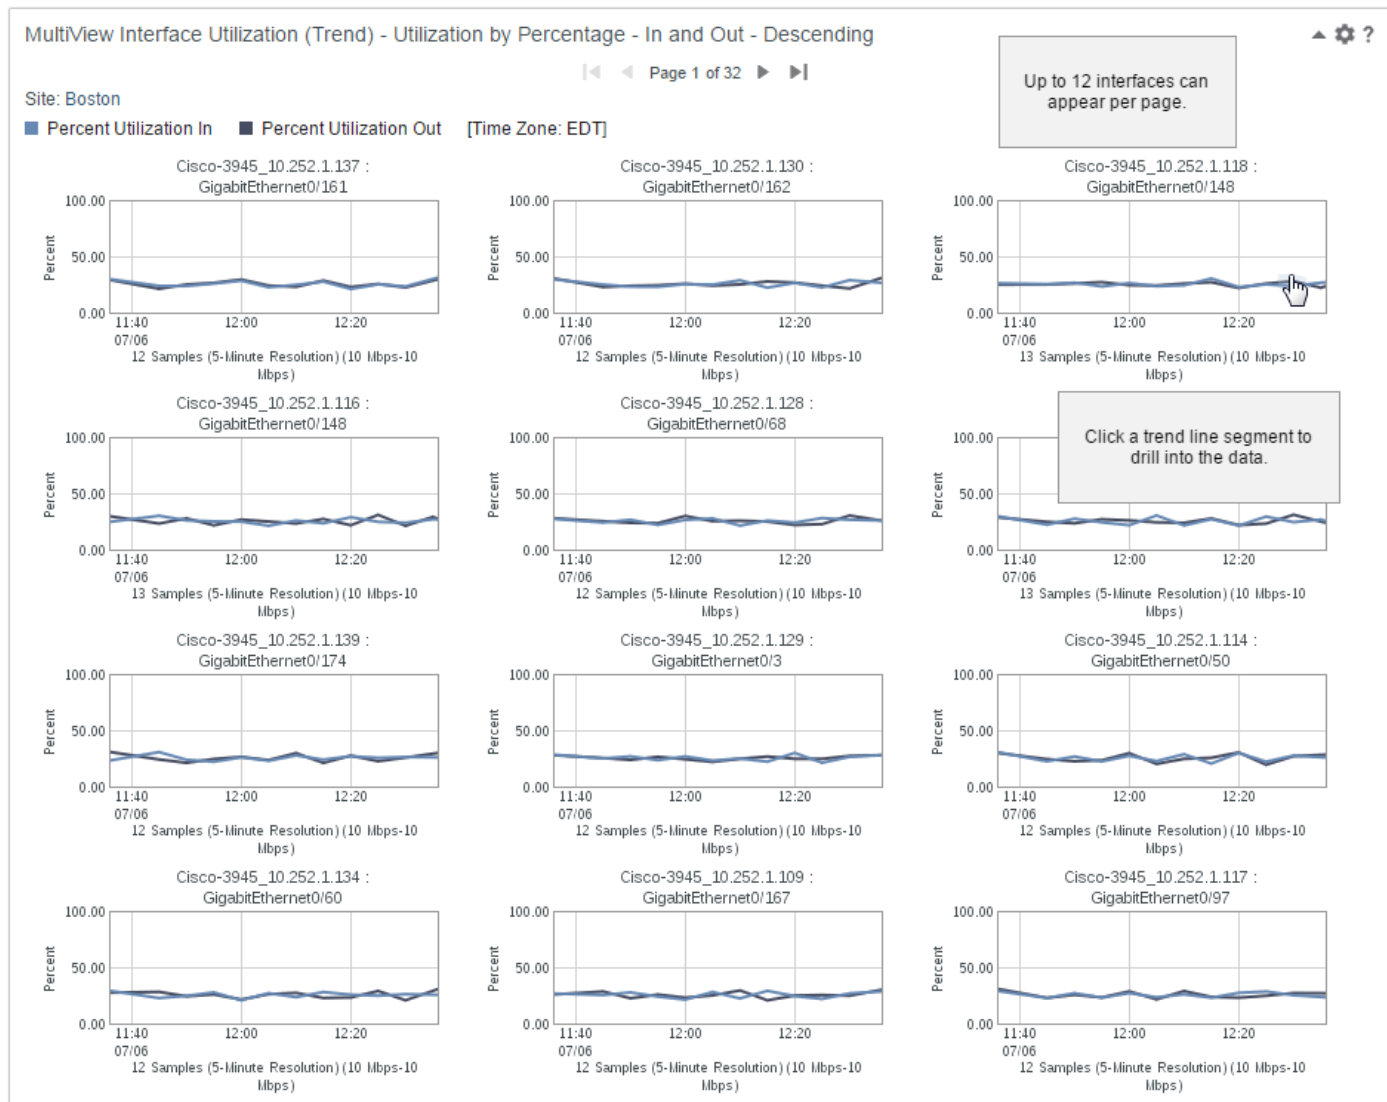
3. In the **Customize Time Range** section, [filter data based on specific time periods by selecting a custom time range for the view](#).
4. In the **Apply Business Hours** section, [assign a business hours filter and a time zone to the view](#).  
Applying a business hours filter to a dynamic trend view displays the same data in the view, but the data with the periods *outside* of the applied business hours are shaded.  
For more information about dynamic trend views, see [Dynamic Trend Views](#).
5. Click **Save**.

The MultiTrend view is configured.

### **Configure a MultiView**

MultiViews combine statistics from multiple interface charts in a single view. Use these views to diagnose performance issues on a device. For example, you can compare the recent performance of all interfaces on a card in a single view. More MultiViews that query for different types of metrics are available in the Device category when you edit a dashboard.

The following image shows the important elements of a MultiView:

**Figure 77: MultiView****Follow these steps:**

- Do one of the following:
  - To edit an existing MultiView, click the **View Settings** (gear) icon on the view, and then select **Edit** from the menu.
  - To configure a new MultiView, click the **Edit** icon (the Pencil icon).

**TIP**

MultiViews require more page space than other types of views. To add a MultiView to a dashboard, add it to a single-column area in the layout. Do not add a MultiView to a small section on a dashboard layout.

MultiViews also require more space for the PDF output.

For more information about how to add a view to a dashboard, see [Manage Dashboards](#).

The view settings dialog opens.

- Complete the following fields:
  - Title**  
Defines the title for the MultiView.

**Required:** Yes

– **Metric Family**

Specifies the metric family for the MultiView. For example, Wireless Access Point.

**Required:** Yes

– **Metrics Fields**

Specifies the metrics to display in reports. You can specify up to 10 metrics to display. Move the metrics that you want displayed in reports from the **Metrics Fields Available** list to the **Metrics Fields Selected** list. The metrics listed in the **Metrics Fields Available** list are those that are related to the specified metric family. In the **Metrics Fields Selected** list, use the arrows to specify an order for the selected metrics.

**Required:** No

– **Metric Value**

Defines the metric value for the MultiView. The options are based on the view that you are configuring.

**Required:** Yes

– **Direction Settings**

Defines the direction settings for the MultiView.

**Options:**

- In
- Out
- In and Out

**Default:** In and Out

**Required:** Yes

– **Metric Sort**

Defines the column for NetOps Portal to use to sort the table by default.

**Default:** Utilization - Average

**Required:** Yes

– **Sort Direction**

Determines which interfaces are reflected in the first pages of trend charts.

**Options:**

- Descending
- Ascending

**Default:** Descending

**Required:** Yes

– **Resolution**

Defines the resolution frequency.

**Options:**

- 1 Minute
- 5 Minute
- 10 Minute
- 15 Minute

**Default:** Use default resolution (5 Minute)

**Required:** Yes

– **Chart Legend Type**

Defines how NetOps Portal displays the information about the trend chart below the chart, either in a table or chart legend.

**Options:**

- **Legend Type Chart:** The trend chart information is displayed in a chart legend.
- **Legend Type Table:** The trend chart information is displayed in a table legend.

**Default:** Legend Type Chart

– **Chart Type**

Defines the type of chart, either lines or area.

**Options:**

- **Trend Chart:** Show trend lines for each metric.
- **Stacked Chart:** Show lines for each metric stacked with area fills under each.

**Default:** Trend Chart

– **Projection**

Defines whether a line that plots projected average values for twice the current time period is displayed.

**Options:**

- **Hide:** A line is not displayed.
- **Show:** A line is displayed.

**Default:** Hide

– **Chart Display Order**

Defines how the interface charts are sorted.

**Options:**

- **Display by Metric:** Sort the charts by metric value, with the highest (most severe) values shown first, top to bottom.
- **Display by Name:** Sort the charts in alphanumeric order, by device name and then item name. If the **Metric Calculate Level** is set to **by Device**, the charts are sorted only by device name.

**Default:** Display by Metric

– **Metric Calculate Level**

Determine the level of aggregation for metrics. Defines whether to show metric values for devices or for the device components that are plotted on the charts in the view.

**Options:**

- **by Component:** Show metric values only for the device components.
- **by Device:** Show metric values only for devices.

**Default:** by Component

– **Standardized Axis**

Defines the options for the Y-axis for the trend chart.

**Options:**

- **Calculated:** Let the Y-axis adjust dynamically, based on the range of metric values that are included. The calculated setting is applied to all charts in the view.
- **Fixed at 0 to 100:** Maintain a static range, 0 through 100, for the Y-axis.
- **Scale per Chart:** Let the Y-axis adjust dynamically, based on the range of metric values for each chart in the view. The Y-axis scaling of one chart in a view does not affect the scaling of any other chart.

**Default:** Calculated

– **Number of Columns on Page**

Defines the number of columns for each page.

**Options:** 1, 2, 3

**Default:** 3

**Required:** Yes

– **Number of Charts on Page**

Defines the number of charts for each page.

**Default:** 15

**Limit:** 48

**Required:** No

– **Maximum Number of Charts**

Defines the maximum number of charts.

**Default:** 60

**Limit:** 1200

**Required:** No

– **Metric Filtering**

Specifies whether to filter metrics.

**Options:**

- **Enabled:** Metric are filtered.
- **Disabled:** Metric are not filtered.

**Default:** Enabled

– **Baseline Metrics**

Determines whether NetOps Portal changes utilization trends using the baseline trend line. The baseline data that NetOps Portal plots in many views shows statistical deviations from normal performance for a given statistic.

**Options:** Selected (enabled), Cleared (disabled)

**Options:**

- **Enabled:** Metric are filtered.
- **Disabled:** Metric are not filtered.

**Default:** Cleared (disabled)

For more information, see [Baseline Calculations](#).

3. (23.3.4 and higher) In the **Set Trend Goal/Threshold Line** section, complete the following fields:

– **Goal/Threshold Line**

Specifies whether to show or hide a goal/threshold line trend line in the view. Goal line trend lines serve as visual indicators of fixed metric values.

**Options:**

- **Disabled:** Do not show a goal or threshold line trend line in the view.
- **Enabled:** Show a goal or threshold line trend line in the view.

**Default:** Enabled

– **Goal Line**

(If you have decided to show a goal/threshold line trend line) Specifies the type or trend line to show in the view, and where NetOps Portal shows it.

**Options:**

- **Hide:** Do not show a goal/threshold line trend line in the view.
- **Show Goal Line:** Show a goal line trend line in the view and in the chart legend (golden color).
- **Show Threshold Line (Above):** Show a threshold line trend line in the view with shading above the line with pink offset of critical color. Show the line in trend charts and in the chart legend.
- **Show Threshold Line (Below):** Show a threshold line trend line in the view with shading below the line with pink offset of critical color. Show the line in trend charts and in the chart legend.
- **Show Threshold Line (Above and Below):** Show a threshold line trend line in the view with shading above the line with pink offset of critical color and shading below the line with green offset of normal color. Show the line in trend charts and in the chart legend.

**Default:** Show Goal Line

– **Goal Line Value**

Specifies a value for the goal line trend line. Goal line trend lines serve as visual indicators of fixed metric values. For example, select a value within a critical threshold range to determine whether performance is approaching a degraded level quickly.

**Default:** 40

– **Value Format**

Specifies whether NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.

**Options:**

- **Scaled:** NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.
- **Unscaled:** NetOps Portal does not scale the goal line or threshold line value. Instead, it uses the raw value on the trend line on the chart Y-axis.

**Default:** Scaled

– **Goal Line Label**

Specifies the label for the goal line trend line.

**Default:** Goal Line

4. In the **Customize Time Range** section, [filter data based on specific time periods by selecting a custom time range for the view](#).
5. In the **Apply Business Hours** section, [assign a business hours filter and a time zone to the view](#).  
Applying a business hours filter to a dynamic trend view displays the same data in the view, but the data with the periods *outside* of the applied business hours are shaded.
6. Click **Save**.

The MultiView is configured.

### **Applied Business Hours Filters**

When applying a business hours filter to a page as part of your view, the business hours override other business hour settings in the following priority order:

1. Business hours and time zone explicitly assigned to a view.
2. Site group that is associated with a business hours filter explicitly applied to a view.
3. Business hours filter applied to a dashboard or to a context page (user-session level business hours).
4. Site group that is associated with a business hours filter specified at the context level.

Applying a business hours filter to a trend view displays the same data in the view, but the data with the periods *outside* of the applied business hours are shaded.

For more information:

- About business hours filtering, including how to define business hours definitions, see [Configure Business Hours Filtering](#).
- About how to apply a business hours filter to a dashboard, see [Dashboards](#).
- About how to apply a business hours filter to a context page, see [Context Pages](#).

## **NetOps Business Reports**

When the NetOps Report Manager Service is installed and the NetOps business reports are enabled, you can run the reports in NetOps Portal.

(23.3.2 and higher)

### **IMPORTANT**

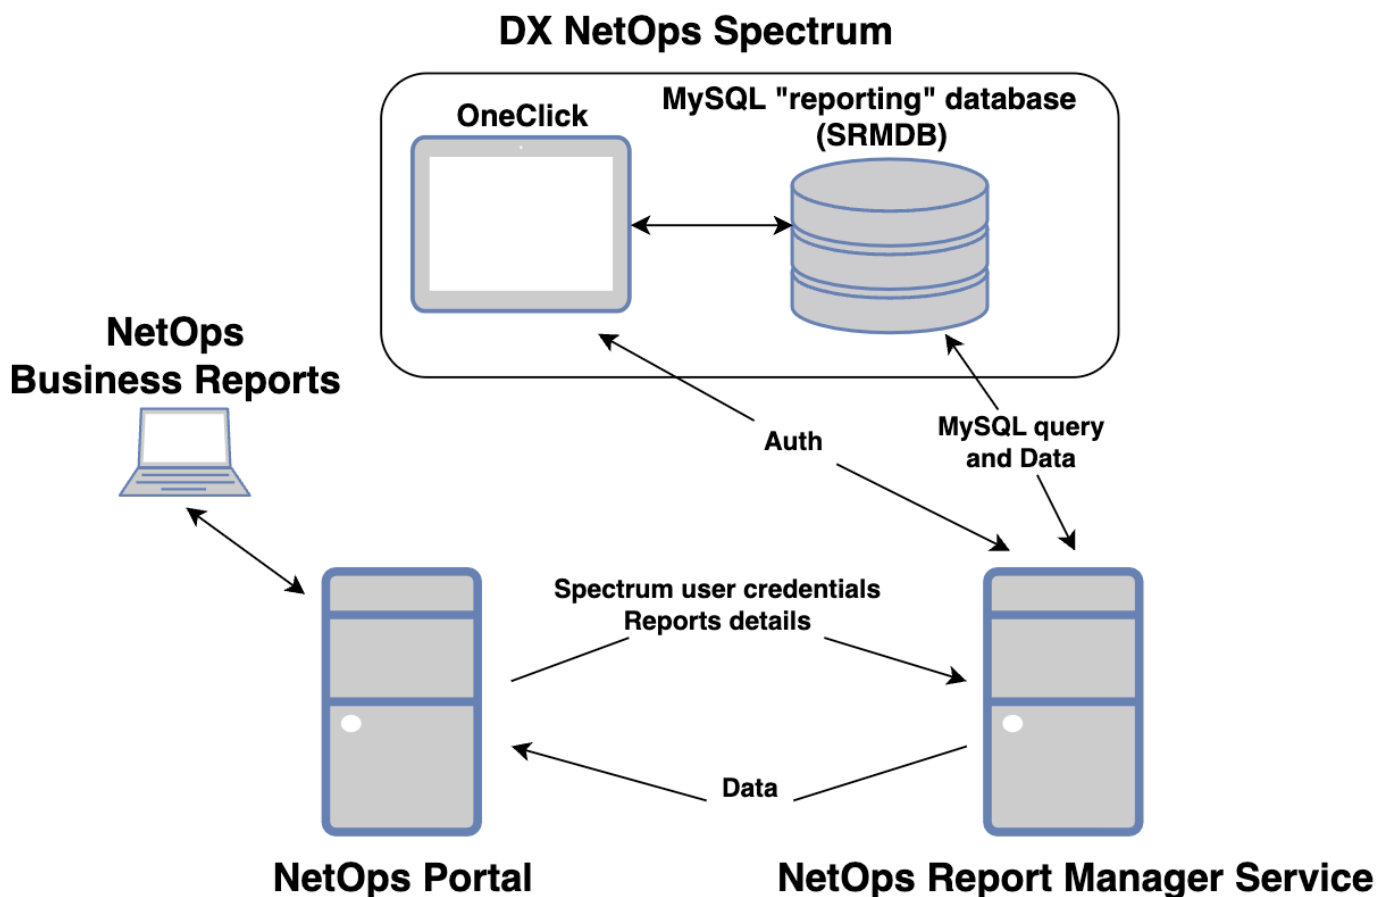
Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

NetOps business reports consist of the following primary components:

- **DX NetOps Spectrum**  
DX NetOps Spectrum (Spectrum) provides inventory and alarm data to the NetOps Report Manager Service, and authenticates the Report Manager Service.
- **The NetOps Report Manager Service**  
The NetOps Report Manager Service is a standalone application that processes requests from NetOps Portal. It retrieves data from Spectrum's MySQL "reporting" database (SRMDB) and forwards it to NetOps Portal. It authenticates against a configured Spectrum OneClick instance.  
For more information about the SRMDB, see [SpectroSERVER and DX NetOps Spectrum Databases Overview](#).
- **NetOps Portal**  
NetOps Portal sends requests to the NetOps Report Manager Service. You configure NetOps business reports and then run them using NetOps Portal.

The following image shows the system architecture for the NetOps Report Manager Service:

**Figure 78: NetOps Report Manager Service Architecture**



For more information:

- About how to install the NetOps Report Manager Service, see [Enable the NetOps Report Manager Service](#).
- About how to run NetOps business reports in NetOps Portal, see [Manage NetOps Business Reports](#).



## Run Business Reports

If you have integrated DX NetOps Spectrum (Spectrum) with CA Business Intelligence (CABI), you can use the following methods to run business reports:

- (Preferred) (23.3.2 and higher) [Run NetOps Business Reports in NetOps Portal](#)

### IMPORTANT

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

- [Run CABI Reports](#)

## Manage NetOps Business Reports

You can run NetOps business reports, edit the setting for NetOps business reports, and view a list of NetOps business reports.

(23.3.2 and higher)

### IMPORTANT

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

For more information, see [NetOps Business Reports](#).

You can manage the NetOps business reports in the following ways:

### NOTE

**Prerequisite:** [You have enabled the NetOps business reports.](#)

- [View a List of NetOps Business Reports](#)
- [Run a NetOps Business Report](#)
- [Edit the Settings for a NetOps Business Report](#)
- [Schedule a NetOps Business Report](#)
- (23.3.4 and higher) [View the Settings for a Scheduled NetOps Business Report](#)
- (23.3.4 and higher) [Edit the Settings for a Scheduled NetOps Business Report](#)
- [Print/Export a NetOps Business Report](#)
- [View a List of Archived Run Now NetOps Business Reports](#)
- [Delete an Archived Run Now NetOps Business Report](#)
- [Check the Report Manager Service Health Status](#)
- [Update the SRMDB Username and Password](#)

### View a List of NetOps Business Reports

The NetOps business reports that NetOps Portal includes display on the **Run NetOps Business Reports** page. To view this page, log in to NetOps Portal as a user with the Run NetOps Business Reports role right, hover over **Reports**, and then click **NetOps Business Reports**.

NetOps Portal includes the following NetOps business reports:

**NOTE**

These reports are based on Spectrum alarms, and include information from Spectrum's MySQL "reporting" database (SRMDB).

- **Alarm Log: All**  
Displays a list of alarms for devices and models in all Global Collection groups that are available for the logged in NetOps Portal user and that satisfy the criteria of the alarm condition and minimum duration threshold. The alarms are displayed in the list in chronological order and are grouped by alarm condition and landscape.
- **Alarm Log: Group**  
Displays a list of alarms for devices and models in a selected Global Collection group that satisfy criteria of alarm condition and minimum duration threshold. The alarms are displayed in the list in chronological order and grouped by alarm condition and landscape.
- **Top N Most Common Alarms: All**  
Displays a list of the most frequently occurring alarm types for devices and models in all Global Collection groups available for the logged in NetOps Portal user. Clicking an alarm report to get details for each of the alarms for a selected alarm type.
- **Top N Most Common Alarms: Group**  
Displays a list of the most frequently occurring alarm types for devices and models in a selected Global Collection group. Clicking an alarm report retrieves the details for each of the alarms for a selected alarm type.
- **Top N Devices and Models with Most Alarms: All**  
Displays a list of managed devices and models in all Global Collection groups available for the logged in NetOps Portal user with the most number of alarms. Clicking an alarm sub-report retrieves the details for each of the alarms for a selected device or model.
- **Top N Devices and Models with Most Alarms: Group**  
Displays a list of managed devices and models in a selected Global Collection group with the most number of alarms. Clicking an alarm sub-report retrieves the details for each of the alarms for a selected device or model.

**Run a NetOps Business Report****Follow these steps:**

1. From the **Run NetOps Business Reports** page, launch a specific report by clicking the name of the report (link).  
The report configuration dialog opens.

**NOTE**

The following fields are an example for the **Alarm Log: All** NetOps business report.

2. In the **Settings** section, complete the following fields:
  - **Title**  
Identifies the name of the dashboard (the NetOps business report).  
**Required:** Yes
  - **Name**  
Specifies the name assigned to the report so that you can find it and run it again.  
**Required:** Yes
  - **Description**  
Describes the report.  
**Required:** Yes
  - **Use Alarm Types**  
Specifies whether to use selected alarm types or to show all.  
**Default:** No, Show All  
**Required:** No
  - **Use Alarm Cause Codes**  
Specifies whether to use selected alarm cause codes or to show all.  
**Default:** No, Show All

**Required:** No

– **Alarm Condition**

Specifies the alarm condition that this report shows. Move the alarm condition that you want shown in the report from the **Alarm Condition Available** list to the **Alarm Condition Selected** list.

**Required:** No

– **Minimum Duration Time (sec)**

Specifies the minimum duration time that this report shows.

**Default:** 300

**Required:** No

– **Results Limit (Max Rows)**

Specifies the total number of rows that the table can display. If the number of results exceeds the max rows, the table shows the top metrics according to the sorted column.

**Default:** 1000

**Required:** No

– **Show Chart**

Specifies whether to show a chart in the NetOps business report.

**Options:**

- **Enabled:** A chart is shown in the NetOps business report.
- **Disabled:** A chart is not shown in the NetOps business report.

**Default:** Enabled

**Required:** No

3. In the **Customize Time Range** section, complete the following fields:

– **Custom Time Range**

Defines whether to run the report on a custom time range.

**Options:**

- **Enabled:** Run the report on a custom time range. You must define a time range in the **Time Range** field.
- **Disabled:** Run the report for all times.

**Default:** Disabled

– **Time Range**

When the custom time range is enabled, defines the custom time range.

**Options:** Last Hour, Last 4 Hours, Last 8 Hours, Last 12 Hours, Last 24 Hours, Last 7 Days, Last 14 Days, Last 30 Days, Last 3 Months, Last 12 Months, Yesterday, Previous Week, Previous Month, Today, Current Week, Current Month

**Default:** Last Hour

4. In the **Global Collections to Include** section, to add or remove global collections to the view, do the following steps:

a. Click **Add / Remove GC**.

The **Add/Remove Global Collections** dialog appears.

b. Move the global collection that you want to add to the view from the **Available** list to the **Selected** list.

c. Click **OK**.

5. To determine which users see view changes, from the **Apply Changes** drop-down, [set the scope when changing the view settings](#).

6. Click **Save**.

The NetOps business report is run.

## **Edit the Settings for a NetOps Business Report**

**Prerequisite:** The NetOps business report for which you want to edit the setting is run.

**Follow these steps:**

1. From the ran NetOps business report, edit the settings of the report by clicking the **Edit the settings of this report template** icon (the Pencil icon).  
The report configuration dialog opens.
2. [Edit the settings for the report](#), and then click **Save**.

Your changes to the settings are saved.

**Schedule a NetOps Business Report**

You schedule a NetOps business report by saving dashboard data (the NetOps business report) as reports. These reports are attached to email messages and are saved as scheduled reports.

(23.3.4 and higher) You can schedule a NetOps business report with different settings by [scheduling the NetOps business report](#), ensuring that the **Scheduled Copy** checkbox is selected, and then scheduling the report again with different settings.

For more information about how to schedule a NetOps business report, including information about this checkbox, see [Manage Scheduled Reports](#).

**View the Settings for a Scheduled NetOps Business Report**

(23.3.4 and higher)

**Prerequisites:**

- You have scheduled the NetOps business report for which you want to view settings, and you have selected the **Scheduled Copy** checkbox.
- You are the current (active) user of the scheduled NetOps business report.

**Follow these steps:**

1. Hover over **Administration, Configuration Settings**, and then click **All Scheduled Reports**.  
The **Manage Scheduled Reports** page opens, and a list of scheduled reports displays.
2. Run the NetOps business report for which you want to view settings by clicking the name of the report (link) in the **Dashboard** column.  
The NetOps business report is run.
3. Select the NetOps business report for which you want to view settings, and then click the **Edit the settings of this report template** icon (the Pencil icon).  
The report configuration dialog opens.
4. [View the settings for the report](#), and then click **Save**.

**Edit the Settings for a Scheduled NetOps Business Report**

(23.3.4 and higher)

**Prerequisites:**

- You have scheduled the NetOps business report for which you want to edit settings, and you have selected the **Scheduled Copy** checkbox.
- You are the current (active) user of the scheduled NetOps business report.

In the report configuration dialog, edit the settings for the report, and then click **Save**. Your changes to the settings are saved. These settings are applied to the newly-scheduled NetOps business report.

## **Print/Export a NetOps Business Report**

You print/export NetOps business reports by downloading the dashboard data (the NetOps business report) as a report in PDF or comma-separated values (CSV) format or preview the dashboard data in PDF or CSV format.

For more information, see [Download Dashboard Data as Reports](#).

## **View a List of Archived Run Now NetOps Business Reports**

### **Prerequisites:**

- You have defined to have NetOps Portal run the NetOps business report, and for the email server to send dashboard data as reports from schedule reports attached to emails, immediately (a *Run Now* report). (You have selected **Run Now** as the **Frequency** at which NetOps Portal saves dashboard data as schedule reports.)  
For more information about how to select the frequency, see [Manage Scheduled Reports](#).
- You have defined to have NetOps Portal archive the report (archive reports is enabled).  
For more information about how to enable archive reports, see [Configure the Email Server](#).
- You are logged in as a user with the Send Reports to Archive role right.

Run Now reports are one-time-only scheduled reports. You can view these archived reports from the **Archived Run Now Reports** dialog.

For more information about how to view these reports, see [Manage Scheduled Reports](#).

## **Delete an Archived Run Now NetOps Business Report**

### **Prerequisites:**

- You have defined to have NetOps Portal run the NetOps business report, and for the email server to send dashboard data as reports from schedule reports attached to emails, immediately (a *Run Now* report). (You have selected **Run Now** as the **Frequency** at which NetOps Portal saves dashboard data as schedule reports.)
- You have defined to have NetOps Portal archives the report (archive reports is enabled).  
For more information about how to enable archive reports, see [Configure the Email Server](#).
- You are logged in as a user with the Send Reports to Archive role right.

You can delete archived Run Now reports from the **Archived Run Now Reports** dialog.

For more information, see [Manage Scheduled Reports](#).

## **Check the Report Manager Service Health Status**

You can view the overall system status of the NetOps Report Manager Service on the **System Status** page.

For more information, see [View System Status](#).

## **Update the SRMDB Username and Password**

For more information about how to upgrade Spectrum's MySQL "reporting" database (SRMDB), see [Upgrade the NetOps Report Manager Service](#).

# **Manage On-Demand Reports**

You can view a static data set from a narrow context by creating, or generating, an on-demand report. Use on-demand reports to investigate and troubleshoot issues.

You manage on-demand reports by adding them to a dashboard, generating them, and by editing them. On-demand reports retrieve data sets from specific sets of items or groups without building a dashboard. To view the same data in different ways, change the view type of the on-demand report.

In this article:

- [Add an On-Demand Trend Report to a Dashboard](#)
- [Generate an On-Demand Report](#)
- [Download Generated On-Demand Report Data](#)
- [Edit an On-Demand Report](#)
- [On-Demand Report View Types](#)
- [Applied Business Hours Filter to Trend Views](#)

## **Add an On-Demand Trend Report to a Dashboard**

You add on-demand trend reports to dashboards as views.

### **Follow these steps:**

1. [Add the on-demand trend report as a view to the dashboard.](#)
2. Complete the following fields:
  - **Title**  
Defines the title for the on-demand trend report.  
**Required:** Yes
  - **View Type**  
Defines the view type.
  - **Resolution**  
Defines the resolution frequency.  
**Options:**
    - **1 Minute**
    - **5 Minute**
    - **10 Minute**
    - **15 Minute****Default:** Use default resolution (5 Minute)
3. (23.3.4 and higher) In the **Set Trend Goal/Threshold Line** section, complete the following fields:
  - **Goal/Threshold Line**  
Specifies whether to show or hide a goal/threshold line trend line in the view. Goal line trend lines serve as visual indicators of fixed metric values.  
**Options:**
    - **Disabled:** You do not want to show a goal or threshold line trend line in the view.
    - **Enabled:** You want to show a goal or threshold line trend line in the view.**Default:** Disabled
  - **Goal Line**  
(If you have decided to show a goal/threshold line trend line) Specifies the type or trend line to show in the view, and where NetOps Portal shows it.  
**Options:**
    - **Hide:** A goal/threshold line trend line is hidden from the view.
    - **Show Goal Line:** A goal line trend line is shown in the view and in the chart legend (golden color).
    - **Show Threshold Line (Above):** A threshold line trend line is shown in the view with shading above the line with pink offset of critical color. The line is shown in trend charts and in the chart legend.
    - **Show Threshold Line (Below):** A threshold line trend line is shown in the view with shading below the line with pink offset of critical color. The line is shown in trend charts and in the chart legend.
    - **Show Threshold Line (Above and Below):** A threshold line trend line is shown in the view with shading above the line with pink offset of critical color and shading below the line with green offset of normal color. The line is shown in trend charts and in the chart legend.

**Default:** Hide

– **Goal Line Value**

Specifies a value for the goal line trend line. Goal line trend lines serve as visual indicators of fixed metric values. For example, select a value within a critical threshold range to determine whether performance is approaching a degraded level quickly.

**Default:** 40

– **Value Format**

Specifies whether NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.

**Options:**

- **Scaled:** NetOps Portal scales the goal line or threshold line value along with the metric values on the trend line on the chart Y-axis.
- **Unscaled:** NetOps Portal does not scale the goal line or threshold line value. Instead, it uses the raw value on the trend line on the chart Y-axis.

**Default:** Scaled

– **Goal Line Label**

Specifies the label for the goal line trend line.

**Default:** Goal Line

4. In the **Customize Time Range** section, [filter data based on specific time periods by selecting a custom time range for the view](#).

5. In the **Apply Business Hours** section, [assign a business hours filter and a time zone to the view](#).

Applying a business hours filter to a dynamic trend view displays the same data in the view, but the data with the periods *outside* of the applied business hours are shaded.

Select whether to **Compress Non Business Hours** by selecting or clearing the **Enabled** checkbox.

**Options:**

- **Selected:** Non-business hours are compressed.
- **Cleared:** Non-business hours are not compressed.

**Default:** Cleared (Disabled)

For more information, see [the "Applied Business Hours Filter to Trend Views" section](#).

6. (If you are adding this view to a dashboard) In the **Items or Groups to Include** section, complete the following field:

– **Context**

Defines the context of the view.

**Options:**

- **Dynamic:** Indicates that the context of the view changes with the context of the dashboard. With this option, the page context is honored for site groups that are associated with a business hours definition.
- **Fixed:** Indicates that the view uses a specified group, device, or component as a context for the data. With this option, the page context is honored only if the view is configured with a single site group that is associated with a business hours definition.

**Default:** Dynamic

The on-demand trend report is added to the dashboard.

### **Generate an On-Demand Report**

Generating an on-demand report launches the report. You can report on different device types and components in a single on-demand report. For example, you can include routers and servers in the same report.

#### **Follow these steps:**

1. From an inventory page, select an item, and then click **On Demand**.  
The **IM On-Demand/Multi-Metric Trend Report** dialog opens.

2. Complete the fields. Select managed items and metric families that complement each other (the **Metrics to Include** field), set the scope for the view settings (the **Apply Changes** field), and then click **Run**.

For more information:

- About the fields, see [Manage On-Demand Report Templates](#).
- About how to set the scope for views, see [Customize Views](#).

The on-demand report is generated.

### **Download Generated On-Demand Report Data**

After you have generated the on-demand report, you can download the data as a report in PDF or comma-separated values (CSV) format.

#### **NOTE**

- On-demand reports reflect interface data, but do not reflect rollups to the router level.
- On-demand reports in PDF format of chart/table views do not show the gauge or chart that is associated with the currently-selected row in the table.

For more information, see [Download Dashboard Data as Reports](#).

### **Edit an On-Demand Report**

You can edit on-demand reports while editing a dashboard that includes this view (from the **Edit Dashboard** page) or by editing the view from the dashboard (the **IM On-Demand/Multi-Metric Trend Report** dialog).

#### **NOTE**

Changing the item context can clear the original selections. For example, if you select three routers for reporting, and then add interfaces, the routers are cleared.

For more information:

- About how to edit dashboards, see [Manage Dashboards](#).
- About how to edit views, see [Customize Views](#).

### **On-Demand Report View Types**

On-demand report view types determine the chart format. For options that require aggregation, the system determines the aggregation method. For example, Bits has the counter metric profile type and uses sum. The Rate metric is profiled as type gauge and uses average.

Many view types provide more configuration options, such as Number of Charts on Page and Maximum Number of Charts.

#### **TIP**

To change the visualization of the data, edit the on-demand report and change the **View Type**. The summary table shows only aggregate metrics. If you change to this view type, the view editor does not remove selected non-aggregate metrics. However, the view shows only the aggregate values. For example, you can configure a Chart with Multiple Metrics to show the CPU Utilization maximum. When you change the view type to the summary table, the view shows the CPU Utilization average.

#### **NOTE**

On-demand reports do not show data for the SD-WAN Tunnel metric family.

The following is list of the view types that you can define for an on-demand report:



- [Chart with Multiple Items and Metrics](#)
- [Chart with Multiple Metrics](#)
- [Chart per Metric](#)
- [Chart per Item with Multiple Metrics](#)
- [Chart per Metric with Multiple Items](#)
- [Chart per Metric by Single Item](#)
- [Table with Multiple Metrics](#)
- [Table per Item with Multiple Metrics](#)
- [Date Table with Multiple Metrics](#)
- [Date Table per Item with Multiple Metrics](#)
- [Summary Table/Chart by Metrics](#)
- [Summary Table/Chart by Items](#)

For more information about how to define the view type for an on-demand report, see [Manage On-Demand Report Templates](#).

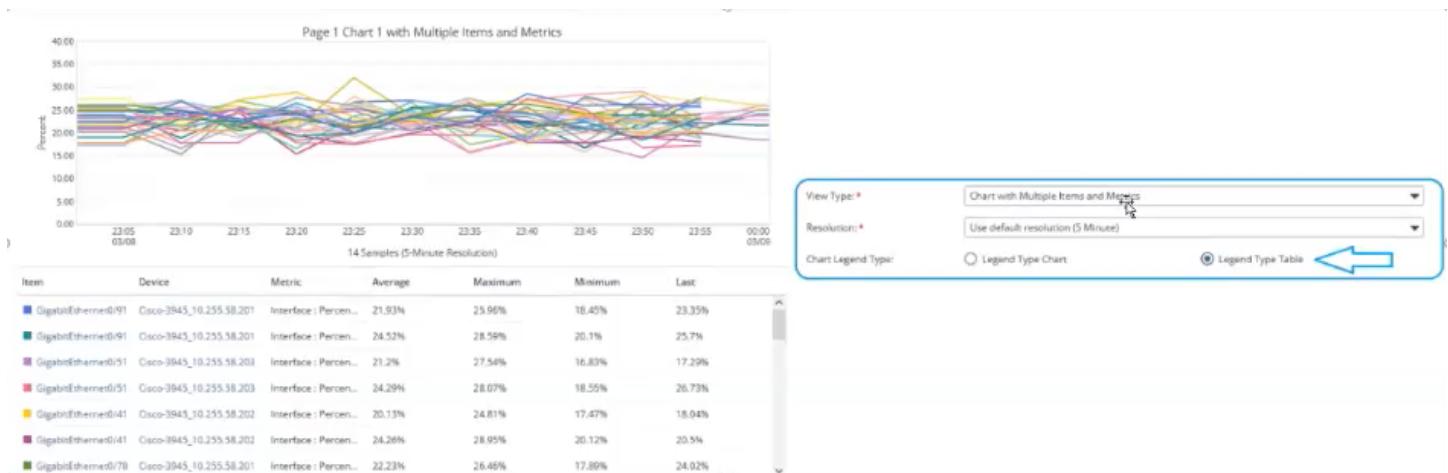
### **Chart with Multiple Items and Metrics**

The **Chart with Multiple Items and Metrics** on-demand report view type consists of a trend chart that consolidates items and metrics for a device on the same chart. This report view displays a trend line for all items and metrics, which is an aggregate for all the selected groups or items.

Below the trend chart, the legend table reports the average, maximum, minimum, and last reported values for each trend line contained in the chart.

The following image shows an example of this on-demand report view type:


**Figure 79: ChartWithMultItemsAndMet**



For more information about trend charts, see [Trend Views](#).

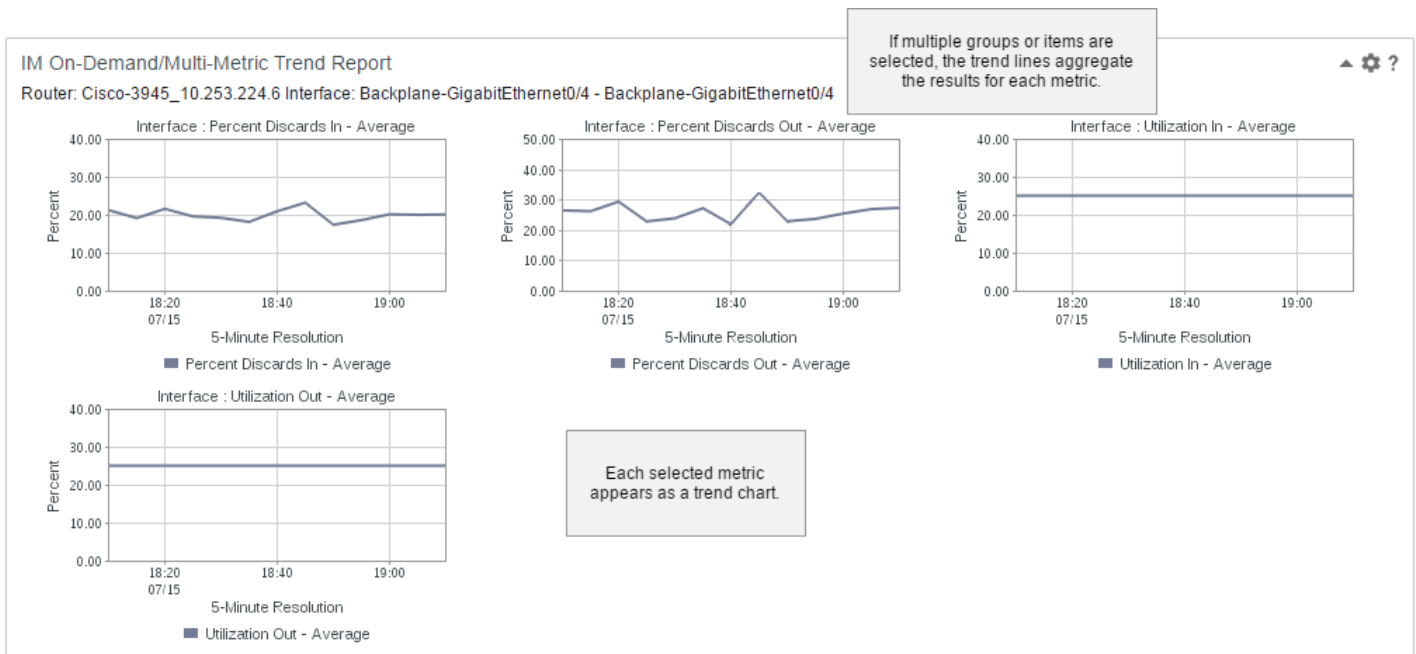
### **Chart with Multiple Metrics**

The **Chart with Multiple Metrics** on-demand report view type consists of a chart that displays a trend line for each selected metric. The trend line is an aggregate for all the selected groups or items. The following image shows an example of this on-demand report view type:

**Figure 80: ChartWithMultMet****Chart per Metric**

The **Chart per Metric** on-demand report view type consists of charts for each selected metric. Each chart displays a trend line for the metric. The trend line is an aggregate for all the selected groups or items.

The following image shows an example of this on-demand report view type:

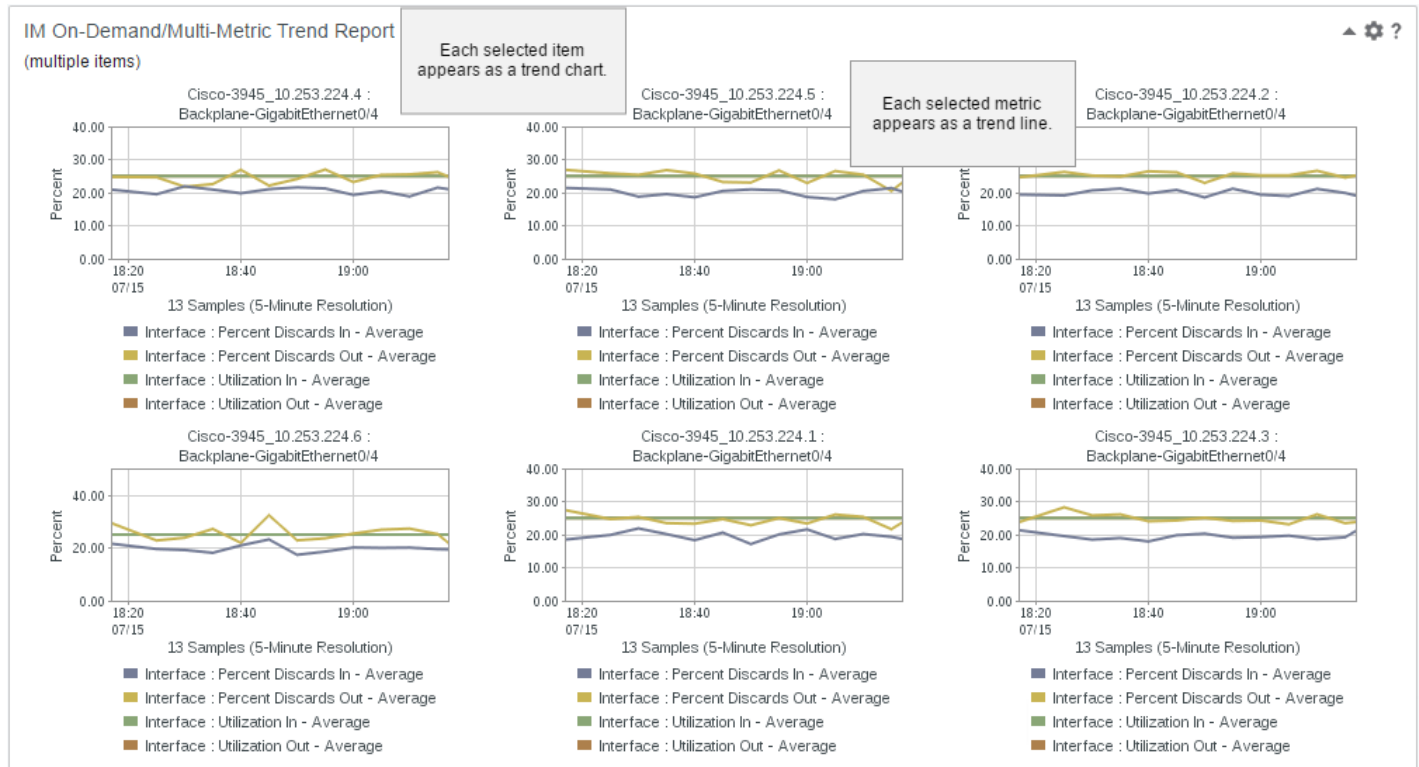
**Figure 81: ChartPerMet**

## Chart per Item with Multiple Metrics

The **Chart with Item with Multiple Metrics** on-demand report view type consists of charts for each selected item or group. Each chart displays trend lines for each selected metric.

The following image shows an example of this on-demand report view type:


**Figure 82: ChartPerItemWithMultMet**



## Chart per Metric with Multiple Items

The **Chart per Metric with Multiple Items** on-demand report view type consists of charts for each selected metric. Each chart displays trend lines for each selected item or group.

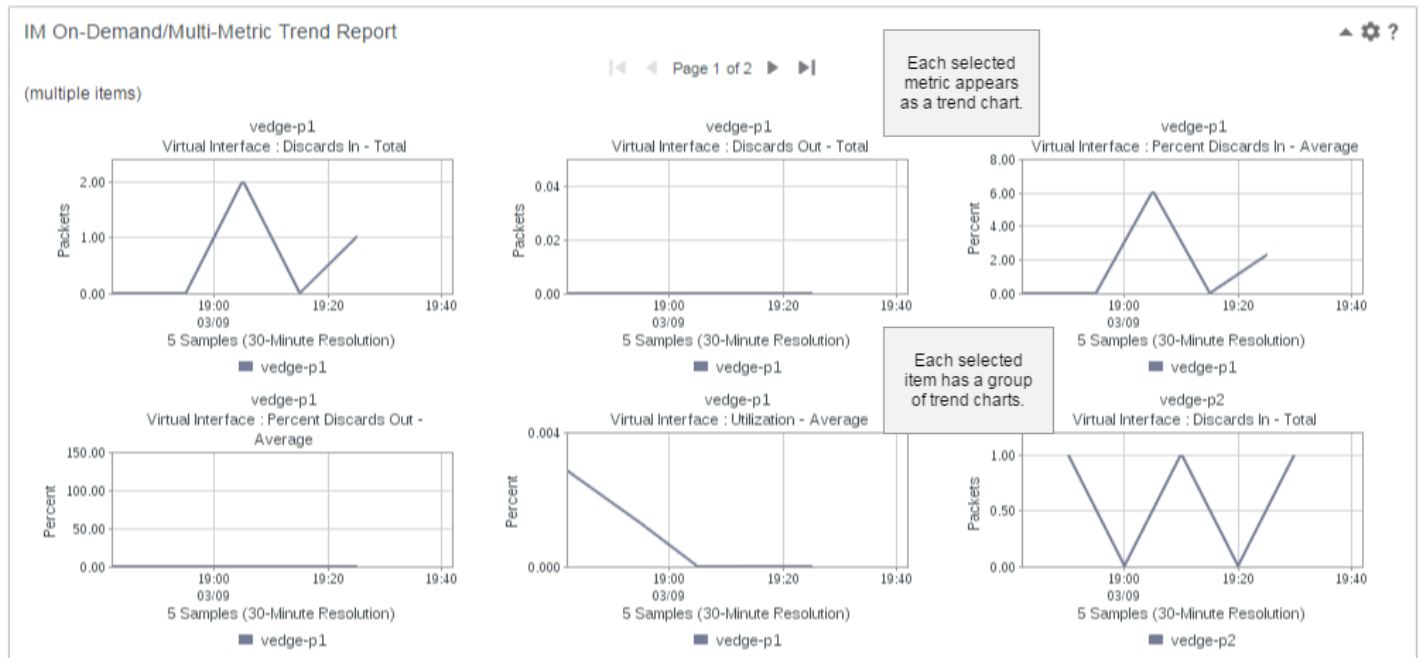
The following image shows an example of this on-demand report view type:

**Figure 83: ChartPerMetWithMultiItems**

### Chart per Metric by Single Item

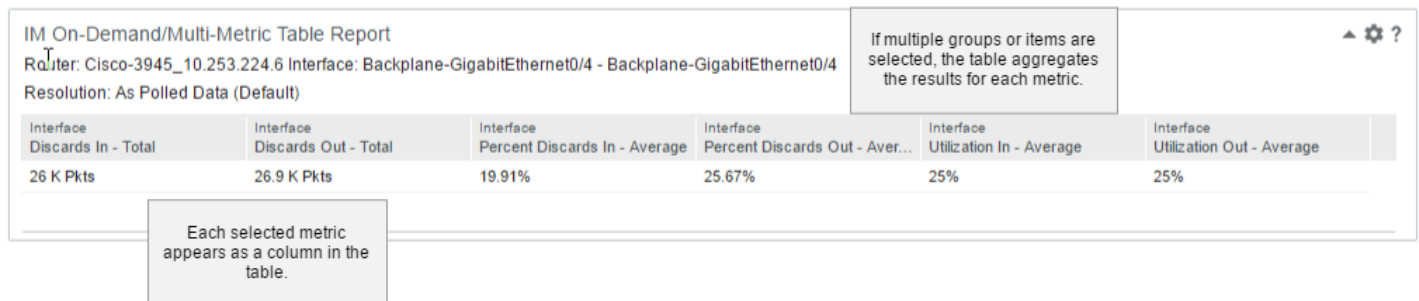
The **Chart with Metric by Single Item** on-demand report view type consists of chart groupings for each selected item or group. Each chart grouping consists of a chart for each selected metric.

The following image shows an example of this on-demand report view type:

**Figure 84: ChartperMetricbySingleItemReport****Table with Multiple Metrics**

The **Table with Multiple Metrics** on-demand report view type consists of a table that displays a list of the selected metrics. The table shows aggregate values for the selected groups or items.

The following image shows an example of this on-demand report view type:

**Figure 85: TableWithMultMet****NOTE**

This on-demand report view type does *not* apply business hours filters.

**Table per Item with Multiple Metrics**

The **Table per Item with Multiple Metrics** on-demand report view type consists of a table that displays a list of the selected metrics. Each row in this table represents a single item.

The following image shows an example of this on-demand report view type:

**Figure 86: TablePerItemWithMultMet**

IM On-Demand/Multi-Metric Table Report  
(multiple items)  
Resolution: As Polled Data (Default)

Each selected metric appears as a column in the table.

Device Name	Name	Interface Discards In - Total	Interface Discards Out - Total	Interface Percent Discards In ...	Interface Percent Discards O...	Interface Utilization In - Average	Interface Utilization Out - Aver...
Cisco-3945_10.253...	Backplane-GigabitEt...	23.9 K Pkts	25.2 K Pkts	20%	26.05%	25%	25%
Cisco-3945_10.253...	Backplane-GigabitEt...	34.9 K Pkts	33.7 K Pkts	19.87%	24.32%	25%	25%
Cisco-3945_10.253...	Backplane-GigabitEt...	35.9 K Pkts	35.4 K Pkts	20.17%	24.71%	25%	25%
Cisco-3945_10.253...	Backplane-GigabitEt...	34.4 K Pkts	35.9 K Pkts	19.73%	25.43%	25%	25%
Cisco-3945_10.253...	Backplane-GigabitEt...	34.3 K Pkts	34.3 K Pkts	19.25%	24.68%	25%	25%
Cisco-3945_10.253...	Backplane-GigabitEt...	35.5 K Pkts	35.7 K Pkts	19.69%	24.68%	25%	25%

Search  Page 1 of 1 Displaying 1 - 6 of 6 Max Per Page 10

Each selected item appears as a row in the table.

**NOTE**

This on-demand report view type does *not* apply business hours filters.

**Date Table with Multiple Metrics**

The **Date Table with Multiple Metrics** on-demand report view type consists of a table that aggregates metric values by time intervals.

The following image shows an example of this on-demand report view type:

**Figure 87: DateTableWithMultMet**

Date Table with Multiple Metrics  
(multiple items)  
View Type: Date Table with Multiple Metrics  
Resolution: As Polled Data (Default) Time Interval: 5 Minute Metric Calculate Level: by Group [Time Zone: EDT]

Quick Filter

End Time	CPU Utilization - Average	CPU Utilization - Baseline Average	Memory Utilization - Average	Memory Utilization - Baseline Average
September 7, 2022 1:05:00 PM	34%	49.66%	54.71%	48.42%
September 7, 2022 1:10:00 PM	9%	49.66%	45.21%	48.42%
September 7, 2022 1:15:00 PM	19%	49.66%	50.61%	48.42%
September 7, 2022 1:20:00 PM	25.67%	49.66%	62.24%	48.42%
September 7, 2022 1:25:00 PM	34%	49.66%	70.05%	48.42%
September 7, 2022 1:30:00 PM	42.67%	49.66%	63.03%	48.42%
September 7, 2022 1:35:00 PM	50.67%	49.66%	59.44%	48.42%
September 7, 2022 1:40:00 PM	59%	49.66%	58.06%	48.42%
September 7, 2022 1:45:00 PM	67.67%	49.66%	57.21%	48.42%
September 7, 2022 1:50:00 PM	75.67%	49.66%	50.21%	48.42%

10 per page Page 1 of 2 Displaying 1 - 10 of 13

**Date Table per Item with Multiple Metrics**

The **Date Table per Item with Multiple Metrics** on-demand report view type consists of a table that aggregates group/device/component metric values by time intervals. This view adds support for time series tabular reporting.

The following image shows an example of this on-demand report view type:

**Figure 88: DataTablePerItemWithMultMet**

**Date Table per Item with Multiple Metrics**

ID: 1000128, Type: RIBTableDateModel - Expand for query details

(multiple items)

View Type: Date Table per Item with Multiple Metrics  
Resolution: As Polled Data (Default) Time Interval: 5 Minute Metric Calculate Level: by Device [Time Zone: EDT]

Quick Filter

End Time	Device Name	CPU Utilization - Average	CPU Utilization - Baseline Average	Memory Utilization - Average	Memory Utilization - Baseline Average
September 7, 2022 1:05:00 PM	Cisco-3945_10.35.58.201	98%	49.66%	35.87%	48.73%
September 7, 2022 1:05:00 PM	Cisco-3945_10.35.58.202	1%	49.12%	79.17%	49.36%
September 7, 2022 1:05:00 PM	Cisco-3945_10.35.58.203	3%	50.21%	49.09%	47.16%
September 7, 2022 1:10:00 PM	Cisco-3945_10.35.58.201	6%	49.66%	20.17%	48.73%
September 7, 2022 1:10:00 PM	Cisco-3945_10.35.58.202	10%	49.12%	70.83%	49.36%
September 7, 2022 1:10:00 PM	Cisco-3945_10.35.58.203	11%	50.21%	44.64%	47.16%
September 7, 2022 1:15:00 PM	Cisco-3945_10.35.58.201	19%	49.66%	41.67%	48.73%
September 7, 2022 1:15:00 PM	Cisco-3945_10.35.58.202	18%	49.12%	62.5%	49.36%
September 7, 2022 1:15:00 PM	Cisco-3945_10.35.58.203	20%	50.21%	47.67%	47.16%
September 7, 2022 1:20:00 PM	Cisco-3945_10.35.58.201	23%	49.66%	66.67%	48.73%
September 7, 2022 1:20:00 PM	Cisco-3945_10.35.58.202	26%	49.12%	57.56%	49.36%
September 7, 2022 1:20:00 PM	Cisco-3945_10.35.58.203	28%	50.21%	62.5%	47.16%
September 7, 2022 1:25:00 PM	Cisco-3945_10.35.58.201	31%	49.66%	81.57%	48.73%
September 7, 2022 1:25:00 PM	Cisco-3945_10.35.58.202	35%	49.12%	55.36%	49.36%
September 7, 2022 1:25:00 PM	Cisco-3945_10.35.58.203	36%	50.21%	73.21%	47.16%
September 7, 2022 1:30:00 PM	Cisco-3945_10.35.58.201	40%	49.66%	78.26%	48.73%
September 7, 2022 1:30:00 PM	Cisco-3945_10.35.58.202	43%	49.12%	57.5%	49.36%
September 7, 2022 1:30:00 PM	Cisco-3945_10.35.58.203	45%	50.21%	53.33%	47.16%
September 7, 2022 1:35:00 PM	Cisco-3945_10.35.58.201	48%	49.66%	67.45%	48.73%
September 7, 2022 1:35:00 PM	Cisco-3945_10.35.58.202	51%	49.12%	62.5%	49.36%

250 per page

Page 1 of 1


Displaying 1 - 39 of 39

### Summary Table/Chart by Metrics

The **Summary Table/Chart by Metrics** on-demand report view type consists of a trend graph that represents one or more items and multiple metrics. The trend graph aggregates the items for each selected metric.

Below the trend graph, the legend table is sorted by metric. The custom percentile values that you have defined for the selected metrics also appear in the legend table.

The following image shows an example of this on-demand report view type:

**Figure 89: SummaryTableChartByMet****NOTE**

This on-demand report view type does *not* apply business hours filters.


**Summary Table/Chart by Items**

The **Summary Table/Chart by Items** on-demand report view type consists of a trend graph that represents multiple items and a single metric. The trend graph aggregates the items for each selected metric.

Below the trend graph, the legend table is sorted by item. The custom percentile values that you have defined for the selected metric also appear in the legend table.

The following image shows an example of this on-demand report view type:



**Figure 90: SummaryTableChartByItems****NOTE**

This on-demand report view type does *not* [apply business hours filters](#).

**Applied Business Hours Filter to Trend Views**

Applying a business hours filter to a trend view in an on-demand report displays the same data in the view, but the data with the periods *outside* of the applied business hours are shaded.

**NOTE**

The on-demand report view types that do not apply business hours filters are noted.

For more information:

- About business hours filtering, including how to define business hours definitions, see [Configure Business Hours Filtering](#).
- About how to apply a business hours filter to a dashboard, see [Dashboards](#).
- About how to apply a business hours filter to a context page, see [Context Pages](#).
- About trend views, see [Trend Views](#).

**Manage On-Demand Report Templates**

View a static data set from a narrow context by creating, or generating, an on-demand report template.

On-demand report templates are reusable report definitions. You can manage the on-demand report templates in your tenant by viewing them, creating/generating them, editing them, copying them, and deleting them.

**Prerequisite:** Before you can manage the on-demand report templates, ensure that [your user account has the appropriate role rights for managing on-demand report templates](#).

In this article:

- [View a List of On-Demand Report Templates](#)
- [Create or Edit an On-Demand Report Template](#)
- [Apply an On-Demand Report Template to Dashboards, Group Context Pages, and Device Context Pages](#)
- [Copy an On-Demand Report Template](#)
- [Run an On-Demand Report Template](#)

## **View a List of On-Demand Report Templates**

### **Follow these steps:**

1. Log in as a user with the Run On-Demand Report Templates role right.
2. Hover over **Reports**, and then click **On-Demand Report Templates**.  
The **Manage On-Demand Report Templates** page opens, and the current list of on-demand report templates appears. Users see only the report templates that are associated with the current tenant. The global administrator sees only reports that are associated with the Default Tenant.
3. (Optional) To filter the list to show only your on-demand report templates, click **My Reports**.  
The **My Reports** page appears. You can manage on-demand reports in your tenant from this page.

## **Create or Edit an On-Demand Report Template**

You can report on different device types and components in a single on-demand report template. For example, you can include routers and servers in the same report. Select managed items and metric families that complement each other.

### **Follow these steps:**

1. Do one of the following from the **On-Demand Report Templates** page:
  - (If you want to create an on-demand report template and you have the Create On-Demand Report Templates and Save On-Demand Report Templates for All Users role rights) Create an on-demand report template by clicking **New**.
  - (If you want to edit an on-demand report template that you created or own) Edit the on-demand report template by selecting it, and then click **Edit**.

#### **NOTE**

To edit an on-demand report that another user or administrator has created and owns, copy the on-demand report template, and then edit it.

For more information, see [the "Copy an On-Demand Report Template" section](#).

The **IM On-Demand/Multi-Metric Trend Report** dialog opens.

2. Specify the following fields:
  - **Title**  
The title appears on the view and in the on-demand report.
  - **Name**  
The name that identifies the report in the **On-Demand Report Templates** list and that appears as a title for the on-demand report.
  - **Description**
  - **View Type**  
Defines the view type for the on-demand report template.

**Options:** Chart with Multiple Metrics, Chart per Metric, Chart per Item with Multiple Metrics, Chart per Metric with Multiple Items, Chart per Metric by Single Item, Table with Multiple Metrics, Table per Item with Multiple Metrics,

Summary Table/Chart by Metrics, Summary Table/Chart by Items, Date Table with Multiple Metrics, Date Table per Item with Multiple Metrics, Chart with Multiple Items and Metrics

**Default:** Chart with Multiple Metrics

For examples of these view types, see [Manage On-Demand Reports](#).

– **Resolution**

Defines the resolution frequency.

**Options:** Use default resolution (5 Minute), 1 Minute, 5 Minute, 10 Minute, 15 Minute, 30 Minute, 1 Hour, 1 Day

**Default:** Use default resolution (5 Minute)

For more information, see [Customize Views](#).

– (Trend charts) **Chart Legend Type**

Defines how the information about the trend chart is displayed below the chart, either in a table or chart legend.

**Options:**

- **Legend Type Chart:** The trend chart information is displayed in a chart legend.
- **Legend Type Table:** The trend chart information is displayed in a table legend.

**Default:** Legend Type Chart

– (Devices) **Baseline Metrics**

To add baseline metrics to the metric value drop-down list or item selector, enable this option. Baseline data helps characterize past performance for the selected monitored parameters, assess present performance, and estimate future performance.

**Default:** Disabled

– (If you are generating a report by adding the **IM On-Demand/Multi-Metric Trend Report** to a dashboard or context page) **Context**

Defines the context for the on-demand report template. Set to one of the following:

- **Dynamic:** A dynamic context indicates that the context of the view changes with the context of the dashboard. With this option, the page context is honored for site groups that are associated with a business hours definition.
- **Fixed:** A fixed context indicates that the view uses a specified group, device, or component as a context for the data. With this option, the page context is honored only if the view is configured with a single site group that is associated with a business hours definition.

**Default:** Fixed

– (If the view supports business hours) In the **Apply Business Hours** section, complete the following fields:

• **Current Context Business Hours**

(If business hours have been applied to the context page) Defines the current context business hours. Otherwise, `No business hours selected` displays.

• **Apply Business Hour Filter**

By default, the view inherits the business hours filter applied to context pages (user-session level business hours). Enable this option to override this filter and assign a business hours filter and time zone to the view.

**Default:** Disabled

• **Compress Non Business Hours**

Defines whether to compress the non-business hours.

**Default:** Disabled

– (If you selected **Fixed** as the context for the on-demand report template) For **Items or Groups to Include**, complete one of the following:

- Select **Add Items**, click **Add / Remove Items**, select a **Context Type**, and then add up to 15 items.
- Select **Add Groups**, click **Add / Remove Groups**, and then add up to 15 groups.

**NOTE**

- Only the groups and items that are included in your permission groups are displayed.
- You can include any group in a report *except* the ones that are in the **Custom Collections** group.

– (If you selected Chart per Item with Multiple Metrics, Chart per Metric with Multiple Items, Table per Item with Multiple Metrics, or Summary Table/Chart by Items as the **View Type**) **Metric Calculate Level**

Determines the level of aggregation for metrics. Set this value to one of the following:

- by Group
- by Device
- by Component

**Default:** by Device

– **Metrics to Include**

Expand from the list of available metric families, and select up to 15 metrics to display in the report. You can select metrics from multiple metric families. You can use the arrows to specify an order for the selected metrics and the order is saved.

**NOTE**

For SD-WAN metric families (Tunnels and Application Paths), percentile and projection metrics are available, but are unsupported and the views render incorrectly.

For more information about this known issue, see [Known Issues](#).

The **Metrics to Include** list is populated with the available metric families and metrics for the selected items or groups.

3. (Optional) To preview the report, complete the following:

a. Click **Run**.

A preview dashboard shows the view format that you have selected.

**NOTE**

If the metric families that you selected do not apply to the selected components, N/A appears in the report.

b. From the preview, click the **Save the settings of this report template** icon, and then click **Save**.

4. (Optional) Share the on-demand report results.

For more information, see [Share Data with Other Users](#).

The on-demand report template is created and appears on the **Manage On-Demand Report Templates** page.

### **Apply an On-Demand Report Template to Dashboards, Group Context Pages, and Device Context Pages**

You can apply on-demand report templates to dashboards and to group and device context pages. For group context pages, the reports use the selected group as the context for the view. For device context pages, the reports use the device as the context.

### **Copy an On-Demand Report Template**

Copy an on-demand report template that another user or administrator has created and owns, so that you can edit it.

**Follow these steps:**

1. Log in as a user with the Copy On-Demand Report Templates from other Users role right.
2. From the **Manage On-Demand Report Templates** page, click **Copy**.  
The **IM On-Demand/Multi-Metric Trend Report** dialog opens.
3. Specify the fields as described in [the "Create or Edit an On-Demand Report Template" section](#).

A copy of the on-demand report template is created, and it is available for you to edit.

### **Run an On-Demand Report Template**

Run an on-demand report template to execute the report.

**Follow these steps:**

1. Log in as a user with the Run On-Demand Report Templates role right.

2. From the **Manage On-Demand Report Templates** page, select the on-demand report that you want to run, and then click **Run**.

The on-demand report is shown in the dashboard.

## Performance Metrics

You can view the performance data that the data sources collect, process, and aggregate from the dashboards in NetOps Portal. In some cases, the dashboard views are rollups of more granular data that is available in the data source, or by drilling down into the details from the dashboard.

## Baseline Calculations

Some views include baseline data as a basis for comparison. The baseline calculation method varies by the registered data source. The baseline data that is plotted in many views shows statistical deviations from "normal" performance for a given statistic. Metrics are considered to be "normal" based on the calculated baseline average. The Standard Deviation is used to gauge the statistical validity of the baseline values. Baseline values are included in charts to help you see places where performance values are changing rapidly.

Baseline data helps to characterize past performance for selected monitored parameters, assess present performance, and estimate future performance. For example, comparing current CPU utilization to a known baseline average level helps to determine whether current utilization is within a typical range. A monitored parameter that exceeds a baseline can indicate additional load on the server from a new application process, an increase in the number of users or sessions, or an increase in the amount of data being processed.

In this article:

### Baseline Averages

Depending on the amount of polled data that is collected, *baseline averages* are calculated in two ways:

- Initially, averages are calculated for the same hour regardless of the day.
- After enough data is collected, averages are calculated for the same day of the week and the same hour.

Baseline averages help to characterize past performance for selected monitored metrics, and helps to assess present performance. NetOps Portal continually calculates baseline averages and related standard deviations as each hour passes. The standard deviation provides a statistical indicator of how much variability exists in the population data that factored into the baseline average calculations.

In the data aggregator, "normal" for a specified duration within a window of time is based on the calculated baseline average.

### Baseline Average Calculations

When a limited amount of data is first collected, the baseline average is calculated for the same hour for every preceding day of the week. For example, after two days worth of history, a baseline average value for the 9:00 AM to 10:00 AM time period is calculated by averaging the hourly rollups for the same time periods for two consecutive days.

Eventually, when more data is available, a switchover in the calculation method occurs automatically and the data aggregator establishes "normal" by averaging hourly samples across available preceding same days of the week. This method, then, considers the day of the week patterns in utilization. This method produces a better approximation of what is "normal", which can lead to a reduction in the number of missed violations and false positive events that are generated. In the same example as above, after three weeks of history, a baseline average is calculated by averaging the 9:00 AM to 10:00 AM hourly rollups for the three Mondays within the three-week period.

**NOTE**

By default, this automatic switchover occurs when at least three same day of the week, same hour data samples are available for the past 12 weeks. The data aggregator switches back to the every day, same hour calculation method automatically when the required number of data points is no longer available. These default settings are configurable.

For more information, see [Configure Data Retention Rates](#).

Baseline averages are calculated for event and report generation purposes.

**Standard Deviation Calculations**

The standard deviation is calculated from the baseline average for rollups, threshold events, and report generation purposes.

Rollups:

- For hourly rollups, the standard deviation is calculated for the polled values.
- For daily rollups, the standard deviation is calculated for hourly averages.
- For weekly rollups and beyond, the standard deviation is calculated for the daily averages.

Threshold Events:

- The standard deviation provides a statistical indicator of how much variability exists in the population data that factored into the baseline average calculations.

Reporting:

- For hourly reporting, the standard deviation is calculated for the polled values.
- For daily reporting, the standard deviation is calculated for hourly averages.
- For weekly reporting and beyond, the standard deviation is calculated for the daily averages.

The formula for calculating this standard deviation is:

population deviation = Square root of (Sum ( X - population mean)/number of data points)

- **X**  
The data point value in the population
- **Population**  
The set of potential values that includes observed cases and potentially observable cases

**Example: Calculate the Same Hour Average and Population Standard Deviation for CPU Utilization**

The following example shows how the "same hour" average (mean) and population standard deviation are calculated for CPU utilization on a specific device, when there are three points of data for 2:00 AM on Monday, Tuesday, and Wednesday.

1. Collect three points of data:

Day:	Monday	Tuesday	Wednesday
Mean (Average) CPU utilization:	76	65	10

2. Calculate the population mean. The formula for calculating the population mean is as follows:

The population mean = sum of data point values in population/number of data points.

The equation for this example is as follows:

$(76+65+10) / 3$

The population mean= 50.33

3. Calculate the difference of each data point from the mean. The differences for this example are:

25.67      14.67      -40.33

4. Calculate the square of the difference for each data point. The squares for this example are:

658.78      215.11      1,626.778

5. Calculate the sum of the squares.

The sum of the squares for this example is 2,500.67.

6. Calculate the sum of the squares, divided by the number of data points in the population.

The result for this example is 833.56.

7. Calculate the square root of the sum of squares of data point value from the population mean.

The square root for this example is 28.87, and the standard deviation is 28.87.

The following table depicts the hourly averages (mean) of rate data by day, the average (mean) of hourly averages, and the population standard deviation of the hourly averages for the same hour:

Time	Monday	Tuesday	Wednesday	...	Mean	Standard Deviation
2:00 AM	76	65	10	...	50.33	28.87
3:00 AM	87	18	32	...	45.67	29.78
4:00 AM	10	56	40	...	35.33	19.07
5:00 AM	60	45	19	...	41.33	16.94
Hour...	...	...	...	...	...	...

#### **Example: Calculate the Same Day of the Week Same Hour Average and Population Standard Deviation for CPU Utilization**

The following example shows how the average (mean) and population standard deviation are calculated for CPU utilization on a specific device, when there are three points of data for three Mondays at 2:00 AM.

1. Collect three points of data.

Monday of Week:	1	2	3
Mean (Averages) CPU utilization:	76	4	6

2. Calculate the population mean. The formula for calculating the population mean is as follows:

The population mean = sum of data point values in population/number of data points.

The equation for this example is as follows:

$$(76+4+6)/3$$

The population mean = 28.67.

3. Calculate the difference of each data point from the mean. The differences for this example are:

47.33      -24.67      -22.67

4. Calculate the square of the difference for each data point. The squares for this example are:

2,240.44      608.44      513.78

5. Calculate the sum of the squares.  
The sum of the squares for this example is 3,362.67.
6. Calculate the sum of the squares, divided by the number of data points in the population.  
The result for this example is 1,120.89.
7. Calculate the square root of the sum of squares of the data point value from the population mean.  
The square root for this example is 33.48, and the standard deviation is 33.48.

The following table depicts the hourly averages (mean) of rate data by day, the average (mean) of hourly averages and the population standard deviation of the hourly averages for the same day of the week, same hour:

Time	Week 1		Week 2		Week 3	Monday			
	Monday	...	Monday	...	Monday	...	Mean	Standard Deviation	
2:00 AM	76	...	4	...	6	...	28.67	33.48	
3:00 AM	87	...	71	...	56	...	71.33	12.66	
4:00 AM	10	...	27	...	58	...	31.67	19.87	
5:00 AM	60	...	3	...	32	...	31.67	23.27	
Hour	...	...	...	...	...	...	...	...	

#### **Example: Deviation from Normal using the Same Day of the Week Same Hour Average and Population Standard Deviation for CPU Utilization**

Assume that the data aggregator is polling CPU utilization data at a 5-minute interval. You define an event rule to generate an event when CPU utilization is greater than one standard deviation above the mean for a single 5-minute poll interval.

In this example, event rule duration and window are both set to 5 minutes.

The formula for calculating when an event is raised is as follows:

```
CPU utilization = mean value + 1(standard deviation value)
```

Therefore, substituting mean and standard deviation values from the preceding same day of the week, same hour for Monday at 2:00 AM is as follows:

```
CPU utilization = 28.67 + 1 (33.48)
CPU utilization = 62.15
```

As a result, if CPU utilization were to exceed 62.15 for a single 5-minute poll interval between 1:05 AM and 2:00 AM on Monday, an event would be raised. This event indicates that the CPU utilization deviated from normal for that timeframe.

#### **Example: Examine CPU Utilization Events in a Trend Chart View**

Assume that the data aggregator is polling CPU utilization data at a 5-minute interval. In this example, you want to be alerted whenever CPU utilization on one of your business critical servers drops below the expected level. You define an event rule to generate an event when CPU utilization is one standard deviation below the mean for a single 5-minute poll interval.

For illustrative purposes only, assume that CPU utilization is 50 percent from Monday, 12:00 AM to Sunday, 12:00 AM. From Sunday, 12:00 AM to Monday, 12:00 AM, CPU utilization drops to 10 percent. You expect this drop in utilization. However, when the data aggregator begins to calculate the baseline average, an event is raised when the CPU utilization drops to 10 percent. The event clears when the CPU utilization goes back up to 50 percent. The erroneous event is raised



because, initially, when a limited amount of data is collected, the baseline average is calculated for the same hour for every day, not taking into account the difference in utilization across days of the week. The data aggregator is expecting the CPU utilization to be 50 percent *always*.

After three weeks pass, three same days of the week, same hour data samples are available, and the baseline average calculation method changes. The data aggregator establishes "normal" by averaging hourly samples across same days of the week. The data aggregator is now expecting the CPU utilization to be 10 percent every Sunday at 12:00 AM to Monday at 12:00 AM. The erroneous event that was raised previously every Sunday at 12:00 AM is no longer raised.

## Rate Metrics

You can select metric families and metrics to display for a selected managed item or group using the Dynamic and Custom view types. For some metrics, the following additional options are available: **Average Rate** and **Total Rate**. These rate variants represent a "per second" value for the metric, achieved by dividing the raw sample value by the number of seconds in the sample's poll period.

The data aggregator aggregates these variants differently across managed items when you have selected a group for reporting. The data aggregator also aggregates them differently when a selected device has child items that are monitored, or when the selected time period requires aggregation. For example, in a table, a single value can represent the aggregated result of one hour of samples: 12 samples from five-minute polling are aggregated into a single value. In this case, using the **Average Rate** option would average out the "rate" values for the 12 samples, where using the **Total Rate** option would simply sum them.

The following calculations illustrate the difference:

```
Average Rate (AvgRate) = ((sample1/duration1) + (sample2/duration2) + (sample3/duration3)) / numberOfSamples
Total Rate (TotalRate) = ((sample1/duration1) + (sample2/duration2) + (sample3/duration3))
```

For most situations, **Average Rate** is the preferable metric to use.

### Examples

Particularly for the Composite Trend view type, selecting the more useful rate metric option can yield more revealing results. **Average Rate** is adequate for most situations. But if you select **Total Rate** for a view that reports on groups, you can see the total for all devices of a similar type, or for all devices in the same region, for example.

Use the **Total Rate** metric to see metrics for a device across all of its interfaces. For a device with child interfaces, you can see an average across all of its interfaces, or you can see totals per interface. Select trend chart settings that display data from each interface as a separate line.

In cases where you are reporting on subinterfaces, consider that each one has a rate of N bytes/sec. The rate per interface is an aggregation of rate data from all subinterfaces. Typically, you would select **Total Rate** to see a metric such as Bytes/second per interface. To see data from all subinterfaces, use **Total Rate** to see a Bytes/second total for all interfaces.

If you do an aggregation of the 12 samples from an hour time period, those samples would be averaged to get the **Average Rate**. They would be added together to get the **Total Rate**.

## Interface Reporting

Digital network interfaces use serial transmission to send data from the transmission port to the receive port on the other end of the communications channel or circuit. Some transmission channels, such as copper Gigabit Ethernet, aggregate serial data across multiple channels to establish their overall circuit capacity. For copper Gigabit Ethernet, 4 channels of 250 Mbps are used to establish the 1-Gbps Ethernet circuit.

Most digital interfaces are *full-duplex*. The term means that they can transmit outbound data at the same moment that they are receiving inbound data. Because data transmission and data reception are independent interface tasks, they are reported separately.

Interface Utilization represents the average amount of data that is transmitted by the interface in a single direction (In or Out), divided by the interface bandwidth, or capacity. Interface utilization can be expressed as a percentage or as a transmission rate in bits per-second (bps).

Interface utilization rates can contribute to network performance issues. For a given interface, monitor whether it is transmitting frames at or below the rate at which it is receiving them. Acceptable interface utilization rates also depend on various SLAs and failover scenarios within your network. For example, two interfaces use a load-sharing algorithm to balance outbound traffic to the next hop. The average interface utilization must remain low enough that a failure of one link does not saturate the remaining available link, which now transmits all data.

## **Interface Utilization**

*Interface utilization* refers to the transmission and reception of data and associated framing of device interfaces. Interface utilization is commonly referred to as "network utilization," "circuit utilization," or "uplink utilization."

The interface utilization percentage metric is calculated from average data because an instantaneous reading of individual interface utilization is either 100% (actively transmitting or receiving a frame) or 0% (not actively transmitting/receiving a frame). The average utilization percentage value includes the amount of time that the interface was in use over the given interval.

The interface utilization rate metric takes into account the interface speed, or its available bandwidth. For a physical interface, the available bandwidth of an interface is defined as the actual clockspeed rate at which the interface is capable of transmitting data. For example, 1536 Kbps, 44.728 Mbps, 100 Mbps, 1 Gbps, and 10 Gbps describe clockspeed rates. For logical interfaces, such as subinterfaces, the available bandwidth is defined as the bandwidth value assigned to the interface by a network administrator. However, the total amount of real bandwidth available to the logical interfaces cannot exceed the physical interface capacity in terms of actual transmission rates.

Full-duplex interfaces have the capability to transmit data independently at the same time that they can receive data. This capability requires independent hardware dedicated in the transmit direction and the receive direction of the interface. Accordingly, the average utilization of an interface is reported separately in either the inbound or outbound direction. For example, separate views show "Average Utilization Out" and "Average Utilization In".

The network utilization can be derived from interface utilization values averaged over time from interfaces that are in use or not in use.

## **Interface Errors and Discards**

Elevated interface error rates usually indicate a problem with the transmission medium. For example, the cable, fiber, or interface hardware can cause errors. Each error indicates that the associated packet was dropped during the attempt to transmit or receive it.

When detected in the inbound direction, errors typically indicate problems with the transmission medium (for example, cable or fiber). Outbound errors indicate problems with the interface hardware. The acceptable rate of errors for any given interface is typically zero (0) errors.

Interface discards typically occur when interface buffers no longer have the capacity to store packets (because, for example, buffer memory is exhausted). Buffer congestion often indicates that the rate at which packets are arriving at the interface exceeds its transmission rate.

Each reported discard is a packet that the reporting interface threw out. The sending host must retransmit such packets if a reliable protocol such as TCP is used to send the data end-to-end. Interface discards are typically the result of congested queues serving the interface. Discards frequently occur in bursts. Elevated discard rates may be the result of either microcongestion or chronic congestion issues.

Acceptable discard rates depend on the applications being served, the transmission protocols, and the SLAs established within your organization. Views of interface error and discard totals use a "k" to indicate thousands. Units labeled 'kErrors' or 'kDiscards' therefore refer to thousands of packet errors or thousands of packet discards.

## Interface Bandwidth

The term "bandwidth" generally refers to the available capacity of an interface to transmit data at a given bit rate. The bandwidth associated with an interface depends on the type of interface.

For a physical interface, the bandwidth is the physical clock rate that the interface uses to transmit data. The clock rate is also typically a function of the type of interface. The following list shows the clock rates for common physical interfaces:

- DS-0: 64 Kbps
- DS-1: 1.536 Mbps
- E-1: 2.048 Mbps
- E-3: 34.368 Mbps
- DS-3: 44.278 Mbps
- OC-3: 155 Mbps
- Fast Ethernet: 100 Mbps
- Gigabit Ethernet: 1000 Mbps
- Ten Gigabit Ethernet: 10 Gbps

## CPU Utilization

The CPU utilization metric is based on average CPU usage over the time period selected for the view.

*Utilization percent* is a term applied to the portion of a time period that a component is doing work, divided by the total amount of time in the time period. The result is multiplied by 100 to obtain a percentage.

For a CPU, the busy time is spent processing program instructions. Here is an example of how to interpret a utilization rate of 70% for a 5-minute time period: "For 70% of the 5 minutes, the CPU was fully utilized."

High rates of CPU utilization can indicate poor application performance. With high CPU utilization, processes must wait in the processor queue for a previous process to complete execution.

## Memory Utilization

In addition to utilization, the data sources send information about memory capacity and usage to NetOps Portal. The **Memory Utilization** metric is an average utilization statistic derived from the percentage of available memory in use at a given moment, averaged over the reporting interval.

High rates of memory utilization may indicate that processes are paging instruction sets into and out of virtual memory. Such paging leads to slower memory read times, and can require the CPU to be interrupted to manage the paging process. The result is decreased performance for the associated application processes.

In addition, a steady increase in memory utilization over time may indicate a 'memory leak' condition. Such a condition exists where memory is allocated by processes as they start, but is not being released as the processes end. Memory leaks degrade device performance over time. Typically, the device becomes unresponsive when memory is no longer available.

## Device Availability and Reachability

Availability measures system uptime and reachability measures device connectivity to NetOps Portal. *Availability* is a percentage of time that the device is powered on and capable of processing data. NetOps Portal uses system uptime to calculate the percentage of time within a poll cycle that the device was available. The metric value is recorded as up (100), or down (0), or a percentage of the poll cycle.

If NetOps Portal cannot reach a device for a particular timeframe, availability is backfilled when the device sends a poll response. NetOps Portal uses the system uptime to calculate which poll cycles to mark as available. If the device was restarted, or if a counter rollover for uptime occurred, NetOps Portal backfills only from the 0 to the current system uptime

value. Poll cycles before the rollover or restart, where no poll responses were received, remain blank. Hourly and daily rollups reflect the backfill data only if the data is received before the rollup occurs. Availability is the only metric that NetOps Portal backfills.

NetOps Portal might not be able to reach a device that is available because of a network or communications failure by another device. *Reachability* refers to whether a data source can reach device. Typically, data sources use ICMP (ping testing) to communicate regularly with the target device. Any communication failures, including the loss of the network path or routing, affect the reachability statistics. If ICMP is blocked, NetOps Portal uses SNMP to determine reachability.

Reachability data comes from regular ping testing of all devices that support ICMP. The reachability value is the percentage ping responses that are received from the device during each reporting interval.

## Reachability Status and Contact Status

In this article:

- [Reachability Status](#)
- [Contact Status](#)
- [View Device Status](#)

### Reachability Status

The reachability status indicates whether NetOps Portal can reach the device during the selected time range of the context page.

The value is based on the polled reachability metric for the device at the last poll cycle during the selected time range:

- **Reachable**  
Indicates that the data collectors reached the device during the last poll cycle during the selected time range.
- **Not Reachable**  
Indicates that the data collectors could not reach the device during the last poll cycle during the selected time range.
- **Reachability Status Unknown**  
Indicates that the data collectors did not poll the device during the last poll cycle during the selected time range. This status might indicate that the device has a life cycle state of "Retired" or is offline for maintenance (the life cycle state is "Maintenance").

### Contact Status

The contact status indicates the live status of a pingable or manageable device and is based on the outcome of active polling:

**Values:**

- **Up**  
Indicates that the data collectors are successfully polling the device.
- **Down**  
Indicates that the data collectors cannot reach the managed item through the management protocol, SNMP, or ping. At least two polls failed to receive a response.
- **DC Connection Lost**  
Indicates that the data aggregator can no longer reach the data collector that monitors the device.
- **Not Monitored**  
Indicates NetOps Portal is not monitoring the device. Polling is disabled, for example, you have paused polling of the monitored device using NetOps Portal or using the Data Aggregator REST web services, by changing the life cycle state of the device from "Active" to "Retired" or "Maintenance", or by disassociating the device from any device collection that has monitoring profiles associated to it.  
For more information:

- About how to pause polling using NetOps Portal, see [Manage Monitored Devices](#).
- About how to disable polling using the `discoverydefaultconfig` or `pollable` Data Aggregator REST web services, see [Manage Default Polling Behavior](#) or [Manage Polling Behavior for Components](#).
- About the life cycle states, see [Manage Device Life Cycles](#).
- **Management Lost**  
Indicates that the data collector that is associated with the device can reach the device through the ICMP protocol, but it cannot through the management protocol, such as SNMP.
- **Unknown**  
Indicates an unexpected condition.

## View Device Status

You can view the following statuses/states on the **Details** tab of a device context page, at the top of the page. To view a device context page, hover over **Inventory**, **Items**, click **Devices**, and then click the device for which you want to view the device context page:

- **Reachability status**, for example, "Reachability Status Unknown"  
For more information about this status, see [the "Reachability Status" section](#).
- (23.3.3 and higher) **Contact status**, for example, "DC Connection Lost"  
For more information about this status, see [the "Contact Status" section](#).

### NOTE

You can also view the contact status on the **Monitored Devices** page, on the **Details** tab. The status is shown in the **Status** field.

For more information about this page, see [Manage Monitored Devices](#).

- **Life cycle state**, for example, "Maintenance"  
For more information about this state, see [Manage Device Life Cycles](#).
- **Current alarm state**, for example, "Maintenance Alarm"

### NOTE

If NetOps Portal is not discovering the device from Spectrum, then the current alarm state shows as "Alarm Status Unknown", and on the **Devices** page, the current alarm status for the device shows as blank.

For more information about this state, see [Monitor Device Inventory with Alarm States](#).

The following image shows an example of **Details** tab for a device showing these statuses:

**Figure 91: Reachability Status for a Device**

The screenshot shows the DX NetOps portal interface. The top navigation bar includes links for Home, Alarms, Performance, Inventory (selected), Reports, System Health, and Administration. The main content area is titled 'Router: Cisco-2800' and 'Switch Cisco-2800'. On the left, a sidebar lists various views: Details (selected), System Health, Interface Health, Top Interface Discards, Top Interface Errors, Top Interface Utilization, Custom View - Infrastructure..., and Alarms. The main panel displays the 'Details' tab for the 'Cisco-2800' device. At the top of this panel, there are status indicators: 'Reachability Status Unknown' (with a red icon), 'DC CONNECTION LOST' (with a red 'X' icon), 'MAINTENANCE' (with a wrench icon), and 'MAINTENANCE ALARM' (with a wrench icon). Below these, a table provides detailed information about the device, categorized into Identity, Device, SNMP, Group Membership, Time Settings, and Custom Attributes.

IDENTITY	DEVICE	SNMP	GROUP MEMBERSHIP	TIME SETTINGS
Context Types: Router, Switch	Vendor: Cisco Systems/Cisco	Profile: default	/All Groups/Inventory/IP Domains/Default Domain	The device is not a member of a site group that includes time settings.
Device Name: Cisco-2800	Model: Cisco2811	Profile: SNMPv1/SNMPv2C	/All Groups/Inventory/Data Sources/Spectrum	
Device Name Alias:	Description: Cisco IOS Software, 2800 Software (C2800NM-ADVIPSERVICESK9-M), Version 12.4(25b), RELEASE SOFTWARE	Timeout: 3000		
Device IP:		Retries: 2	/All Groups/cisco	
SysName: Cisco-2800-10.255.0.101			/All Groups/bbbCopy1/bbbCopy1/cccCopy1	
				CUSTOM ATTRIBUTES
				The device does not have any custom attributes.

For more information about these statuses on the **Devices** page, see [Inventory Pages and Views](#).

## Scorecard Projections

Scorecard projections use a customizable set of data points to predict future values for metrics. DX NetOps Performance Management calculates projected values when it renders the view. These values are based on the historical time frame of the view. You can add these values to the IM Group Scorecard Trend View.

### TIP

Do not use scorecard projections for error metrics. Each error is a discrete event that is not affected by historical errors.

The scorecard view includes the following methods to calculate projections:

- **Approximation**  
This method uses the average from each time frame in the view to calculate the projection values. DX NetOps Performance Management calculates a least squares regression on the averages, then uses the line equation to project future values. This calculation method is faster than the Detailed Data method.
- **Detailed Data**  
This method uses the polled data for the entire time frame of the view. DX NetOps Performance Management calculates a least squares regression for the entire set of data points. This calculation is more statistically accurate than the **Approximation** method, and provides extra columns in the view.

### NOTE

Detailed data scorecard projections are supported only for gauge metrics (for example, Bits Out - Average Rate). Detailed data scorecard projections are *not* supported for counter metrics (for example, Bits Out - Total). Projection values are calculated on the As Polled (rate) data to ensure precision.

The following columns are hidden by default:

- **Slope**  
Indicates the slope of the line equation.
- **Intercept**  
Indicates the intercept of the line equation.
- **Degrees**  
Degrees of freedom, which indicates the sample size.
- **Linear Fit**  
Indicates the confidence level of the projected values as related to the sample data.
- **Days to Threshold**  
Indicates the projected number of days before the specified critical threshold is reached.

## Percentiles

A *percentile* is the value of a variable below which a certain percent of observations fall. For example, the 95th percentile is the value (or score) below which 95 percent of the observations are found. NetOps Portal calculates percentiles for rollups, dashboards, and report generation purposes.

Percentile monitoring is useful in measuring throughput data. This statistic more accurately reflects the required capacity of the monitored link for applications that are bandwidth sensitive. For example, the 95th percentile says that 95 percent of the time, the bandwidth usage is below this amount. The remaining 5 percent of the time, the bandwidth usage is above that amount.

NetOps Portal aggregates metric values during rollups.

By default, NetOps Portal calculates the 95th percentile for some metrics. You can include up to two configurable metric properties (percentile calculations) in a metric family. *Percentile calculations* use the Microsoft Excel method. Hourly and daily rollup performs the percentile calculation on the polled rate data. Weekly rollup performs the percentile calculation on the results of the daily rollups. NetOps Portal calculates percentiles and exposes them to the reports and dashboards.

**IMPORTANT**

To minimize the affect on system performance, include only those percentile calculations in a metric family that you require.

The following is a list of the metric properties (percentile calculations) that you can include in a metric family:

- **Percentile**  
**Values:**
  - 0: Disabled
  - 95**Default:** Disabled or 95
- **Percentile2**  
**Values:**
  - 0: Disabled
  - 1-99**Default:** Disabled
- **Percentile3**  
**Values:**
  - 0: Disabled
  - 1-99**Default:** Disabled

For more information about how to edit metrics, see [Edit a Metric](#).

Changes to these metric properties (percentile calculations) generate administration events. After you update the metric family, the percentile data is available for reporting within several poll cycles. For trend views, changes cause a gap in the trend line. For table views, changes affect the value of the percentile for time ranges when the change occurred. For example, in a daily time range, the value is inaccurate for one day.

For more information about administration events, see [Event Types](#).

**TIP**

If you disable a percentile calculation, avoid errors in views that include the percentile by modifying them.

Three percentile values appear in the **Metric Families** table for a data source. A dash in the **Percentiles** column of the **Metrics** tab indicates that all three percentile values are equal to zero (disabled).

For more information about how to view this table, see [Manage Metric Families](#).

## Metric Projection

To calculate future values based on historical metric data, use metric projection. Metric projection is useful for capacity planning. For example, to verify that the interface bandwidth is sufficient for a specific time in the future, calculate the projected interface utilization.

To see future trends, metric projection supports up to three configurable intervals. For example, you can project to 20, 60, and 180 days in the future for the metric. Projection shows an overall trend. Typically, the longer the projection interval, the less accurate the exact value.

### Scorecard Projections

- Scorecard views provide line-of-business owners a group-level summary of how key metrics perform over time. Performance is based on a set of user-defined thresholds.
- The scorecard view displays historical time intervals, and up to three projected values.
- Projected values are calculated when the view is rendered, and are based on the historical time frame of the view.

## Metric Projections

- Metric projection is designed for network and capacity planners that want the system to calculate and store projections.
- Projections are configured for individual metrics. Up to three projection intervals can be specified per metric.
- Projected values are based on up to 90 days of historical data. Once configured, projected metric values can be included in custom table views.

For more information, see [Scorecard Projections](#).

You can add the projection values to custom table views.

### TIP

Table views are useful for capacity planning.


## Metric Projection Calculations

Projection values are calculated for a specified interval and are based on a configurable daily percentile calculation on the metric. For capacity planning, 95 percent is a typical percentile value. The system uses the following process to calculate the projection:

1. Calculates and stores the daily percentile value using the Microsoft Excel method from the as polled data.
2. Calculates a linear regression line from the daily percentile values. The calculation uses a simple linear regression (least squares regression).  
The calculation uses all the available daily percentile values from the last 90 days as input data. Projection requires at least two days of daily percentile values. The accuracy of the projection typically increases with the available data points.
3. Calculates the future value for the interval from the linear equation.

The following diagram illustrates the calculation method:



**Figure 92: Metric Projection Calculation****Configure Metric Projection**

To configure metric projection, edit the metric through the UI. For more information, see [Edit a Metric](#).

**TIP**

When you disable a projection, modify views that include the projection to avoid an error in the view.

**WARNING**

Calculate only projections that you need. Each projection that you calculate may significantly affect system performance. For more information, see the [DX NetOps Performance Management Sizing Tool](#).

To configure metric projection using REST web services, execute a PUT operation on the target metric. Configure **ProjectionPercentile** and **ProjectionInterval** for the metric. For more information, see [Metric Family XML Structure](#).

**WARNING**

Changes to **ProjectionPercentile** cause inaccurate projections for up to 90 days. When you change the value of **ProjectionPercentile**, the percentile values for days before the change are not recalculated.

Changes to these properties generate Administrative events.

**Example:**

This example shows the XML for interface utilization. This projection is calculated based on the 95th percentile. The projection is calculated for 20, 30, and 90 days.

**Endpoint:**

`http://DA_host:8581/typecatalog/metricfamilies/extension/NormalizedPortInfo`

```
<DataModel xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" namespace="http://
im.ca.com/normalizer" xsi:noNamespaceSchemaLocation="IMDBCertificationFacet.xsd">
  <Author>CA</Author>
  <Version>1.2</Version>
  <FacetType name="NormalizedPortInfo"
descriptorClass="com.ca.im.core.datamodel.certs.NormalizedFacetDescriptorImpl">
    <Documentation/>
    <FacetOf namespace="http://im.ca.com/core" name="Item"/>
    <AttributeGroup external="true" list="true" name="PortInfoPollable">
      <Documentation/>
      <Attribute name="Utilization"
type="double">
        <ProjectionPercentile>95</ProjectionPercentile>
      </Attribute>
    </AttributeGroup>
    <BaselineDefinitions>
      <Baseline name="DailyBaseline">
        <ID>26</ID>
        <PerformanceMetric>Utilization</PerformanceMetric>
        <Period>1 Day</
Period>
        <ProjectionInterval>20</
ProjectionInterval>
        <ProjectionInterval2>30</
ProjectionInterval2>
        <ProjectionInterval3>90</ProjectionInterval3>
        <Window>90 Days</Window>
        <StartDate>0</StartDate>
        <EndDate>0</EndDate>
        <DaysOfWeek>0</DaysOfWeek>
      </Baseline>
    </BaselineDefinitions>
  </FacetType>
</DataModel>
```

## Total, Average, Minimum, and Maximum Values

The total, average, minimum, and maximum values are calculated for rollups and for reporting purposes. These values let you observe the upper and lower bounds of performance across a given time interval.

A *rollup* is the process during which metric values are aggregated. In an hourly rollup, the 1 minute, 5 minute, 15 minute, 30 minute, and 60 minute polled values for metrics are aggregated every hour. In a daily rollup, as-polled values for metrics are aggregated once a day. In a weekly rollup, daily values for metrics are aggregated once a week. A

given metric is aggregated by either averaging or summing data points during the rollup (not both) based on the "RollupStrategy" configured for that metric.

Hourly rollups:

- Total: Sum of all the as-pollled data points for the hour.
- Average: Sum of all the as-pollled data points divided by the number of data points for the hour.
- Minimum: The lowest single poll value of the hour.
- Maximum: The highest single poll value of the hour.

Daily rollups:

- Total: Sum of all the as-pollled data points for the day.
- Average: Sum of all the as-pollled data points for the day divided by the number of data points.
- Minimum: The lowest single value of the hourly minimums.
- Maximum: The highest single value of the hourly maximums.

Weekly rollups and beyond:

- Total: Sum of all the daily data points for the week.
- Average: Sum of all the daily data points divided by the number of data points for the week.
- Minimum: The lowest single value of the daily minimums for the week.
- Maximum: The highest single value of the daily maximums for the week.

Five-minute resolution reporting:

- Total: Sum of all the as-pollled data points for the 5 minute increment.
- Average: Sum of all the as-pollled data points divided by the number of data points for the 5-minute increment.
- Minimum: The lowest single poll value of the five-minute increment.
- Maximum: The highest single poll value of the five-minute increment.

One hour resolution reporting:

- Total: Sum of all the as-pollled data points or the hourly rollup value for the hour.
- Average: Sum of all the as-pollled data points divided by the number of data points or the hourly rollup value for the hour.
- Minimum: The lowest single value of the as-pollled data or hourly rollup minimum for the hour.
- Maximum: The highest single value of the as-pollled data or hourly rollup maximum for the hour.

Day resolution reporting:

- Total: Sum of all the hourly data points or the daily rollup value for the day.
- Average: Sum of all the hourly data points divided by the number of data points or the daily rollup value for the day.
- Minimum: The lowest single value of the hourly data or daily rollup minimum for the day.
- Maximum: The highest single value of the hourly data or daily rollup maximum for the day.

## Events

You can monitor the health and status of your system and your environment using events. *Events* are messages that provide information about what is happening in NetOps Portal. Events include basic information, such as related devices and the time of the occurrence that triggered the event.

### TIP

You can integrate with DX NetOps Spectrum (Spectrum), and then in Spectrum, you can configure events that generate alarms.

For more information, see the [DX NetOps Spectrum documentation](#).

In this article:

- [Events Views on the Events Display Dashboard](#)
- [View and Manage Events](#)

## **Events Views on the Events Display Dashboard**

The following views are included in the **Events Display** dashboard:

- **Events**  
This view displays the events that occurred in the selected time range for the dashboard. You can filter this view for a specific group. This view is the default view in the **Events Display** dashboard.
- **Filtered Events**  
This view includes filters for data source, severity, event type, event subtype, and threshold profile.

### **TIP**

To see the complete event properties, select an event in any view, and then click **Details**.

## **View and Manage Events**

The following is a list of scenarios for viewing and managing events:

- **Troubleshoot performance issues**  
Troubleshoot performance issues with a specific device by filtering the events to that device. The **Events** view filters all events to display only those for the selected device.
- **Monitor thresholds**  
Configure threshold profiles that include event rules that define the conditions that trigger them to raise or clear a violation for specified devices, components, or groups. Use these events to track when the monitored devices do not meet operating requirements.  
For more information, see [Configure Threshold Profiles](#) and [Threshold Event Processing Self-Monitoring Metrics](#).
- **Track configuration changes**  
If you have chosen to manually control when changes to monitoring occur (the **Automatically Update Metric Families** option is cleared for a custom monitoring profile), view the events log to view configuration changes. Reconfiguration events with the **Detected** subtype indicate detected changes on the monitored devices. Reconfiguration events with the **Changed** subtype indicate changes to device items.  
For more information about this option, see [Manage Monitoring Profiles](#).
- **Audit changes to the system**  
Administration events provide useful information about changes to the system configuration, such as the time of a change.

## **Event Types**

The event type determines what information an event contains. Events have an event type, and many events have an event subtype.

The following are a list of events, categorized:

### **NOTE**

When you have integrated other products with DX NetOps Performance Management, there can be more event types than those listed here.

## **Administration**

Administration events include system status notifications and audit events. These events include the following event subtypes:

- **ActiveMQ Status**  
Occurs when the ActiveMQ status changes.
- **Administration Message**  
Administrative tasks that are related to discovery and tenants can trigger this event type.
- **Certification Change**  
Occurs when a vendor certification or metric family is extended or updated.
- **Data Collector Status**  
Occurs when the status of a data collector changes.
- **Monitoring Profile**  
Occurs when the monitored metrics selected for a monitoring profile are modified.
- **Percentile Change**  
Occurs when percentile calculations are changed for a metric family. When the percentile is changed in the UI, the event includes the user.
- **Projection Change**  
Occurs when projection calculations are changed for a metric family. When the projection is changed in the UI, the event includes the user.
- **Threshold Monitoring Status**  
Provides information about the threshold monitoring engine.  
For more information, see [Threshold Monitoring and Threshold Limiter Behavior](#).
- **Vendor Certification**  
Occurs when the vendor certification priority order is modified for a metric family.

### **Data Collection**

Data collection events provide information about contact with monitored devices. These events include the following subtypes:

- **Contact Established**  
Occurs when the data collector established contact with a device.
- **Contact Lost**  
Occurs when the data collector loses contact with a device.
- **Management Established**  
Occurs when the data collector established contact with the management protocol of the device, such as SNMP.
- **Management Lost**  
Occurs when the data collector loses contact with the management protocol of the device, such as SNMP.
- **Polling Disabled**  
Occurs when polling for a device stops.
- **Polling Enabled**  
Occurs when polling for a device starts.

### **Data Repository State**

Data repository state events provide information about the status of the data repository. Data repository state events include information about the data repository host. These events include the following subtypes:

- **Connected**  
Occurs when the data aggregator connects to a data repository node after another node goes down. These events include information about the connection.
- **Data Repository Failover**  
Occurs when the data repository cluster changes the primary host. These events include the new host name and the previous host name.
- **Data Repository Failure**

Occurs when a data repository cluster fails. These events include the name of the host that failed.

- **Degraded**  
Occurs when the data repository node response time to the data aggregator heartbeat is over 20 seconds.
- **Down**  
Occurs when a data repository node is down.
- **Start Up**  
Occurs when a data repository node starts up. These events include availability information about all the data repository nodes.

For more information about these events, see [View System Status](#).

### **Inventory Discovery**

Inventory discovery events provide information about discovery operations. These events include the following subtypes:

- **Unreachable Devices Report**  
Occurs when a discovery operation results in one or more unreachable devices. The event details include the count of unreachable devices. If the count is 10 or fewer, a list of the IP addresses of each unreachable device is included.
- **Pingable Devices Report**  
Occurs when a discovery operation results in one or more pingable devices. The event details include the count of pingable devices. If the count is 10 or fewer, a list of the IP addresses of each pingable device is included.
- **Discovery Completion Report**  
Occurs when a discovery operation completes successfully. The event details include the start and end time of the operation, as well as counts of new manageable devices, new pingable devices, changed devices, unreachable devices and existing devices.
- **Discovery Failure Report**  
Occurs when a discovery operation fails. The event details include the start time of the operation, as well as any available details about the failure.

For more information about these events, see [Discovery](#).

### **Life Cycle**

Life cycle events track changes to the device life cycle. These events include the following subtype:

- **State Change**  
Occurs when the life cycle state of a device changes. These events include the current state, the previous state, and the user who changed the state.

For more information about these events, see [Manage Device Life Cycles](#).

### **Override**

Override events provide information about manual changes to metric values. These events include the following subtypes:

- **Override Cleared**  
Occurs when a user clears the override for the Speed In and Speed Out values for an interface.
- **Override Set**  
Occurs when a user manually updates the Speed In and Speed Out values for an interface.

For more information about these events, see [Override Speed In and Speed Out Values on Interfaces](#).

### **Polling State**

Polling state events track whether polling is disabled or enabled on an interface. These events include the following subtype:

- **State Change**

Occurs when the polling state of an interface changes. The event includes the current state, the previous state, and the user who changed the state.

For more information about these events, see [Manage Interface Polling Behavior](#).

## **Reconfiguration**

Reconfiguration events provide information about change detection. These events include the following subtypes:

- **Changed**  
Occurs when the system applies a change to a component item.
- **Detected**  
Occurs when the system detects a change to a component.
- **Rebooted**  
Occurs when the system detects that a device has rebooted.

For more information about these events, see [Device Reconfiguration](#).

## **Syslog**

**Prerequisite:** To receive events of type syslog (Syslog events) from DX Operational Intelligence, log analytics for Insights must be installed and configured. For more information, see [Install and Configure Log Analytics for Insights](#).

Syslog events occur when Syslog pattern match rules defined in DX Operational Intelligence have triggered. These events include the following subtypes:

- **Syslog Info**  
Syslog events of this type are notification events only, and do not result in alarms.
- **Syslog Alarm**  
Syslog events of this type result in alarms.

By default, triggered rules result in Syslog events of subtype **Syslog Info**, which do not result in alarms. You can configure the OI Connector to generate Syslog events of subtype **Syslog Alarm**, which result in alarms, or you can configure it to ignore the notifications it receives from DX Operational Intelligence. You can configure the disposition of Syslog events for each pattern match rule defined in DX Operational Intelligence.

For more information:

- About how to configure the disposition of Syslog events, see [Configure the OI Connector](#).
- About how to monitor logs for critical log patterns, see [Use Log Analytics for Insights](#).
- About log analytics for Insights, see [Insights](#).

## **Threshold Violation**

Threshold violation events occur when monitored groups or devices violate the configured criteria in assigned threshold profiles. These events include the following subtypes:

- **Cleared**  
Occurs when a device or group that is in violation of a threshold clears the violation. These events include detailed information about the violation and the reason the violation was cleared.
- **Raised**  
Occurs when the monitored device or group violates the event rule in an assigned threshold profile. These events include detailed information about the violation.

For more information about these events, see [Configure Threshold Profiles](#).

## Change the Event Manager Properties

You can configure the Event Manager service (`caperfcenter_eventmanager`) by changing its properties (the `em.properties` file).

For example, you can configure the following:

- How long the Event Manager stores, or retains, events.
- The maximum time (in seconds) that a script can take to complete execution before it stops the execution.
- Whether to publish events to a Kafka topic (enable the Event Manager-Kafka integration), and whether to secure the integration.

### Follow these steps:

1. Open the `<installation_directory>/PerformanceCenter/EM/webapps/EventManager/WEB-INF/em.properties` Event Manager properties file.

#### Example:

```
/opt/CA/PerformanceCenter/EM/webapps/EventManager/WEB-INF/em.properties
```

#### – **installation\_directory**

The installation directory for NetOps Portal.

**Default:** `/opt/CA`

2. Edit one or more of the following properties, and then save your changes:

```
#
# Database properties
#
```

```
db.driverClassName=com.mysql.cj.jdbc.Driver
db.url=jdbc:mysql://localhost:3306/em?
useUnicode=true&characterEncoding=UTF-8&character_set_server=utf8mb4&useSSL=true&verifyServer
db.username=netqos
db.password=netqos
db.passwordEncrypted=false
db.maxActive=40
db.timeout=120
db.validationQuery=SELECT 1 FROM DUAL
db.validationQueryTimeout=3000
db.testWhileIdle=true
db.testOnBorrow=false
db.testOnReturn=false
db.timeBetweenEvictionRunsMillis=300000
```

```
# event retention period in days
# Increasing the retention period will increase
# the Event Manager database size/space requirements
Event.Retention=30
# Event table analysis happens periodically after the following number of events
em.event.analysisThreshold=10000
```

```
#
# Webservice properties
```



```
#

# Override the Event Manager location with the following properties
nqevents.dbHost=
em.web.port=8281
em.ws.maxqueue=50
em.ws.queueTimeout=180
# number of seconds to wait before any event RIB queries should be timed out
# and errors returned
rib.queryTimeout=100

#
# Notification settings
#

# The maximum size of the notification processing thread pool before notifications
  are dropped
em.notification.notification_max_pool_size=20
# The maximum # of notification processing task
em.notification.notification_max_task_size=5000
# The constant size of the notification processing thread pool
em.notification.notification_pool_size=10
# The maximum size of the notification action handler processing thread pool before
  handlers are dropped
em.notification.handler_max_pool_size=20
# The constant size of the notification action handler processing thread pool
em.notification.handler_pool_size=5
# The maximum # of notification action handler task
em.notification.handler_max_task_size=5000

#
# Script-based notifications
#

# Set to 0 to disable script notifications
em.notification.script_notification_enable=1

# The maximum size of the script notification action handler processing thread pool
  before handlers are dropped
em.notification.script_notification_handler_max_pool_size=1
# The constant size of the script notification action handler processing thread pool
em.notification.script_notification_handler_pool_size = 1
# The script notification queue size
em.notification.script_notification_handler_queue_size=5000
# The script notification execution time in seconds
em.notification.script_notification_execution_time=300
```

```
# Working directory for notification scripts (default is $PC_HOME/
NotificationScripts)
# em.notification.script_notification_execution_dir=

#
# Trap-based notifications
#

# Whether we will send NetQOS (Event Manager) traps using the existing format, or the
  format that matches the variables clause in the MIB
# Setting this value to false will use the existing format, true will use new format
  that matches the variables clause in the MIB
em.notification.trap.eventmanager_trap_use_new_format=false

# When encountering "Operation not permitted (sendto failed)", how many times to
  retry send trap and how long to wait between retries.
em.notification.trap.retry_count=2
em.notification.trap.retry_sleepms=1000

# Whether to send device IP as trap agentAddr in v1/v2 traps.
em.notification.trap.trap_send_device_ip_as_trap_address=false

# Whether to use "DELETE QUICK" followed by "OPTIMIZE" when deleting items in DB.
em.min.event.count.for.delete.quick=500000

# Kafka Configuration
em.kafka.enabled=false
em.kafka.bootstrap.servers=
#em.kafka.security.protocol=
#em.kafka.ssl.truststore.location=
#em.kafka.ssl.truststore.password=
#em.kafka.ssl.keystore.location=
#em.kafka.ssl.keystore.password=
#em.kafka.ssl.key.password=

# Consul for ServiceQueryIf
consul.host=localhost
consul.port=8900

# Consul Service Query
ServiceQuery.enabled=false
# service poll interval (in seconds)
ServiceQuery.pollInterval=60

# Default DM URL
dm.url=http://localhost:8481/dm/
# Default EM URL
```

`em.url=http://localhost:8281/EventManager/`

– **Event.Retention**

Defines how long (in days) the Event Manager stores, or retains, events.

**IMPORTANT**

Increasing the retention period increases the Event Manager database size and space requirements.

**Default:** 30

– **em.event.analysisThreshold**

Defines the number of events that must occur before NetOps Portal performs event table analysis.

**Default:** 10000

– **ngevents.dbHost**

Defines the Event Manager host to override.

– **em.web.port**

Defines the Event Manager port to override.

**Default:** 8281

– **em.ws.maxqueue**

Defines the Event Manager max queue to override.

**Default:** 50

– **em.ws.queuetimeout**

Defines the number of seconds to wait before timing out the queue.

**Default:** 180

– **rib.queryTimeout**

Defines the number of seconds to wait before timing out event RIB queries and returning errors.

**Default:** 100

– **em.notification.script\_notification\_execution\_time**

Defines the maximum time (in seconds) that a script can take to complete execution before it stops the execution. For more information about these notification scripts, see [Configure Notifications](#).

**Default:** 300

– **em.notification.notification\_max\_pool\_size**

The maximum size of the notification processing thread pool before notifications are dropped.

**Default:** 20

– **em.notification.notification\_max\_task\_size**

The maximum number of notification processing tasks.

**Default:** 5000

– **em.notification.notification\_pool\_size**

The constant size of the notification processing thread pool.

**Default:** 10

– **em.notification.handler\_max\_pool\_size**

The maximum size of the notification action handler processing thread pool before handlers are dropped.

**Default:** 20

– **em.notification.handler\_pool\_size**

The constant size of the notification action handler processing thread pool.

**Default:** 5

– **em.notification.handler\_max\_task\_size**

The maximum # of notification action handler task.

**Default:** 5000

– **em.notification.script\_notification\_enable**

Specifies if script notifications are enabled or disabled.

**Values:** 1 (Enabled), 0 (Disabled)

**Default:** 1 (Enabled)

- **em.notification.script\_notification\_handler\_max\_pool\_size**  
The maximum size of the script notification action handler processing thread pool before handlers are dropped.  
**Default:** 1
- **em.notification.script\_notification\_handler\_pool\_size**  
The constant size of the script notification action handler processing thread pool.  
**Default:** 1
- **em.notification.script\_notification\_handler\_queue\_size**  
The script notification queue size.  
**Default:** 5000
- **em.notification.script\_notification\_execution\_time**  
The script notification execution time in seconds.  
**Default:** 300
- **em.notification.script\_notification\_execution\_dir**  
The default directory for the notification scripts.  
**Default:** \$PC\_HOME/NotificationScripts
- **em.notification.trap.eventmanager\_trap\_use\_new\_format**  
Defines whether NetOps Portal will send NetQOS (Event Manager) traps using the existing format, or the format that matches the variables clause in the MIB. Setting this value to false will use the existing format, true will use new format that matches the variables clause in the MIB.  
**Default:** false
- **em.notification.trap.retry\_count**  
When encountering "Operation not permitted (sendto failed)", how many times to retry send traps.  
**Default:** 2
- **em.notification.trap.retry\_sleepms**  
When encountering "Operation not permitted (sendto failed)", how long to wait between retries to send traps.  
**Default:** 1000
- **em.notification.trap.trap\_send\_device\_ip\_as\_trap\_address**  
Defines whether to send device IP as trap agentAddr in v1/v2 traps.  
**Default:** false
- **em.min.event.count.for.delete.quick**  
Defines whether to use "DELETE QUICK" followed by "OPTIMIZE" when deleting items in the database.  
**Default:** 500000
- **em.kafka.enabled**  
Specifies if the Event Manager service publishes events to a Kafka topic (whether the Event Manager-Kafka integration is enabled).  
**Values:**
  - **true:** The Event Manager service publishes events to a Kafka topic (the Event Manager-Kafka integration is enabled).
  - **false:** The Event Manager service does not publish events to a Kafka topic (the Event Manager-Kafka integration is not enabled).**Default:** false
- **em.kafka.bootstrap.servers**  
Defines the hostname and port (hostname :port ) for Kafka.
- **em.kafka.security.protocol**  
Defines the security protocol for Kafka.
- **em.kafka.ssl.truststore.location**  
Defines the Kafka SSL truststore location for the Event Manager.
- **em.kafka.ssl.truststore.password**

Defines the password for the Kafka SSL truststore.

- **em.kafka.ssl.keystore.location**  
Defines the Kafka SSL keystore location for the Event Manager.
- **em.kafka.ssl.keystore.password**  
Defines the password for the Kafka SSL keystore.
- **em.kafka.ssl.key.password**  
Defines the password for the Kafka SSL key.
- **consul.host**  
Defines the host for Consul for ServiceQueryIf.  
**Default:** localhost
- **consul.port**  
Defines the port for Consul for ServiceQueryIf.  
**Default:** 8900
- **ServiceQuery.enabled**  
Specifies whether the Consul Service Query is enabled.  
**Options:**
  - **false:** The Consul Service Query is disabled.
  - **true:** The Consul Service Query is enabled.**Default:** false
- **ServiceQuery.pollInterval**  
Defines the service poll interval (in seconds).  
**Default:** 60
- **dm.url**  
Defines the URL for the Device Manager.  
**Default:** http://localhost:8481/dm/
- **em.url**  
Defines the URL for the Event Manager.  
**Default:** http://localhost:8281/EventManager/

3. Using the command line, stop and then start the Event Manager service by issuing the following commands:

```
systemctl stop caperfcenter_eventmanager
systemctl start caperfcenter_eventmanager
```

The Event Manager service starts and uses the new values to, for example, determine the retention period.

## Threshold Monitoring and Threshold Limiter Behavior

The threshold limiter monitors how long the evaluation engine takes to process rules in the data aggregator.

If the threshold monitoring exceeds the specified percentage of the poll cycle, the evaluation engine enters a DEGRADED state. In the DEGRADED state, the evaluation engine waits for the monitoring to drop below the specified percentage. After a specified time, the threshold monitoring engine reassesses whether to suspend threshold evaluations. If threshold violations continue to exceed the percentage during the time period, the evaluation engine is suspended. Threshold evaluations will not resume, even if you restart the Data Aggregator. If the threshold violation does not exceed the specified percentage, the evaluation engine returns to normal operation.

### WARNING

The default settings provide protection against potential polled data loss. Only modify the threshold limiter settings with the direction from Broadcom Support.

In this article:

- [View Changing Trends Over Time](#)
- [Threshold Monitoring Engine Status Events](#)
- [Take Action If Threshold Evaluations Are Suspended](#)
- [Threshold Limiter Behavior](#)
- [Resume Threshold Evaluations](#)
- [Change the Default Behavior of the Threshold Limiter](#)
- [Disable the Limiter](#)

### **View Changing Trends Over Time**

Use the **Threshold Monitoring** dashboard to view changing trends over time. This dashboard provides information about the state of the threshold monitoring engine, and includes the following views:

- **The Number of Event Rules Evaluated - Total**  
This view displays the number of actual rule evaluations that have occurred for an associated set of polled items.
- **Percentage of Poll Cycle to Complete Event Processing**  
This view displays the percentage of the poll cycle that the threshold monitoring engine takes to complete event processing.

#### **Follow these steps:**

1. Hover over **Administration**, and then click a data aggregator data source.
2. In the **Tree View** tab, expand the **All Data Aggregators** collection, and then select the same data aggregator data source.
3. In the **Details** tab, click the name of the data aggregator.
4. Click the **Threshold Monitoring** tab in the **Data Aggregator Pages** view.

### **Threshold Monitoring Engine Status Events**

Threshold monitoring engine status events describe the status of threshold evaluations. You can see these events in the **Data Aggregator Events** tab. The following table shows the possible threshold monitoring engine status events:

Event Type	Event Subtype	
Administration event	Threshold monitoring engine status	Threshold evalua
Administration event	Threshold monitoring engine status	The Threshold M in {X} minutes if t
Administration event	Threshold monitoring engine status	Threshold evalua
Administration event	Threshold monitoring engine status	Threshold evalua evaluate the mor
Administration event	Threshold monitoring engine status	The Threshold M
Administration event	Threshold monitoring engine status	The Threshold M
Administration event	Threshold monitoring engine status	The Threshold M
Administration event	Threshold monitoring engine status	Threshold evalua evaluate the mor

### **Take Action If Threshold Evaluations Are Suspended**

If threshold evaluations are suspended, consider the following options before you resume evaluations:

- Try to correlate the change in performance to configuration changes in NetOps Portal.
- Reduce the overall number of active event rules. Turn off event rules one at a time. Check the performance after you turn off each rule before turning off another rule.
- Reduce the overall number of active event rules that have windows greater than 300 seconds.
- Reduce the number of Violation event conditions within event rules.
- Reduce the number of event rules that use a condition type of Standard Deviation.
- Verify that only *required* collections are applied to the monitoring profile or threshold profiles that contains event rules.
- Verify that only *required* devices are contained within collections that are associated with these monitoring profiles or threshold profiles.

### **Threshold Limiter Behavior**

To determine whether to suspend threshold evaluations, the limiter looks at how long the engine takes to evaluate thresholds as a percentage of the poll cycle time:

- **Percentage of Poll Cycle Threshold**

Specifies the percentage of the poll cycle that can be used to monitor thresholds. By default, the Percentage of Poll Cycle Threshold attribute value is 80 percent.

For example, 4 minutes for items that are polled at a 5-minute rate. The threshold monitoring engine becomes DEGRADED when the engine takes more than 240 seconds to complete threshold evaluations. An event is generated on the Data Aggregator item when the threshold monitoring engine becomes DEGRADED.

- **Recovery Interval**

Specifies how long the threshold monitoring engine remains in the DEGRADED state. By default, the Recovery Interval attribute is 15 minutes. If the processing time does not drop below the specified percentage with the recover interval, threshold evaluations are suspended. An event is generated on the Data Aggregator item when threshold evaluations are suspended.

### **Resume Threshold Evaluations**

Manually resume suspended threshold evaluations.

#### **IMPORTANT**

If threshold evaluations are suspended frequently, contact Broadcom Support.

#### **Follow these steps:**

1. Navigate to the `http://DA_host:port/rest/thresholdmonitoring/config` URL.
2. Take note of the ID value of the `ThresholdMonitoringConfiguration` item.

#### **Example:**

```
<ThresholdMonitoringConfigurationList>
  <ThresholdMonitoringConfiguration version="1.0.0">
    <ID>16</ID>
    <ThresholdMonitoringEnabled>true</ThresholdMonitoringEnabled>
    <PercentOfPollCycleThreshold>80</PercentOfPollCycleThreshold>
    <ThresholdMonitoringLimiterEnabled>true</ThresholdMonitoringLimiterEnabled>
    <RecoveryIntervalInMinutes>15</RecoveryIntervalInMinutes>
  </ThresholdMonitoringConfiguration>
</ThresholdMonitoringConfigurationList>
```

3. Open a REST client editor or HTTP tool that sends requests and gets responses.
4. Set the Content-type to application/xml.
5. Enter the following filter criteria:

**URL:** `http://DA_host:port/rest/thresholdmonitoring/config/ID`

– **ID**

The identification number that is assigned to the ThresholdMonitoringConfiguration item.

**HTTP method = PUT**

Resume threshold evaluations by setting the value of the ThresholdMonitoringEnabled parameter to true on the **Body** tab of the HTTP Request pane:

```
<ThresholdMonitoringConfiguration version="1.0.0">
  <ThresholdMonitoringEnabled>true</ThresholdMonitoringEnabled>
</ThresholdMonitoringConfiguration>
```

Threshold evaluations resume.

### **Change the Default Behavior of the Threshold Limiter**

Modify the behavior of the threshold limiter only when instructed to do so by Broadcom Support.

**Follow these steps:**

1. Navigate to the `http://DA_host:port/rest/thresholdmonitoring/config` URL.
2. Note the ID value of the ThresholdMonitoringConfiguration item.
3. Open a REST client editor or HTTP tool that sends requests and gets responses.
4. Set the Content-type to application/xml.
5. Enter the following filter criteria:

**URL:** `http://DA_host:port/rest/thresholdmonitoring/config/ID`

– **ID**

The identification number that is assigned to the ThresholdMonitoringConfiguration item.

**HTTP method = PUT**

Increase the Percentage of Poll Cycle Threshold value (PercentOfPollCycleThreshold and RecoveryIntervalInMinutes ) on the **Body** tab of the HTTP Request pane:

```
<ThresholdMonitoringConfiguration version="1.0.0">
  <PercentOfPollCycleThreshold>
```

**percent**

```
</PercentOfPollCycleThreshold>
```

```
<RecoveryIntervalInMinutes>
```

**minutes**

```
</RecoveryIntervalInMinutes >
```

```
</ThresholdMonitoringConfiguration>
```

– **percent**

Specifies the percentage value.

– **minutes**

Specifies the number of minutes to wait before reassessing the threshold monitoring engine.

The threshold limiter runs with the updated values.

### **Disable the Limiter**

Disable the threshold limiter only when instructed to do so by Broadcom Support.

**Follow these steps:**

1. Navigate to the `http://DA_host:port/rest/thresholdmonitoring/config` URL.
2. Note of the ID value of the ThresholdMonitoringConfiguration item.
3. Open a REST client editor or HTTP tool that sends requests and gets responses.



4. Set the Content-type to application/xml.

5. Enter the following filter criteria:

**URL:** `http://DA_host:port/rest/thresholdmonitoring/config/ID`

– **ID**

The identification number that is assigned to the `ThresholdMonitoringConfiguration` item.

**HTTP method** = PUT

Disable the limiter by setting the value of the `ThresholdMonitoringLimiterEnabled` parameter to `false` on the **Body** tab of the HTTP Request pane:

```
<ThresholdMonitoringConfiguration version="1.0.0">
  <ThresholdMonitoringLimiterEnabled>false</ThresholdMonitoringLimiterEnabled>
</ThresholdMonitoringConfiguration>
```

The limiter is disabled and the evaluation engine cannot enter a DEGRADED state.

## Threshold Event Processing Self-Monitoring Metrics

To determine if you are doing too much eventing, monitor the key performance indicators in the data aggregator.

Eventing in the data aggregator is performed in batches, for example, events are simultaneously evaluated and generated for large groups of items. You can assess the health of the data aggregator using the self-monitoring metrics.

In general, steady values for these self-monitoring metrics indicate a healthy system. Some intensive database jobs cause fluctuation in these metrics. Typically, these jobs run between 2 AM and 4 AM UTC. Turn on eventing slowly and judge the system health before moving forward with different rules. Monitor the health of the system over 24 hours after each subsequent change.

Errors in the Karaf log on the data aggregator can also indicate that your system is under stress.

For more information about the threshold best practices, see [Configure Threshold Profiles](#).

To view these metrics, add a custom **IM Device MultiTrend** view to a dashboard. Edit the dashboard, using the following metrics from the **Data Aggregator Event Calculation Times** metric family:

- **Event Process Queue Size**  
This metric shows the size of the event processing queue. An increase in queue size without a subsequent recovery (trending downward) indicates that eventing is backed up.
- **Count of Cleared Events**  
This metric indicates the number of cleared events that are in the reporting resolution window.
- **Count of Created Events**  
This metric indicates the number of raised events that are in the reporting resolution window.  
A continuously large number of events that are raised or cleared can affect the Event Manager database. These metrics can indicate when your system has exceeded the recommended event generation rate. Event generation/clear bursts are acceptable.
- **Count of Processed Event Rule Evaluations**  
This metric indicates the sum of event rules multiplied by the number of items to which those rules are applied. The higher the number of evaluations, the more work your system is doing. Some evaluations are more expensive than others. For example, evaluations with more conditions, more standard deviation conditions, or longer duration and window are more expensive. The total acceptable number of evaluations depends on your event rules.
- **Total Time to Calculate Events**  
This metric indicates the total amount of time that was spent processing events for this metric family. If the value of this metric exceeds the number of seconds in the reporting resolution window, the eventing was delayed or backlogged at that point in time.

## Insights

Log Analytics in DX NetOps Insights is a purpose-built analytics capability that collects analyses and visualizes the log data generated by your applications and IT infrastructure to gain operational insights. With Log Analytics in DX NetOps Insights, you can diagnose and inspect the events that lead to a particular outage event using the logs in context of a device or alarm. You can stream system logs from devices into DX Operational Intelligence using the log collector, and then inspect the log lines contextually from alarms and inventory within NetOps Portal.

The following diagram shows the data flow of system logs into DX Operational Intelligence with log analytics for Insights:

You can forward, or stream, system logs to a DX Operational Intelligence log analytics cluster using the syslog (rsyslog) logging standard on Linux systems, and then serve syslog on-demand within NetOps Portal in context of a network device or an alarm. Syslog data includes messages with different kinds of information and includes an in-built severity level from 0 (Emergency) to 7 (Debug).

You can perform post-mortem after an incident using syslog monitoring. For example, you can proactively monitor syslog events, which can speed up the damage control process, saving minutes or even hours of downtime.

Get started by installing and configuring log analytics for Insights.

For more information, see [Install and Configure Log Analytics for Insights](#).

For more information about log analytics and the log collector, see Log Analytics in [the DX Operational Intelligence documentation](#).

## Use Log Analytics for Insights

### IMPORTANT

To use log analytics for Insights, the DX Operational Intelligence log analytics capabilities must be enabled (by installing and configuring log analytics for Insights) with NetOps Portal.

For more information, see [Install and Configure Log Analytics for Insights](#).

With log analytics for Insights, you can do the following:

- Forward, or stream, system logs from devices into DX Operational Intelligence, and then inspect the log lines contextually from alarms and inventory within NetOps Portal.  
For more information about the log collector, see Log Analytics in [the DX Operational Intelligence documentation](#).
- Test log analytics in the context of a device or an alarm.
- [View the log for a specific device in the context of an alarm](#).
- [View device log events](#).
- [Monitor Syslog events for specified log patterns](#).
- Filter logs (quick filter).
- Print logs to PDF or CSV.

### View the Log for a Device in Context of an Alarm

You can view the log for a specific device in the context of an alarm from the Alarm Console. NetOps Portal harvests the log events that were collected for the 15 minutes prior to the alarm and displays those in the log events table.

The following image shows the Alarm Console with an alarm selected, and the corresponding log showing for the last 15 minutes for an affected device relative to the alarm creation time in the **Log Events** view:

## View Device Logs Events

You can inspect the last 15 minutes of logs in the context of a device. On the device details page, click the **Log Events** tab. The **Log Events** view displays.

The following image shows an example of this view:

## Monitor Syslog Events for Specified Log Patterns

**Prerequisite:** For NetOps Portal to receive events of type syslog (Syslog events) from DX Operational Intelligence, log analytics for Insights must be installed and configured. For more information, see [Install and Configure Log Analytics for Insights](#).

You can monitor Syslog events for specified log patterns in NetOps Portal. You configure Syslog pattern match rules in DX Operational Intelligence. These rules create notifications that the OI Connector polls. The OI Connector creates Syslog events from these notifications and sends them to the event manager in NetOps Portal.

Syslog events display in the **Events** view, which you can add to dashboards.

The following image shows an example of this view:

For more information:

- About how to configure notifications for Syslog events, see [Configure Notifications](#).
- About how to define and configure log patterns, see Log Analytics in [the DX Operational Intelligence documentation](#).
- About how to configure the disposition of Syslog events, see [Configure the OI Connector](#).
- About the Syslog event type, see [Event Types](#).

## NetOps Flow

You can view flow from the NetOps Flow (flow) dashboards.

The NetOps Flow dashboards capture flow data using the Internet Protocol Flow Information Export (IPFIX) protocol. NetOps Flow stores data alongside performance data, allowing you to combine performance and flow in custom charts.

## Flow Dashboards

Flow is the consumption of data that comes from devices, such as routers and switches. You can view flow from the flow dashboards.

In this article:

- [Make the Flow Menu Available](#)
- [View the Dashboards](#)

### Make the Flow Menu Available

The **Flow** menu and the flow dashboards are hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

#### Follow these steps:

1. Hover over **Administration**, **User Settings**, and then click **Roles**.  
The **Manage Roles** page appears.
2. Select the role for which you want to show the **Flow** menu, and then click **Edit**.

The **Edit Role** page appears. The **Menu Set** is shown.

3. Click **Edit** to edit the menu set.  
The **Edit Menu Set** dialog opens.
4. Move the **Flow** menu from the **Available Menus** column to the **Selected Menus** column by selecting the menu, and then clicking the **Select item** arrow.
5. Click **OK**.

The **Flow** menu and the flow dashboards are now available.

### View the Dashboards

After you have made the **Flow** menu available, you can view flow from the flow dashboards. Access the dashboards by hovering over **Performance, Flow**.

The following dashboards are available:

- **Flow Dashboard**

This dashboard includes a Sankey diagram that shows the following:

- One-minute-aggregate flow data using the NetOps Portal default time range.
- Application mapping of flow traffic (by way of NBAR2 or user-defined application mapping), with the highest-bandwidth consumption from top to bottom.

**TIP**

You can customize the flow name displayed in the Sankey diagram.

For more information, see [Manage Application Mappings](#).

- Internet Protocol Flow Information Export (IPFIX) flow and the passing of data, showing device data stored in the data repository using the data aggregator inventory. Each flow shows the direction of data traffic. The line thickness is based on the amount of traffic across the port.

Use this dashboard to quickly identify abnormal traffic flows by grouping/ungrouping source, application, port, and direction. Use the toggles at the top of the dashboard to turn on or off the groupings based on what you want to see.

The following image shows an example of this dashboard:

- **Flow Statistics**

This dashboard shows the same information as the **Network Interface Performance** dashboard, which is an DX NetOps Network Flow Analysis-driven dashboard.

## View Network Flow

You can view the data that passes from a source to a destination associated for devices that export flow to NetOps Flow.

You can view network flow from a device, such as a router or switch, or interface, or interface's context page, on the **Network Flow** tab.

**Prerequisite:** NetOps Flow is installed and is collecting flow.

**Follow these steps:**

1. Complete the following based on the inventory type for which you want to view network flow:
  - **Device**
    - a. Hover over **Inventory, Items**, and then click **Devices**.  
The **Devices** page appears.
    - b. Click the device for which you want to view network flow. The device's context page opens to the **Details** tab.
  - **Interface**
    - a. Hover over **Inventory, Items**, and then click **Interfaces**.  
The **Interfaces** page appears.

- b. Click the interface for which you want to view network flow. The interface's context page opens to the **Details** tab.
2. Click the **Network Flow** tab.  
This tab includes the following views:
  - **Top Network Flows**  
This view includes a Sankey diagram depicting the flows that are going through the device/interface, by default, sorted by protocol. Each flow data record represents an application. The thickness of each flow reflects the amount of data (bytes) that is passing in the flow.  
You can control what is shown in the diagram using the following checkboxes:
    - **Show Source:** Shows the source IP address and the interface on the device/interface receiving the flow for each flow.
    - **Show Protocol:** Shows the flows by the type of protocol (for example, UDP, TCP, and IP) that the flow data contains. Selected by default.
    - **Show Application Name:** Shows the application name for each flow.
    - **Show Port:** Shows the destination port for each flow.
    - **Show In/Out:** Shows the flow from the destination to the source (in) and from the source to the destination (out) for each flow.
 Clicking a flow in this view shows the details for that flow in the **Top Flows by Protocol** view.
  - **Top Flows by Protocol**  
This view includes a table that displays the details about the flows that are going through the device/interface, such as the protocol type and the total amount of data passing in the flows (bytes). The information displayed in this view is dependent on the checkboxes that you have selected for the **Top Network Flows** view. For example, if you have selected only the **Show Protocol** checkbox for the **Top Network Flows** view, the protocol type, the out interface, the destination IP address, and the total bytes for the flow are shown. Selecting all checkboxes shows the source IP address, the in interface, the protocol type, the application name, the destination port, the out interface, the destination IP address, and the total bytes for the flow.  
Clicking a flow in this view shows the details for that flow in the **Selected Flow Over Time** view.
  - **Selected Flow Over Time**  
This view shows a line chart depicting the recorded total number of bytes for a flow over time for a specific source and destination. The source and destination IP addresses for the flow are shown at the top of the view.  
The following image shows an example of these views:

## Update the NetOps Flow Configuration

When required, such as when a credential to access the Data Aggregator/NetOps Portal and Data Repository services changes, update the NetOps Flow configuration.

You can update the NetOps Flow configurations that you used with the `helm install` commands post-deployment. For example, you can update the configuration to change the credentials to access those external services, such as the data aggregator/NetOps Portal and the data repository.

### NOTE

You can also enable HTTPS for NetOps Flow using this procedure.  
For more information, see [Enable HTTPS for NetOps Flow](#).

### Example:

The following example updates the NetOps Flow configuration with new credentials:

```
helm -n <NAMESPACE> upgrade --reuse-values \
--set global.pc.adminUser=<PC_USER> \
--set global.pc.adminPassword=<PC_PASSWORD> \
```

```
--set global.dr.dbUser=<DR_DB_USER> \
--set global.dr.dbPassword=<DR_DB_PASSWORD>
```

### Example:

```
helm -n default upgrade --reuse-values \
--set global.pc.adminUser=admin2 \
--set global.pc.adminPassword=mypass2 \
--set global.dr.dbUser=user2 \
--set global.dr.dbPassword=pass2
```

- **NAMESPACE**  
The Kubernetes namespace to where the NetOps Flow components will be deployed. To use the default namespace, omit the `-n <NAMESPACE>` option.  
**Default:** default
- **PC\_USER**  
The name of the NetOps Portal user.  
**Default:** admin
- **PC\_PASSWORD**  
The password for the NetOps Portal user.  
**Default:** admin
- **DR\_DB\_USER**  
The name of the data repository database user.  
**Default:** dauser
- **DR\_DB\_PASSWORD**  
The password for the data repository database user.  
**Default:** dapass

The components are restarted, and their state change back to "Running" within a few minutes.

## Modern Network Monitoring

DX NetOps Virtual Network Assurance (VNA) provides modern network monitoring for software-defined architectures and hybrid cloud platforms.

VNA collects data from traditional, software-defined networking (SDN), software-defined data center (SDDC), software-defined Wide Area Network (SD-WAN), network functions virtualization (NFV), and hybrid-cloud network architectures, and delivers that information to DX NetOps Performance Management. You can monitor the VNA data related to the relevant technologies and architectures of your environment, as well as the other VNA data and performance information from the dashboards in NetOps Portal. These dashboards are populated after you have integrated with VNA and configured the relevant plug-ins.

For more information:

- About VNA, see [the DX NetOps Virtual Network Assurance documentation](#).
- About how to integrate with VNA, see [Integrate with Virtual Network Assurance](#).

## Monitor SDN/NFV Virtual Inventory

You can view your virtual inventory on a dashboard in NetOps Portal.

In software-defined networking (SDN) and network functions virtualization (NFV) environments, your virtual inventory changes according to the requirements of your services. Modern network monitoring provides insight into the trends in virtual inventory. Inventory metrics provide assurance that the programmatic processes that manage the network are functioning properly. These metrics also highlight anomalies in provisioning.

The **SDN/NFV Virtual Inventory Overview** dashboard shows your virtual inventory, and provides the following information:

- Virtual network function (VNF) count by type over time as a trend chart and a stacked chart.
- Service chain and virtual network inventory trends.
- VNF inventory trends by type.

#### NOTE

The **Other VNFs Count** view on this dashboard shows VNFs where NetOps Portal cannot determine the type.

To view this dashboard, go to **Performance, SDN/NFV Reporting, SDN/NFV Virtual Inventory Overview**.

#### NOTE

The **SDN/NFV Reporting** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate a menu with a role, see [Manage User Account Roles](#).

## Monitor SDN/NFV Virtual Resource Usage

Investigate performance degradation and perform capacity planning by monitoring the resource usage of your virtual environment using these dashboards.

You can monitor SDN/NFV virtual resource usage using the following dashboards:

- **SDN/NFV Virtual Compute Usage Overview**

The views on this dashboard show virtual CPU and memory usage for virtual network functions (VNFs). The views use different visualizations to show aggregated and individual resource utilization. By default, views show the VNFs with the highest utilization.

- **SDN/NFV Virtual Storage Usage Overview**

The views on this dashboard show virtual storage, reads, and writes for VNFs. The views use different visualizations to show aggregated and individual resource utilization. By default, views show the VNFs with the highest utilization. The views use virtual disk usage data on existing SDN devices only from the SDN Devices Metrics metric family that it receives from VNA.

To view these dashboards, go to **Performance, SDN/NFV Reporting, SDN/NFV Virtual Compute Usage Overview | SDN/NFV Virtual Storage Usage Overview**.

#### NOTE

The **SDN/NFV Reporting** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate a menu with a role, see [Manage User Account Roles](#).

The following video highlights the key capabilities of these dashboards:

## Monitor SDN/NFV Physical Host Resource Usage

Investigate performance degradation and for capacity planning by monitoring the physical resource usage of your virtual environment.

In a virtual network, the virtual network functions (VNFs) run on physical hosts. If the load on the physical hosts is too high, the VNFs compete for resources, leading to performance degradation. The **SDN/NFV Physical Server Usage Overview** dashboard shows the status of the physical infrastructure that supports the virtual environment.

This dashboard includes the following information:



- Aggregated CPU, memory, and disk utilization for physical servers
- Availability and Reachability information
- Individual servers with the highest CPU, memory, and disk utilization
- Inventory of all servers in the SDN/NFV environment

To view this dashboard, go to **Performance, SDN/NFV Reporting, SDN/NFV Physical Server Usage Overview**.

#### NOTE

The **SDN/NFV Reporting** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate a menu with a role, see [Manage User Account Roles](#).

The following video highlights the key features of this dashboard:

## Monitor SDN/NFV vSwitch Performance

Identify bottlenecks in your SDN or NFV environment by monitoring vSwitch performance.

Because virtual switches (vSwitches) connect virtualized network functions (VNFs) to each other and to physical interfaces, the performance of the vSwitch determines the overall bandwidth of the connected VNFs. With the transition from physical to virtual networking, the generalized architecture of the x86 box replaces the highly-specialized hardware of the physical switch. Because this architecture is not optimized for switch performance, vSwitch performance is an important indicator of overall virtual network health.

#### NOTE

In NetOps Portal, virtual routers (vRouters) appear as vSwitches.

In this article:

- [Analyze the Overall Performance of a vSwitch](#)
- [Get Detailed Information about a vSwitch](#)
- [Get Performance Information for Interfaces on a vSwitch](#)
- If you have integrated with DX NetOps Virtual Network Assurance (VNA) and you have configured the VMware NSX-T plug-in, [View NSX-T BGP Sessions for Tier-0 Virtual Switches](#)

### Analyze the Overall Performance of a vSwitch

You can analyze the overall performance of the vSwitches in your environment using the **SDN/NFV vSwitch Performance Overview** dashboard. This dashboard uses different visualizations to show the following information about vSwitches:

- The inventory of vSwitches in the virtual network
- The vSwitch throughput in bits and packets
- The minimum, maximum, and average throughput in bits and packets
- The vSwitch utilization and throughput in bits and packets aggregated to the device level
- The CPU and memory utilization of vSwitch service

#### TIP

The views that show vSwitch network performance, such as throughput, errors, and discards, show the top vSwitches for these metrics. You can get detail information about a vSwitch by clicking a vSwitch from these views.



### **Get Detailed Information about a vSwitch**

The **vSwitch** context page provides detailed information about a specific vSwitch. To access this page, click a vSwitch from the inventory list (on the **SDN/NFV vSwitch Performance Overview** dashboard), in another view, or in the service chain visualization.

From the **vSwitch** context page, you can access the information from the following tabs:

- **Details tab**  
Shows basic information about the vSwitch, a list of events, and availability information. Use this tab for basic troubleshooting.
- **Network tab**  
Shows network performance information for the vSwitch. Use this tab to identify network performance bottlenecks that are related to the vSwitch. The tab includes the following information:
  - Average interface utilization
  - Errors and discards for interfaces on the vSwitch
  - Aggregated throughput in bits and packets
  - Interface throughput in bits and packets
  - Inventory of the virtual interfaces

High dropped packets might indicate that the vSwitch service is overburdened.
- **Resources tab**  
Shows resource use for the vSwitch service. Use this tab to ensure that the vSwitch is not consuming too much of the host resources. The tab includes the following information:
  - Current CPU and memory utilization

**NOTE**  
CPU and memory utilization for the vSwitch service are compared against the total resources available to the host.

  - Trend chart for CPU and memory utilization

High utilization for the service means that other VMs on the host might have insufficient resources.

### **Get Performance Information for Interfaces on a vSwitch**

The **Virtual Interface** context page provides performance information for interfaces on a vSwitch. To access this page, go to **Inventory, Items, Virtual Interfaces**, and then click a virtual interface from the **Virtual Interfaces** list.

This context page includes the following tabs:

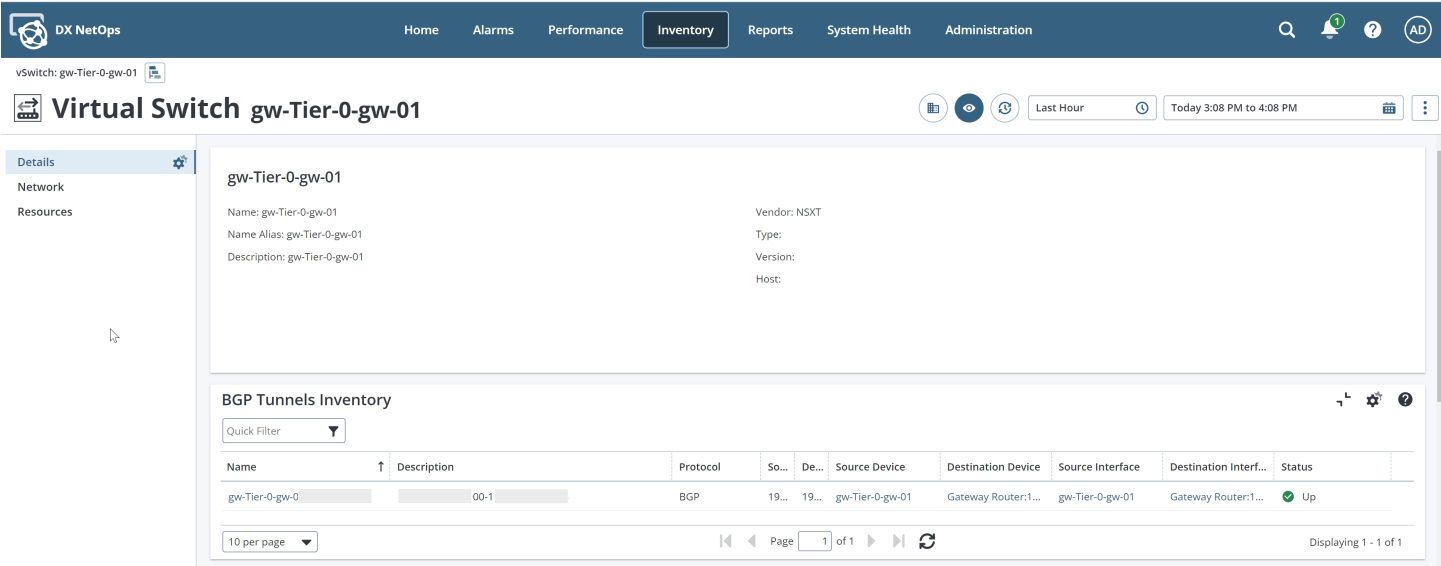
- **Details Tab**  
Shows basic information about the interface and a list of events for to the interface.
- **Health Tab**  
Shows performance information for the virtual interface, such as discards, errors, interface utilization, and throughput in bits and packets.

### **View NSX-T BGP Sessions for Tier-0 Virtual Switches**

If you have integrated with DX NetOps Virtual Network Assurance (VNA) and you have configured the VMware NSX-T plug-in, you can view a list of NSX-T BGP sessions and their statuses for Tier-0 virtual switches (vSwitches) from the **Virtual Switch** context page. To view the list, open a context page for a Tier-0 vSwitch that an NSX-T plug-in monitors.

The following image shows an example of the **Virtual Switch** context page:

Figure 93: Virtual Switch context page



Monitor Service Chains

In software-defined networking (SDN)/network functions virtualization (NFV) environments, services are delivered dynamically through virtual network functions (VNF). The service chain represents the required VNFs and the stack of physical and virtual building blocks that support the VNF.

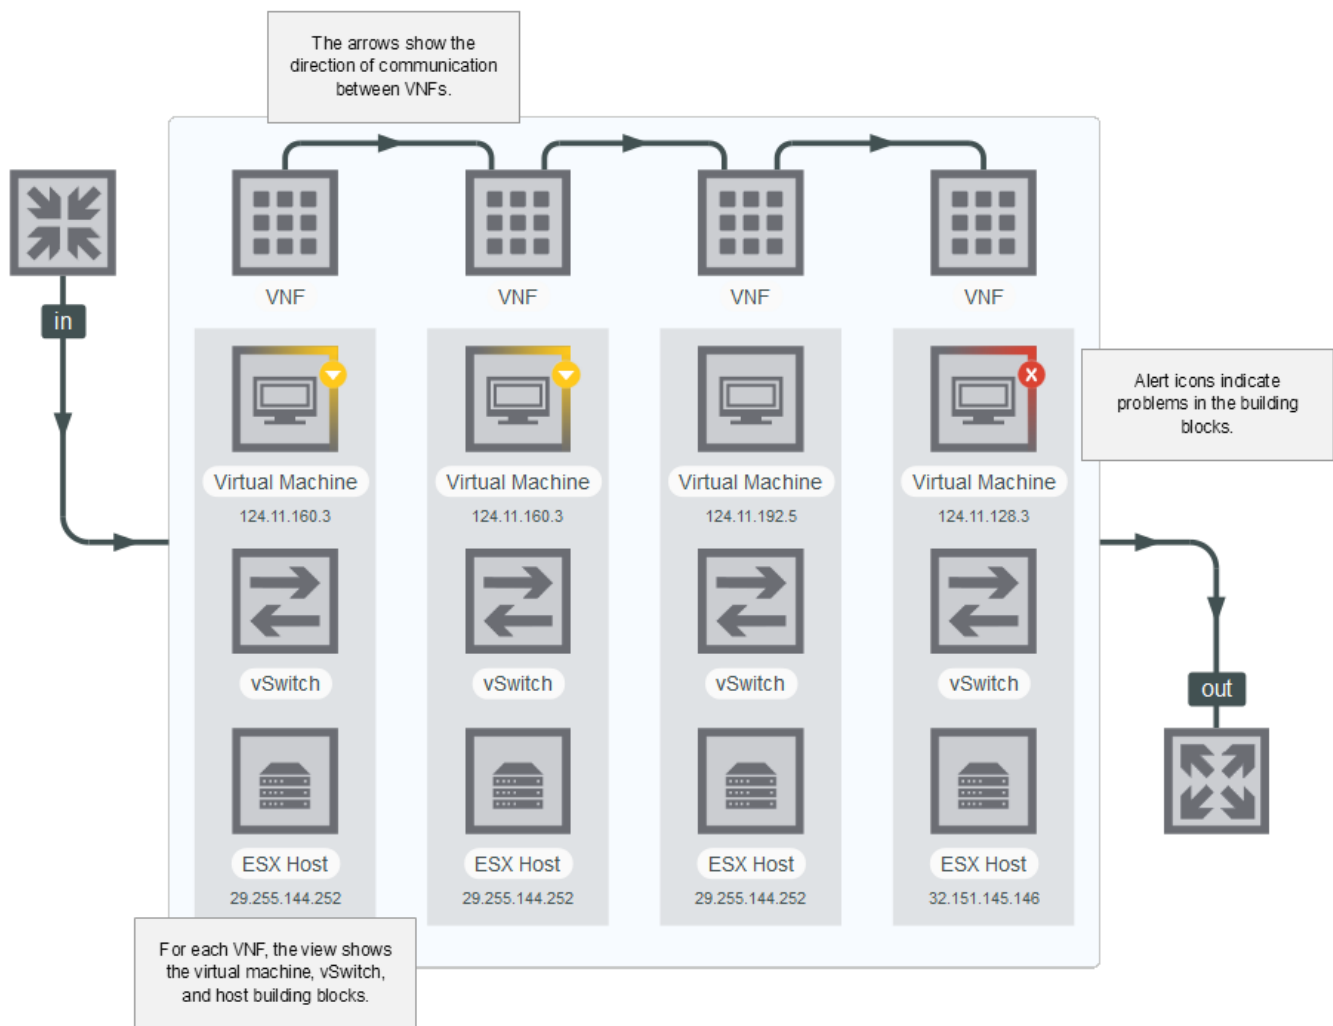
The **Service Chain** page provides information about the status of service chains. To access this page, go to **Inventory**, **Service Chains**, and then click a service chain in the list.

NOTE

The **Service Chains** menu is a shown option only if there are service chains in NetOps Portal.

The main view shows the VNFs in the service chain and includes the virtual and physical building blocks. The view indicates the type of each VNF with an abbreviation. The arrows show whether communication is one-way or bidirectional.

The following image shows the important elements of the service chain visualization:

**Figure 94: Service Chain Details Elements**

The view provides the following functionality:

- Show name, type, IP address, and performance information by hovering over a building block in the service chain.
- Alert icons indicate problems in the building blocks. Hover over the icon for information about the violated threshold.
- Load the detailed context page for a building block by right-clicking a building block, and then selecting a context.

### **Configure the Thresholds for the Alert Icons**

Click the **View Settings** (gear) icon, and then click **Edit**. Assign Critical (red), Major (orange), and Minor (yellow) threshold for the following metrics:

- CPU utilization
- Memory utilization

### **VNF Type Abbreviations**

The view uses the following abbreviations to identify the currently defined VNF types:

- vFW Firewall
- vLB Load Balancer
- vNAT NAT
- vOPT Optimizer
- vCMF Content Filter
- vDPI DPI
- vWOL WAN Accelerator
- vADC ADC
- vCACHE Cache
- vROUTER Router
- Other

## Monitor AWS

You can use the custom dashboards and context pages in NetOps Portal to monitor Amazon Web Services (AWS).

DX NetOps Virtual Network Assurance (VNA) verifies that the AWS Cloud environment is operating as expected and highlights problem areas by monitoring the environment.

You can monitor the following items from the dashboards and context pages:

- **Tunnel**  
Represents the connection between two devices and has the following polled statistics:
  - TunnelState
  - TunnelDataIn
  - TunnelDataOut
- **Ec2**  
Represents a virtual machine and has the following polled metrics:
  - CPU
    - Utilization
    - Aggregated Network Incoming Bytes
    - Aggregated Network Outgoing Bytes
    - Aggregated Network Incoming Packets
    - Aggregated Network Outgoing Packets
  - Disk
    - Disk IOPS
    - Disk Read Bytes
    - Disk Write Bytes
    - Disk Capacity

### AWS Context Pages

From the inventory, you can navigate to context pages for more information about the following AWS items:

- **Devices**  
View the device context page for any of your AWS devices. To access the device context page, hover over **Inventory**, **Items**, click **Devices**, and then click a device from the list.
- **Tunnels**

**NOTE**

The context page for AWS tunnels do not include AWS cloud metrics. Create custom dashboards to view tunnel performance metrics.

- **EC2s (Virtual Machines)**

From the **Summary** context tab, you can view the details for the selected EC2 with the corresponding event list. To access the context tabs for a virtual machine, hover over **Inventory**, **Items**, click **Devices**, and then click a virtual machine from the list.

From the **VMware Virtual machine** context tab, you can view trend charts for CPU Usage, Active Memory Usage, and Power State Connection state.

**TIP**

You can create custom dashboards to view the supported EC2 performance metrics.

- **Virtual Interfaces**

From the **Details** context tab, you can view the details for the selected virtual interface from the **Event List** view. To access the **Details** context tab, hover over **Inventory**, **Items**, click **Virtual Interfaces**, and then click a virtual interface from the list.

## Monitor Cisco ACI

Cisco Application Centric Infrastructure (ACI) uses Nexus 9000 switches to create a dynamic virtual environment that hosts and serves applications. Monitor the Cisco ACI environment to verify that everything is operating as expected and to highlight problem areas. NetOps Portal provides dashboards and context pages that support Cisco ACI monitoring.

In this article:

- [ACI Console](#)
- [ACI Health Dashboard](#)
- [Switches Overview Dashboard](#)
- [ACI Workflows](#)

### **ACI Console**

The **ACI Console** provides a searchable inventory list that shows relationships between items in the Cisco ACI environment. Where relevant, the inventory shows the most recent ACI health score:

- To limit the list to a particular node in the hierarchy, double-click that node.
- To return to a higher level of the hierarchy, click the breadcrumb links.
- To filter the inventory by a health-score threshold, use the slider at the top of the pane. The inventory shows nodes with children with health scores equal to or lower than the threshold.

The **ACI Console** also provides relationship diagrams for various aspects of the ACI inventory. The diagrams show alert icons on items that exceed customizable thresholds. Use the diagrams to understand how particular problems with the ACI infrastructure affect your applications.

The **ACI Console** includes diagrams for the following inventory item types:

- **Application Profile Diagram**  
Shows the relationships between endpoint groups (EPGs) that support the application profiles (APs). Provider EPGs are connected to contracts connected to consumer EPGs. The diagram also shows the underlying leafs that support the endpoints in the EPGs.
- **EPG Diagram**  
Shows the AP parent and each end point. For each endpoint, the diagram shows the associated leaf.
- **Endpoint Diagram**

Shows the EPG parent. If the endpoint is a vLAN, the diagram shows the connection to the leaf. If the endpoint is an application, the diagram shows the full supporting technology stack.

- **Leaf Diagram**

Shows application profiles that contain endpoints which are connected to the leaf.

- **APIC Diagram**

Shows leaves that are connected to the APIC controller.

**NOTE**

Because each spine is connected to each leaf and has no other relationships, the **ACI Console** does not include a diagram for spines.

You can do the following from the diagrams:

- To open a diagram, click an item in the inventory.
- To show more details for an item in the relationship map, hover over the icon for the item.
- To show which thresholds are exceeded, hover over the alert icon in the relationship map.
- To open the relationship map for another item, click the icon for an item in the current relationship map. The previous relationship map becomes a breadcrumb link on the right side of the diagram.
- To open the context page for an item, click the item name in the relationship map.
- To go to the **ACI Health** dashboard or the **Switches Overview** dashboard, click links in the upper left corner of the dashboard.

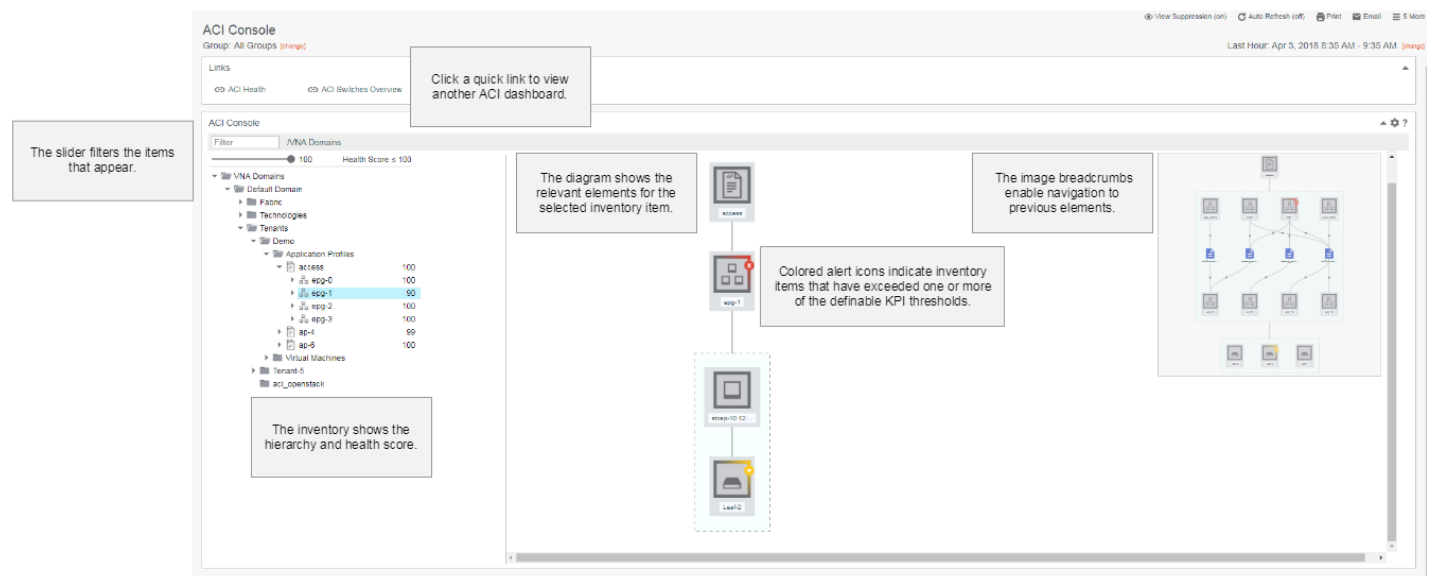
## ACI Console Thresholds

To configure the thresholds for the alert icons, click the **View Settings** (gear) icon, and then click **Edit**. Assign Critical (red), Major (orange), and Minor (yellow) threshold for the following metrics:

- CPU utilization
- Health score
- Interface utilization
- Memory utilization

The following image shows the important elements of the **ACI Console**:

**Figure 95: ACI Console Elements**



## **ACI Health Dashboard**

The **ACI Health** dashboard provides focused information about your Cisco ACI environment. To highlight problem areas, the dashboard focuses on health scores and faults. To narrow the area of troubleshooting, change the context of the dashboard:

- To find problematic switches, scope to the fabric.
- To find a problem in a tenant, scope to a tenant.
- To find a problem in an application, scope to an AP.

This dashboard shows the following information:

- Counts of items in the Cisco ACI environment
- Count of critical faults
- ACI health scores for items in the environment
- Aggregated fault trends by severity
- AP and EPG inventory

## **Switches Overview Dashboard**

The **Switches Overview** dashboard provides information about the status of switches. If you select the appropriate group of Nexus 9000 switches, you can monitor the infrastructure behind an ACI environment. For example, you can analyze critical faults, CPU, memory, and interface utilization for these switches.

### **NOTE**

Most metrics on this dashboard require SNMP data collection from the Nexus 9000 switches through DX NetOps Performance Management.

This dashboard shows the following information:

- Switches with the highest CPU and memory utilization
- Switches with the highest interface utilization
- Switches with the highest policy CAM utilization
- Switches with the most critical faults

## **ACI Workflows**

The dashboards apply to the following role-based workflows:

- ACI System Administrators are responsible for maintaining the Cisco ACI environment.
  - a. Use the **ACI Console** to look at the fabric and VMs:
    - View the relationship of the fabric to APs and EPGs.
    - View the relationship of fabric nodes to compute resources.
    - View the relationship of VMs to APs and EPGs.
    - View the relationship of VMs to fabric nodes and compute resources.
  - b. Use the ACI inventor to track the number of entities in the ACI environment.
  - c. Track the performance of APs and EPGs on context pages.
- Tenant Administrators are responsible for managing tenants in the ACI environment.
  - a. Scope the **ACI Console** to the tenant, and look at VM relationships to APs, EPGs, compute resources, and the fabric.
  - b. Use the **ACI Health** dashboard to look at the health score and faults within the tenant.
  - c. Track the performance of APs and EPGs on context pages.
- Application Owners are responsible for a single application that is hosted in the ACI environment.
  - a. Scope the **ACI Console** to the AP, and look at VM relationships to the AP, EPGs, compute resources, and the fabric.

- b. Use the AP context page to monitor computer and storage utilization by VM and in aggregate and top VM utilization.
- Fabric Administrators are responsible for maintaining the health of the network in the ACI environment.
  - a. Scope the **ACI Console** to the fabric, and look at fabric relationships to APs and EPGs and the relationships of fabric nodes to compute resources.
  - b. Use the **ACI Health** dashboard to look at the health score and faults within the fabric.
  - c. Track the performance of the Nexus 9000 switches on the **Switches Overview** dashboard and the context pages for the switches.

## Monitor Cisco DNA Center

If you have integrated with DX NetOps Virtual Network Assurance (VNA) and you have configured the Cisco DNAC plug-in, you can monitor Cisco DNA Center (DNAC) performance metrics using the dashboards, views, and context pages in NetOps Portal.

Monitor DNAC performance metrics from the following dashboards:

- [Cisco DNA Center Console](#)
- [Cisco DNA Center Summary](#)

To view these dashboards, go to **Performance, Cisco DNA Center Reporting**.

### NOTE

The **Cisco DNA Center Reporting** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate a menu with a role, see [Manage User Account Roles](#).

For more information:

- About the DNAC performance metrics that VNA collects using the Cisco DNAC plug-in, see [the VNA documentation](#).
- About Cisco DNA Center, see [the Cisco DNA Center documentation](#).

### Cisco DNA Center Console

The **Cisco DNA Center Console** dashboard provides the following information:

- A list of the monitored VNA domains (for example, from DNAC environments).

### NOTE

The view on this dashboard shows only VNA domains related to Cisco DNAC plug-in instances.

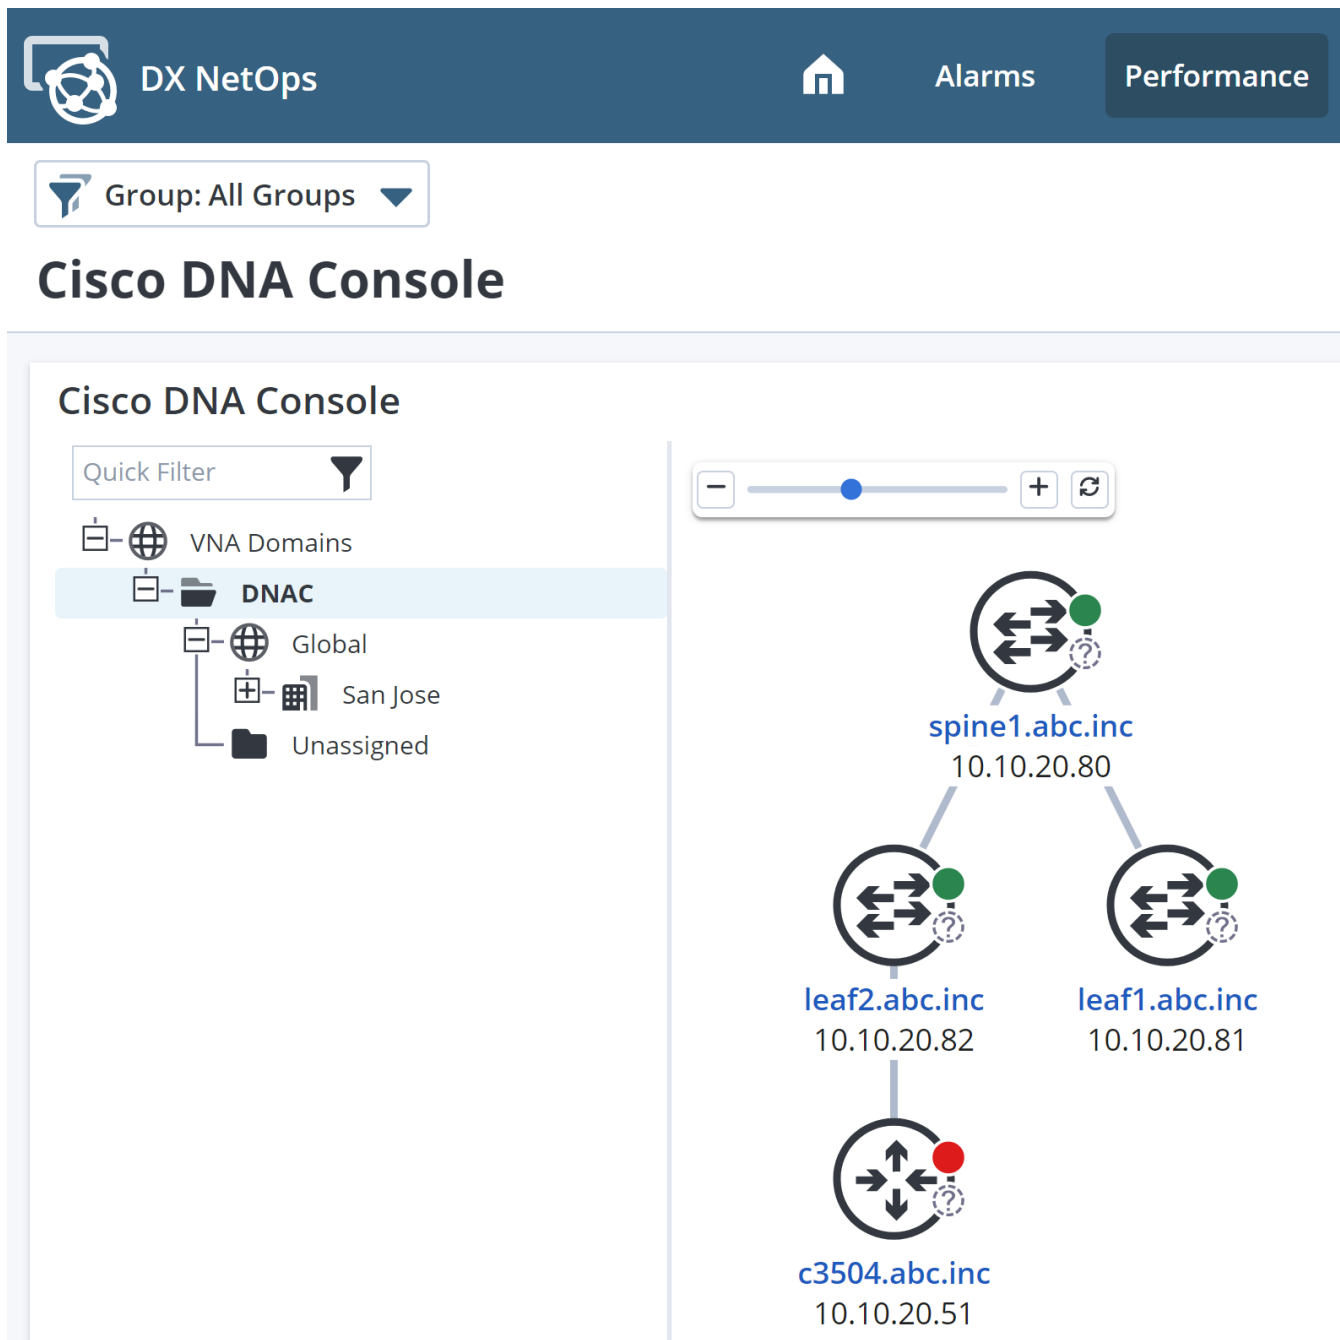
**Prerequisite:** To view VNA domains in the view, you have access to the groups where the DNAC VNA domains exist.

- An overview of the DNAC entities and the relationship of the network fabric.

The following image shows an example of the **Cisco DNA Center Console** dashboard showing the topology of monitored Cisco DNAC devices:



Figure 96: Cisco DNA Center Console page



### Cisco DNA Center Summary

The **Cisco DNA Center Summary** dashboard displays views for monitoring DNAC inventory.

**Prerequisite:** To view the content in the views on this dashboard, you have selected a group/domain that has DNAC inventory.

This dashboard includes the following views:

- **Cisco DNA (Inventory)**

This view shows DNAC inventory for the selected group, classified into the following table views:

### – Devices

The following image shows an example of the **Devices** view within the **Cisco DNA (Inventory)** view:

**Figure 97: Devices View**

Cisco DNA (Inventory)

Devices

Automatic Group: DNAC Domain

Quick Filter ▼ On Demand

<input type="checkbox"/>	Name	Type	Domain	Address	Description	Current Al...	Life Cycle S...	Context Types	DevCustom1
<input type="checkbox"/>		Router	Default Domain		Description - Cisco Catalyst 4500 Router - 5	✓ Active	Active	Router	
<input type="checkbox"/>		Router	Default Domain		Description - Cisco Catalyst 6800 Router - 2	✓ Active	Active	Router	
<input type="checkbox"/>		Switch	Default Domain		Description - Cisco Catalyst 9300 Switch - 2	✓ Active	Active	Switch	
<input type="checkbox"/>		Router	Default Domain		Description - Cisco Catalyst 9500 Router - 3	✓ Active	Active	Router	
<input type="checkbox"/>		Switch	Default Domain		Description - Cisco Catalyst 9600 Switch - 4	✓ Active	Active	Switch	
<input type="checkbox"/>		Switch	Default Domain		Description - Cisco Catalyst Digital Building Switch - 1	✓ Active	Active	Switch	
<input type="checkbox"/>		Switch	Default Domain		Description - Cisco Catalyst Micro Switch - 5	✓ Active	Active	Switch	
<input type="checkbox"/>		Server	Default Domain			✓ Active	Active	Server	

Page 1 of 1

Displaying 1 - 16 of 16

### – Wi-Fi Devices

The following image shows an example of the **Wi-Fi Devices** view within the **Cisco DNA (Inventory)** view:

**Figure 98: Wi-Fi Devices View**

DX NetOps

All Groups > VNA Domains > DNAC Domain

Cisco DNA Center Summary

☐ Cisco Catalyst 2960-X Switch - 3
 Switch
 Default Domain
 10.214.27.224
 Description - Cisco Catalyst 2960-X Switch - 3
 Critical
 Active
 Switch

Page 1 of 1

Displaying 1 - 16 of 16

Wi-Fi Devices

Automatic Group: DNAC Domain

Quick Filter ▼ On Demand

<input type="checkbox"/>	Name	Domain	Address	Description	Current Alar...	Life Cycle State	Context Types	DevCustom1
<input type="checkbox"/>	Catalyst 9105 Unified Access Point - 11	Default Domain		Description - Catalyst 9105 Unified Access Point - 11	▲ Critical	✓ Active	Device, Wireless A...	
<input type="checkbox"/>	Catalyst 9105 Unified Access Point - 9	Default Domain		Description - Catalyst 9105 Unified Access Point - 9	▲ Critical	✓ Active	Device, Wireless A...	
<input type="checkbox"/>	Catalyst 9120 Unified Access Point - 13	Default Domain		Description - Catalyst 9120 Unified Access Point - 13	▲ Critical	✓ Active	Device, Wireless A...	
<input type="checkbox"/>	Catalyst 9120 Unified Access Point - 2	Default Domain		Description - Catalyst 9120 Unified Access Point - 2	▲ Critical	✓ Active	Device, Wireless A...	
<input type="checkbox"/>	Catalyst 9120 Unified Access Point - 23	Default Domain		Description - Catalyst 9120 Unified Access Point - 23	▲ Critical	✓ Active	Device, Wireless A...	
<input type="checkbox"/>	Catalyst 9130 Unified Access Point - 1	Default Domain		Description - Catalyst 9130 Unified Access Point - 1	▲ Critical	✓ Active	Device, Wireless A...	
<input type="checkbox"/>	Catalyst 9130 Unified Access Point - 15	Default Domain		Description - Catalyst 9130 Unified Access Point - 15	▲ Critical	✓ Active	Device, Wireless A...	
<input type="checkbox"/>	Catalyst 9136 Unified Access Point - 12	Default Domain		Description - Catalyst 9136 Unified Access Point - 12	▲ Critical	✓ Active	Device, Wireless A...	

Page 1 of 1

Displaying 1 - 30 of 30

### – Virtual Interfaces

The following images show an example of the **Virtual Interfaces** view within the **Cisco DNA (Inventory)** view:

**Figure 99: Virtual Interfaces View**

**Virtual Interfaces**

Automatic Group: DNAC Domain

Quick Filter On Demand

<input type="checkbox"/>	Name	Description	Device	Life Cycle State	Domain	IP Address
<input type="checkbox"/>	router Interface 683071/3/31	router Interface 683071/3/31	Cisco Catalyst 9500 Router - 3	Active	Default Domain	
<input type="checkbox"/>	router Interface 762905/25/1	router Interface 762905/25/1	Cisco Catalyst 6800 Router - 2	Active	Default Domain	
<input type="checkbox"/>	router Interface 765686/6/16	router Interface 765686/6/16	Cisco Catalyst 4500 Router - 5	Active	Default Domain	
<input type="checkbox"/>	router Interface 770351/26/20	router Interface 770351/26/20	Cisco 2900 Router - 4	Active	Default Domain	
<input type="checkbox"/>	router Interface 774409/8/22	router Interface 774409/8/22	Cisco Catalyst 9500 Router - 3	Active	Default Domain	
<input type="checkbox"/>	switch Interface 201575/31/21	switch Interface 201575/31/21	Cisco Catalyst Micro Switch - 5	Active	Default Domain	
<input type="checkbox"/>	switch Interface 266238/20/13	switch Interface 266238/20/13	Cisco Catalyst 9300 Switch - 2	Active	Default Domain	
<input type="checkbox"/>	switch Interface 295723/20/5	switch Interface 295723/20/5	Cisco Catalyst Micro Switch - 5	Active	Default Domain	

Page 1 of 1

Displaying 1 - 30 of 30

- Cisco DNA (Alarm Table View)**

This view shows the alarms corresponding to DNAC inventory for the selected group.

The following image shows an example of this view:

**Figure 100: Cisco DNA (Alarm Table View)**

**Cisco DNA (Alarm Table View)**

Automatic Group: DNA Central Time Range: No Time Range

No Filter Quick Filter Acknowledge Unacknowledge Clear Troubleshooter Poll Ping Traceroute On Demand Manage Life Cycle

<input type="checkbox"/>	Severity	Date/Time	Item Name	Model Type	IP Ad...	Alarm Title	Impact	Number of Occ...	Acknowledged	Troubleshooter	Trouble Ticket ID
<input type="checkbox"/>	Critical	Aug 1, 2023 3:4...	Catalyst 9136 U...	IP Device		SDN Model OPERATIONAL STATUS IS REPORTED AS DOWN BY SDN GATEWAY	0	1			Ticket Number L...
<input type="checkbox"/>	Critical	Aug 1, 2023 3:4...	Catalyst 9136 U...	IP Device		SDN Model OPERATIONAL STATUS IS REPORTED AS DOWN BY SDN GATEWAY	0	1			Ticket Number L...
<input type="checkbox"/>	Critical	Aug 1, 2023 3:4...	Catalyst 9120 U...	IP Device		SDN Model OPERATIONAL STATUS IS REPORTED AS DOWN BY SDN GATEWAY	0	1	Yes		Ticket Number L...
<input type="checkbox"/>	Critical	Aug 1, 2023 3:4...	Catalyst 9120 U...	IP Device		SDN Model OPERATIONAL STATUS IS REPORTED AS DOWN BY SDN GATEWAY	0	1			Ticket Number L...

100 per page Page 1 of 34

Displaying 1 - 100 of 3380

Select an alarm to view its details.

**Alarm Details**

- Impact: Management Lost
- Impact: Symptoms
- Neighbor Topology
- Interfaces
- Events
- Log Events

## Monitor SD-WAN Devices

SD-WAN solutions, such as Viptela and Cisco IWAN, route traffic according to predefined performance requirements. If you have integrated with DX NetOps Virtual Network Assurance (VNA) and you have configured a VNA plug-in, you can monitor SD-WAN devices using the dashboards and context pages in NetOps Portal.

Monitor the health of your SD-WAN tunnels or applications using the **SD-WAN Tunnel Statistics** dashboard. This dashboard reports metrics for jitter, latency, and packet loss. You can also monitor the health of your SD-WAN tunnels or applications, as well as identify issues impacting the delivery of services and applications according to service level

agreements (SLAs), using the **SD-WAN App Path Statistics** dashboard. This dashboard reports metrics for Percentage of Jitter SLA Threshold, Percentage of Packet Loss SLA Threshold, and Percentage of Latency SLA Threshold.

To view these dashboards, go to **Performance, SD-WAN Reporting, SD-WAN Tunnel Statistics | SD-WAN Application Statistics**.

#### NOTE

The **SD-WAN Reporting** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate a menu with a role, see [Manage User Account Roles](#).

In this article:

- [Monitor SD-WAN Tunnels and Applications](#)
- [View the Aggregator Health Counts](#)
- [View the Location of a Site](#)
- [View Metrics from Time Bar Charts](#)
- [View Metrics from the Scorecard](#)
- [View Metrics from Tables with Color-Coded Bars](#)
- [View Metrics from Gauge/Table Views](#)
- [View Metrics from Trend/Table Views](#)
- [View Side-by-Side Comparison of Metrics from Trend Charts](#)
- [View Source and Destination Endpoint](#)
- [View Errors and Discards in Separate Tables](#)
- [Navigate to More Details from the Dashboards](#)

## Monitor SD-WAN Tunnels and Applications

You can monitor the tunnels and application/SLA paths using the dashboards and context pages in NetOps Portal. Tunnels represent the connection between two devices and has polled statistics (jitter, latency, and packet loss). Application/SLA paths represent tunnels in relation to service level agreement (SLA) thresholds. The measured SLA metrics (jitter, latency, packet loss) are reported as percentages of the related SLA thresholds. Rather than reporting tunnel traffic, application/SLA paths show the ability of a tunnel to meet the SLAs for different types of traffic (for example, voice or video). These traffic types are referred to as SLA classes.

#### NOTE

For integrations with DX NetOps Virtual Network Assurance (VNA) and a configured Viptela plug-in, Viptela sites map to site groups in NetOps Portal. You can update site names by managing the site groups in NetOps Portal.

For more information, see [Manage Groups](#).

For more information about the metrics that Viptela collects, including the VNA metrics, SNMP metric families, and NFA metrics, see [the DX NetOps Virtual Network Assurance documentation](#).

The following video examines how to monitor Viptela SD-WAN devices using the dashboards and context pages in NetOps Portal:

#### NOTE

If you have integrated with DX NetOps Virtual Network Assurance and you have configured the Fortinet plug-in, you can view virtual site groups (Gateway sites) and monitor the health and location of these sites from the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards. The tunnel and application statistics geographic maps and cards on these dashboards display Gateway sites.

The following video examines SD-WAN application/SLA paths and why SD-WAN monitoring is important to delivering a reliable digital experience:

The following workflow illustrates how an Operations Engineer can use the SD-WAN dashboards and context pages for SD-WAN monitoring:

1. One of the following events drives you to the relevant SD-WAN dashboard:
  - For an event impacting site connectivity, go to the **SD-WAN Tunnel Statistics** dashboard.
  - For an event that could impact service level agreements (SLAs), go to the **SD-WAN App Path Statistics** dashboard.
2. Use the health counts to view the site connections or application/SLA paths with issues.
3. Use the geographic map and scorecard to view site details.
4. Review the remaining trend data.
5. Drill into edge devices.

### View the Aggregator Health Counts

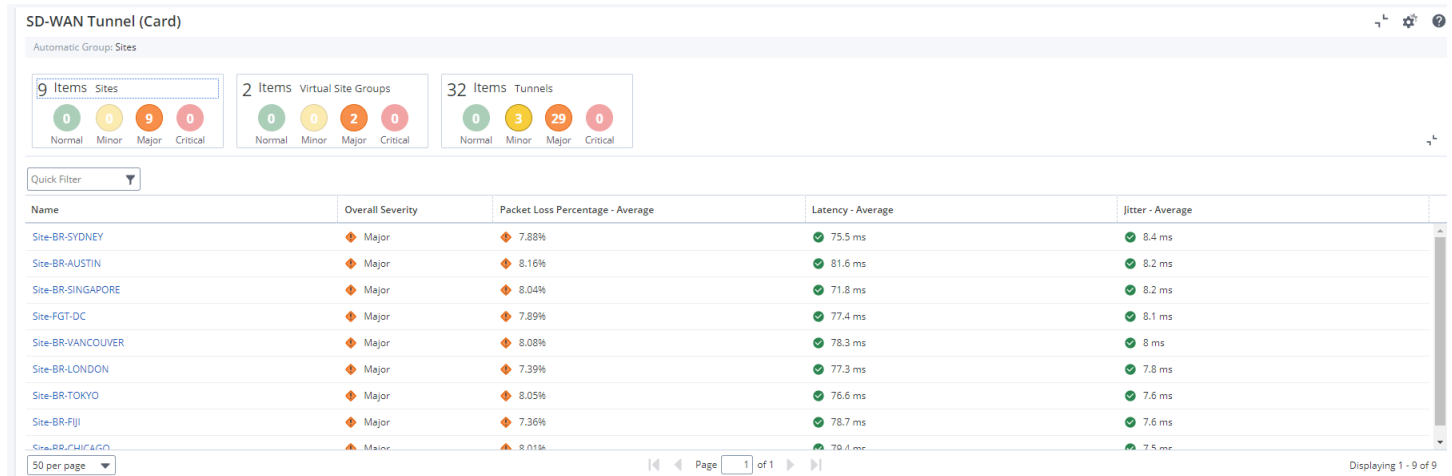
You can view the aggregated health counts of the following items from the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards:

- Sites
- Edge Routers
- Application Paths
- Applications
- Tunnels

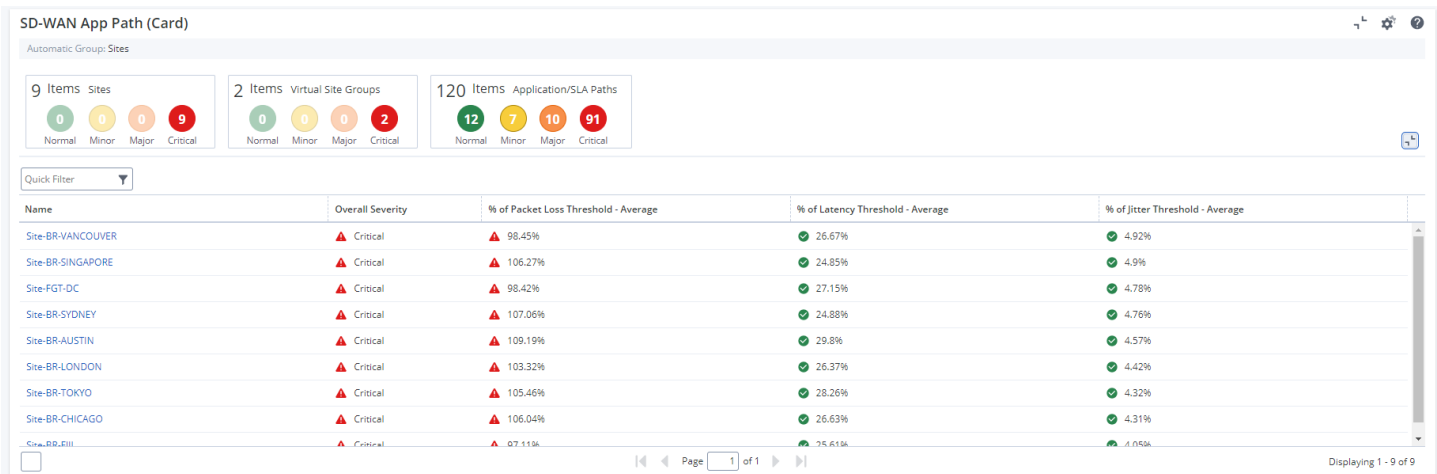
Counts appear in color-coded circles for each threshold range. To expand the view, click the arrow in the lower-right corner. Click a card title or a circle to view the corresponding details in the expanded view.

The following image shows an example of an SD-WAN Tunnel card showing SD-WAN tunnel statistics:

**Figure 101: SD-WAN Tunnel Card**



The following image shows an example of an SD-WAN App Path card showing SD-WAN application statistics:

**Figure 102: SD-WAN App Path Card****View the Location of a Site**

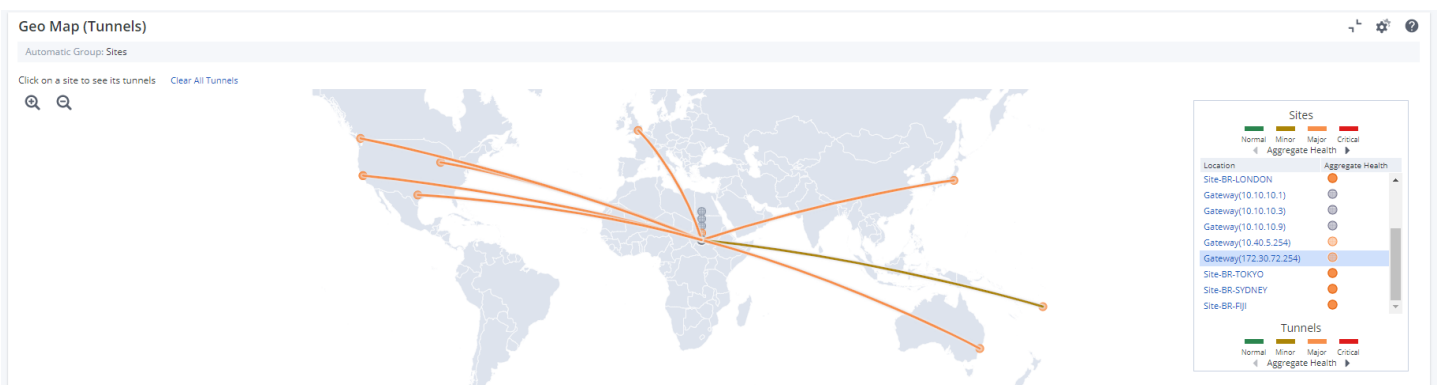
You can view the location of each site from the geographic map on the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards.

Select a site to display the connections to and from other sites. The connection lines are color-coded based on health metrics. Site router details appear when you hover over a site. Tunnel or application/SLA path details appear when you hover over a connection.

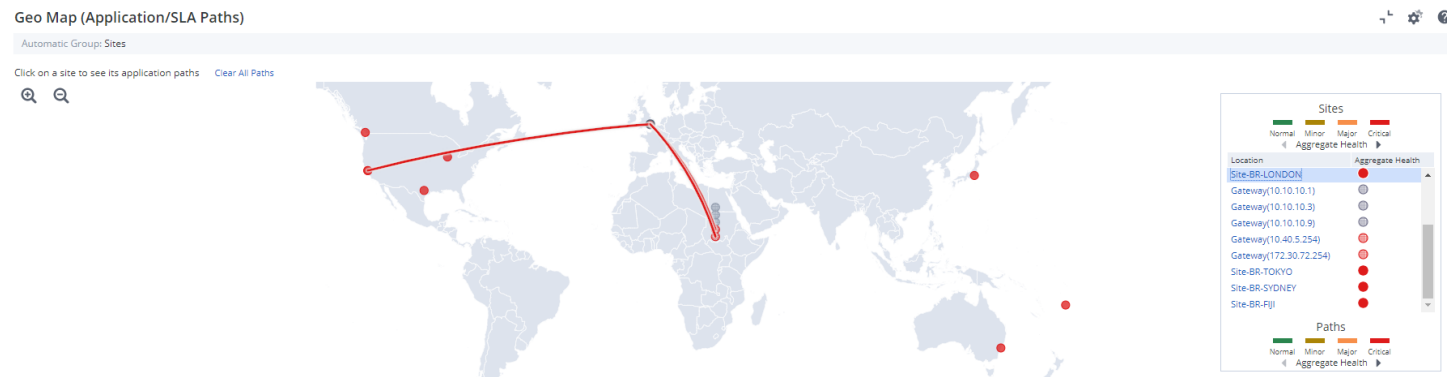
**NOTE**

Only sites within the page-level group context are shown even when a tunnel or application path runs between an included and excluded site.

The following image shows an example of a geographic map showing SD-WAN tunnel statistics:

**Figure 103: Geographic Map for SD-WAN Tunnels**

The following image shows an example of a geographic map showing SD-WAN application statistics:

**Figure 104: Geographic Map for SD-WAN Applications****View Metrics from Time Bar Charts**

You can view aggregate packet loss, latency, and jitter for each metric from time bar charts that display on the **SD-WAN App Path Statistics** dashboard.

**NOTE**

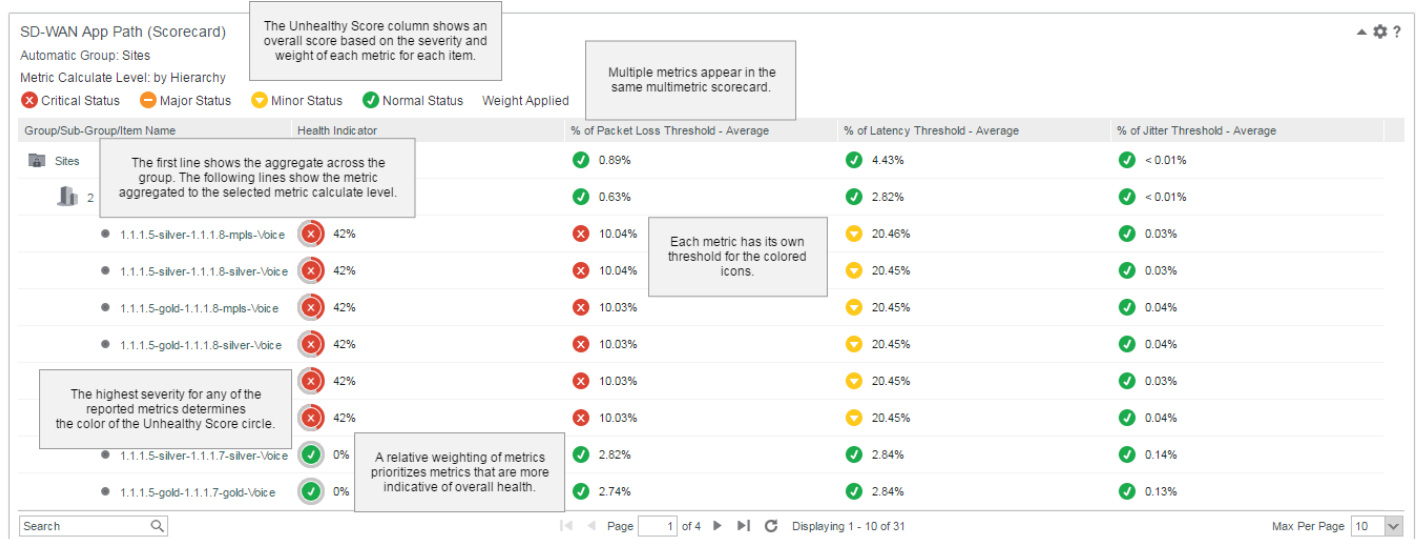
The time bar charts only shows sites within the page-level group context even when a tunnel or application path runs between an included and excluded site. The selected site in the map view filters this view.

The following image shows an example of a time bar chart showing SD-WAN application statistics:

**Figure 105: SD-WAN Timebar****View Metrics from the Scorecard**

You can view tunnel or application/SLA path metrics by subgroup and component for the selected group from the scorecards that display on the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards. Subgroups and the items in those groups appear in a hierarchical format. Colored icons indicate performance levels for metrics and an overall health indicator.

The following image shows an example of a scorecard showing SD-WAN application statistics:

**Figure 106: SD-WAN Scorecard****View Metrics from Tables with Color-Coded Bars**

You can view tunnel or application metrics in the tables that display on the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards. These tables display with color-coded bars for percentage metrics.

**View Metrics from Gauge/Table Views**

You can view view packet loss from the a gauge/table views that display on the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards.

**View Metrics from Trend/Table Views**

You can view packet loss, latency, and jitter metrics in the trend/table views that display on the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards.

**View Side-by-Side Comparison of Metrics from Trend Charts**

View latency and jitter metrics side-by-side for comparison in the trend charts that display on the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards.

**View Source and Destination Endpoint**

View the source and destination endpoints of your tunnels or applications from the charts and tables on the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards.

**View Errors and Discards in Separate Tables**

View SD-WAN virtual interface errors in one table and SD-WAN virtual interface discards in another table from the **SD-WAN Tunnel Statistics** and **SD-WAN App Path Statistics** dashboards.

**Navigate to More Details from the Dashboards**

From the inventory or the dashboards, you can navigate to context pages containing more details for the following SD-WAN items:



- **Devices**

View the typical device context page for any of your SD-WAN devices.

- **Application/SLA Paths**

On the **Details** tab, view the details for the selected application/SLA path. View its statistics, packet loss in a gauge, heat chart, and trend chart, and its traffic.

On the **Class Health** tab, view the percentage of packet loss, percentage of jitter, and the percentage of latency.

These metrics appear in gauges, heat charts, side-by-side bar charts, trend charts, and a bar chart table.

- **Tunnels**

On the **Details** tab, view the details for the selected tunnel. View its statistics, packet loss in a gauge, heat chart, and trend chart, and its traffic.

- **Virtual Interfaces**

On the **Details** tab, view the details for the selected virtual interface with its event list.

On the **Health** tab, view trend charts for discards, utilization, errors, packet rate, and bit rate.

## Monitor VMware NSX-T

If you have integrated with DX NetOps Virtual Network Assurance (VNA) and you have configured the VMware NSX-T plug-in, you can monitor the NSX-T item types, alarms, and performance metrics that VNA supports from NetOps Portal.

You can monitor NSX-T item types, alarms, and performance metrics from the following dashboards:

- [NSX-T Console dashboard](#)

This dashboard shows your virtual inventory.

- [NSX-T Health dashboard](#)

This dashboard provides health information for key NSX-T item types, inventory, alarms, and performance metrics, such as transports nodes.

Also in this topic:

- [Configure to Show NSX-T BGP Sessions and their Statuses For a Device](#)
- [View NSX-T BGP Sessions for Gateway Routers](#)
- [Adjust the Number of Devices Shown in the Alarm State Indicators for a Card](#)

For more information about the VMware NSX-T plug-in, see [the DX NetOps Virtual Network Assurance documentation](#).

### NSX-T Console Dashboard

The **NSX-T Console** dashboard provides the following information:

- An overview of the NSX-T environments that NetOps Portal monitors, including topology, items' alarm state, and drill-down links.
- A list of the monitored VNA domains, for example, NSX-T environments.

**NOTE**

The **NSX-T Console** dashboard shows only VNA domains related to NSX-T plug-in instances.

**Prerequisite:** To view VNA domains on the dashboard, you have access to the groups where the NSX-T VNA domains exist.

- An overview of the NSX-T entities and relationship of the underlying virtual network fabric.
- An overview of NSX-T items that belong to the same transport zone.
- A mapping of tier-0-routers-to-border-gateway routers to which they are connected, and a view of the BGP sessions that connect those items.
- A list of the host transport node items.

To view this dashboard, go to **Performance**, **NSX-T Reporting**, **NSX-T Console**.

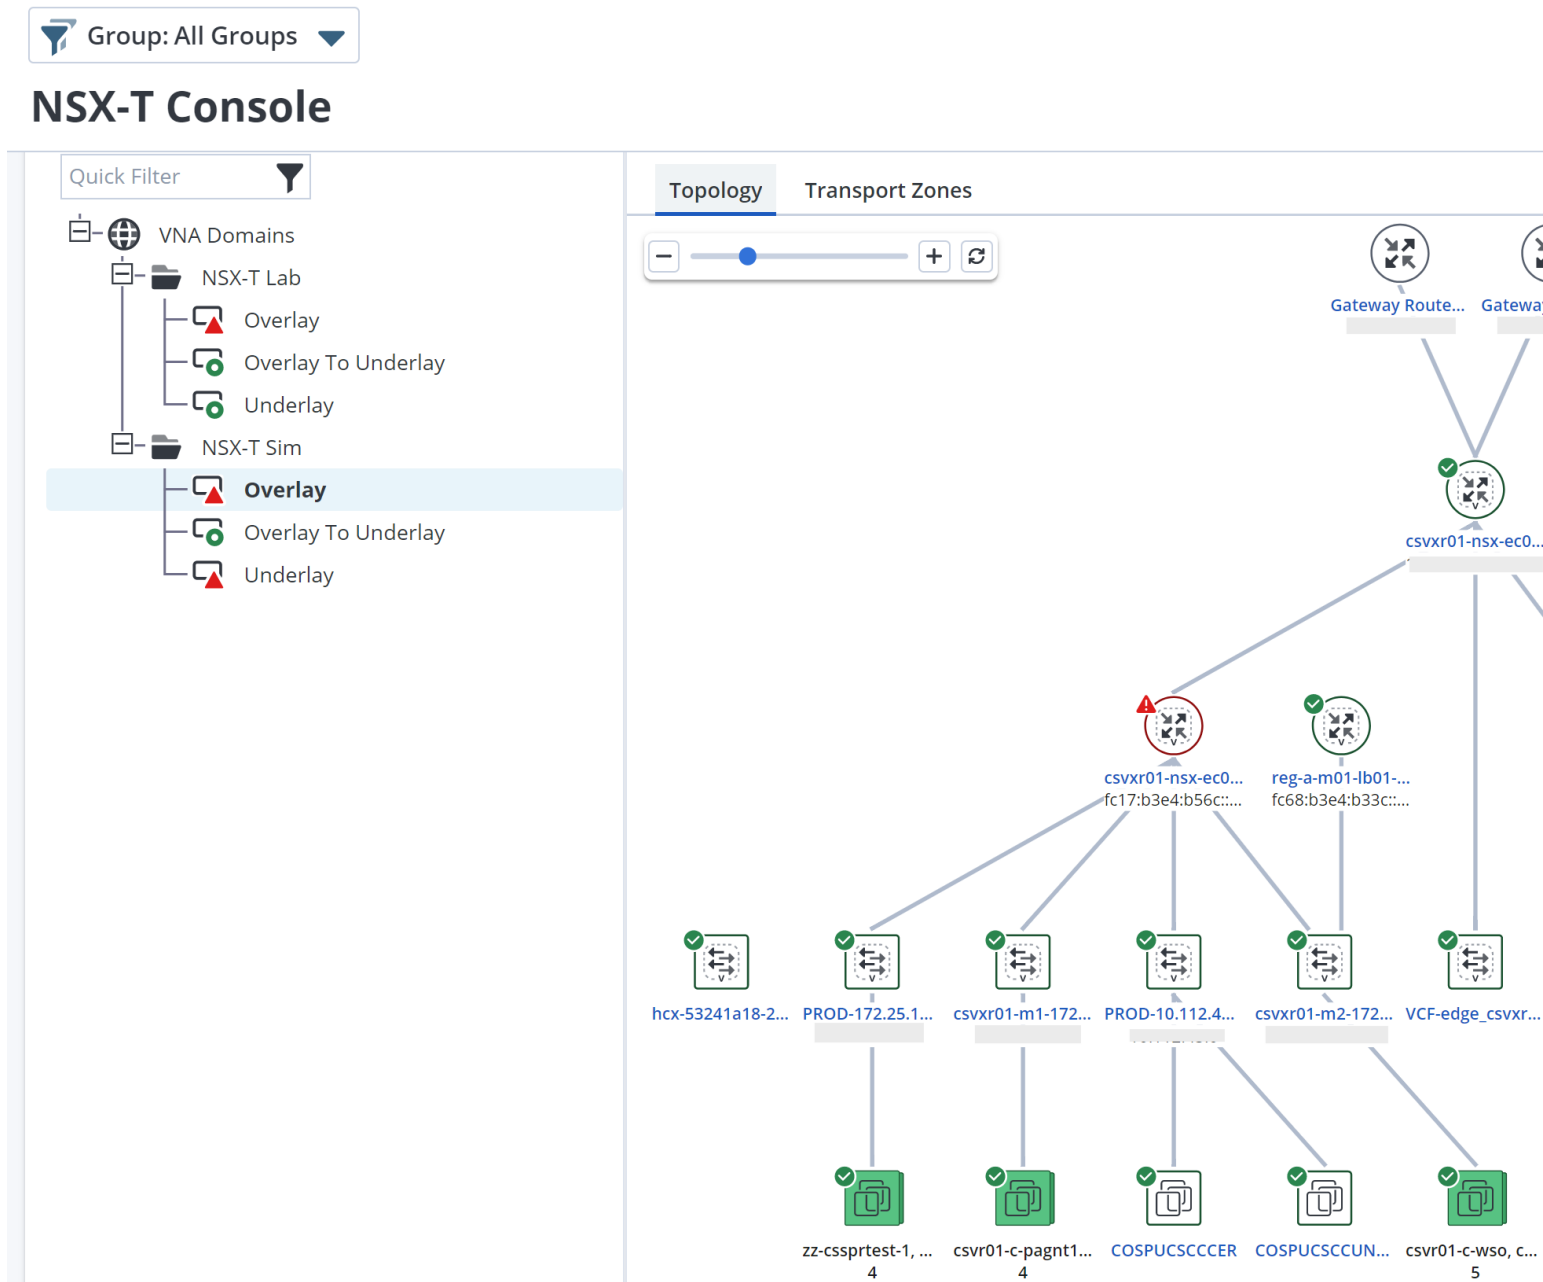
**NOTE**

The **NSX-T Reporting** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate menus with roles, see [Manage User Account Roles](#).

The following image shows the **NSX-T Console** dashboard:

**Figure 107: NSX-T Console Dashboard**



## NSX-T Health Dashboard

Use the **NSX-T Health** dashboard to quickly identify issues and to drill down to the related items for the further troubleshooting.

### NOTE

This dashboard requires that you specify a unique domain name for each configured NSX-T plug-in for it to function properly. Ensure that no other plug-in is using the same domain name. You give plug-ins domain names when you create the domain during the plug-in configuration process.

For more information about the plug-in configuration process and how to manage domains, see [the DX NetOps Virtual Network Assurance documentation](#).

To view this dashboard, go to **Performance, NSX-T Reporting, NSX-T Health**.

### NOTE

The **NSX-T Reporting** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate menus with roles, see [Manage User Account Roles](#).

The **NSX-T Health** dashboard consists of the following views:

- **NSX-T Devices Alarm State By Device Type - Card**

### NOTE

This view displays only if you have integrated with DX NetOps Spectrum (Spectrum).

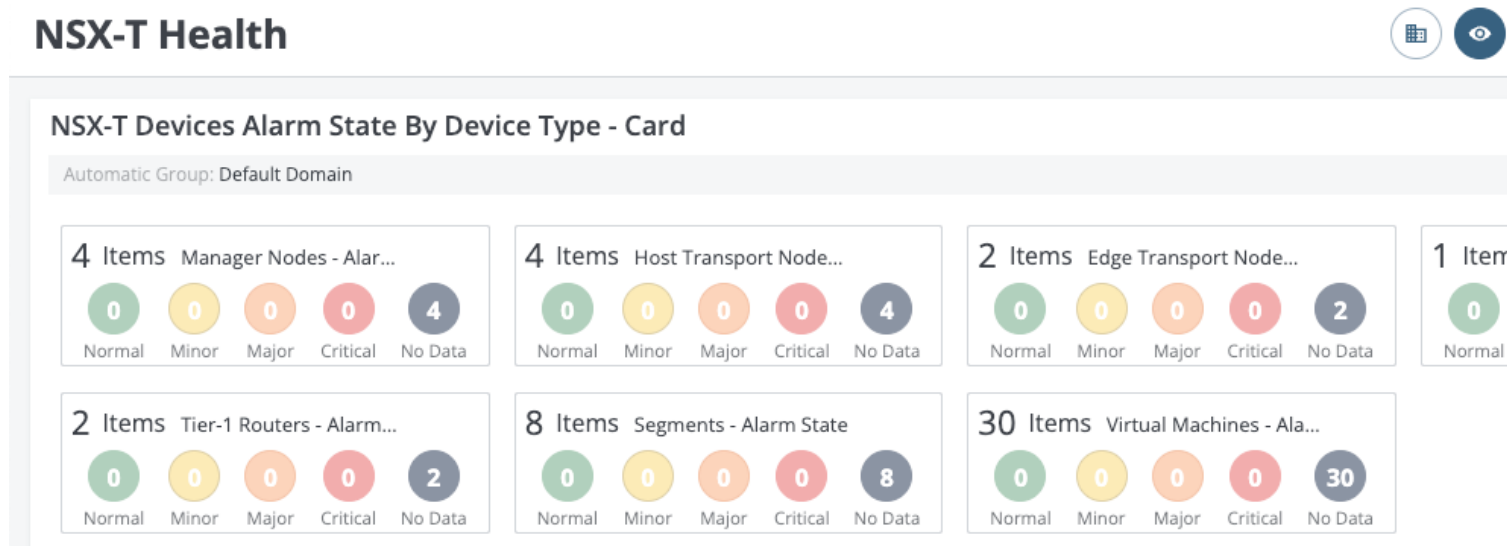
This view displays a card for each NSX-T device type (NSX-T Manager nodes, Host transport nodes, Edge transport nodes, Tier-0 routers, Tier-1 routers, logical switches (Segments), and virtual machines). The total number of NSX-T devices of that particular device type (the total item count) is shown at the top of each card, along with the number (counter) of NSX-T devices per alarm severity level/state (alarm state indicators) per device type.

### TIP

By default, up to 5000 devices are shown in the alarm state indicator for each card. Improve performance by [limiting the number of NSX-T devices shown in the alarm state indicators for each device type card](#).

The following image shows this view on the **NSX-T Health** dashboard:

**Figure 108: The NSX-T Devices Alarm State By Device Type - Card views on the NSX-T Health Dashboard**



The cards include alarm state indicators for only the following alarm severity levels/states:

- **Normal**

Indicates the following:

- NetOps Portal has made contact with the device, and the device is operating normally.
- A Normal alarm has been generated. If an event generated an alarm but a severity for the alarm was not specified (for example, if you have created the supporting EventDisp configuration file manually and inadvertently omitted a severity), Spectrum assigns the "Normal" severity status to the alarm.

– **Minor**

Indicates that the device is experiencing an abnormal situation, but that no immediate action is required. This level of severity is also used for alarms created only to convey information, such as "Duplicate IP."

– **Major**

Indicates that the device is experiencing a loss of service or a loss of service is impending, and that you should take action.

– **Critical**

Indicates that the device is experiencing a loss of service, and that immediate action is required.

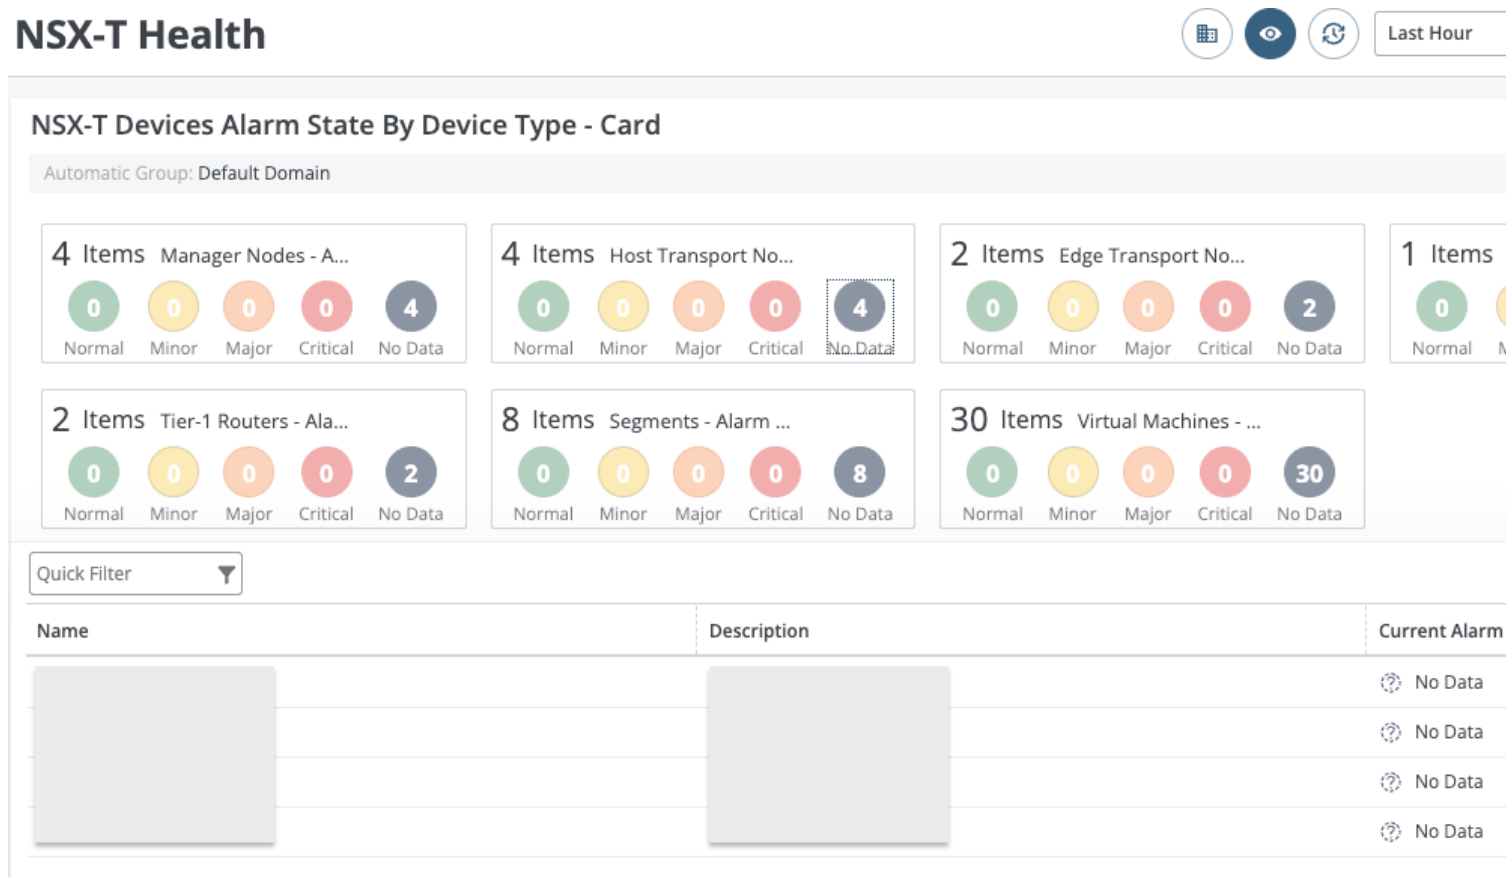
– **No Data**

Indicates that the device does not have an alarm state.

Clicking an *alarm state indicator* in a card displays a list of NSX-T devices of that type with that particular alarm state. Clicking a *total item count* for a card of a particular device type displays a list of NSX-T devices of that type with their alarm state.

The following image shows an example of the table that displays when you click an alarm state indicator or a total item count:

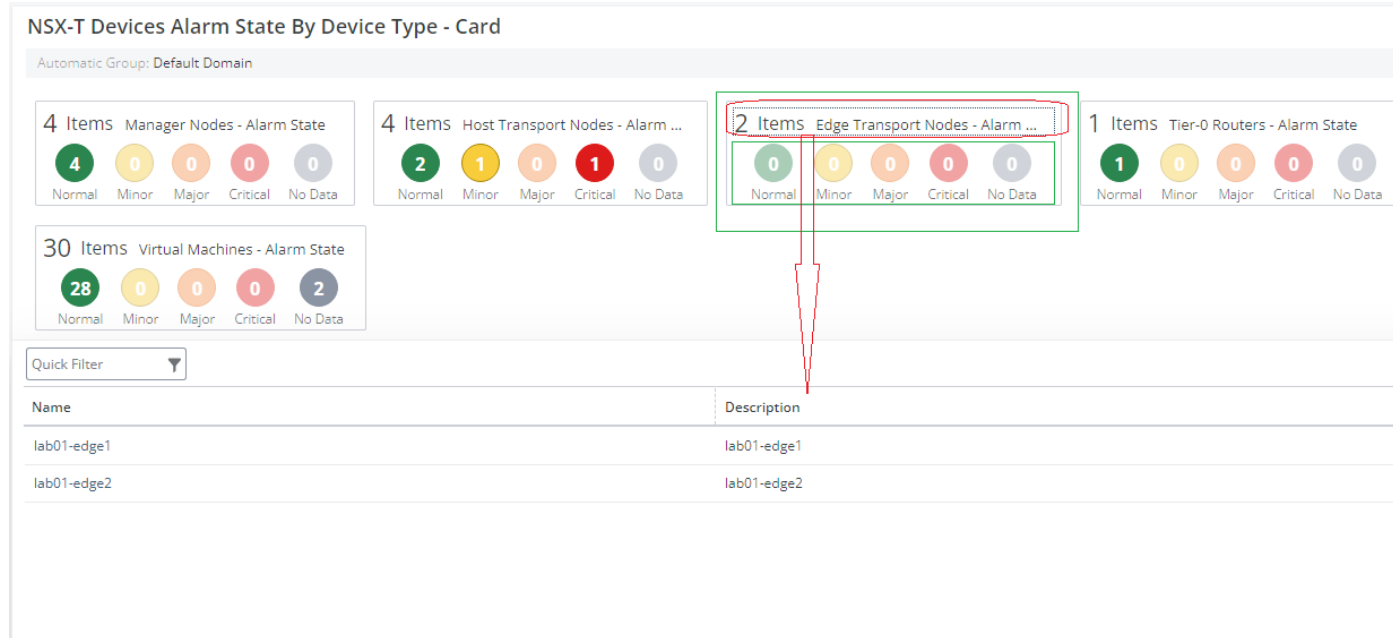
**Figure 109: The table with the NSX-T devices with a specific alarm state on NSX-T Health Dashboard**



**NOTE**

If the sum of the number (counter) of NSX-T devices for all alarm severity levels/states for a particular device type differs from the total item count for that device type, then there are devices with an alarm severity level/state value other than those included in the cards. For example, these devices might have a "Suppressed", "Maintenance", "Initial", or "Unknown" alarm severity level/state. To display a list of the alarm severity level/state for the NSX-T devices of a particular device type, click the *total item count* for the card. The following image shows an example of a list of devices with the "Suppressed" and "Maintenance" alarm severity levels/states, which are states not included as alarm state indicators on the cards:

**Figure 110: The table with the NSX-T devices with an alarm state not shown in the cards on NSX-T Health Dashboard**



- NSX-T Alarms**

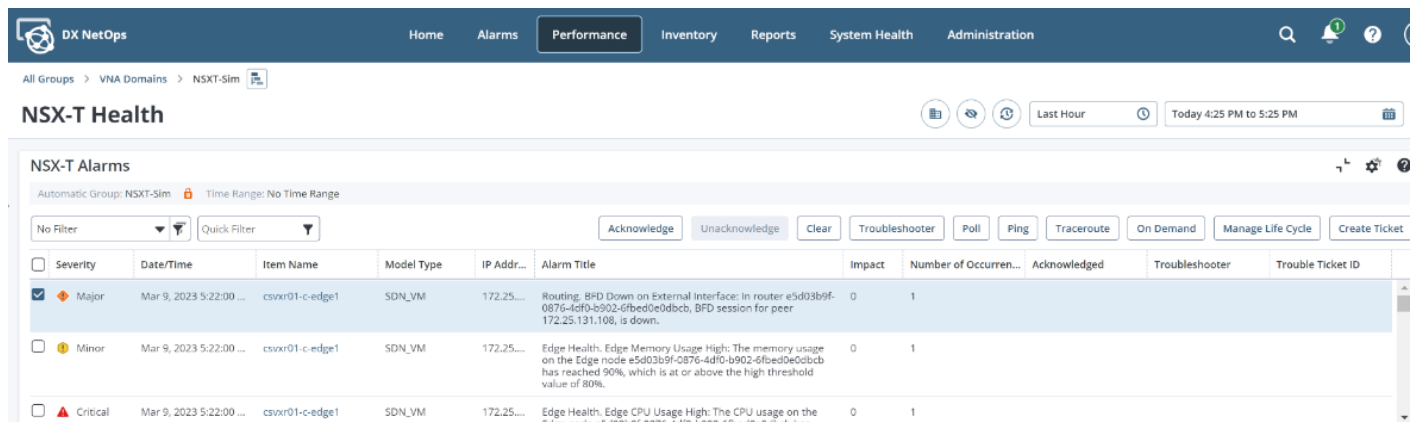
**NOTE**

This view displays only if you have integrated with Spectrum.

This view shows alarms for the items in the selected NSX-T domain group.

The following image shows this view on the **NSX-T Health** dashboard:

**Figure 111: The NSX-T Alarms view on NSX-T Health Dashboard**



Clicking an alarm (the link) displays the details for that alarm.

- **NSX-T Inventory**

This view shows devices that belong to the selected NSX-T domain group.

The following image shows this view on the **NSX-T Health** dashboard:

**Figure 112: The NSX-T Inventory view on the NSX-T Health Dashboard**

NSX-T Health

NSX-T Inventory

Devices

Quick Filter

On Demand

<input type="checkbox"/>	Name	Type	Address	Description	Current Alarm S...	Life Cycle State	Context Types
<input type="checkbox"/>	COSPUSCCER	Virtual Machine		COSPUSCCER		Active	Device
<input type="checkbox"/>	COSPUSCCPU8	Virtual Machine		COSPUSCCPU8		Active	Device
<input type="checkbox"/>	COSPUSCCSUB	Virtual Machine		COSPUSCCSUB		Active	Device
<input type="checkbox"/>	COSPUSCCCTFTP	Virtual Machine		COSPUSCCCTFTP		Active	Device
<input type="checkbox"/>	COSPUSCCUNITY	Virtual Machine		COSPUSCCUNITY		Active	Device
<input type="checkbox"/>	csvr01-c-pagnt1	Virtual Machine		csvr01-c-pagnt1		Active	Device
<input type="checkbox"/>	csvr01-c-vrli1	Virtual Machine		csvr01-c-vrli1		Active	Device
<input type="checkbox"/>	csvr01-c-vrli2	Virtual Machine		csvr01-c-vrli2		Active	Device

Clicking a device name (the link) opens the context page for that device.

**TIP**

You can configure to show NSX-T BGP sessions and their statuses in the context view on the **NSX-T Health** dashboard for a device that you select in the **NSX-T Inventory** view.

For more information, see the ["Configure to Show NSX-T BGP Sessions and their Statuses For a Device" section](#).

- **Virtual Interface**

This view shows interfaces for the items in the selected NSX-T domain group.

The following image shows this view on the **NSX-T Health** dashboard:

**Figure 113: The Virtual Interface view on the NSX-T Health Dashboard**

Virtual Interface

Quick Filter

On Demand

<input type="checkbox"/>	Name	Description	Device	Life Cycle State	IP Address
<input type="checkbox"/>	docker0	docker0	lab01-edge1	Active	
<input type="checkbox"/>	docker0	docker0	lab01-edge2	Active	
<input type="checkbox"/>	eth0	eth0	lab01-edge1	Active	
<input type="checkbox"/>	eth0	eth0	lab01-edge2	Active	
<input type="checkbox"/>	eth1/1		leaf-2	Active	
<input type="checkbox"/>	eth1/1		leaf-1	Active	
<input type="checkbox"/>	eth1/1			Active	
<input type="checkbox"/>	eth1/10		leaf-2	Active	

Page 1 of 7

Displaying 1 - 100 of 661

Clicking an interface name (the link) opens the context page for that virtual interface.

- **Trend charts**

The **NSX-T Health** dashboard includes the following views:

- **Top Host Transport Nodes by Device Utilization - Trend/Table**

This view shows a trend chart of the top host transport nodes by device utilization item types and metrics.

- **Top Host Transport Nodes by Memory Utilization - Trend/Table**

This view shows a trend chart of the top host transport nodes by memory utilization item types and metrics.

- **Top Edge Transport Nodes by CPU Utilization - Trend/Table**

This view shows a trend chart of the top edge transport nodes by CPU utilization item types and metrics.

- **Top Edge Transport Nodes by Memory Utilization - Trend/Table**

This view shows a trend chart of the top edge transport nodes by memory utilization item types and metrics.

- **Top Virtual Interfaces by Incoming Discards - Trend/Table**

This view shows a trend chart of the top virtual interfaces by incoming bits item types and metrics.

- **Top Virtual Interfaces by Outgoing Discards - Trend/Table**

This view shows a trend chart of the top virtual interfaces by outgoing bits item types and metrics.

- **Top Virtual Interfaces by Incoming Errors - Trend/Table**

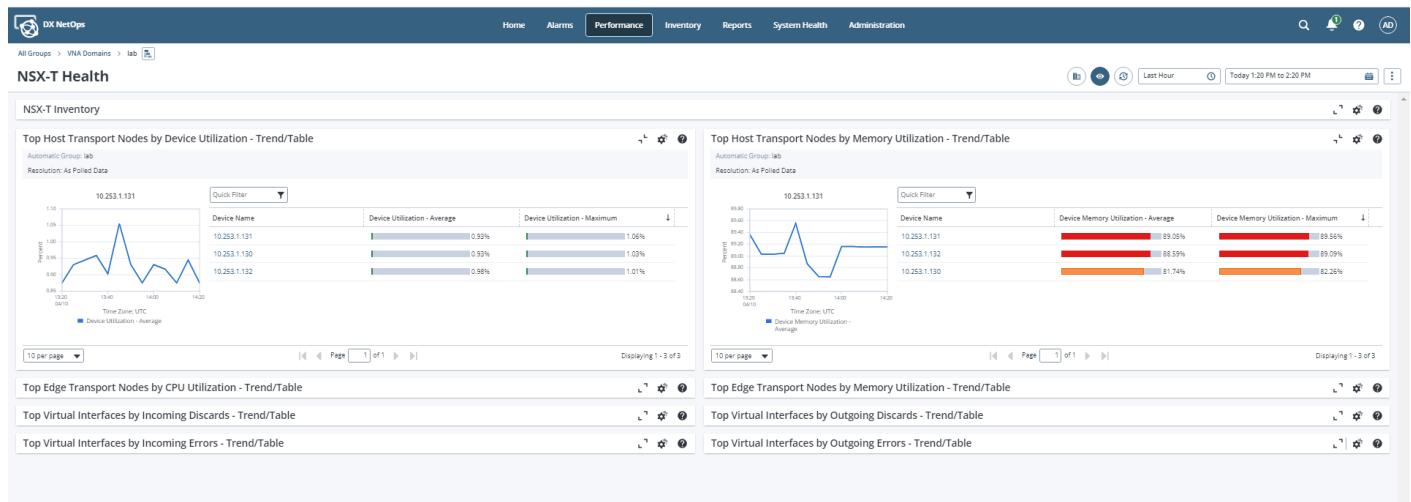
This view shows a trend chart of the top virtual interfaces by incoming errors item types and metrics.

- **Top Virtual Interfaces by Outgoing Errors - Trend/Table**

This view shows a trend chart of the top virtual interfaces by outgoing errors item types and metrics.

The following image shows these views on the **NSX-T Health** dashboard:

**Figure 114: Views on the NSX-T Health dashboard**



## Configure to Show NSX-T BGP Sessions and their Statuses For a Device

You can show NSX-T BGP sessions and their statuses in the context view on the **NSX-T Health** dashboard for a device selected in the **NSX-T Inventory** view. This context view can show various item types. To have this sub-view show NSX-T BGP sessions and their statuses, configure the data context for the view.

The following image shows an example of the context view configured to show NSX-T BGP sessions and their statuses. The example shows a selected device in the **NSX-T Inventory** view, and the NSX-T BGP sessions and their statuses for this device in the **BGP Tunnels** sub-view:

**Figure 115: BGP Tunnels sub-view on the NSX-T Health dashboard**

**NSX-T Health**

Name	Description	Protocol	Source Device	Destination Device	Source Interface	Destination Interface	Status
gw-Tier-0-gw-01-192.100.1.1	45	BGP	192...	192...	gw-Tier-0-gw-01	Gateway Router:192.1...	Up

**Follow these steps:**

- On the **NSX-T Health** dashboard, edit the **NSX-T Inventory** view by clicking the **View Settings** (gear) icon for the view, and then click **Edit**.  
The view settings dialog for the **NSX-T Inventory** view opens.
- From the **Context Type** field, select the following option, and then save your changes:
  - **BGP Tunnels**  
Specifies the context to be NSX-T BGP sessions and their statuses.  
**Default:** Virtual Interface  
For more information about how to change the data context for views, see [Customize Views](#).

**View NSX-T BGP Sessions for Gateway Routers**

You can view a list of NSX-T BGP sessions and their statuses for Gateway routers from the (Gateway) **Router** context page. To view the list, open a context page for an NSX-T Gateway router that an NSX-T plug-in monitors.

The following image shows an example of the (Gateway) **Router** context page:



**Figure 116: (Gateway) Router context page**

Router: Gateway Router: [ID]

**Router Gateway Router:** [ID]

Details

Interface Health

Custom View - Infrastructu...

Log Events

SDN Device Metrics

**Gateway Router:** [ID] Reachability Status Unknown ACTIVE

**IDENTITY**  
Context Types: Router  
Device Name: Gateway Router  
Device Name Alias: Gateway Router  
Device IP: [IP]

**DEVICE**  
Description: Gateway Router: :connected to TIER0  
gw-Tier-0-gw-01

**SNMP**  
Profile: N/A

**GROUP MEMBERSHIP**  
/All Groups/Inventory/IP Domains/Default Domain  
/All Groups/Inventory/All Items/Routers

**TIME SETTINGS**  
The device is not a member of a site group that includes time settings.

**CUSTOM ATTRIBUTES**  
The device does not have any custom attributes.

**BGP Tunnels Inventory**

Quick Filter [v]

Name	Description	Protocol	So...	De...	Source Device	Destination Device	Source Interface	Destination Inter...	Status
gw-Tier-0-gw-01	:5000-	BGP	19...	19...	gw-Tier-0-gw-01	Gateway Router:1...	gw-Tier-0-gw-01	Gateway Router:1...	Up

10 per page

Page 1 of 1

Displaying 1 - 1 of 1

**Adjust the Number of Devices Shown in the Alarm State Indicators for a Card**

In the view settings, to improve the performance of the device type cards in the **NSX-T Devices Alarm State By Device Type - Card** view on the **NSX-T Health** dashboard, you can adjust the number of NSX-T devices that are shown in the alarm state indicators for each device type card. By default, up to 5000 devices are shown in the alarm state indicators.

If the number of NSX-T devices for a device type is greater than the set limit, then at the top of the card view, a warning message displays, and a list of cards with truncated results are displayed.

**Example:**

In the following example, the limit the number of NSX-T devices shown in the alarm state indicators for each device type card (the **Items Limit** field) is set to 8. Because more than eight NSX-T devices of device type Segments and Virtual Machines have an alarm state of "Normal", the "Normal" indicator shows truncated results (7) and at the top of the view, and the following warning message appears:

Showing first 7 results for Segments, Virtual Machines card

The following image shows an example of the message:

**Figure 117: The NSX-T Devices Alarm State By Device Type - Card view with a warning**

NSX-T Devices Alarm State By Device Type - Card

Showing first 7 results for Segments, Virtual Machines card.

Automatic Group: Default Domain

4 Items Manager Nodes - Alarm State

4 Items Host Transport Nodes - Alarm ...

2 Items Edge Transport Nodes - Alarm ...

1 Items Tier-0 Routers - Alarm State

2 Items Tier-1 Routers - Alarm State

7 Items Segments - Alarm State

7 Items Virtual Machines - Alarm State

**Follow these steps:**

1. On the **NSX-T Health** dashboard, edit the **NSX-T Devices Alarm State By Device Type - Card** view by clicking the **View Settings** (gear) icon for the view, and then click **Edit**.

The view settings dialog for the **NSX-T Devices Alarm State By Device Type - Card** view opens.

2. Edit the following field, and then save your changes:

– **Items Limit**

Specifies the limit for the number of NSX-T devices (counter) shown in the alarm state indicator for each card. If the number of NSX-T devices for a device type is equal to or greater than the set limit, then at the top of the view, a warning message displays with a list of cards with truncated results.

**Default:** 5000

The number of devices shown in the alarm state indicators for a card are adjusted.

## Monitor Wi-Fi Device Inventory

If you have integrated with DX NetOps Virtual Network Assurance (VNA) and you have configured a plug-in that monitors Wi-Fi devices, such as the Cisco Meraki and Aruba Central plug-ins, you can view the Wi-Fi devices that VNA supports and that NetOps Portal monitors, such as wireless access points (APs) and wireless LAN controllers (WLCs), from NetOps Portal.

Monitor these Wi-Fi devices from the **Wi-Fi Devices** page and the **Wi-Fi Health** dashboard.

In this article:

- [View a List of Wi-Fi Devices](#)
- [View Client Count and Bandwidth for a Wi-Fi Device](#)
- [Monitor the Client Devices for a Specific Wireless Access Point](#)
- [Monitor Performance Baseline of Wi-Fi Devices](#)
- [View AP Radio Details for Wi-Fi Devices](#)
- [View a List of Inventory and Performance Metrics](#)
- [Filter Wi-Fi Alarms](#)

### View a List of Wi-Fi Devices

You can get an overview of the Wi-Fi devices that NetOps Portal monitors from the **Wi-Fi Devices** page. To view this page, hover over **Inventory**, **Items**, and then click **Wi-Fi Devices**.

**NOTE**

The **Wi-Fi Devices** menu is hidden (not available) by default. Your user account role determines the menus that you can access. For this menu to be available to user accounts with a particular role, you must make it available for that role.

For more information about how to associate a menu with a role, see [Manage User Account Roles](#).

The following image shows an example of a list of Wi-Fi devices on the **Wi-Fi Devices** page:

**Figure 118: Wi-Fi Devices page****Wi-Fi Devices**

Wi-Fi Devices

⌵ ⚙️ ?

On Demand Manage Life Cycle

<input type="checkbox"/>	Name	↑ Domain	Address	Description	Current Al...	Life Cycle St...	VNA SubType	Context Types	DevCustom1
<input type="checkbox"/>		Default Domain			🕒 Initial	✅ Active		Device, Wireless...	
<input type="checkbox"/>		Default Domain		Description - Cisco 3504 Wireless Controller - 2	✅ Normal	✅ Active	distribution	Device, Wireless...	
<input type="checkbox"/>		Default Domain		Description - Cisco 5520 Wireless Controller - 964		✅ Active	access	Device, Wireless...	
<input type="checkbox"/>		Default Domain			✅ Normal	✅ Active		Device, Wireless...	
<input type="checkbox"/>		Default Domain				✅ Active		Device, Wireless...	
<input type="checkbox"/>		Default Domain		Description - Cisco 5520 Wireless Controller - 608	✅ Normal	✅ Active	border_router	Device, Wireless...	
<input type="checkbox"/>		Default Domain		Description - Cisco 5520 Wireless Controller - 1838	✅ Normal	✅ Active	access	Device, Wireless...	
<input type="checkbox"/>		Default Domain		Description - Cisco 5508 Wireless Controller - 803		✅ Active	access	Device, Wireless...	
<input type="checkbox"/>		Default Domain		Description - Cisco 5508 Wireless Controller - 66	✅ Normal	✅ Active	border_router	Device, Wireless...	
<input type="checkbox"/>		Default Domain		Description - Cisco 8540 Wireless Controller - 1186	✅ Normal	✅ Active	border_router	Device, Wireless...	

**NOTE**

In NetOps Portal views, for Cisco Meraki APs, ensure that you have set a **Resolution** that is based on the poll rates (intervals) that are configured for the Meraki plug-in.

For more information:

- About the poll intervals, see [Cisco Meraki](#).
- About the resolution for table views, see [Set the Resolution for Reported Data](#).

To view the details for an AP or WLC, click the name of the AP or WLC. The **Summary** tab is selected by default.

You can view details for the AP or WLC from the following context tabs:

- **Summary**

Displays details for the AP or WLC, including a description of the AP or WLC and group membership. This context tab includes the following views:

- **Calendar Heat Chart - Availability**

- **Event List**

- **Client Count**

[View the count of network devices \(clients\) connected to an AP or WLC over a selected time range/interval time from this view.](#)

- **Bandwidth**

[View the cumulative value of bytes of data that have flown into an AP \(Rx\) and the bytes of data that have flown out of an AP \(Tx\) from this view.](#)

- **Active Access Points**

- **Clients Associated**

- **CPU Utilization**

- **Memory Utilization**

- **Latency Trend**

- **Reachability Trend**

- **Radio Details**

[Troubleshoot AP Radio performance from this table view.](#)

- **Interfaces** (23.3.1 or lower)

- **Wireless Access Context** (23.3.1 or lower)

This context tab includes the following views:


- **Wireless Access Table**
- **Wireless Access Gauge Chart**
- **Wireless Access Pie Chart**
- **Wireless Access Horizontal Bar Chart**
- **Wireless Access Trend Chart**
- **Clients**  
This context tab includes the **Clients** view.
- **Wireless Controller Context** (23.3.1 or lower)  
This context tab includes the following views:
  - **Wireless Controller Table**
  - **Wireless Controller Gauge Chart**
  - **Wireless Controller Pie Chart**
  - **Wireless Controller Horizontal Bar Chart**
  - **Wireless Controller Trend Chart**
- **Log Events**  
This context tab includes the **Log Events** view.
- **Custom View - Infrastructure Management**
- **SDN Device Metrics**  
This context tab includes the following views:
  - **Health Score-Issue Count (Horizontal Bar Chart)**
  - **AP Interface-Clients Count (Trend Chart)**
  - **CPU Utilization (Gauge Chart)**
  - **Memory Utilization (Gauge Chart)**
  - **CPU/Memory Utilization (Trend Chart)**
  - **Health Score-Issue Count (Trend Chart)**
- **Radio Metrics - Trend Chart**  
[Monitor performance baselining of Wi-Fi devices from this view.](#)
- **Alarms**  
Displays the alarms for the AP or WLC in the **Alarms** view.

For more information about AppNeta, see [the DX NetOps Virtual Network Assurance documentation](#).


### **View Client Count and Bandwidth for a Wi-Fi Device**

You can view the count of network devices (clients) connected to an AP or WLC over a selected time range/interval time and the cumulative value of bytes of data that have flown into an AP (Rx) and the bytes of data that have flown out of an AP (Tx) from the **Client Count** and **Bandwidth** views on the **Summary** context page for an AP or WLC. To view these views, from the **Wi-Fi Devices** page, select the AP for which you want to view client count and bandwidth. By default, the **Summary** context tab is selected and the **Summary** page displays.

The following image shows an example of the **Client Count** view on the **Summary** page:

**Figure 119: The Client Count view**

The following image shows an example of the **Bandwidth** view on the **Summary** page:

**Figure 120: The Bandwidth view**

### Monitor Performance Baselining of Wi-Fi Devices

You can monitor performance baselining of Wi-Fi devices from the **Radio Metrics - Trend Chart** context page for an AP or WLC. APs have radio interfaces, and are modeled with the Wi-Fi device. This page shows a list of metrics of radio interfaces for the AP or WLC.

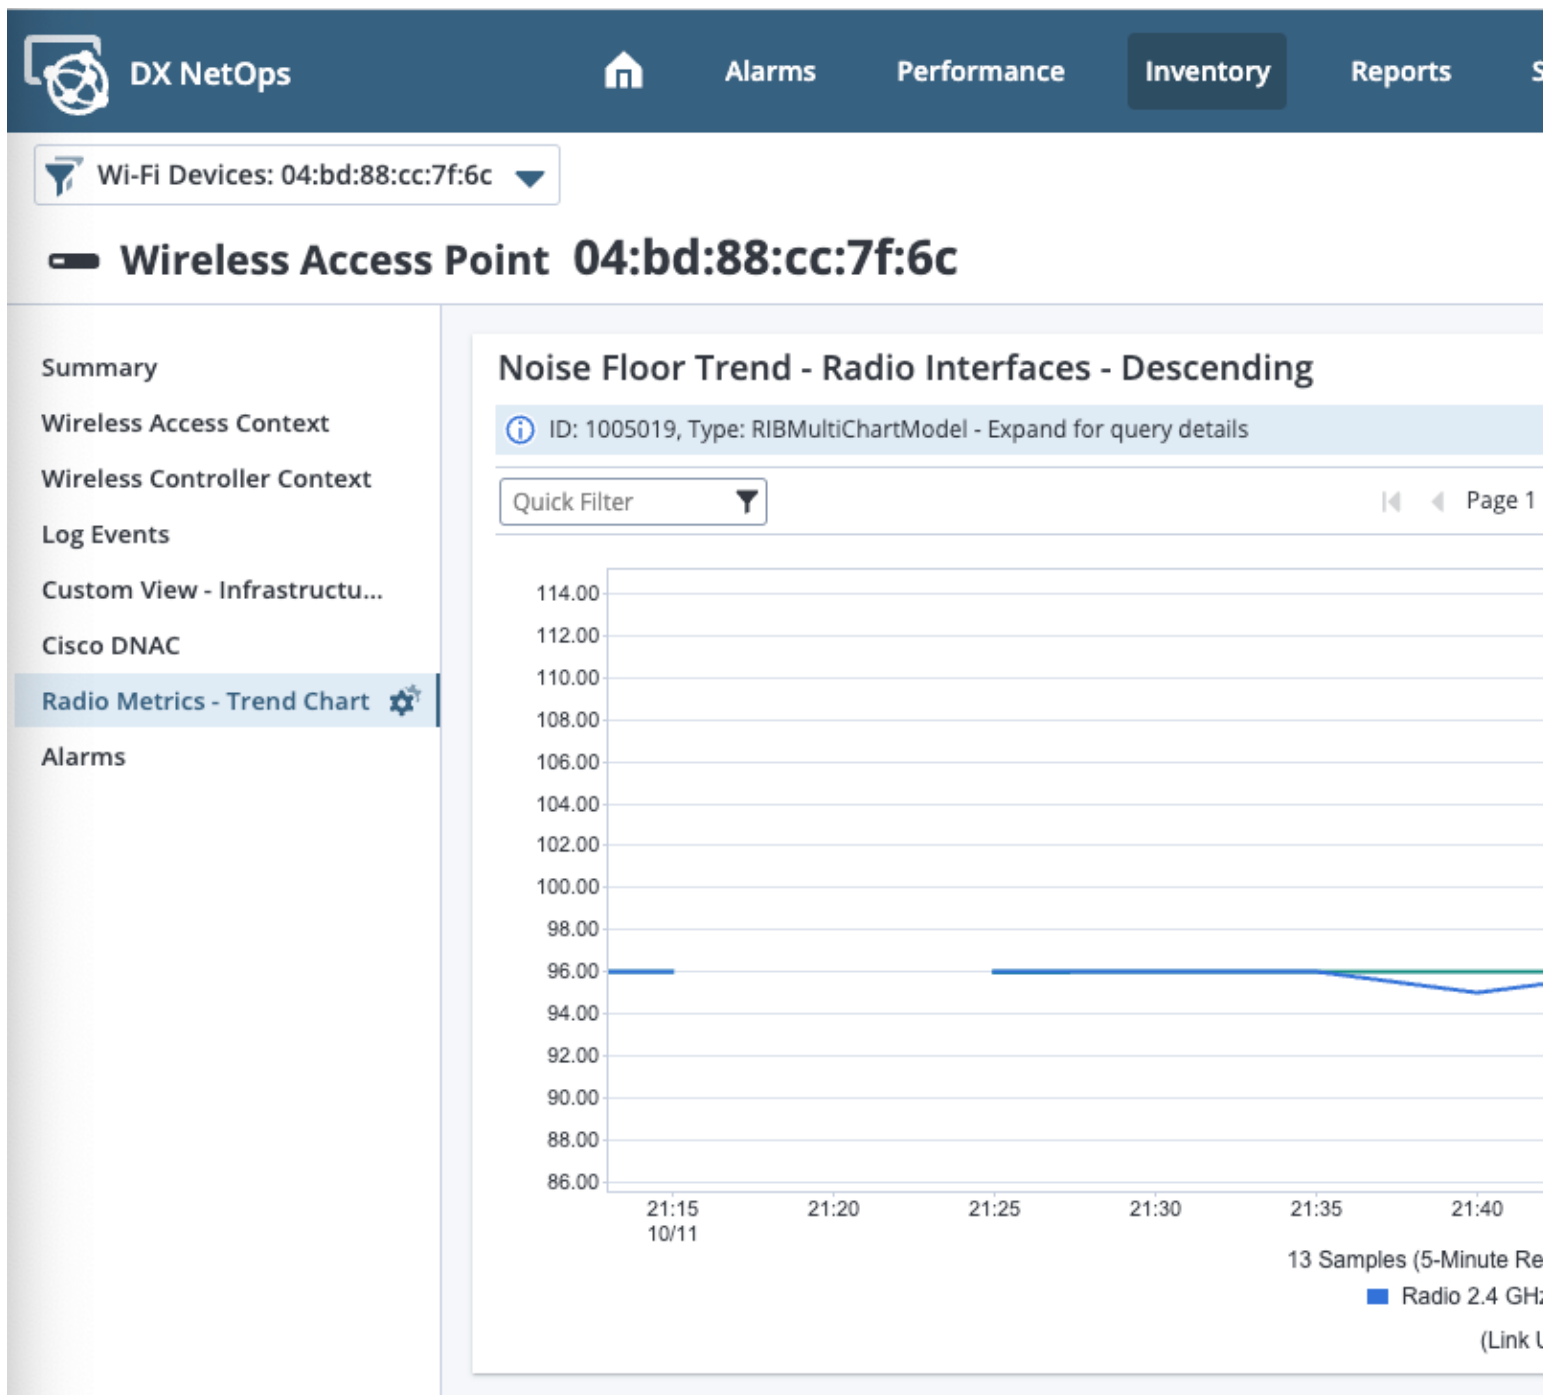
To view this page, from the **Wi-Fi Devices** page, select an AP or WLC, and then click the **Radio Metrics - Trend Chart** context tab.

The **Radio Metrics - Trend Chart** page includes the following views, which show MultiTrend view charts:

- **Noise Floor Trend - Radio Interfaces - Descending**
- **Utilization Percentage Trend - Radio Interfaces - Descending**
- **TxBytes Trend - Radio Interfaces - Descending**
- **TxPower Trend - Radio Interfaces - Descending**
- **TxDropped Packets Trend - Radio Interfaces - Descending**

The following image shows an example of the **Noise Floor Trend - Radio Interferences - Descending** view on the **Radio Metrics - Trend Chart** page:

Figure 121: Radio Metrics - Trend Chart page

**View AP Radio Details for Wi-Fi Devices**

You can view the following radio details:

- [AP Radio Details \(with Frequency Range \(FR\)\)](#)
- [AP Radio Interference](#)

**AP Radio Details (with FR)**

You can troubleshoot AP Radio performance from the **Radio Details** table view on the **Summary** context tab for an AP.

The following image shows an example of the table in the **Radio Details** table view:

**Figure 122: Radio Details table view**

Radio Band	Status	MAC Address	Mode	Channel
Radio 5 GHz	UP	f4:2e:7f:66:2e:f0	0	124+
Radio 2.4 GHz	UP	f4:2e:7f:66:2e:e0	0	11

The +/- character in the **Channel** column for a radio band represents channel bonding with the adjacent (higher or lower) channel having the same width (40 MHz). The 'E' character (for example, 149E) shows channel bonding with 80 MHz width.

#### TIP

You can also view Radio frequency range from this table view, but you must first configure it to add the **Frequency Range(MHz)** column to the table view. This column is hidden (not show) by default. You can analyze the situation, optimize the network configuration, and improve the Wi-Fi client experience.

For more information about how to add columns to table views, see [Table Views](#).

### AP Radio Interference

You can troubleshoot AP Radio interference from the **Wireless Access Point Devices** table view on the out-of-the-box **Wi-Fi Health** dashboard, but you must first configure the table view to add the **2.4 GHz Channel Frequency Range**, **5 GHz Channel Frequency Range**, **6 GHz Channel Frequency Range**, and **SSID** columns to the view. These columns are hidden (not shown) by default. You can analyze the situation, optimize the network configuration, and improve the Wi-Fi client experience.

To view the **Wi-Fi Health** dashboard, hover over **Performance**, **Infrastructure Health**, and then click **Wi-Fi Health**.

The following image shows an example of the table in the **Wireless Access Point Devices** table view configured with the columns:

**Figure 123: Wireless Access Point Devices table view**

Quick Filter											On Demand	Manage Life Cycle
<input type="checkbox"/>	Name	↑	Domain	Address	Description	Curre...	Life Cyc...	Context T...	2.4 GHz Chann...	5 GHz Channel ...	6 GHz Channel ...	SSID
<input type="checkbox"/>			Default D...			Initial	Active	Device, Wi...				
<input type="checkbox"/>			Default D...			Majoi	Active	Device, Wi...				BCLMT-Person...
<input type="checkbox"/>			Default D...			Majoi	Active	Device, Wi...				BCLMT-Person...
<input type="checkbox"/>			Default D...		Description - Cisco Aironet 3800 Unified Ac...	Mino	Active	Device, Wi...				
<input type="checkbox"/>			Default D...		Description - Catalyst 9120 Unified Access ...		Active	Device, Wi...				
<input type="checkbox"/>			Default D...		Description - Cisco Business 100 Unified Ac...	Mino	Active	Device, Wi...				
<input type="checkbox"/>			Default D...		Description - Cisco Catalyst 9120AX Unified ...		Active	Device, Wi...				
<input type="checkbox"/>			Default D...		Description - Catalyst 9120 Unified Access ...		Active	Device, Wi...				
<input type="checkbox"/>			Default D...		Description - Catalyst 9105 Unified Access ...	Mino	Active	Device, Wi...				
<input type="checkbox"/>			Default D...		Description - Cisco Aironet 2800 Unified Ac...		Active	Device, Wi...				

For more information about how to add columns to table views, see [Table Views](#).



## View a List of Inventory and Performance Metrics

You can view a list of inventory—such as APs and WLCs—and performance metrics—such as most and least utilized APs and WLCs by CPU and memory usage—from the **Wi-Fi Health** dashboard.

This dashboard includes the following views:

- **Wireless Access Points Pie Table (By Severity)**

The pie chart shows the health states of the wireless APs.

- **Wireless Controllers Pie Table (By Severity)**

The pie chart shows the health states of the WLCs.

**TIP**

You can get the aggregated severity status of the alarms for a device by adding the device to the global collection in DX NetOps Spectrum (Spectrum).

- **Wi-Fi Alarms**

Displays the Wi-Fi alarms.

- **Wireless Controller Devices**

Displays the WLC devices.

- **Wireless Access Point Devices**

Displays the AP devices. You can [troubleshoot AP Radio interference from this table view](#).

- **Least Available APs**

- **Least Available WLCs**

- **Top CPU Utilization - APs**

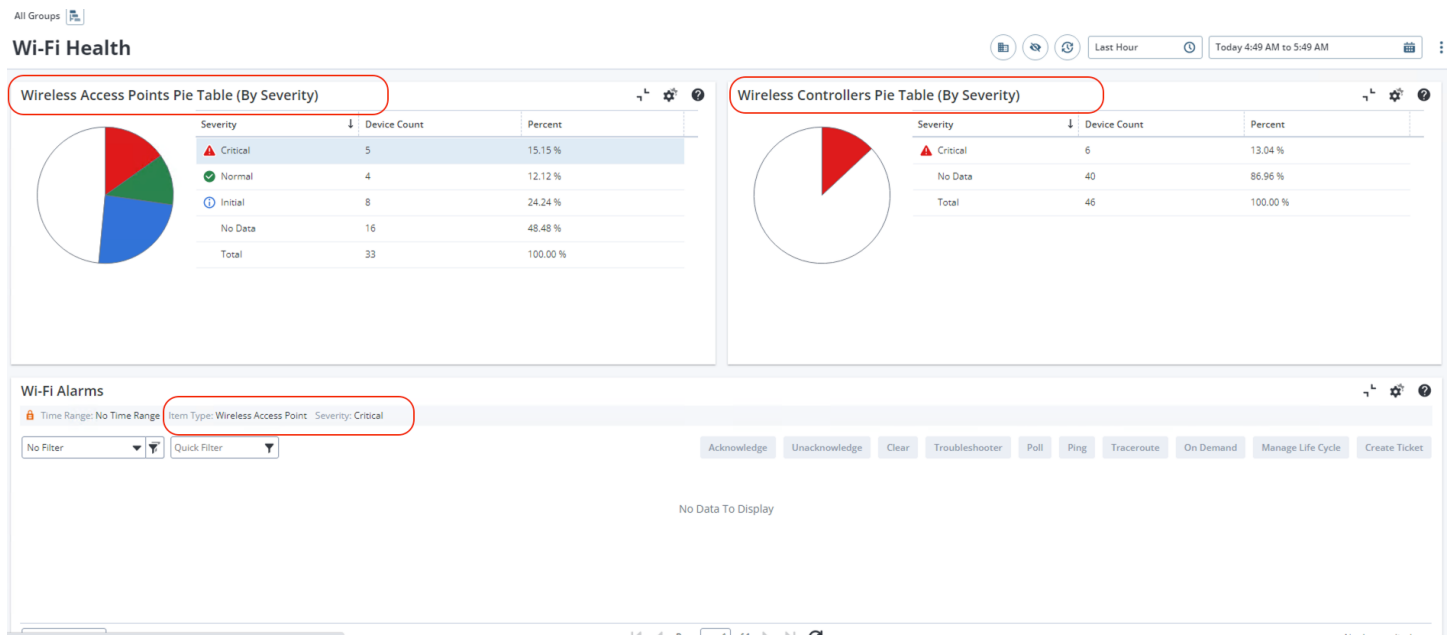
- **Top CPU Utilization - WLCs**

- **Top Memory Utilization - APs**

- **Top Memory Utilization - WLCs**

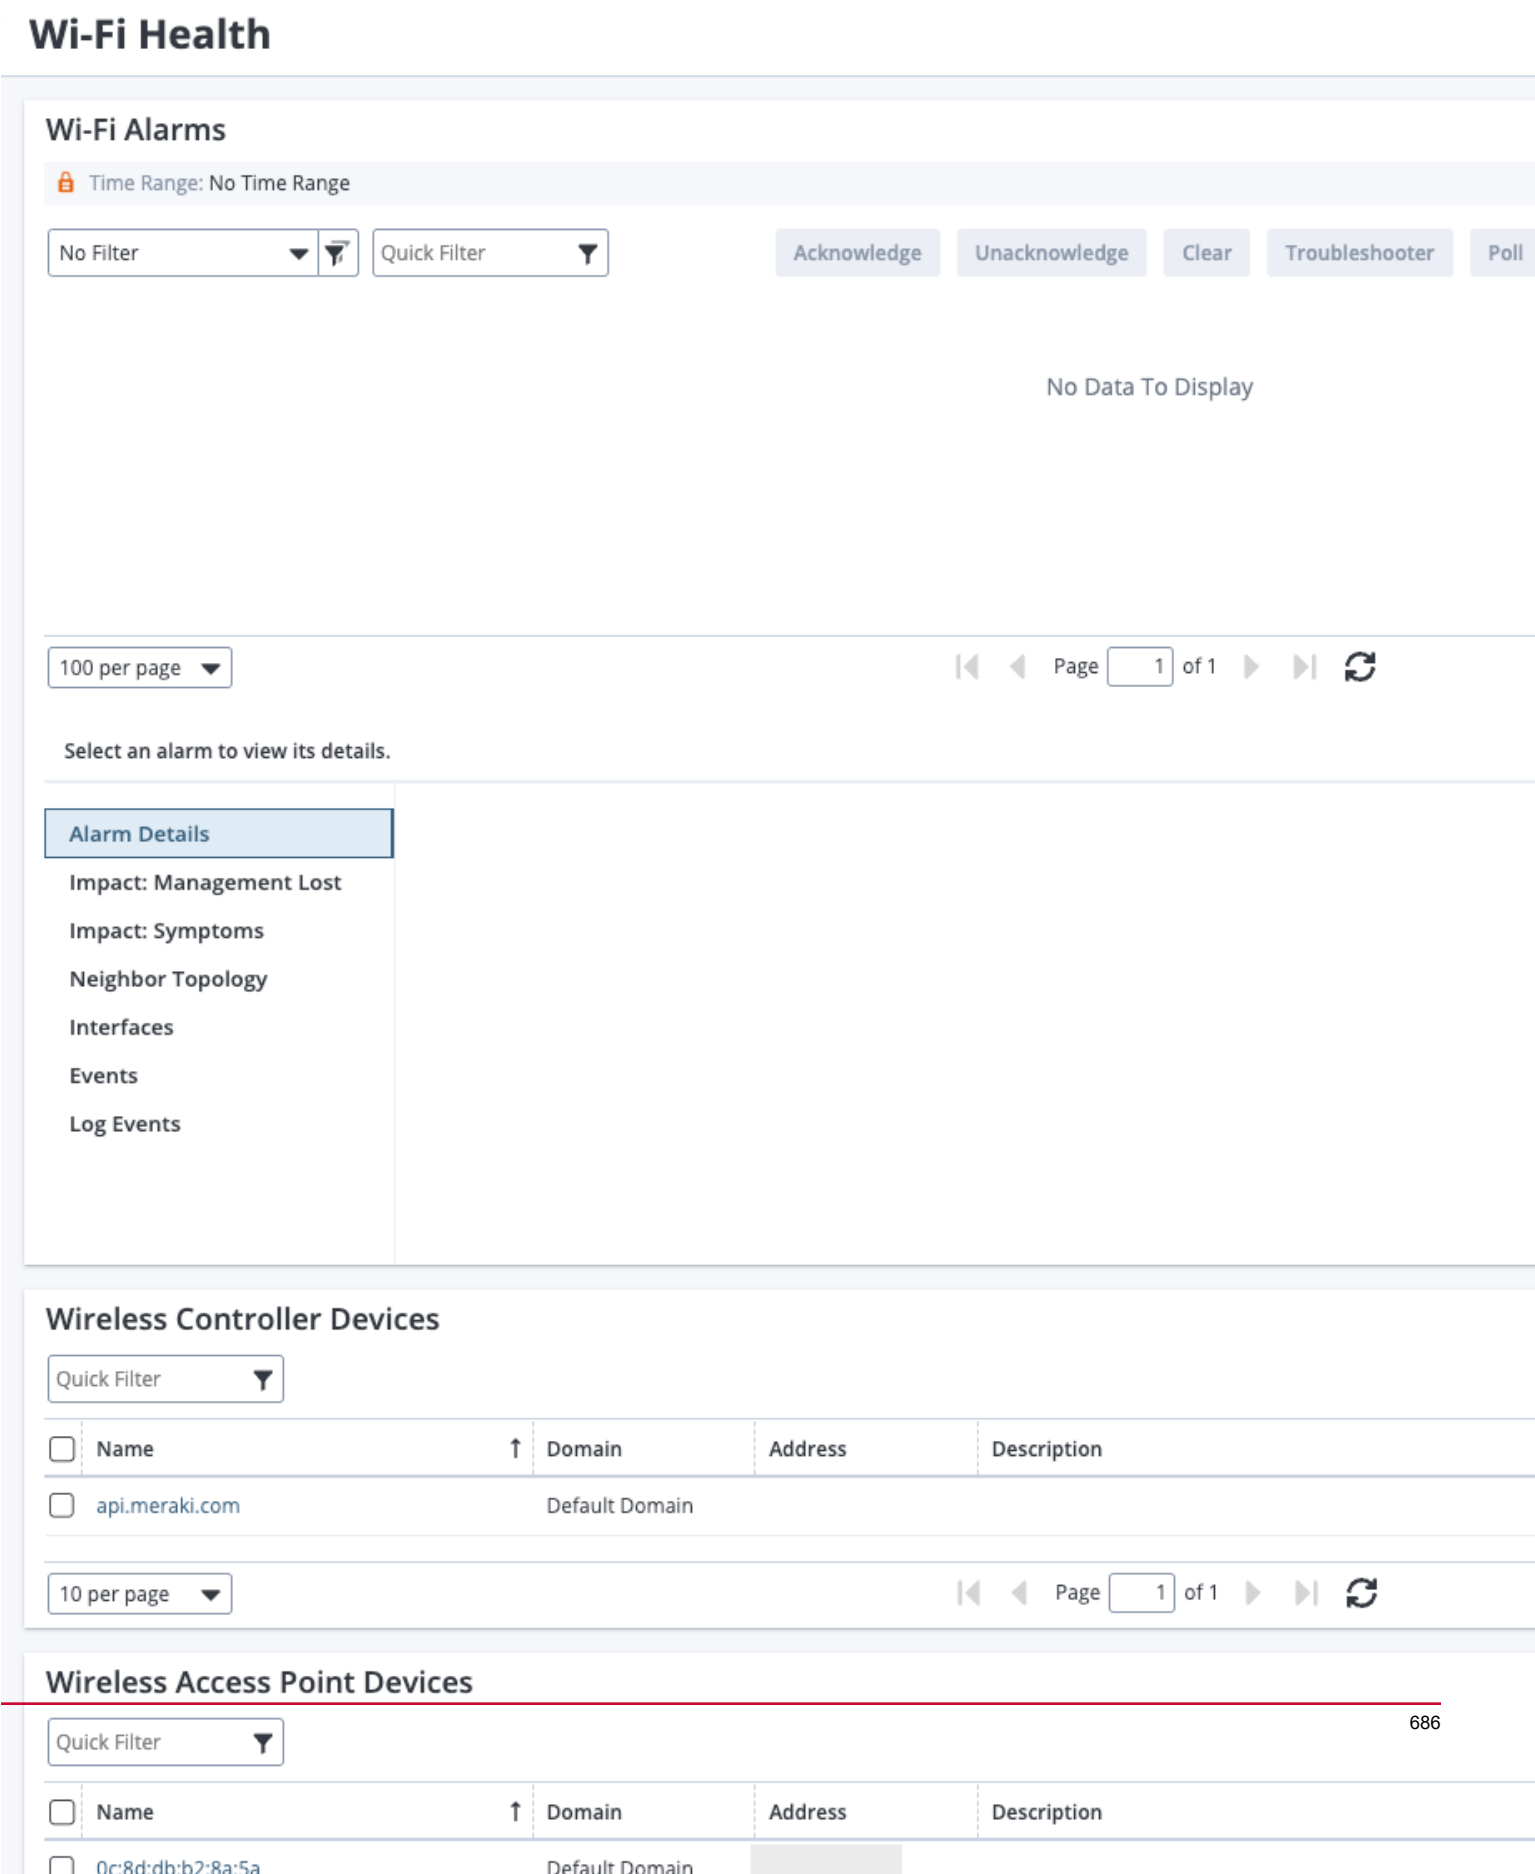
The following image shows an example of the **Wi-Fi Health** dashboard with the pie charts on the **Wireless Access Points Pie Table (By Severity)** and **Wireless Controllers Pie Table (By Severity)** views, and the **Wi-Fi Alarms** view:

**Figure 124: Wi-Fi Health dashboard with the views that include pie charts**



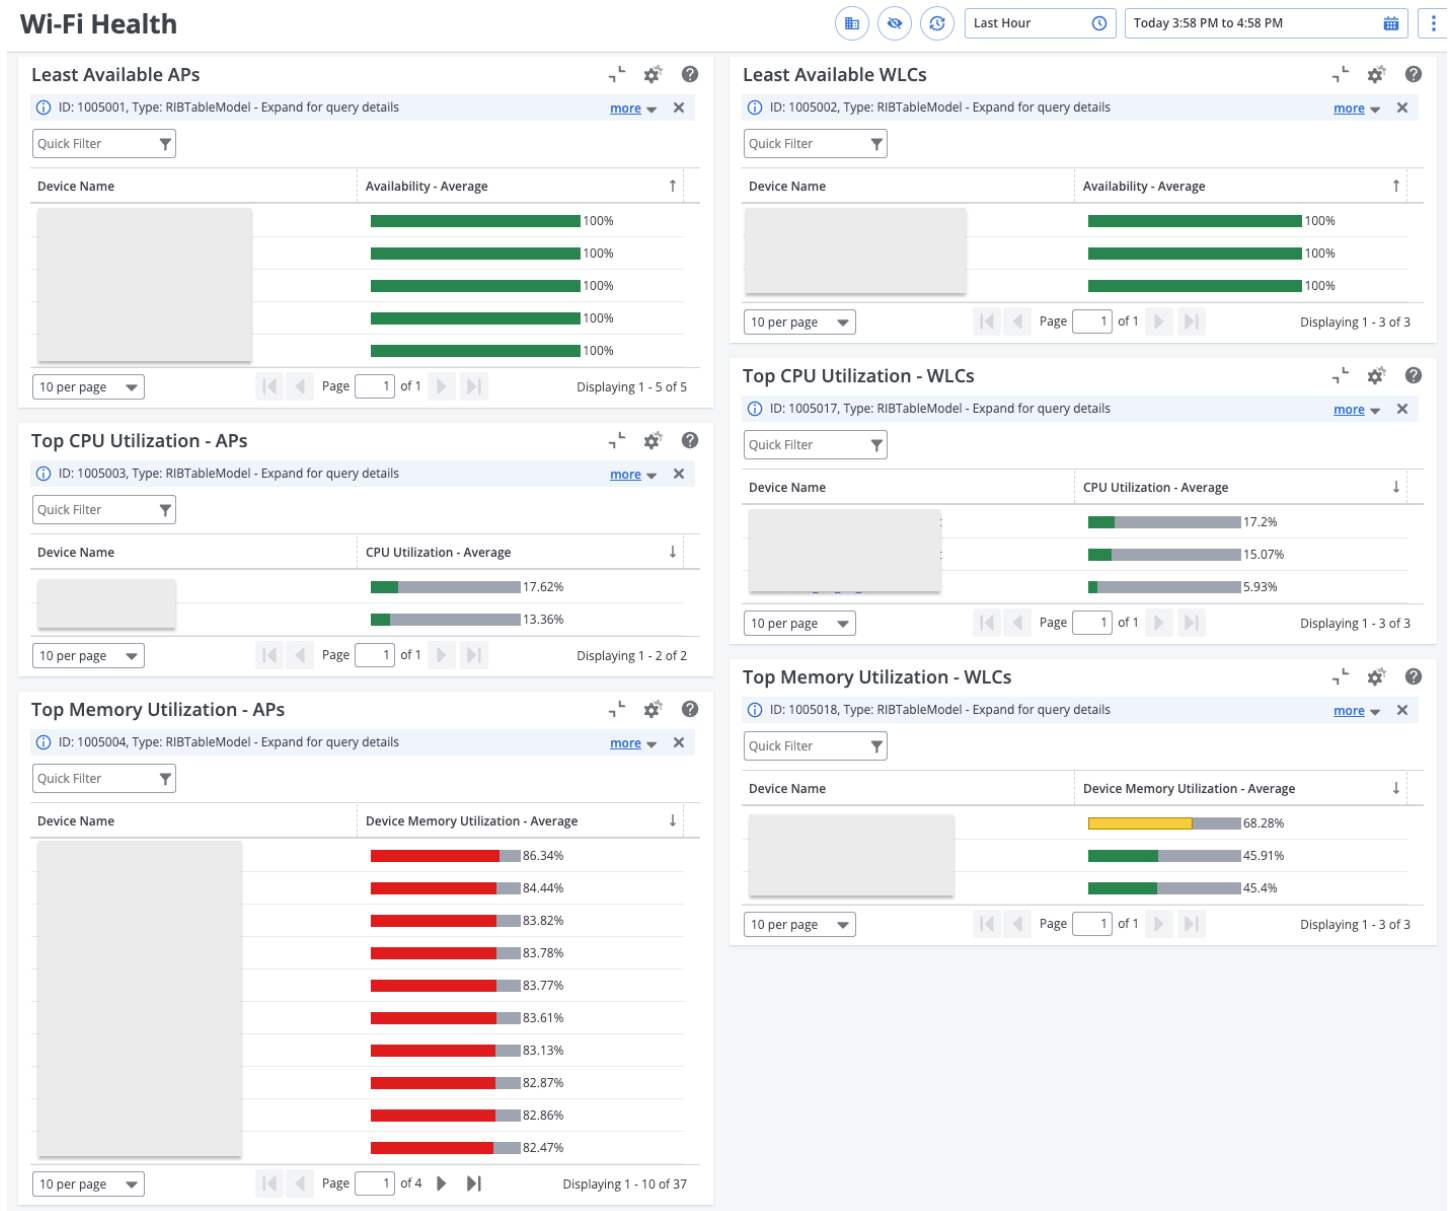
The following image shows an example of the **Wi-Fi Health** dashboard with the **Wi-Fi Alarms**, **Wireless Controller Devices**, and **Wireless Access Point Devices** views:

Figure 125: Wi-Fi Health dashboard with views



The following image shows an example of the **Wi-Fi Health** dashboard with the **Least Available APs**, **Least Available WLCs**, **Top CPU Utilization - APs**, **Top CPU Utilization - WLCs**, **Top Memory Utilization - APs**, **Top Memory Utilization - WLCs**, and **Wireless Access Point Devices** views:

**Figure 126: Wi-Fi Health dashboard with views**

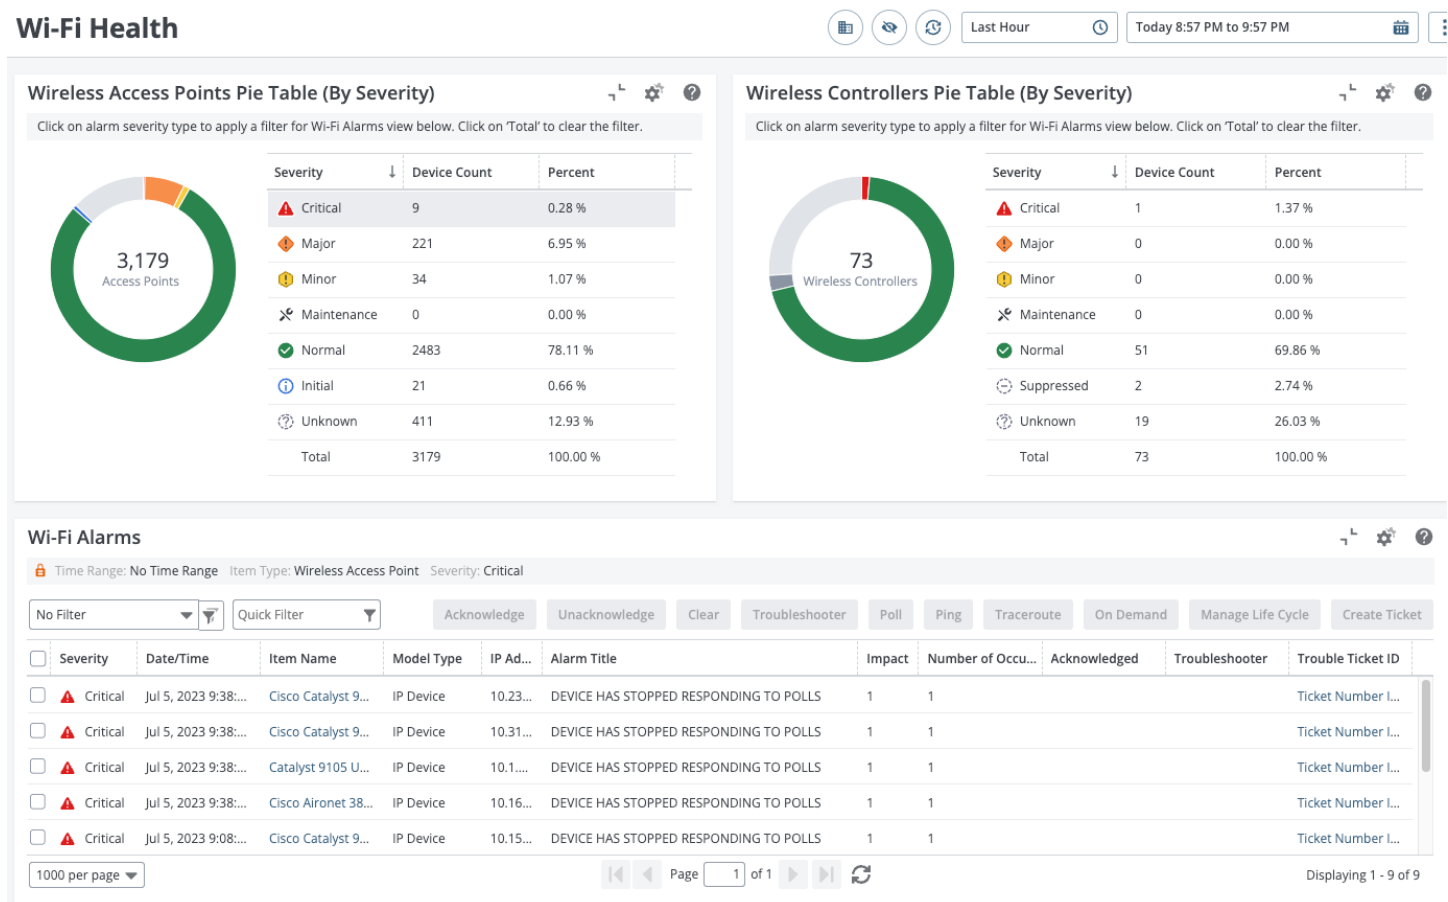


### Filter Wi-Fi Alarms

You can filter the Wi-Fi alarms that are listed in the **Wi-Fi Alarms** view on the **Wi-Fi Health** dashboard by severity and device type (AP or WLC) by clicking a row (a severity) in the legend table of the **Wireless Access Points Pie Table (By Severity)** and **Wireless Controllers Pie Table (By Severity)** views. When a filter is applied, the selected severity and the item type are shown as highlighted at the top of this view. You can clear an applied filter by clicking the **Total** row in a legend table.

The following image shows an example of the **Wi-Fi Health** dashboard with the Wi-Fi alarms filtered by Critical severity and by AP device type:

**Figure 127: Wi-Fi Health dashboard filtered**



## Monitor Monitoring Point Devices

If DX NetOps Virtual Network Assurance (VNA) is registered as a data source and the AppNeta plug-in is configured, you can monitor these devices using the dashboards and context pages in NetOps Portal.

Monitor the health of your monitoring points from the **Monitoring Point** context page in NetOps Portal. To view a list of monitoring points, go to **Inventory, Items, Devices**, and then filter for monitoring points. And then, to view the details for a monitoring point, click the name of the monitoring point. The **Details** tab is selected by default.

You can view details for the monitoring point from the following context tabs:

- **Details**  
Displays details for the monitoring point, including a description of the monitoring point and group membership. This context tab includes the **Events** view.
- **Network Paths Health**  
Displays data metrics for the network paths that have this monitoring point as its source, such as utilization and performance. This context tab includes the following views:

- **Utilization and Performance**
- **Data**
- **Voice**
- **Network Paths Health Scorecard**
- **Network Paths Trend**  
Displays data metrics for the network paths that have this monitoring point as its source, such as data loss and data jitter. This context tab includes the following views:
  - **Network Paths - Utilization, Latency and Round Trip Time**
  - **Network Paths - Data Jitter and Data Loss**
  - **Network Paths - Voice MOS, Voice Jitter and Voice Loss**
- **Custom View - Monitoring Point**
- **Alarms**  
This context tab includes the the **Monitoring Point Alarms** view which displays the alarms for the monitoring point.

For more information about AppNeta, see [the DX NetOps Virtual Network Assurance documentation](#).

## Monitor Client Device Inventory

You can monitor the network devices (clients) that DX NetOps Virtual Network Assurance (VNA) supports in NetOps Portal.

If you have integrated with VNA and you have configured a plug-in that monitors client devices, such as the Aruba Central plug-in, you can monitor client device inventory from NetOps Portal using the following methods:

- [Monitor Client Devices](#)
- [Access the Details and Performance Data for a Specific Client Device](#)
- [Monitor the Client Devices for a Specific Wireless Access Point](#)

### Monitor Client Devices

You can monitor client devices from the **Clients** page. To view this page, hover over **Inventory**, **Items**, and then click **Clients**. This page provides an overview of the client devices that NetOps Portal monitors.

The following image shows an example of this page:

**Figure 128: Clients page**

#### Clients

Clients					<span>On Demand</span> <span>Manage Life Cycle</span>	
Quick Filter <span>▼</span>						
<input type="checkbox"/>	Name	↑	Address	Description	Life Cycle State	Context Types
<input type="checkbox"/>	00:10:40:bf:2c:b0				✓ Active	Device, Wireless Client
<input type="checkbox"/>	00:10:40:bf:2f:6b				✓ Active	Device, Wireless Client
<input type="checkbox"/>	00:10:40:bf:32:23				✓ Active	Device, Wireless Client
<input type="checkbox"/>	00:10:40:bf:81:7c				✓ Active	Device, Wireless Client
<input type="checkbox"/>	00:10:40:bf:89:d7				✓ Active	Device, Wireless Client
<input type="checkbox"/>	00:22:58:1f:6d:31				✓ Active	Device, Wireless Client
<input type="checkbox"/>	00:80:92:92:82:04				✓ Active	Device, Wireless Client
<input type="checkbox"/>	00:cd:fe:ca:bb:61				✓ Active	Device, Wireless Client
<input type="checkbox"/>	02:05:2e:ee:fe:ea				✓ Active	Device, Wireless Client
<input type="checkbox"/>	02:9a:0a:85:87:18				✓ Active	Device, Wireless Client

## Access the Details and Performance Data for a Specific Client Device

From the **Clients** page, you can access the details and performance data for a specific client device on the **Details** context page for that client device by selecting that client from the list.

The following image shows an example of this context page for a client device:

**Figure 129: Details context page**

The screenshot displays the 'Details' context page for Client 315020. The page header includes the DX NetOps logo and navigation tabs: Home, Alarms, Performance, Inventory (selected), Reports, System Health, and Administration. The client's status is 'Reachability Status Unknown' and 'ACTIVE'. The page is divided into several sections:

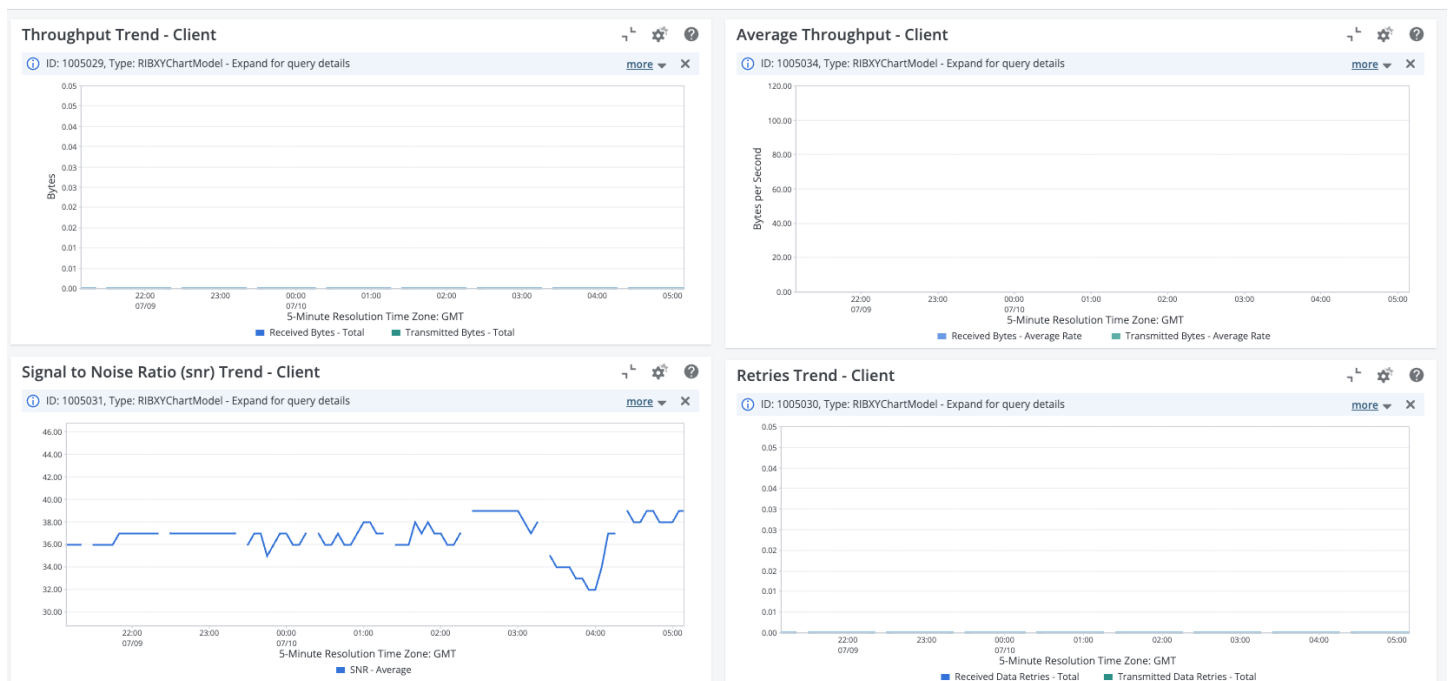
- IDENTITY:** Context Types: Device, Wireless Client; Device Name: Client 315020; Device Name Alias: Client 315020; Device IP: ; MAC Address: .
- DEVICE:** Description: ; SSID: BCLMT-Handheld.
- SNMP:** Profile: N/A.
- GROUP MEMBERSHIP:** /All Groups/Inventory/IP Domains/Default Domain; /All Groups/Inventory/All Items/Clients; /All Groups/VNA Domains/ArubaDomain1/Sites/e/Mobile Clients; /All Groups/VNA Domains/ArubaDomain1/Sites/e.
- TIME SETTINGS:** The device is not a member of a site group that includes time settings.
- CUSTOM ATTRIBUTES:** The device does not have any custom attributes.

Below these sections is the **Event List**, which includes a search filter and a 'No Data To Display' message. At the bottom, there are two trend charts: **Throughput Trend - Client** and **Retries Trend - Client**, both showing 'No Data To Display'.

The **Details** context page includes the following views:

- **Client device summary**  
Displays the details and performance data for the client device.
- **Event List**  
Displays the history of the wireless access points (APs) with which the client device is associated over time.
- **Throughput Trend - Client**  
Displays the total received and total transferred bytes from the client device.
- **Average Throughput - Client**  
Displays the average speed of the data transfer for the client device.
- **Signal to Noise Ratio (snr) Trend - Client**  
Displays the average signal-to-noise (snr) ratio trend for the client device.
- **Retries Trend - Client**  
Displays the retries that have happened while receiving and transmitting data from the client device.

The following image shows an example of the views:

**Figure 130: Views****Monitor the Client Devices for a Specific Wireless Access Point**

You can monitor client devices in context of an AP from the **Clients** context page for that AP. This page displays a list of the client devices that are connected to the AP.

For more information about APs, see [Monitor Wi-Fi Device Inventory](#).

**Follow these steps:**

1. Hover over **Inventory**, **Items**, and then click **Wi-Fi Devices**.  
The **Wi-Fi Devices** page appears.
2. Select the AP to which the client devices that you want to monitor are associated.


**TIP**

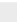






Filter the list for APs, if necessary.

The **Summary** page (the **Summary** context tab) for the AP appears.


3. Click the **Clients** context tab.  
The client devices that are associated to the AP are listed on the **Clients** context page.  
The following image shows an example of this context page for an AP:

**Figure 131: Clients context page**




Wi-Fi Devices: 

**Wireless Access Point**     Last Hour  Today 9:40 AM to 10:40 AM  

Summary

- Wireless Access Context
- Clients** 
- Wireless Controller Context
- Log Events
- Custom View - Infrastructu...
- SDN Device Metrics
- Radio Metrics - Trend Chart
- Alarms

Client Name	OS	IP Address	Radio	SSID	Max Speed(Mbps)	Protocol
Client 315020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 297020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 324020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 306020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 360020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 306020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 369020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 351020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 279020			E4:22:83:72:86:11	BCLMT-Handheld	0	
Client 56582			bc:4b:a9:aa:69:03	Testing	0	

10 per page  Page 1 of 2   Displaying 1 - 10 of 14

## Monitor VNA System Health

You can monitor DX NetOps Virtual Network Assurance (VNA) system usage and identify performance trends and degradation using the following VNA system health dashboards:

### [Virtual Network Assurance Health / Polling Dashboard](#)

VNA sends collected self-monitoring performance metrics to NetOps Portal. NetOps Portal uses these metrics to populate the **Virtual Network Assurance Health** and **Virtual Network Assurance Polling** dashboards.

For more information about the self-monitoring performance metrics that VNA sends to NetOps Portal, see [the DX NetOps Virtual Network Assurance documentation](#).

## Virtual Network Assurance Health / Polling Dashboard

You can monitor DX NetOps Virtual Network Assurance (VNA) system usage and identify performance trends and degradation.

You can do the following:

- [Track the health of VNA](#) using the **Virtual Network Assurance Health** dashboard.
- [Determine the polling load on VNA](#) using the **Virtual Network Assurance Polling** dashboard.

To view these dashboards, hover over **System Health**, and then click **Virtual Network Assurance Health/Virtual Network Assurance Polling**.

### Track the Health of VNA

Track the health of VNA using following views on the **Virtual Network Assurance Health** dashboard:

- **VNA System Statistics**  
This view shows CPU, memory, and disk utilization for the VM where the VNA is installed.
- **VNA System Memory Utilization (Trend)**  
This view shows a trend chart of the memory utilization for the VM where the VNA is installed.
- **VNA System CPU Utilization (Trend)**  
This view shows a trend chart of the CPU utilization for the VM where the VNA is installed.
- **VNA WildFly Statistics**  
This view shows CPU and memory utilization for WildFly.
- **VNA WildFly Memory Utilization (Trend)**



This view shows a trend chart of the memory utilization for WildFly.

- **VNA WildFly CPU Utilization (Trend)**

This view shows a trend chart of the CPU utilization for WildFly.

### **Determine the Polling Load on VNA**

Determine the polling load on VNA using following views on the **Virtual Network Assurance Polling** dashboard:

- **VNA Engine Statistics**

This view shows the total API requests processed and the failed API requests for the VM where the VNA is installed.

- **VNA Engine Total API Requests Processed (Trend)**

This view shows a trend chart of the total API requests processed for the VM where the VNA is installed.

- **VNA Engine Failed API Requests (Trend)**

This view shows a trend chart of the failed API requests for the VM where the VNA is installed.

- **VNA ActiveMQ Statistics**

This view shows the total messages processed and the queued messages for ActiveMQ.

- **VNA ActiveMQ Total Messages Processed (Trend)**

This view shows a trend chart of the total messages processed for ActiveMQ.

- **VNA ActiveMQ Queued Messages (Trend)**

This view shows a trend chart of the queued messages for ActiveMQ.

## **Fault Monitoring**

DX NetOps Spectrum (Spectrum) provides fault monitoring of managed elements, including devices, applications, host systems, and connections.

You can share models, Global Collections, and events between Spectrum and DX NetOps Performance Management.

For more information:

- About Spectrum, see [the DX NetOps Spectrum documentation](#).
- About how to integrate with Spectrum, see [Integrate with DX NetOps Spectrum](#).

## **Customize Event Integration**

By default, DX NetOps Spectrum polls threshold violation and clear events only. You can also configure DX NetOps Spectrum to poll for events in the Event Manager database.

### **NOTE**

For DX NetOps Spectrum to process an event, the device or port must be modeled in DX NetOps Spectrum and included in the synchronization process. To process events that are not associated with an item that is modeled in DX NetOps Spectrum, configure a trap notification in DX NetOps Performance Management, and use the DX NetOps Spectrum South Bound Gateway.

To determine the event types available to add to DX NetOps Spectrum, see the `em.event_types` table in the Event Manager database.

Use the following process to configure DX NetOps Spectrum to poll for specific events:

1. [Review the configuration example](#).
2. [Obtain a developer ID](#).
3. [Update the `netqos-integration-application-config.xml` file](#).
4. [Update the event disposition file](#).
5. [Create event format files](#).
6. [Create probable cause files](#).

## Review the Configuration Example

This example shows how to configure DX NetOps Spectrum to poll for a specific event in the Event Manager database. The event in this example identifies when a router device experiences high memory usage.

1. Identify a device or port for which you want DX NetOps Spectrum to poll for in the Event Manager database. If the device or port is not modeled in DX NetOps Spectrum, model the element. For example, to monitor specific events for a particular router, the router must be modeled in the DX NetOps Spectrum database.
2. Obtain a developer ID from Broadcom Support for use with the integration with DX NetOps Spectrum. This example uses the default developer ID value, 0xffff.
3. Identify the events for which DX NetOps Spectrum polls. For example, you can identify Incident events that originates from CA Application Delivery Analysis.
4. Define the event by completing the following steps:
  - a. Copy the `<${SPECROOT}>\tomcat\webapps\spectrum\WEB-INF\netqos\config\container\netqos-integration-application-config.xml` file to the `<${SPECROOT}>/custom/netqos/config/container` directory.
  - b. Open the `/custom/netqos/config/container/netqos-integration-application-config.xml` file for editing.
  - c. Define the custom event. Update the existing `eventTypeManager` element as follows:
    - a. Add the Incident event to the list of events for which to poll.
    - b. Establish an alarm map value.
    - c. Specify a default alarm clear code.

The following code shows these changes. Note that the alarm clear code uses a developer ID:

```
<bean id="eventTypeManager"
  class="com.ca.im.netqos.integration.event.type.EventTypeManager">
  <property name="interestingEventTypes">
    <map>
      <entry key="ThresholdViolation" value-ref="thresholdViolationAlarmCodes" />
      <entry key="Incident" value-ref="IncidentAlarmCodes" />
    </map>
  </property>
  <property name="alarmClearCodes">
    <map>
      <entry key="ThresholdViolation" value="0x5c40009" />
      <entry key="Incident" value="0xffff0004" />
    </map>
  </property>
</bean>
```

- d. Define the alarm map by adding the following bean element:

```
<bean id="IncidentAlarmCodes"
  class="org.springframework.beans.factory.config.MapFactoryBean">
  <property name="sourceMap">
    <map>
      <entry key="1" value="0xffff0001" />
      <entry key="2" value="0xffff0002" />
      <entry key="3" value="0xffff0003" />
    </map>
  </property>
</bean>
```

- e. Save and close the file.

5. Specify how DX NetOps Spectrum processes the encountered event by completing the following steps:
  - a. Open the `<${SPECROOT}>\SS\CsVendor\netqos\EventDisp` file for editing.

- b. Add the following map entries for the Incident event:

```
#Incident Event
0xffff0001E 50 A 1, 0xffff0001,107
0xffff0002E 50 A 2, 0xffff0002,107
0xffff0003E 50 A 3, 0xffff0003,107
0xffff0004E 50 C 0xffff0001,107 C 0xffff0002,107 C 0xffff0003,107
```

- c. Save and close the file.

- d. Back up this file in case the contents are changed during a DX NetOps Spectrum upgrade.

6. Complete the following steps:

- a. Create an event format file for each of the alarm codes using the following naming convention (*AlarmCode - EventFormatFile*):

- 0xffff0001 - Eventffff0001
- 0xffff0002 - Eventffff0002
- 0xffff0003 - Eventffff0003
- 0xffff0004 - Eventffff0004

#### NOTE

When creating the Eventffff0004 file, use appropriate wording for clearing an alarm.

- b. Create a text file containing content similar to the following text:

```
{d "%w- %d %m-, %Y - %T"} - {S 109} is reporting a minor threshold violation.
Detail of Threshold Violation:
  1) Incident Start Time: {D 111}
  2) Event ID: {S 107}
  3) Event Source: {S 113}
  4) Alert Message: {S 76620}
A corresponding minor Threshold Violation Alarm will be generated.
(event [{e}])
```

- c. Save the file to the <\$SPECROOT>\SG-Support\CsEvFormat directory.

- d. Repeat steps a and b for each alarm code.

7. Complete the following steps for each alarm code:

- a. Create a probable cause file for each of the alarm codes using the following naming convention (*AlarmCode - ProbableCauseFile*):

- 0xffff0001 - Probffff0001
- 0xffff0002 - Probffff0002
- 0xffff0003 - Probffff0003
- 0xffff0004 - Probffff0004

- b. Create a text file containing content similar to the following text:

```
A minor threshold violation has occurred.
SYMPTOMS:
The monitored threshold has been exceeded.
PROBABLE CAUSES:
RECOMMENDED ACTIONS:
Launch the "Performance View" to see incident details.
```

- c. Save the file to the <\$SPECROOT>\SG-Support\CsPCause directory.

8. Restart the SpectroSERVER and OneClick servers.

DX NetOps Spectrum is configured to poll for a specific event in the Event Manager database. DX NetOps Spectrum uses the updated files to poll for the Incident event, generating events and alarms as specified.

## Obtain a Developer ID

When defining events for the integration with DX NetOps Spectrum, you use identifying event codes. The first 2 bytes of any event code contain a developer ID. You can obtain a registered developer ID from Broadcom so that you can specify unique codes for your events. Using a unique developer ID lets you easily recognize your new codes in OneClick. Doing so also prevents potential conflicts with other DX NetOps Spectrum event codes.

Obtain a developer ID from Broadcom by contacting Broadcom Support.

## Update the netqos-integration-application-config.xml File

DX NetOps Spectrum determines the events for which to poll using the `netqos-integration-application-config.xml` file. DX NetOps Spectrum polls for `ThresholdViolation` events by default. To poll for more events, modify this file to define the [event codes](#) and [associated alarms](#) for each event. This file is located in the `$SPECROOT\tomcat\webapps\spectrum\WEB-INF\netqos\config\container` directory.

## Define Events

The `eventTypeManager` bean defines the events for which DX NetOps Spectrum polls. The entries for `ThresholdViolation` events appear in the file by default. You can manually add more events.

```
<bean id="eventTypeManager"
  class="com.ca.im.netqos.integration.event.type.EventTypeManager">
  <property name="interestingEventTypes">
    <map>
      <entry key="ThresholdViolation" value-ref="thresholdViolationAlarmCodes" />
      <entry key="TestEvent" value-ref="TestEventAlarmCodes" />
    </map>
  </property>
  <property name="alarmClearCodes">
    <map>
      <entry key="ThresholdViolation" value="0x5c40009" />
      <entry key="TestEvent" value="TestEventAlarmClearCode" />
    </map>
  </property>
</bean>
```

Add events that DX NetOps Spectrum can include in polling by updating the following property elements:

- **interestingEventTypes**  
Specifies the types of events to include in polling. Each entry element identifies a specific event type and an alarm code map value. The `ThresholdViolation` entry is included by default. Add an entry element, as follows:

```
<entry key="TestEvent" value-ref="TestEventAlarmCodes" />
```

- **TestEvent**  
Specifies the name of an event in the Event Manager database.
- **TestEventAlarmCodes**  
Specifies the value of the map that identifies the alarms for this event.

### NOTE

The alarm code map is described in the next section.

- **alarmClearCodes**  
Specifies the alarm clear codes for polled events. The default alarm clear code for the `ThresholdViolation` event is `0x5c40009`. For each event, add an entry element, as follows:

```
<entry key="TestEvent" value="TestEventAlarmClearCode" />
```

- **TestEvent**

Specifies the name of the event that was added for polling.

– **TestEventAlarmClearCode**

Specifies the alarm clear code for the event.

## Define Alarms

An alarm map defines the alarm code values that are associated with a particular event. For each polled event (or, each `interestingEventTypes` entry), define a corresponding alarm map. The alarm map for the `ThresholdViolation` event appears in the file by default. Manually add an alarm map for each custom event.

```
<bean id="thresholdViolationAlarmCodes"
  class="org.springframework.beans.factory.config.MapFactoryBean">
  <property name="sourceMap">
    <map>
      <entry key="1" value="0x5c40010" />
      <entry key="2" value="0x5c40011" />
      <entry key="3" value="0x5c40012" />
    </map>
  </property>
</bean>
<bean id="testEventAlarmCodes"
  class="org.springframework.beans.factory.config.MapFactoryBean">
  <property name="sourceMap">
    <map>
      <entry key="alarmSev1" value="alarmCode1" />
      <entry key="alarmSev2" value="alarmCode2" />
      <entry key="alarmSev3" value="alarmCode3" />
    </map>
  </property>
</bean>
```

To add alarm maps for custom events, add a bean element for each event, and then update the following values:

- **testEventAlarmCodes**  
Specifies the alarm code map value for a particular event. This value is established on the `interestingEventTypes` entry and must match that value.
- **alarmSev1 - alarmCode1, alarmSev2 - alarmCode2, alarmSev3 - alarmCode3**  
Specifies the `alarmSeverity` - `alarmCode` pairs for a particular event. For example, for the default `ThresholdViolation` event, the Minor (1), Major (2), and Critical (3) alarm codes are 0x5c40010, 0x5c40011, and 0x5c40012, respectively.

## Update the Event Disposition File

The Event Disposition (`EventDisp`) file is used to determine how to process the events that are configured in the `netqos-integration-application-config.xml` file. Each event entry maps an event to a DX NetOps Spectrum event file.

The `EventDisp` file for the integration with DX NetOps Spectrum is located in the `<${SPECROOT}>\SS\CsVendor\netqos` directory.

For the default `ThresholdViolation` event, the following entries map the alarm codes to individual DX NetOps Spectrum event files:

```
text#PC Threshold 0x5c40010 E 50 A 1,0x5c40010,107 0x5c40011 E 50 A 2,0x5c40011,107
0x5c40012 E 50 A 3,0x5c40012,107 0x5c40009 E 50 C 0x5c40010,107 C 0x5c40011,107 C
0x5c40012,107
```

For each custom event, add new event map entries to the file. The following example shows syntax that generates or clears alarms that are based on the event code:

```
text#New Event alarmCode1E 50 A
1, alarmCode1_filename,107 alarmCode2E 50 A
2, alarmCode2_filename,107 alarmCode3E 50 A
3, alarmCode3_filename,107 alarmClearCode4E 50 C alarmCode1,107 C alarmCode2,107 C
alarmCode3,107
```

### Create the Event Format Files

Event format files contain the messages about the event that is displayed to users on the **Events** tab in OneClick. The events that are defined in the `netqos-integration-application-config.xml` file require an event format file. The file enables the event to display correctly in the OneClick **Events** view. The file name must match the alarm code (for example, alarm code 0x5c40010 uses the file "Event05c40010"), and it must exist in the `<$SPECROOT>\SG-Support\CsEvFormat` directory.

The following example shows the file format:

```
text{d "%w- %d %m-, %Y - %T"} - {S 109} is reporting a minor threshold violation.
Detail of Threshold Violation:
  1) Incident Start Time: {D 111}
  2) Event ID: {S 107}
  3) Event Source: {S 113}
  4) Alert Message: {S 76620}
A corresponding minor Threshold Violation Alarm will be generated.
(event [{e}])
```

### Create Probable Cause Files

Probable-cause files define the symptoms, probable causes, and recommended corrective actions for alarms. Alarm codes require a probable-cause file so that the alarm displays correctly in the OneClick **Alarms** view. The file name must match the alarm code (for example, alarm code 0x5c40010 uses the `Prob05c40010` file), and it must exist in the `<$SPECROOT>\SG-Support\CsPCause` directory.

The following example shows the file format:

```
textA minor threshold violation has occurred.
SYMPTOMS:
The monitored threshold has been exceeded.
PROBABLE CAUSES:
RECOMMENDED ACTIONS:
Launch the "Performance View" to see incident details.
```

### Next Steps

After you have configured DX NetOps Spectrum to poll for specific events, to have polling reflect the changes, deploy these changes by restarting the SpectroSERVER and OneClick servers.

**NOTE**

If you have multiple SpectroSERVERs, distribute the event file configuration changes to all servers.

## Device Configuration View

View DX NetOps Spectrum (Spectrum) device configuration in NetOps Portal.

If you have integrated with Spectrum, you can view current (running) and previously-captured Spectrum device configurations.

Use the following process to view device configurations:

1. [Verify the prerequisites.](#)
2. [View device configurations.](#)

### Verify the Prerequisites

Before you can see the **Device Configuration** view on a device's context page, verify the following prerequisites:

- You have integrated DX NetOps Performance Management with Spectrum, which includes enabling the data aggregator to discover Spectrum devices.  
For more information:
  - About how to enable synchronized discovery in DX NetOps Spectrum, see [the DX NetOps Spectrum documentation](#).
  - [Integrate with Spectrum](#).
- You have configured Network Configuration Manager in DX NetOps Spectrum.  
For more information about Network Configuration Manager and the devices that Network Configuration Manager supports, see [the DX NetOps Spectrum documentation](#).

**TIP**

You can configure an unsupported device by creating a custom device family using the Network Configuration Manager Extension utility.

For more information about this utility, see [the DX NetOps Spectrum documentation](#).

### View Device Configurations

You can view device configurations in the **Device Configuration** pane on a device's context page.

Follow these steps:

1. In NetOps Portal, hover over **Inventory**, **Items**, and then click **Devices**.  
The **Devices** page appears.
2. Click the device for which you want to view device configurations.  
The device's context page opens to the **Details** tab.
3. Click the **Device Configuration** tab.  
Spectrum device configuration data, such as routers, displays in the **Device Configuration** pane. It shows the device configuration history table that lists the configurations that Spectrum has captured for the device.  
By default, this pane displays the following attributes:
  - **Capture Time**  
The date and time when Spectrum captured the configuration for the device.
  - **Line Changes**  
The number of lines that have changed when Spectrum compares a configuration file to the previous configuration file in the configuration history table.
  - **Reference**

Indicates if the configuration has been designated as the reference for the device. If the current configuration is different from the reference, you can optionally generate an alarm.

- **Running vs. Startup**

Indicates whether a difference between the device's startup and running configuration is available.

- **Last Verified Time**

The time Spectrum last verified this configuration to present on the device.

4. (Optional) To show the contents of the device configuration text file for a device configuration, click the device configuration.

The contents display in the **Configuration Details** pane that opens below the **Device Configuration** pane.

Sensitive device configuration information, such as passwords, only displays for Administrators. For all other users, this information is masked out (gray boxes are displayed) unless the user has the **Show Unmasked Device Configurations** role right.

For more information about this role right, see [Role Rights](#).

Sensitive information is determined by the device family mask, which you specify in Spectrum.

For more information, see [the DX NetOps Spectrum documentation](#).

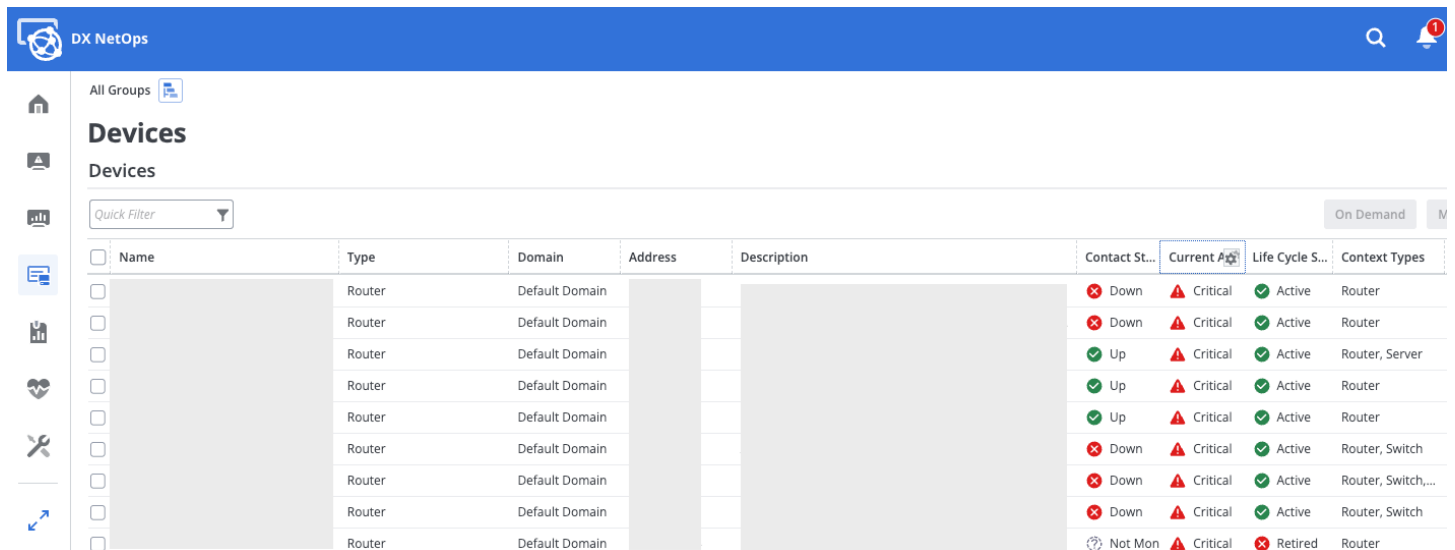
## Monitor Device Inventory with Alarm States

If you have integrated with DX NetOps Spectrum (Spectrum), you can monitor device inventory with alarm states from NetOps Portal.

Monitor device inventory with alarm states from the **Devices** page. To view this page, hover over **Inventory**, **Items**, and then click **Devices**. This page provides an overview of the device inventory with alarm states.

The following images show examples of this page based on your version:


**Figure 132: (23.3.3 and higher) Devices Page**






	Name	Type	Domain	Address	Description	Contact St...	Current Alarm	Life Cycle S...	Context Types
<input type="checkbox"/>		Router	Default Domain			Down	Critical	Active	Router
<input type="checkbox"/>		Router	Default Domain			Down	Critical	Active	Router
<input type="checkbox"/>		Router	Default Domain			Up	Critical	Active	Router, Server
<input type="checkbox"/>		Router	Default Domain			Up	Critical	Active	Router
<input type="checkbox"/>		Router	Default Domain			Up	Critical	Active	Router
<input type="checkbox"/>		Router	Default Domain			Down	Critical	Active	Router, Switch
<input type="checkbox"/>		Router	Default Domain			Down	Critical	Active	Router, Switch,...
<input type="checkbox"/>		Router	Default Domain			Down	Critical	Active	Router, Switch
<input type="checkbox"/>		Router	Default Domain			Not Mon	Critical	Retired	Router




**Figure 133: (23.3.2 and lower) Devices Page**

















All Groups 

## Devices

Devices   

spectrum 

On Demand Manage Life Cycle

<input type="checkbox"/>	Name	Type	Domain	Address	Description	Current Alarm State 	Life Cycle...	Context Types	DevCusto...
<input type="checkbox"/>		Router	Default Dom...			 Active		Router, Switch	
<input type="checkbox"/>		Virtual Machine	Default Dom...			 Active		Device	
<input type="checkbox"/>		Virtual Machine	Default Dom...			 Active		Device	
<input type="checkbox"/>		Virtual Machine	Default Dom...			 Active		Device	
<input type="checkbox"/>		Virtual Machine	Default Dom...			 Active		Device	
<input type="checkbox"/>		Router	Default Dom...			 Normal	 Active	Router, Switch	
<input type="checkbox"/>		Router	Default Dom...			 Normal	 Active	Router, Switch	
<input type="checkbox"/>		Virtual Machine	Default Dom...			 Normal	 Active	Device	
<input type="checkbox"/>		Virtual Machine	Default Dom...			 Normal	 Active	Device	
<input type="checkbox"/>		Router	Default Dom...			 Minor	 Active	Router, Switch	

The following columns on the **Devices** page contain alarm state information:

- **Current Alarm State**

Displays the current alarm state of the device, which is the alarm severity level in Spectrum. This column is displayed in the view by default.

**NOTE**

This column displays a state only if NetOps Portal is discovering the device from Spectrum. Otherwise, it is blank.

**Values:**

- Normal

Indicates the following:

- Spectrum OneClick has made contact with the device, and the device is operating typically.
- A "Normal" alarm is generated.
- An event generated an alarm but a severity for the alarm is not specified. For example, you created the supporting `EventDisp` configuration file manually and inadvertently omitted a severity.

- Minor

Indicates that an abnormal situation exists, but no immediate action is required.

**NOTE**

On the **Details** tab of a device context page, this status shows as "Minor Alarm".

- Major

Indicates that a loss of service has occurred or is impending. Action is required within a short period.

**NOTE**

On the **Details** tab of a device context page, this status shows as "Major Alarm".

- Critical

Indicates that a loss of service has occurred and immediate action is required.

**NOTE**

On the **Details** tab of a device context page, this status shows as "Critical Alarm".

- Maintenance

Indicates that the device has been taken offline for maintenance purposes.

**NOTE**

On the **Details** tab of a device context page, this status shows as "Maintenance Alarm".

- Suppressed

Indicates that Spectrum OneClick cannot reach the device due to a known error condition that exists on another device.

**NOTE**

On the **Details** tab of a device context page, this status shows as "Suppressed Alarm".

- Initial

Indicates that Spectrum OneClick has not yet established contact with the device.

**NOTE**

On the **Details** tab of a device context page, this status shows as "Initializing Alarm".

- Unknown

Indicates that NetOps Portal is not discovering the device from Spectrum and that the alarm state is unknown. In the column, the current alarm state for the device is shown as blank.

**NOTE**

On the **Details** tab of a device context page, this status shows as "Alarm Status Unknown".

For more information:

- About these alarm states, see [the DX NetOps Spectrum documentation](#).
- About how to view Spectrum alarm data in NetOps Portal, see [Alarms View](#).

- **Last Alarm State Change**

Displays the date and time that the alarm state last changed. This column is hidden from the view by default.

**NOTE**

You can also:

- View device status, including the current alarm state, from the **Details** tab of a device context page, by clicking the name of the device from the **Devices** page.  
For more information, see [Reachability Status and Contact Status](#).
- Monitor device inventory with alarm states from DX NetOps Spectrum OneClick (go to **Inventory**, **Consoles**, and then click **Spectrum OneClick WebApp**).
- Add this column to views and device context pages.  
For more information, see [Customize Views](#).

For more information about how to integrate with Spectrum, see [Integrate with DX NetOps Spectrum for Fault Management](#).

## Configure Notifications

Configure notifications for events that come from a data source to the event manager.

The event manager evaluates incoming events against the conditions that you configure for the notification criteria. Only when the criteria are met does it take a notification action. If an event does not trigger a notification, you can still display the event in the **Event** list.

Users configure and receive notifications only for events for items in groups to which they have access.

Consider the following information:

- Users can only see their own notifications.
- Deleting a notification for an event does not affect the event or future events.

In this article:

- [Manage Notifications](#)
- [Create a Notification that Sends Messages Related to Events Automatically](#)
- [Notification Actions](#)

## **Manage Notifications**

Administrators can manage notifications by viewing, creating, or deleting them. Manage notifications on the **Manage Notifications** page.

**Prerequisite:** Event Manager is enabled and in the Available synchronized state.

### **NOTE**

As a default tenant administrator, you can work in a real-user context to create notifications for a tenant administrator or tenant user. Log in as a tenant administrator or tenant user. The default tenant administrator can also administrate the tenant, and then create a tenant-scoped notification by proxying to the user.

### **Follow these steps:**

1. Do *one* of the following tasks:
  - Hover over **Administration**, and **Configuration Settings**, and then click **Notifications**.
  - Click the name of your user account in the upper-right corner, and then click **Manage Notifications**.
 The **Manage Notifications** page appears.

## **Create a Notification that Sends Messages Related to Events Automatically**

You can create notifications that sends messages that are related to events automatically.

**Prerequisite:** The Event Manager is enabled and in the Available synchronized state.

The following video examines how to set up notifications for threshold violation events in NetOps Portal to notify teams when threshold violation events occur or change:

### **Follow these steps:**

1. From the **Manage Notifications** page, click **New**.  
The **Create Notification** dialog opens.
2. Specify a name and description, and then click **Next**.
3. Select the groups that generate events to trigger the notification, and then click **Next**.
4. Select conditions for the notification, and then click **Next**.
5. Specify how the event manager should execute the notification, and then click **Next**:

- **Email**

Send email notifications.

### **NOTE**

To include the device name for events that are triggered on components, use the `Item Parent Name` property.

### **TIP**

To create or update a notification email template, select **Save** or **Update Email Template**. Changes to templates do not affect existing messages.

For more information about this notification action, see [the "Email Notification Action" section](#).

- **Trap**

Send trap notifications. Multiple destinations are supported, but the first destination is required.

Two MIB choices are available in the Notifications wizard to provide compatibility for existing customers.

For more information about this notification action, see [the "Trap Notification Action" section](#).

- **Script**

Users with the Create Notifications role right can create or edit script notification actions. Specify the script file name. All scripts must be executable. Store scripts in the `/<installation_directory>/PerformanceCenter/NotificationScripts` directory. You can use scripts in another language such as Python or Perl.

- ***installation\_directory***

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

For more information about this notification action, see [the "Script Notification Action" section](#).

NetOps Portal saves the notification and sends messages when the selected conditions occur.

## **Notification Actions**

When you create a notification, you specify one of the following actions:

### **Email Notification Action**

The event manager sends email notifications to one or more recipients when an event is raised or cleared. The email provides a link to the context page for the device or component that triggered the alarm.

#### **NOTE**

To use the hostname in the URL instead of the IP address of the NetOps Portal host, configure the **Web Site Host**.

For more information, see [Configure the DX NetOps Security Settings Using the SSO Configuration Tool](#).

**Supported roles:** Users with a role that contains the Create Notifications role right and Event Manager access can configure email notifications. However, the Administrator role must first specify an SMTP server.

In the **Email** tab, select **Enable**, and then configure the email notification settings.

### **Trap Notification Action**

The event manager sends trap notifications to fault or network management system (NMS) in your environment, such as DX NetOps Spectrum (Spectrum).

#### **NOTE**

If you have integrated with Spectrum, while Spectrum can receive traps, this method is not the preferred method for the integration.

For more information, see [the DX NetOps Spectrum documentation](#).

#### **IMPORTANT**

Create an SNMP profile with the outgoing trap port (typically 162) before creating the notification.

**Supported roles:** Users with the Administrator role (global administrators) can configure trap notifications. Administrators must also have product privileges to Event Manager and data sources that create events.

The trap receivers must be preconfigured to receive traps. Each destination can have its own configuration regarding SNMP community and IPV4 destination. To receive and decrypt SNMPv3 traps, the SNMP profile for the trap receiver should match this notification configuration.

For more information, see [SNMP Profiles](#).

For more information about trap formats, see the corresponding NMS documentation for your trap receiver.

### **Script Notification Action**

Scripts can store events to a database, forward notifications to multiple systems, and send specific types of notifications to some specific system. You can log output from scripts that you own. Script return codes are logged in the in the NetOps Portal log file. The event manager executes script notification actions serially to ensure that it processes sets before it processes clears. If the queue of unprocessed script notification actions exceeds a certain size (5000 by default), the

event manager drops any new incoming script notification actions. It generates an event when events start dropping and when the processing of new script notification actions restarts. If a script takes too long to complete execution (300 seconds by default), the event manager kills the script, and then it generates an event. The event manager generates script notification action for events on the data aggregator item in NetOps Portal.

For more information about how to change event properties that are related to script notification actions, see [Change Event Properties](#).

#### NOTE

Store scripts in the `/<installation_directory>/PerformanceCenter/NotificationScripts` directory.

- ***installation\_directory***  
The default installation directory for NetOps Portal.  
**Default:** `/opt/CA`

The event manager automatically passes the following parameters to script notification actions:

- CAPM\_EventDataSource
- CAPM\_EventCategory
- CAPM\_EventType
- CAPM\_EventSubType
- CAPM\_EventState
- CAPM\_EventSeverity
- CAPM\_EventOccurredOn
- CAPM\_EventDesc
- CAPM\_ItemParentName
- CAPM\_ItemName
- CAPM\_ItemNameAlias
- CAPM\_ItemDesc
- CAPM\_IPAddress
- CAPM\_ItemUrl
- CAPM\_ItemType
- CAPM\_ItemSubtype
- CAPM\_ItemId
- CAPM\_ItemParentId

#### NOTE

Event-specific properties are also available and prefixed with `CAPM_EvProp_`. The `CAPM_ItemId` and `CAPM_ItemParentId` parameters are NetOps Portal IDs. The script parameters and notifications are always in English regardless of the NetOps Portal language.

### Example: Define a Script that CPU Utilization Threshold Events Can Trigger

In this example, you define a script that prints environment variables to a file. CPU utilization threshold events can trigger this script.

1. Define the following `printenv.sh` script in the `/<installation_directory>/PerformanceCenter/NotificationScripts/` directory by issuing the following command:  

```
printenv > /<installation_directory>/PerformanceCenter/NotificationScripts/out.txt
```

  - ***installation\_directory***  
The default installation directory for NetOps Portal.  
**Default:** `/opt/CA`
2. Make the script executable by issuing the following command:

```
chmod u+x printenv.sh
```

3. Configure notifications so that a CPU utilization threshold event triggers the `printenv.sh` script.  
The script creates the following output:

```
CAPM_EventCategory=PERFORMANCE
CAPM_EvProp__Severity=1
CAPM_IPAddress=10.253.223.1
CAPM_ItemNameAlias=cisco2621-10.253.223.1
CAPM_EvProp_AlarmRuleID=5,313
CAPM_EvProp_AlarmAggregationMethod=No Aggregation
CAPM_ItemId=118
CAPM_EvProp_AlarmProfileName=Test cpu
CAPM_ItemDesc=Cisco Internetwork Operating System Software ^M
IOS (tm) C2600 Software (C2600-IK903S3-M), Version 12.3(9), RELEASE SOFTWARE
(fc2)^M
Copyright (c) 1986-2004 by cisco Systems, Inc.^M
Compiled Fri 14-May-04 14:37 by dchih
CAPM_EventState=OPENED
CAPM_EventType=ThresholdViolation
PWD=/<installation_directory>/bin
CAPM_EventDesc=A Threshold Violation event has been raised. (Profile Name: Test
cpu, Rule Name: test cpu)
CAPM_ItemName=CPU 2
CAPM_ItemParentId=118
CAPM_EvProp_ThresholdProfileFolderId=5,311
CAPM_EvProp_AlarmDuration=60
CAPM_EvProp_AlarmProfileId=5,314
CAPM_EvProp_AlarmViolationRuleDetail=Utilization > 50.0
SHLVL=1
CAPM_EvProp__Alarm_ID=1500
CAPM_ItemType=DEVICE
CAPM_ItemParentName=cisco2621-10.253.223.1
CAPM_EventSeverity=MAJOR
CAPM_EvProp_AlarmMetricFamilyName=CPU
CAPM_EventOccurredOn=Thu May 24 10:39:00 EDT 2018
CAPM_EvProp_AlarmRuleName=test cpu
CAPM_EventDataSource=Data Aggregator@fergi04-dev-da
CAPM_ItemUrl=http://10.237.15.180:8181/pc/desktop/page?
pg=r&DeviceID=118&timeRange=-1&startTime=2018-05-24+10%3A09+America
%2FNew_York&endTime=2018-05-24+11%3A09+America%2FNew_York
CAPM_ItemSubtype=router
CAPM_EvProp_AlarmWindow=60
CAPM_EventSubType=Raised
_=/usr/bin/printenv
```

– **`em_installation_directory`**

The default installation directory for the NetOps Portal Event Manager service.

**Default:** `/opt/CA/PerformanceCenter/EM`

**NOTE**

If an error occurs during the script execution, the errors are logged in the `<em_installation_directory>/logs/EMService.log` file. For example, if the script does not exist or the script is not executable, an error is logged. The exit code of the script is also logged.

**Example: Define a Script that Device Life Cycle Changes Can Trigger**

In this example, you define a script that prints environment variables to a file. Device life cycle changes can trigger this script.

1. Define the following `printenv.sh` script in the `<installation_directory>/PerformanceCenter/NotificationScripts/` directory by issuing the following command:

```
printenv > <installation_directory>/PerformanceCenter/NotificationScripts/env.txt
```

- **installation\_directory**

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

2. Make the script executable by issuing the following command:

```
chmod u+x printenv.sh
```

3. Configure notifications so that device life cycle change events trigger the `printenv.sh` script.

The script creates the following output:

```
CAPM_EventCategory=CONFIG
CAPM_EvProp__Severity=Unknown
CAPM_IPAddress=138.42.96.2
CAPM_ItemNameAlias=138.42.96.2 - alias
CAPM_EvProp_User=admin
CAPM_ItemId=110
CAPM_EvProp_CurrState=RETIRED
CAPM_EvProp_PrevState=ACTIVE
CAPM_ItemDesc=RS 38000 - Riverstone Networks, Inc. Firmware Version: 9.4.1.1 PROM
Version: prom-2.0.1.8
CAPM_EventState=CLOSED
CAPM_EventType=LifeCycle
PWD=<em_installation_directory>/bin
CAPM_EventDesc=LifeCycle - Change
CAPM_ItemName=rs38000-96.2
CAPM_ItemParentId=110
SHLVL=1
CAPM_ItemType=DEVICE
CAPM_ItemParentName=rs38000-96.2
CAPM_EventSeverity=
CAPM_EventOccurredOn=Mon Dec 18 16:20:58 EST 2017
CAPM_EventDataSource=CA Performance Center
CAPM_ItemUrl=http://10.237.10.219:8181/pc/desktop/page?
pg=r&DeviceID=110&timeRange=-1&startTime=2017-12-18+15%3A50+America
%2FNew_York&endTime=2017-12-18+16%3A50+America%2FNew_York
CAPM_ItemSubtype=router
CAPM_EventSubType=Change
_=/usr/bin/printenv
```

- **em\_installation\_directory**

The default installation directory for the NetOps Portal Event Manager service.

**Default:** /opt/CA/PerformanceCenter/EM

**NOTE**

If an error occurs during the script execution, the errors are logged in the `<em_installation_directory>/logs/EMService.log` file. For example, if the script does not exist or the script is not executable, an error is logged. The exit code of the script is also logged.

## Traps Usage

This topic describes traps usage information in DX NetOps Performance Management.

This topic describes traps usage information in DX NetOps Performance Management.

### Get the SNMPengineID

The trap receiver (for example, DX NetOps Spectrum), using SNMPv3, allows encrypted traps from a specific sender for decryption using the `SNMPengineID`.

Get the `SNMPengineID` using the following REST endpoint:

**URL:**

`http://<PC_host>:<port>/EventManager/webservice/notifications/engineId`

- **PC\_host**  
Specifies the NetOps Portal host name.
- **port**  
Specifies the NetOps Portal port number for communication between NetOps Portal and the Event Manager.  
**Default:** 8281

**Method:** GET

**Return:**

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
  <EngineID>
    <ID>
      80:00:13:70:01:8a:2a:f8:41
    </ID>
  </EngineID>
```

The ID string is the `SNMPengineID`.

### EventManager Format Usage

The EventManager MIB is supported for trap notifications.

**TIP**

You can find the MIB files in the `<installation_directory>/PerformanceCenter/PC/webapps/pc/mibs/netqos-em-mib` directory.

- **installation\_directory**  
The directory where NetOps Portal is installed.  
**Default:** /opt/CA

When you select the EventManager format, the trap is sent out with the following variables:

- **netQosEventId**



Specifies an identifier that Event Manager assigned to the event.

- **netQoSEventType**  
Specifies the type of event.
- **netQoSEventCategory**  
Categorizes the event.  
**Values:** 0 Unknown, 1 Fault, 2 Config, 3 Accounting, 4 Performance, 5 Security
- **netQoSEventSeverity**  
Specifies the severity of the event.  
**Values:** 0 Normal, 1 Unknown, 2 Minor, 3 Major, 4 Critical, 5 Unavailable
- **netQoSEventDescription**  
Describes the event.
- **netQoSEventState**  
Specifies the current state of the event. Each state has its own notification.  
**Values:** 0 opened, 1 acknowledged, 2 closed, 3 cleared
- **netQoSEventOpenTime**  
Specifies the UTC timestamp (from the eventState timestamp).
- **netQoSEventMapURL**  
No value is available. The "" string is sent.
- **netQoSEventDetailsURL**  
No value is available. The "" string is sent.
- **netQoSEventAssociatedItemURL**  
Specifies the URL to the item web page.
- **netQoSEventItemName**  
Specifies the item name. There is one notification per item.  
**Maximum length:** 127 bytes
- **netQoSEventItemType**  
Specifies the item type.  
**Maximum length:** 32 bytes
- **netQoSEventItemSubtype**  
Specifies the item subtype.  
**Maximum length:** 32 bytes
- **netQoSEventItemIpAddress**  
Specifies an IP address for the item or an empty string.
- **netQoSEventPropertyName**  
Specifies one name set for each property. A PropertyName exists for each property in the event. (The properties vary by the event type.)  
**Maximum length:** 128 bytes
- **netQoSEventPropertyValue**  
Specifies the property value for the event. A PropertyValue exists for each property in the event. (The properties vary by the event type.)

## Identify Volatility in Network Performance

You can identify volatility in network performance using the **Interface Volatility** dashboard.

### NOTE

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We fully intend to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per capability basis.

The **Interface Volatility** dashboard captures performance of interfaces specific to their volatility. It displays a list of the most volatile interface components for a specific group. You can compare the interface components to other interface components in the group across multiple metrics of varying unit types using the interface component rankings. You can also determine where performance is variable and unpredictable, and then accurately evaluate, predict, and plan resource availability and performance.

To view the **Interface Volatility** dashboard, hover over **Performance**, **Analytics**, and then click **Interface Volatility**.

#### NOTE

The **Analytics** menu and the **Interface Volatility** dashboard are not be visible by default. An Administrator must add the **Analytics** menu to a user account role to which your user account is assigned.

The **Interface Volatility** dashboard includes the **Volatile Interfaces** view.

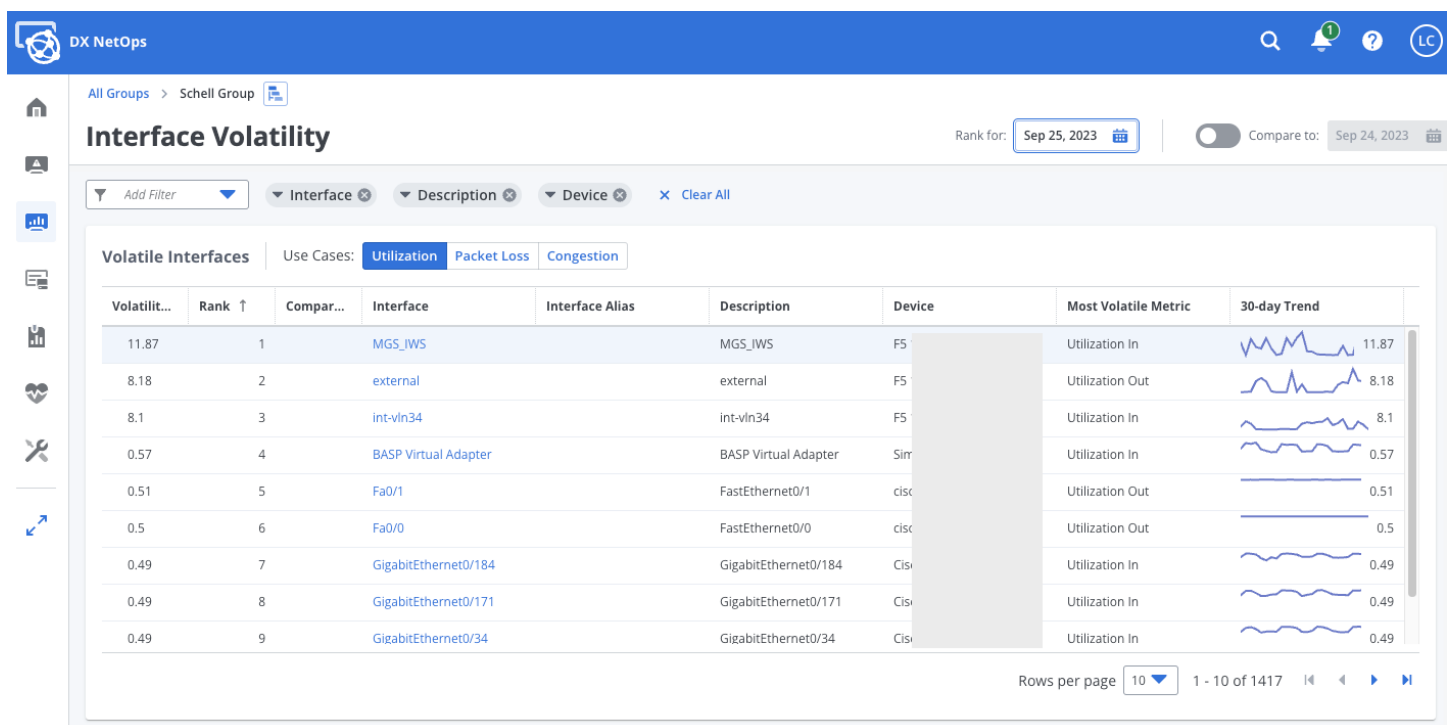
From this view, you can:

- View rapid and unpredictable change in measured performance for interface components in a group.
- Predict and plan resource availability and performance.
- Identify volatile interface components across a range of use cases. The use cases are aligned to metrics that are applicable to those use cases. The weighting indicator drives priority of each associated metric within the volatility calculations.

For more information, see the ["Refine Volatility Analysis" section](#).

The following image shows an example of the **Interface Volatility** dashboard with the **Utilization** use case selected:

**Figure 134: Interface Volatility**



### Refine Volatility Analysis

On the **Interface Volatility** dashboard, refine volatility analysis by focusing on one of the following use cases:

- **Utilization:** Utilization In (50%) and Utilization Out (50%)
- **Packet Loss:** Bits In (25%), Bits Out (25%), Discards In (25%), Discards Out (25%)
- **Congestion:** Utilization In (10%), Utilization Out (10%), Percent Discards In (20%), Percent Discards Out (20%), Errors In (20%), Errors Out (20%)

## Generate a REST API Token

You can protect your password and not send it in clear text by generating a REST API token.

Clients using a REST API to access NetOps Portal must include a time-limited token as part of its REST calls. The token expires after the selected time period.

Generate a REST API token using the following methods:

- [Use NetOps Portal](#)
- [Use a REST client](#)

You can then use the generated REST API token for the following circumstances:

- [When connecting a REST client to the NetOps Portal REST web services.](#)
- (23.3.5 and higher) [When identifying empty and unused groups](#)

Also in this article: [Refresh an Expired Token from a REST client](#)

### Generate the Token from NetOps Portal

Follow these steps:

1. Log in to NetOps Portal.
2. Click the name of your user account in the upper-right corner, and then click **Generate REST API Token**. The **Generate REST API Token** page opens, with the generated REST API token displaying in the **REST API Token** field.

**Example:**

```
NvKlaq6v9LJRgj_nbWgQo3JwSgPrcoXnBJjjQyI3y1zSoNyhkXL0A3ipNmytYDIT8rjy6CHnI6jcTGn08wspHg**
```

3. (Optional) Define the time period at which the token expires in the **Token Expiration** field.

**Values:** 15 Minutes, 30 Minutes, 1 Hour, and 2 Hours

**Default:** 1 Hour

A new REST API token is generated and displays in the **REST API Token** field.

4. Copy the token by clicking the **Copy** icon.

The REST API token is copied to your clipboard.

### Generate the Token from a REST Client

By default, the REST API token using this method expire after 15 minutes.

Follow these steps:

1. Launch a REST client.
2. Enter the following URL for the `BuildToken` endpoint for the NetOps Portal RESTful web services API in the **URL** field, selecting **PUT** for the **HTTP Method**:

```
http://<PC_host>:8381/sso/webservices/BuildToken
```

**XML Format:**

```
<BuildToken xmlns="http://netqos.com/singlesignon/">
  <ssoProductCode>pc</ssoProductCode>
  <username>...</username>
```

```
<password>...</password>
</BuildToken>
```

### Sample results (REST API token):

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<BuildTokenResponse xmlns="http://netqos.com/singlesignon/">

  <BuildTokenResult>NvKlaq6v9LJRgj_nbWgQo3JwSgPrcoXnBJjjQyI3y1zSoNyhkXL0A3ipNmytYDIT8rjy6CHnI6jcTGn08wspHg**</
BuildTokenResult>
</BuildTokenResponse>
```

3. From the **Body**, copy the generated REST API token.

### Refresh an Expired Token from a REST Client

#### Follow these steps:

1. Launch a REST client.
2. Enter the following URL for the `RefreshToken` endpoint for the NetOps Portal RESTful web service API in the **URL** field, selecting **PUT** for the **HTTP Method**:

```
http://<PC_host>:8381/sso/webservices/RefreshToken
```

#### XML Format:

```
<RefreshToken xmlns="http://netqos.com/singlesignon/" >
  <ssoProductCode>pc</ssoProductCode>
  <token>...</token>
</RefreshToken>
```

The token is refreshed.

## Log Out of NetOps Portal

NetOps Portal is the web UI for DX NetOps Performance Management. While logged in to NetOps Portal, to log out, click the name of your user account in the upper-right corner, and then click **Log Out**.

# Administrating

---

This section includes information about how to monitor and maintain the system components, and how to manage users, roles, and tenants, and system configuration options.

## Onboard a New Product Operator

As an Administrator, you can give a coworker with a unique role in the IT organization permissions to use DX NetOps Performance Management.

Administrators can give coworkers in the IT organization permissions to use DX NetOps Performance Management. They give permissions by creating custom user accounts for those users, or operators. Usually, the Administrator assigns predefined menus and roles to these accounts. However, to create a new user with a unique organizational role, the Administrator must define these custom menus and roles.

Onboarding an operator involves creating a menu and populating it with either predefined or custom dashboards. You also create a custom role for the new user, and assign that role to a new user account.

Use the following process to onboard a new product operator:

1. [Set up custom groups for your organization.](#)  
Custom user accounts are best deployed in a well-planned system that includes custom groups. Administrators assign custom groups as permissions to user accounts to restrict these user accounts to viewing only the data, menus, and dashboards that they require to perform their daily tasks.
2. [Add a menu for the operator.](#)  
Administrators can create a custom menu for the operator to reflect this person's organizational responsibilities.
3. [Create a custom dashboard for the operator.](#)
4. [Edit the menu to add the dashboard.](#)  
Administrators and designers can customize menus to meet the requirements of each operator. When you edit custom or out-of-the-box menus, you can add new dashboards, remove dashboards, and change the order of the dashboards in the menu.
5. [Add a custom role.](#)  
You can add a custom user role for each operator. The user account roles let operators perform their job responsibilities. Assign new roles to any menus that you have customized for the intended operator. Roles are disabled until they are assigned to user accounts.
6. [Add a custom user account.](#)  
Add a user account for each operator. For security purposes, do not share user accounts with multiple people.  
**NOTE**  
You can create user accounts with basic parameters, and then edit them as a separate step to assign permissions. This workflow lets you carefully consider the groups that each operator must access.
7. [Assign permissions to the user account.](#)  
Individual operators require data access permissions to monitor data, which are based on groups. You can assign access permissions according to your plan for custom groups. To assign permissions, edit user accounts. Make sure that all operators see only the data that they require for their role.
8. [Log in to test the new user account.](#)

### **Log in to Test the New User Account**

To test a new user account, log in to the account. You can test user accounts while logged in as the administrator using a Proxy. However, user account proxying does not let you test roles or role rights.

**Follow these steps:**

1. Log in using the username and password that you assigned to the new user account.
2. Click the **Inventory** tab.
3. Click item links to verify that the user can see monitored items in the inventory.
4. Select a few dashboards, and verify that the views are populated with data.
5. Test your ability to select a new group or item context for a dashboard by taking the following steps:
  - a. Click the [change] link above the time period selectors.  
A dialog opens with filtering options.
  - b. Click to select another managed item.
  - c. Expand nodes in the Groups tree to select a group context.
  - d. Click **OK**.
6. Verify that the user can see data from the new item or group.
7. Test any special role rights that you assigned to the user, such as saving changes to views.
8. Log out when you are satisfied that the user account meets the requirements of the intended operator.

You are now ready to contact the new operator and provide the username and password.

## Manage Data Sources

Data sources are the products and components that provide performance and configuration data in NetOps Portal.

Data aggregator data source administration is integrated into NetOps Portal. CA Application Delivery Analysis, DX NetOps Network Flow Analysis, and CA Unified Communications Monitor collect and aggregate data independently. However, after you have registered these data sources in NetOps Portal, NetOps Portal controls many administrative functions. NetOps Portal does not administrate DX NetOps Spectrum.

You manage data sources by registering them, testing the data source connections, editing them, and deleting them.

For more information, see [Configure a Data Source](#).

### Maximum Registered Data Source Instances

You can register up to the following number of data sources for each type in NetOps Portal:

Data Source	Maximum Number Supported
CA Application Delivery Analysis	10
DX NetOps Network Flow Analysis (NFA)	10
CA Unified Communications Monitor	4
DX Application Performance Management	1
CA Catalyst Connector	1
Event Manager	1
DX NetOps Spectrum (Spectrum)	1

### Conflict or Duplicate Handling

For user accounts and SNMP profiles, NetOps Portal uses the following processes to handle conflicts or duplicates:

- [Redundant User Accounts](#).
- [Redundant SNMP Profiles](#).

## **Redundant User Accounts**

Users can have two different accounts with the same name in different data sources. The resulting user account retains the password of the first account that is synchronized. The unique role rights and permissions from other accounts are added to the account as you register more data sources.

While multiple user accounts sometimes share a username in different data sources, some account parameters differ. You must edit these account parameters.

### **Example:**

A user named Robert uses DX NetOps Network Flow Analysis, and a different user, also named Robert, uses CA Application Delivery Analysis. In this case, NetOps Portal creates one account that is named Robert. NetOps Portal merges the role rights and permissions from both data sources into a single new account. To preserve the distinct role rights of the two accounts, create an account with a unique username.

## **Redundant SNMP Profiles**

Registering a data source that contains SNMP profile definitions automatically adds the profiles to NetOps Portal. The profiles are distributed to the other registered data sources during the next synchronization.

When a data source is added, NetOps Portal minimizes duplication of SNMP profiles by comparing the following values to existing profiles:

- User for SNMPv3
- Community String for SNMPv1 and SNMPv2

If the compared values match, NetOps Portal retains the SNMP profile with the most recent timestamp.

If the compared values do not match, NetOps Portal appends a number to duplicate profile names. For example, the first profile that is named Boston remains Boston. The second profile becomes Boston(1).

## **Configure a Data Source**

You can view a list of registered data sources in NetOps Portal, register a new data source, test them, edit them, and delete them.

NetOps Portal represents sources of information as registered data sources. To load performance data, register these as data sources in NetOps Portal. You can configure your data source for secure communication using Hypertext Transfer Protocol Secure (HTTPS).

You can do the following to configure a data source:

- [Register a Data Source](#)
- (If data source registration process does not complete successfully) [Test the Data Source Connections](#)
- [Edit a Data Source](#)
- [Delete a Data Source](#)

These procedures require the Administrator role.

### **Register a Data Source**

To view collected data on NetOps Portal dashboards, add, or register, the data sources.

Share the monitoring parameters, user accounts, and other definitions with DX NetOps Performance Management and other registered data sources by registering CA Application Delivery Analysis, DX NetOps Network Flow Analysis, and CA Unified Communications Monitor as data sources in DX NetOps Performance Management. During registration, DX NetOps Performance Management imports shared monitoring data, such as user accounts, SNMP profiles, and other administrative data from these data sources. DX NetOps Performance Management resolves conflicts and eliminates

duplication. At the next synchronization, DX NetOps Performance Management sends updated administrative data to registered data sources.

The registration process includes a "binding" step that prevents further modifications to shared administrative data in individual data sources. As a result, the data source administrator can only modify shared monitoring parameters after registration.

The following video shows how to bind the data aggregator to NetOps Portal and confirm that the data collector is properly connected to the data aggregator:

### Follow these steps:

1. Hover over **Administration**, **Data Sources**, and then click **Data Sources**.  
The **Manage Data Sources** page appears.
2. Click **Add**.  
The **Add Data Source** dialog opens.
3. Complete the following fields:
  - **Source Type**  
Select the type of data source.  
This list includes all possible data sources. The list does not show data sources that already have reached the registration limit.  
**Values:** Application Delivery Analysis, Application Performance Management, Catalyst Connector, Data Aggregator, Network Flow Analysis, Spectrum Infrastructure Manager, Unified Communications Monitor  
**Default:** Application Delivery Analysis
  - **Status**  
Defines the status of the data source.  
**Values:**
    - **Enabled:** NetOps Portal synchronizes the data source.
    - **Disabled:** NetOps Portal delays synchronization with the data source.**Default:** Enabled
  - **Host Name**  
Specifies the hostname for the data source. If you are registering a data aggregator data source, specify the IP address or hostname of the data aggregator host. If you are registering a fault-tolerant data aggregator, specify the proxy server. For other data sources, specify the IP address or hostname of the management console.  
**IMPORTANT**  
If you are planning to enable HTTPS for this data source, enter the hostname that the certificate uses. Typically this is the fully qualified domain name (FQDN), but this depends on the form of the name that was used when the certificate was created.  
**NOTE**  
**Prerequisite:** HTTPS is enabled for NFA and Spectrum.  
For more information, see [the DX NetOps Spectrum documentation](#) and [the DX NetOps Network Flow Analysis documentation](#).  
For more information about fault-tolerant data aggregators, see [Fault Tolerance](#).  
**Character limit:** 255  
**Required:** Yes
  - **Port**  
Specifies the port for the data source.  
**Default:** 80  
**Required:** Yes
  - **Protocol**  
Defines the communication protocol.  
**Options:**



- **http:** The communication between NetOps Portal and this data source is unsecured.
- **https:** The communication between NetOps Portal and this data source is secure.

**NOTE****Prerequisites:**

- (If you are registering an NFA or Spectrum data source ) HTTPS is enabled for [NFA](#) or [Spectrum](#).
- [HTTPS is enabled for NetOps Portal](#).
- The data source self-signed certificate or the CA intermediate/root certificate(s), if signed, has been imported. If NetOps Portal does not restart after importing the certificate, restart all NetOps Portal services.

**Default:** http– **Display Name**

Specifies the display name for the data source. By default, DX NetOps Performance Management combines the data source type and the hostname to create the display name. You can specify another name. For example, instead of NetworkFlowAnalysis@192.168.10.22, specify the name for the data source as NFA\_NewYork.

**Character limit:** 255**Required:** Yes– **Contribute inventory to Data Aggregator****Options:**

- **Selected:** Contribute new data source inventory to the data aggregator.

**NOTE**

Automatic synchronization includes only devices that NetOps Portal discovers after you select this option. To discover previously discovered devices, perform a full synchronization of the data aggregator. DX NetOps Performance Management creates a discovery profile that includes the IP addresses of the devices. This discovery profile attempts discovery 1 minute after new IP addresses are pushed into the IP domain discovery profile for the data aggregator. Otherwise, the discovery profile attempts discovery once per day.

For more information, see [Configure a Data Source](#) and [Synchronize Data Sources](#).

- **Cleared:** The existing inventory remains the same, and the new data source inventory is not contributed to the data aggregator.

**Default:** Cleared**Required:** No– **Web Console**

If the Web Console is on a different system than the data source, clear the **Same as data source** check box, and then specify another **Host Name** and **Port** for the data source console.

**Default:** Selected**Required:** No**Character limit:** For **Host Name**, 255**NOTE**

Use this parameter in cases when network address translation is deployed.

– (Data aggregator data sources only) **Synchronize component items that are not currently present on the monitored device**

Select to have the data aggregator change the status of device components that are no longer present in the environment to "Not Present". By default, the data aggregator does not synchronize device components that have the "Not Present" status because it cannot collect data for them. You can still access historical data that has not reached the data retention limit for these items.

**Default:** Cleared**IMPORTANT**

If the properties of an active component match the identifying properties of the not-present component, the components are indistinguishable. Data from the not-present component might contribute to group-

based dashboards instead of data from the active component. Select this option *only* under the following circumstances:

- You want to report on historical data for no-longer-present items in the environment.
- You want to ensure that not-present items do not conflict with an actively-monitored item.
- You want to identify not-present items so that you can exclude them from groups.

– (DX NetOps Spectrum data sources only) **Synchronize device life cycle state from Spectrum**

When selected, DX NetOps Spectrum controls the life cycle state of devices in NetOps Portal. When a device is deleted in DX NetOps Spectrum, DX NetOps Spectrum changes the life cycle state of that device in NetOps Portal to "Retired".

**Default:** Cleared

For more information:

- About the DX NetOps Spectrum integration behavior, see [Integrate with DX NetOps Spectrum](#).
- About device life cycle states, see [Manage Device Life Cycles](#).

4. (Optional) [Test the connection by clicking Test](#).

5. Click **Save**.

If the data source is enabled (the status of the data source is "Enabled"), data appears in NetOps Portal after the next synchronization.

**Next step:** (Optional) [Configure the data source for secure communication using Hypertext Transfer Protocol Secure \(HTTPS\)](#).

### **Test the Data Source Connections**

If data source registration does not complete successfully, determine the reason for the failure by testing the data source connection. The test validates version compatibility and verifies that the data source is not registered with a different instance of NetOps Portal.

To confirm the proper registration and connection of a data source, on the **Manage Data Sources** page, select the data source that you want to test, and click **Test**.

If the test fails, verify that the server name or IP address is accurate for the source type.

For more information, see [Data Source Test Fails](#).

### **Edit a Data Source**

The changes that you make to a data source apply to the system after the next synchronization.

#### **Follow these steps:**

1. On the **Manage Data Sources** page, select the data source that you want to edit, and then click **Edit**.  
The **Edit Data Source** dialog opens.
2. Complete the fields on the dialog, and then click **Save**.

Your changes are saved.

### **Delete a Data Source**

Users with the Administrator role and the Delete Data Sources role right can delete, or unregister, data sources. This role right is not granted by default. To delete a data source, assign this role right to the Administrator role.

If you are planning to delete a DX NetOps Network Flow Analysis data source, see [Results of Unregistering](#).

**NOTE**

- Deleting a data source removes only certain information from DX NetOps Performance Management.
- Attempting to register a deleted (unregistered) data source at a later time to the *same NetOps Portal instance* can result in unexpected behavior. You can, however, register deleted data sources to another NetOps Portal instance. Instead, consider removing a data source temporarily by setting its **Status** to "Disabled".
- If your user account is configured to hide, or suppress, all views on dashboards (view suppression is enabled), NetOps Portal suppresses those views that are associated with the deleted data source. As a result, deleting a data source can cause some menus and dashboards to become unavailable. Dashboards, context tabs, and custom menus must contain at least one displayed view. Otherwise, the menu item does not appear in NetOps Portal.  
For more information about view suppression, see [Customize Your User Settings](#).

**Follow these steps:**

1. On the **Manage Data Sources** page, select the data source that you want to delete, and then click **Delete**. The **Delete Data Source** dialog opens.
2. Click **Yes**.

The data source is deleted (unregistered from NetOps Portal).

## Synchronize Data Sources

NetOps Portal periodically sends configuration information and retrieves data by synchronizing, or replicating, group configuration, authentication settings, SNMP profiles, users, and roles with registered data sources.

Registered data sources send inventory to NetOps Portal, and NetOps Portal uses the inventory data to request performance metrics from the data sources. To determine whether items from multiple data sources are the same item, NetOps Portal runs **global synchronization**. During global synchronization, NetOps Portal reconciles devices and interfaces.

Devices are reconciled in the following scenarios:

- The primary IP addresses match.
- The name of a device without an IP address matches another device regardless of whether the other device has an IP address.
- The primary IP address for a new device is in the IP address list of an existing device, and the primary IP address of the existing device is in the IP address list of the new device.

**NOTE**

The primary IP address is the IP address that NetOps Portal uses to monitor a device. When a device is first discovered with the IP ranges discovery profile, NetOps Portal tries to use the IP address that maps to the hostname as the primary IP address.

After the devices are consolidated, interfaces are reconciled when the parent device and the `ifIndex` match. Reconciled items are a single item in NetOps Portal. If the reconciled item has different properties in different data sources, NetOps Portal uses the value from the data source with the highest priority. The data aggregator has the highest priority by default. First, the item type priority determines priority. Then the data source priority determines priority. Last, the data source assigned `SourceID` determines priority.

**Incremental synchronization** occurs regularly every five minutes with all registered data sources, or five minutes after the previous synchronization cycle completes. Incremental synchronization also occurs each time an SNMP profile is created. Only the records that are new or changed as of the last synchronization timestamp are included. The **Last Polled On** column shows the completion time of the last successful synchronization for each data source. If this timestamp shows a time more than 5-10 minutes, synchronization for that data source has not completed successfully. Changes might not be reflected in the system.

A **full synchronization** occurs when you first register a data source to NetOps Portal. It does not recur automatically. During a full synchronization, NetOps Portal receives information about all managed items in that data source.

#### NOTE

After the initial data source registration, a full synchronization is not typically required.

In this article:

- [Synchronize a Data Source](#)
- [View Data Source Synchronization Status](#)
- [View the Data Source Log](#)

## **Synchronize a Data Source**

You can propagate a configuration change immediately by manually synchronizing a data source. For example, if you add a group, you can send the change to the data sources immediately.

#### NOTE

If synchronization is in progress, the process is not interrupted and new changes are not applied until the *next* synchronization cycle. A new synchronization begins immediately after the in-progress synchronization is complete.

#### **Follow these steps:**

1. As a user with the Administrator role, hover over **Administration**, **Data Sources**, and then click **Data Sources**. The **Manage Data Sources** page opens.
2. Select the data source that you want to resynchronize, and then click **Resync**. The **Resync Data Source** dialog opens.
3. (Optional) Select **Perform a Full Resynchronization**.

#### NOTE

**Best Practice:** Have NetOps Portal perform a full synchronization *only* when Broadcom Support requests this type of resynchronization.

**Default:** cleared

4. Click **Resync**.

The data source is synchronized. NetOps Portal collects inventory data and sends configuration changes to the selected data source. Only the records that are new as of the last synchronization timestamp are included.

## **View Data Source Synchronization Status**

You can view synchronization status in the table on the **Manage Data Sources** page. The System Health icon at the top of the page indicates when a synchronization failure occurs. To view more information about the failure, click the icon.

#### TIP

You can also determine if there are data source synchronization issues from the **System Status** page. For more information, see [View System Status](#).

As a user with the Administrator role, review the information in the **Global Synchronization Status** section and the **Data Sources** section.

#### IMPORTANT

If the **Last Run Status** in the **Global Synchronization Status** section displays "Failed", contact Broadcom Support.

The **Data Sources** section displays the status of all registered data sources. The following messages describe possible data source status conditions

- **Awaiting Poll**

The data source has never been contacted, and is waiting for the Device Manager to poll it. The data source is polled quickly unless the Device Manager is busy performing another poll.

- **Awaiting Bind**  
Inventory data has been retrieved from the data source. The data source is waiting for NetOps Portal to transmit configuration information, and to lock corresponding administrative features.
- **Available**  
The data source is available for reporting. Registration has succeeded.
- **Polling**  
The Device Manager is in the process of retrieving inventory data.
- **Registering**  
The Device Manager is in the process of registering the data source.
- **Binding**  
The device manager is locking the users, roles, and groups that are defined in the data source. For some data sources, binding prevents further changes to configuration within the data source.
- **Synchronizing**  
The device manager is in the process of synchronizing with the data source.
- **Polling Failure**  
A failure occurred during polling.
- **Synchronization Failure**  
A failure occurred during synchronization.
- **Registration Failure**  
A failure occurred during registration.
- **Bind Failure**  
A failure occurred during the binding of users, groups, and roles.
- **Unable to Contact**  
NetOps Portal is unable to contact the data source due to communication problems.
- **Version Incompatible**  
The version of the data source is not compatible with NetOps Portal.
- **Requires Upgrade**  
The data source requires a software upgrade. Contact CA Support.
- **Requires Registration**  
The data source is waiting for registration.
- **Requires Migration**  
The data source requires migration and is waiting for the Device Manager.
- **Under Maintenance**  
The data source is under maintenance.
- **Disabled**  
The Administrator has disabled the data source.

### **View the Data Source Log**

Failures and detailed statuses are included in the **Data Source Log**. Only the initial synchronization and any failures that occur during subsequent synchronizations are included in the log. You can investigate suspected errors with the data source synchronization using the data source log. Use this information to find events in the device manager logs. The log shows only events that are related to synchronization for the data source that you selected. For more information, see [Logs](#).

As a user with the Administrator role, on the **Manage Data Sources** page, select the data source for which you want to view the log, and then click **Log**. The **Data Source Log** page opens.

## Update the Configuration of the Alarm Service

Use this troubleshooting topic to address DX NetOps Spectrum (Spectrum) alarm issues.

When you add Spectrum as a data source, the alarm service process in the Event Manager creates and maintains an alarm subscription with the configured Spectrum data source. You can control the behavior of the alarm service with regards to the creation and consumption of Spectrum alarm subscriptions by updating the properties in the alarm service configuration settings.

If you are experiencing issues, you can update the properties in the alarm service configuration settings based on the issue that you are experiencing. For example, if the alarm service process in the Event Manager cannot create an alarm subscription, increase the value for the `RequestTimeoutSec` setting to 300 or 900 . Update the default properties in the alarm service configuration settings using the `alarmservice` web service. For example, to have alarms show up faster in NetOps Portal, update the default properties.

**Prerequisite:** Before updating the default properties, ensure that you have reviewed the symptoms and solutions in [Spectrum Alarms Are Missing from the Alarms View](#).

### Follow these steps:

1. Set up a REST client with a connection to the Event Manager server.
2. Display the configuration settings by issuing a GET request to the `alarmservice` endpoint for the Event Manager web service:

```
http://<PC_host>:<EM_port>/EventManager/webservice/alarmservice/config
```

### Example:

```
http://PC_host:8281/EventManager/webservice/alarmservice/config
```

- **PC\_host**  
Specifies the NetOps Portal host name.
- **EM\_port**  
Specifies the required port number for the Event Manager.

**Default:** 8281

For more information about the NetOps Portal server ports that should be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).

The following response is expected:

```
<AlarmServiceConfig>
  <Enabled>
    true
  </Enabled>
  <Parallelism>
    1
  </Parallelism>
  <RequestTimeoutSec>
    120
  </RequestTimeoutSec>
  <SubscriptionRetryDelaySec>
    300
  </SubscriptionRetryDelaySec>
  <SubscriptionBatchSize>
    1000
  </SubscriptionBatchSize>
  <SubscriptionHeartbeatIntervalSec>
```

```

    30
  </SubscriptionHeartbeatIntervalSec>
  <SubscriptionTimeoutMs>
    300000
  </SubscriptionTimeoutMs>
</AlarmServiceConfig>

```

– **Enabled**

Controls whether the alarm service process in the Event Manager subscribes to Spectrum alarms.

**Default:** true

**Values:**

- **true:** Spectrum alarms display in the NetOps Portal Alarm Console.
- **false:** Spectrum alarms do not display in the NetOps Portal Alarm Console.

– **Parallelism**

Defines the number of active alarm subscriptions that the Event Manager can process concurrently.

**IMPORTANT**

Do not modify this setting.

**Default:** 1

– **RequestTimeoutSec**

Controls the timeout (in seconds) of HTTP requests to Spectrum.

**Values:** 60 – 900

**Default:** 120

– **SubscriptionRetryDelaySec**

Controls the amount of time (in seconds) that the Event Manager waits before it attempts to resubscribe for alarms if the alarm subscription request fails.

**Values:** 30 – 900

**Default:** 300

– **SubscriptionBatchSize**

Controls the number of alarms the Event Manager can receive from Spectrum with a single request. The default is generally sufficient.

**IMPORTANT**

Do not modify this setting.

**Default:** 1000

– **SubscriptionHeartbeatIntervalSec**

Controls the frequency (in seconds) at which the Event Manager requests new alarms from Spectrum.

**Values:** 10 – 30

**Default:** 30

– **SubscriptionTimeoutMs**

Defines the maximum amount of time (in milliseconds) that Spectrum keeps the alarm subscription open while waiting for requests from the Event Manager. You must set this value (when converted to seconds) to a value higher than the `SubscriptionHeartbeatIntervalSec`.

**Default:** 300000 (milliseconds)

3. Update the relevant properties in the response body:

**Example:**

The following example response body increases the timeout (in seconds) of HTTP requests to Spectrum to 300 seconds (`RequestTimeoutSec`):

```

<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<AlarmServiceConfig>
  <Enabled>

```

```

        true
    </Enabled>
    <Parallelism>
        1
    </Parallelism>
    <RequestTimeoutSec>
        300
    </RequestTimeoutSec>
    <SubscriptionRetryDelaySec>
        300
    </SubscriptionRetryDelaySec>
    <SubscriptionBatchSize>
        1000
    </SubscriptionBatchSize>
    <SubscriptionHeartbeatIntervalSec>
        30
    </SubscriptionHeartbeatIntervalSec>
    <SubscriptionTimeoutMs>
        300000
    </SubscriptionTimeoutMs>
</AlarmServiceConfig>

```

4. Update the configuration settings by issuing a PUT request to the `alarmservice` endpoint for the Event Manager web service, with the updated response body

```
http://<PC_host>:<EM_port>/EventManager/webservice/alarmservice/config
```

**Example:**

```
http://PC_host:8281/EventManager/webservice/alarmservice/config
```

- **PC\_host**  
Specifies the NetOps Portal host name.
  - **EM\_port**  
Specifies the required port number for the Event Manager.
- Default:** 8281

**Content-Type:** application/xml

The configuration settings are updated.

## Manage Roles and User Accounts

Manage user access with rights assigned to roles. Then refine access by assigning roles, access permissions, administer groups, and product privileges to each specific user account.

- **Roles**

Roles, or user account roles, control user access to the functionality and dashboard pages when assigned to user accounts. Based on job functions, roles grant administrative access to product configuration using role rights. Role rights also determine menu visibility (availability). You can grant access to selected custom and predefined menus by editing these role rights. Roles allow users access data and product features that they require to perform their duties. Roles restrict access to functionality that they do not require.

Roles are shared with registered data sources. Roles determine what users can access in the data source interface when following a drilldown path to a data source. When you add a user, you select a role for the user account. You can



edit roles to include new role rights, and disable roles to prevent users with those role assignments from using NetOps Portal. Use the set of predefined roles to add new users quickly while determining the required customizations.

For more information, see [Manage User Account Roles](#).

- **User Accounts**

Custom user accounts let operators view the data, menus, and dashboards that they require to perform their daily tasks. Operators with the administrator role rights can create user accounts and can manage existing accounts. Tenant administrators can manage user accounts only for their own tenant.

For more information, see [Manage User Accounts](#).

## **Predefined User Account Roles**

By default, DX NetOps Performance Management includes predefined user account roles. If the predefined user account roles do not fit your requirements, add a custom user account role.

The following lists the predefined user account roles:

- **Administrator**  
Performs all administrative tasks in DX NetOps.
- **Designer**  
Performs limited NetOps Portal administration, such as creating dashboards, menus, and roles.
- **IT Architect**  
Designs infrastructure changes and performs capacity planning and Tier-3 troubleshooting and analysis.
- **IT Director**  
Directs IT strategy, manages IT staff, and reviews high-level performance reports.
- **IT Engineer**  
Receives event notifications, implements technology, and performs Tier-2 troubleshooting and analysis.
- **IT Manager**  
Manages IT staff, ensures operation and maintenance of infrastructure, and reviews performance reports.
- **IT Operator**  
Monitors infrastructure, opens and escalates trouble tickets, and performs Tier-1 troubleshooting.
- **Operations Center Manager**  
Manages Operations Center facilities and staff and reviews performance reports.
- **Tenant Administrator**  
Can perform all administrative tasks for tenants in NetOps Portal. Tenant administrators only have access to the user account roles associated with their tenant.
- **User Administrator**  
Can perform all user administrative tasks, with some limitations. Users assigned to this role cannot delete users that are assigned the Administrator role.  
For more information, see [Manage User Accounts](#).
- **VP of IT**  
Provides long-term strategic direction of IT infrastructure and reviews high-level reports.

## **Predefined User Accounts**

By default, DX NetOps Performance Management includes the following predefined user accounts. They are common to all DX NetOps Performance Management installations, and are not substitutes for custom user accounts. You can only edit these user accounts. They are useful for performing initial setup. For example, you can use them to allocate LDAP access with minimal role rights, or use them as templates for creating custom user accounts.

**IMPORTANT**

- Predefined user accounts are less secure.
- **Best Practice:** For improved security, change the default passwords for these user accounts immediately after installing DX NetOps Performance Management.
- **admin**  
Grants all administrative privileges.  
**Role:** Administrator  
**Special Role Rights:** All (the global administrator or Default Tenant administrator)  
**Permission Groups:** Can view data from all groups  
**Default password:** admin
- **user**  
Specifies typical operator privileges, such as viewing data.  
**Role:** IT Operator  
**Special Role Rights:** None  
**Permission Groups:** Can view data from all groups  
**Default password:** user

## Role Rights

Role rights determine the types of dashboards and views that users can see, the administrative functionality that they can change, and whether they can export data.

Administrators can grant role rights to users by editing their role or by assigning a role to their user account.

For more information, see [Manage User Accounts](#).

The following video examines how to configure the settings for user roles to determine the menu rights, NetOps Portal functionality rights, and data source rights for users assigned to that specific role:

### Administrative Role Rights

The following role rights grant users access to administrative functionality. For increased security, limit the number of users with these role rights:

**NOTE**

Some role rights are limited to the Administrator role. Copying the Administrator role does not give the same role rights to the new role.

- **Administer Advanced Flow Application Mapping**  
Allows the Administrator to manage application mappings.  
For more information, see [Manage Application Mappings](#).
- **Administer Business Hours**  
Allows the Administrator to manage business hours definitions.  
For more information, see [Configure Business Hours Filtering](#).
- **Administer DA Threshold Profile**  
Allows users to edit threshold profiles, to configure profiles that other users created, or to transfer ownership of profiles.
- **Administer Data Sources**  
Allows users to register new data sources, test data source connections, view data source status, view the data source log, change data source parameters, and remove data sources.
- **Administer Groups Owned by You**  
Allows users without full administrative rights to manage a specific branch of the Groups tree. Allows users to manage (create, change, and delete) groups only in the specified branch.  
By default, the Administrator role and the Tenant Administrator role have this role right, allowing administration of All Groups and the Tenant root group, respectively.

Only the Administrator and the owner (creator) of the groups in the administered branch can delete and modify groups in that branch. When an administered group is a child of another group, the administered group is deleted when the parent group is deleted. Administered groups are not deleted when the user account of the owner is deleted.

**NOTE**

Assign this role right to users who require limited, branch-specific administrator rights and should not have full administrator rights to the Groups tree. In some organizations, this user is a power user or a super user.

- **Administer Groups Owned by You and Others**

Allows users without full administrative rights to manage the branch of the Groups tree that they have the rights to administer. Allows users to manage (create, change, and delete) groups in the branch that they or other users created. By default, the Administrator role and the Tenant Administrator role have this role right, allowing administration of All Groups and the Tenant root group, respectively.

- **Administer IP Domains**

Allows users to manage IP domains.

For more information, see [Manage IP Domains](#).

- **Administer Life Cycle**

Allows users to manage device state.

**NOTE**

To avoid conflicts in device life cycle management, if you have integrated with DX NetOps Spectrum (Spectrum), and you have enabled Spectrum to control the life cycle state of devices in NetOps Portal (the **Synchronize device life cycle state from Spectrum** checkbox is selected for that data source), do not grant this role to users. In this case, changes to device states in Spectrum trigger changes to device states in NetOps Portal.

For more information, see [Integrate with DX NetOps Spectrum for Fault Management](#) and [Manage Device Life Cycles](#).

- **Administer Maintenance Indicators**

Allows users to manage maintenance indicators.

- **Administer Menus**

Allows users to manage (create, edit, and delete) menus, and to assign dashboards to menus. To assign menus to user accounts, the Administer Roles role right is required.

- **Administer Roles**

Allows users to manage (create, edit, and delete) user account roles, and to assign new menus to user accounts by editing roles.

- **Administer Shared Dashboards**

Allows users to manage (create, edit, and assign) their own dashboards and the dashboards of other users. Users with this role right can edit an existing dashboard page and can save changes that are visible to other users.

For more information, see [Manage Dashboards](#).

- **Administer SNMP Profiles**

Allows users to manage SNMP profiles.

For more information, see [Manage SNMP Profiles](#).

- **Administer Tenants**

Grants users administrative rights over the tenants that are selected in the user wizard. Users with this role can administer certain tenants, but have limited access to the default tenant. NetOps Portal uses this role only in multi-tenant environments. Tenant administration includes the ability to manage users, menus, dashboards, and views.

For more information, see [Administer Tenants](#).

- **Administer Users**

Allows users to manage (create, edit, and delete) user accounts, and to assign new roles to user accounts.

- **Allow Access to Managing the User Administrator Role**

Allows users to edit the User Administrator role.

- **Allow Access to Managing the Administrator Role**

Allows users to edit the Administrator role.

- **Allow Access to REST Services**  
When combined with the Administer Users right, allows users to access REST services for users. When combined with the Administer Roles right, allows users to access REST services for roles.
- **Allow Alarm Filter Creation**  
Allows users to create alarm filters.  
For more information, see [Alarms View](#).
- **Allow Alarm Filter Management**  
Allows users to manage and assign alarm filters to other users.  
For more information, see [Alarms View](#).
- **Allow Alarm Modification Actions**  
Allows users to acknowledge, clear, or assign a troubleshooter to alarms.  
For more information, see [Alarms View](#).
- **Allow Alarm Triage Actions**  
Allows users to triage alarms.  
For more information, see [Alarms View](#).
- **Copy On-Demand Report Templates from other Users**  
Allows users to copy an on-demand report template that another user or administrator has created and owns, and then edit it. This role right is always assigned together with the Create On-Demand Report Templates and Run On-Demand Report Templates role rights.  
For more information about how to copy an on-demand report template, see [Manage On-Demand Report Templates](#).
- **Create a Dashboard**  
Allows users to create dashboards and populate them with views. Other users cannot see these dashboards. To create dashboards for other users, the Administer Shared Dashboards role right is required. To copy a dashboard from the **My Dashboards** menu to another menu, this role right as well as the Edit Context Pages, Edit Shared Views, Generate URLs from views, and the Save Changes to Shared Views role rights are required.
- **Create DA Threshold Profiles**  
Allows users to define and configure threshold profiles. Users can only edit the profiles that they created.
- **Create Notifications**  
Allows users to configure email notifications using the **Create/Edit Notifications** wizard from the **Administration, Notifications** menu. Notifications are not supported for all data sources.  
**NOTE**  
To create notifications, the user also requires access to the Event Manager data source.
- **Create On-Demand Report Templates**  
Allows users to manage (create, edit, and delete) on-demand report templates. This role right is always assigned together with the Run On-Demand Report Templates role right. Users can save on-demand report templates at the user level, which allows only the user to view the templates.  
For more information about how to create, or generate, an on-demand report template, see [Manage On-Demand Report Templates](#).
- **Delete Data Sources**  
Allows users with the Administrator role to delete (unregister) a data source. By default, no users or roles have this role right. You can assign this role right only to the Administrator role.  
For more information, see [Configure a Data Source](#).
- **Drill from Views into DA Admin Page**  
Allows user to access the data aggregator administrator page directly from a page that is associated with the data aggregator. This role right requires the Administer Data Sources role right. The ability to access the data aggregator administrator page is limited to views for data aggregator devices, interfaces, and components. Selecting a data aggregator interface or component causes the administrator page for the associated parent device to appear when clicking the gear icon and Device Admin.
- **Save On-Demand Report Templates for All Users**

Allows users to save on-demand report templates that are visible to all users. This role right is always assigned together with the Create On-Demand Report Templates and Run On-Demand Report Templates role rights. For more information, see [Manage On-Demand Report Templates](#).

- **Manage OI Connector Status**

Allows users to access the **Manage Connector to DX Operational Intelligence** page. From this page, users can view the OI Connector status and can download and install the OI Connector. By default, the Tenant Administrator user account role has this role right.

For more information, see [Install and Configure Log Analytics for Insights](#).

- **Manage Security Settings**

Allows users to manage the security settings for NetOps Portal.

- **Manage Usage Data**

Allows users to manage the usage data for NetOps Portal.

- **Modify Component Polling**

Allows users to modify component polling.

- **Modify Device Alias**

Allows users to modify the alias property for devices.

For more information, see [Manage Context Pages](#).

- **Modify Interface Alias**

Allows users to modify the alias property for interfaces.

For more information, see [Manage Context Pages](#).

- **Modify Device IP Address**

Allows users to modify the IP address property for devices.

For more information, see [Manage Context Pages](#).

- **Modify Interface Speed Overrides**

Allows users to modify the speed override properties for interfaces.

For more information, see [Manage Context Pages](#).

- **Proxy Users**

Allows users to log in as a selected user to view and verify user account settings. It also allows users to edit the **My Dashboards** menu for another user by proxying that user account.

For more information, see [Organize Dashboards in Menus](#).

- **Save Changes to Shared Views**

Allows users to save edits that they have made to the views on a shared page. Other users who can see these views can see the changes if they are applied as a **Default for All Users**. You can save the changes to the user account so that they persist after logout. To copy a dashboard from the **My Dashboards** menu to another menu, this role right as well as the Create a Dashboard, Edit Context Pages, Edit Shared Views, and the Generate URLs from views role rights are required.

For more information, see [Organize Dashboards in Menus](#) and [Manage Dashboards](#).

- **SNMP Clear Text**

Allows users to troubleshoot SNMP profiles and view security information that is typically masked in clear text.

For more information, see [Manage SNMP Profiles](#).

## **Role Rights for Dashboard and View Access**

The following role rights give users access to reporting functionality. Most user accounts require these rights:

- **Drill into Data Sources**

Allows users to navigate to the data source interface during drill-down to see detailed data from a selected item.

- **Drill into Views**

Allows users to drill into a context view to see detailed data from a selected item. The Edit Context Pages role right requires this role right.

- **Edit Context Pages**

Allows users to manage (add, edit, delete, and reorder) tabs on context pages. This role right requires the Drill into Views role right. To copy a dashboard from the **My Dashboards** menu to another menu, this role right as well as the Create a Dashboard, Edit Shared Views, Generate URLs from views, and the Save Changes to Shared Views role rights are required. By default, only the Designer and Administrator roles have this role right.

For more information, see [Manage Context Pages](#).

- **Edit Shared Views**

Allows users to edit the views on a shared page. Other users can see these views, but cannot see the changes. You can apply changes to the current login session or save them to the current user account. To copy a dashboard from the **My Dashboards** menu to another menu, this role right as well as the Create a Dashboard, Edit Context Pages, Generate URLs from views, and the Save Changes to Shared Views role rights are

For more information, see [Organize Dashboards in Menus](#).

- **Edit Time Zone**

Allows users to edit their own time zone setting for data that are displayed in dashboards.

- **Flow**

Allows users to view the **Flow Dashboard** and the **Flow Statistics** dashboard.

For more information, see [Flow Dashboards](#).

- **Run NetOps Business Reports**

(23.3.2 and higher)

**IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

Allows users to run NetOps business reports. By default, the Administrator role has this role right.

For more information about how to run a NetOps business report, see [Manage On-Demand Report Templates](#).

- **Run On-Demand Report Templates**

Allows users to run on-demand report templates. This role right is always assigned together with the Create On-Demand Report Templates right. However, if the Create On-Demand Report Templates right is removed from the user, the user can still edit and delete their dashboards. With this role right alone (without the Create On-Demand Templates role right), users can run on-demand report templates at the tenant level.

For more information about how to run an on-demand report template, see [Manage On-Demand Report Templates](#).

- **Run Dashboards at Higher Resolution**

Allows users to select higher resolutions when viewing dashboards. By default, no users or roles have this role right. To set and save the resolution to higher values than are typically allowed when reporting for longer time ranges, this role right is required. When users save the higher resolution at the tenant level, it is only visible to users with this role right.

For more information about report resolution settings, see [Set the Resolution for Reported Data](#).

- **Show Unmasked Device Configurations**

Allows users to view sensitive Spectrum device configuration information, such as passwords, in NetOps Portal (the information is unmasked). Sensitive device configuration information only displays for Administrators. For all other users, this information is masked out (gray boxes are displayed) unless the Administrator has delegated user access to the information by assigning this role right to those user accounts.

For more information about how to view Spectrum device configurations, see [Device Configuration View](#).

- **View Conversations**

Allows users to see specific client conversations.

- **View Groups Change Log**

Allows users to see the table view that shows changes that are made to groups.

- **View Hosts**

Allows users to see specific client host information.

- **View Item Display Name or Name Alias**

Allows users to see the display names or the aliases for items.

**NOTE**

Users who are given this role right can select the name that appears in their dashboards and views in the **My Settings, Display Settings** menu item.

- **View Item Name Alias Only**  
Allows users to see only the aliases for items.
- **View Inventory and Search**  
Determines whether users can access the Inventory tab and Search field to find items.
- **View Protocols**  
Allows users to view protocol information, where available.
- **View System Health Dashboards**  
Allows users to view system health dashboards that show information about the performance of DX NetOps Performance Management.
- **View ToS**  
Allows users to view the Type of Service information in applicable views.

### **Role Rights to Export and Print Dashboard Data**

The following role rights allow users to export dashboard data in various formats:

- **Export/Print Dashboard Views** (23.3.4 and higher)  
Allows users to export view data as comma-separated values (CSV) files or to print view data as Portable Document Format (PDF) files.  
For more information, see [Download View Data](#) and [Print View Data](#).
- **Export to CSV** (23.3.3 and lower)  
Allows users to download view data as comma-separated values (CSV) files.  
For more information, see [Download View Data](#).
- **Generate URLs from views**  
Allows users to share views with users who do not have access to a dashboard by generating a URL for the view. To copy a dashboard from the **My Dashboards** menu to another menu, this role right as well as the Create a Dashboard, Edit Context Pages, Edit Shared Views, and the Save Changes to Shared Views role rights are required.  
For more information, see [Generate a URL for a View](#) and [Organize Dashboards in Menus](#).
- **Print a Dashboard**  
Allows users to export dashboard data as a report in PDF or CSV format, and then send it to a printer. This role right also allows users to export dashboard data as CSV files.  
For more information, see [Share Data with Other Users](#).
- **Send Reports by Email**  
Allows users to run reports of dashboard data and send them as attachments to email messages. These users can then send the reports to others. The Send Reports on a Schedule role right requires this role right.  
For more information, see [Manage Scheduled Reports](#) and [Share Data with Other Users](#).
- **Send Reports on a Schedule**  
Allows users to create a schedule to run reports of dashboard data, and send them by email on a recurring basis automatically. This role right requires the Send Reports by Email role right.  
For more information, see [Share Data with Other Users](#).
- **Send Reports to Archive**  
Allows users to archive, or save, dashboard data as reports from scheduled reports. The reports are saved to the report repository.  
For more information, see [Share Data with Other Users](#).



## Data Source Role Rights

Each registered data source has its own set of roles with unique rights that let users access features and data within that interface. Administrators can assign rights for a role within that data source. These data source rights apply when users follow a drilldown path from a data view to that particular data source. However, any rights that are granted in this manner are specific to a data source instance. For example, if more than one CA Application Delivery Analysis (ADA) data source is registered, the rights for each management console are managed separately.

For example, an administrator can grant the right to generate reports in a DX NetOps Network Flow Analysis (NFA) data source, but withhold the right to edit dashboards in NetOps Portal. The individual data source Administrator Guides provide detailed information about the application of role rights.

Individual data source administrators can create user accounts and grant users role rights to access features within that data source. After registration, those rights are synchronized with NetOps Portal. The rights are then displayed in NetOps Portal when data source administrators edit a role.

### NOTE

Role rights to individual data sources are distinct from rights to access DX NetOps Performance Management features, but they frequently have the same names.

In this article:

- [Data Aggregator Role Rights](#)
- [DX NetOps Network Flow Analysis Role Rights](#)
- [CA Application Delivery Analysis Role Rights](#)
- [CA Unified Communications Monitor Role Rights](#)

### Data Aggregator Role Rights

- **Drill from Views into DA Admin Page**  
Drill down from data aggregator views to the **Monitored Devices Admin** page to troubleshoot a view that does not display data.
- **Administer Tenants**  
Administer tenants, including user accounts, discover and delete devices for the data aggregator.
- **Administer DA Threshold Profiles**  
Administer threshold profiles, including creating threshold profiles, editing any threshold profiles, and changing the ownership of all threshold profiles.
- **Create DA Threshold Profiles**  
Create threshold profiles, where you can create and manage event profiles. Event profiles contain event rules, and are associated with groups. You can create reports on the events that are generated with these profiles. This role allows you to create threshold profiles, edit your own profiles, and view all threshold profiles.

### DX NetOps Network Flow Analysis Role Rights

The following role rights are applicable to the NFA console:

- **View ToS**  
View Type of Service data.
- **Manage Reports**  
Create, modify, delete, and execute reports.
- **Run Reports**  
Execute defined reports.
- **View Conversations**  
View conversation data.
- **View Hosts**



View host data.

- **View Protocols**  
View protocol data.

### **CA Application Delivery Analysis Role Rights**

The following role rights are applicable to the ADA management console:

- **Engineering**  
Navigate the Engineering section and create Engineering reports.
- **Operations**  
Navigate to the Operations section and create Operations reports.
- **Management**  
Navigate the Management section and create Management reports.
- **Incidents**  
Navigate the Incidents section and view Incidents reports.
- **Investigations**  
Launch Investigations and drill into data from Investigations.

Role rights do not give a ADA user:

- Permission to access the Administration page of the ADA management console.  
To give a user access to the Administration page, give the user the Administrator or Power User product privilege on the ADA data source.
- Access to actual report data in the ADA management console.  
To enable a user to see report data, assign the appropriate groups to the user.

### **CA Unified Communications Monitor Role Rights**

The following role rights are applicable to the CA Unified Communications Monitor management console:

- **Call Details**  
Export call details to a CSV file.
- **Call Performance**  
Access Call Performance reports.
- **Call Quality and Volume**  
Access Call Quality and Volume reports.
- **Call Watch**  
Access Call Watch reports.
- **Call Watch Setup**  
Set up and launch a Call Watch on a selected phone.
- **Collector Incidents**  
Access Collector Incident reports.
- **Incidents**  
Access Incident reports.
- **Investigations**  
Access Investigation reports.
- **Launch Investigation**  
Launch an investigation and view the resulting data.
- **Phone Details**  
Access Phone Details reports.
- **Quality**

- Access Quality reports.
- **Trunk Groups**  
Access Trunk Group reports.
- **Voice Interface**  
Access Voice Interface reports.
- **Midstream Devices**  
Access midstream device and midstream legs reports.

## Manage User Account Roles

User account roles control user access in NetOps Portal, such as dashboards. Based on job functions, user account roles grant administrative access to NetOps Portal configuration. User account roles control user access to those NetOps Portal functions they require to perform their duties, and restrict access to functionality that they do not require.

### View a List of User Account Roles

You can view a list of the user account roles on the **Manage Roles** page. This page includes summary information about the roles and a list of the custom user account roles that you have created. To view this page, hover over **Administration**, **User Settings**, and then click **Roles**.

### View the User Accounts Assigned to a Role

You can access a list of the user accounts assigned to a user account role from the **Manage Roles** page. To view this list, on the **Manage Roles** page, select the role for which you want to view the assigned users, and then click **Users**. The **Users** page appears.

From the **Users** page you can do the following:

- Return to the list of user account roles (the **Manage Roles** page). Click **Roles**.
- Create a user account assigned to the role. Click **New**.  
For more information about how to create a user account, see [Manage User Accounts](#).

### View the Role Rights Assigned to a User Account Role

You can view a list of the role rights that are assigned to a role from the **Add Role** or **Edit Role** dialogs. To view these dialogs, from the **Manage Roles** page, click **New**, or select the user account role that you want to view role rights, and then click **Edit**.

For more information, see [Role Rights](#) and [Data Source Role Rights](#).

### Add or Edit a User Account Role

You can add custom user account roles. Create the user account roles that each unique product operator requires to perform job responsibilities.

Custom user account roles work best within a system of custom groups. You can precisely grant access to dashboards and DX NetOps Performance Management features while restricting access to sensitive data using custom groups. The groups that you create to organize data can serve as permission groups when you set up user account permissions.

Global administrators and users with the required role rights can edit predefined and custom user account roles.

#### **Follow these steps:**

1. From the **Add Role** or **Edit Role** dialogs, enter or edit the required information, and make selections in the provided fields:
  - **Name**  
Specifies the name for the role.

**Maximum Characters:** 45– **Description**

Specifies a description for the role.

– **Role Status**

Specify whether the user account role is active. When enabled, user accounts with this role have the access granted by its role rights. A role can be disabled for security purposes. When a role is disabled, users who are assigned that role are no longer allowed to log in.

**Options:** Enabled or Disabled

**Default:** Enabled

2. To select menus that are available for this role, complete the following:
  - a. Select **Menu Set** (selected by default), and then click **Edit**.  
The **Edit Menu Set** dialog opens.
  - b. Move the menus to which you want this user account role to have access to from the **Available Menus** list to the **Selected Menus** list.
  - c. In the **Selected Menus** list, determine the order that you want the menus to display on the **Performance** tab by ordering them, and then click **OK**.
3. To select role rights for this user account role for NetOps Portal, complete the following:
  - a. Select **DX NetOps**, and then click **Edit**.  
The **Edit Role Rights** dialog opens.
  - b. Move the role rights that you want this user account role to have from the **Available Rights** list to the **Selected Rights** list.
  - c. In the **Selected Rights** menu, determine the priority of the role rights in cases where rights overlap by ordering them, and then click **OK**.

To select role rights for this user account role for the other registered data sources, repeat this step for those data sources.
4. Click **Save** or **Save and Add Another**.

The user account role is added or edited.

**NOTE**

This user account role is not active until you assign it to a user account.

**Next step:** [Assign this user account role to user accounts.](#)

**Delete a User Account Role**

You can delete the custom user account roles that you have created. You cannot delete or disable the Administrator user account role.

**Prerequisite:** Before you delete a custom user account role, verify whether a user account is using the user account role from the **Users** column. Manage the user accounts that are associated with the user account role that you want to delete to remove the associations with the user account role.

**Product Privilege**

The *product privilege* is a type of permission set associated with a user account that grants user access to functionality in selected data sources. It does not apply to DX NetOps Performance Management functionality.

Individual data sources allocate access to the data source. Administrators create user accounts with user access to functionality in selected registered data sources by applying the *product privilege* setting to the data source. For example, a user account that does not have access to administration can have an Administrator product privilege to a specific instance of DX NetOps Network Flow Analysis (NFA). That user has full administrative privileges to that data source when following a drill-down path for NFA managed items.

**NOTE**

To follow a drill-down path to a data source, a user requires the appropriate role right and a product privilege for that data source.

The following are types of privileges that might be available in the data sources and synchronized to NetOps Portal:

- **Administrator**  
Performs all functions, including creating and editing SNMP profiles and other configuration.
- **Power User**  
Creates menus and dashboards, and edits and create roles.
- **User**  
Views menus and dashboards designated by an administrator or power user.
- **None**  
Does not have access to a data source. This setting prevents the user from following a drilldown path from a view to the data source's user interface. By default, all users have this product privilege setting for all data sources. A user can be denied access to a particular data source while having access to others.

The **admin** predefined user account has administrative privileges to the registered data sources. The **user** predefined user account has limited (user-level) privileges for registered data sources.

For more information about how to grant product privileges to registered data sources to user accounts, see [Manage User Accounts](#).

## Data Source Product Privileges

Each registered data source has its own product privilege with unique privileges within that interface. Administrators can assign product privileges to data sources. These privileges apply when users follow a drill-down path from a data view to that particular data source, and are specific to that data source instance. For example, if you have more than one CA Application Delivery Analysis data source, you manage the product privileges for each management console separately.

The default Administrator account (admin) is locked to prevent changes to product privileges. This account is required to have Administrator privileges for all registered data sources. Selecting a group of accounts that includes the admin account prevents you from editing the product privileges for any of the selected accounts.

In this topic:

### CA Application Delivery Analysis Product Privileges

The following list summarizes the product privileges that are applicable to the CA Application Delivery Analysis management console:

A user must have CA Application Delivery Analysis data source product privileges to log in to the management console in CA Application Delivery Analysis.

Product privileges:

- **User**  
Gives access to all pages of the management console, except the **Administration** page.
- **Administrator**  
Gives access to all pages of the management console, including the **Administration** page.
- **Power User**  
Gives user-level product privilege, and **Show Me** menu access to the SNMP profiles, network devices, and device groups on the **Administration** page.

**NOTE**

If a user cannot log in to the management console, verify that the user has the CA Application Delivery Analysis data source product privilege.

## **DX NetOps Network Flow Analysis Product Privileges**

A user must have DX NetOps Network Flow Analysis (NFA) data source product privileges to log in to the NFA console. Product privileges also determine whether a user can access the **Administration** page, and can perform certain functions.

Product privileges:

- **Administrator**  
Gives access to the **Administration** page in the NFA console, and to all functions. Functions include creating and managing user accounts, roles, groups, SNMP profiles, and scheduling for reports.
- **Power User**  
Gives user-level access and any additional abilities that the role setting grants. In DX NetOps Network Flow Analysis, this privilege is equivalent to the Administrator privilege.
- **User**  
Gives access to Top Interfaces reports and Interface Utilization reports on the **Enterprise Overview** page. A User with the appropriate Permission Group settings also has access to the following reports:
  - Top Hosts and Top Protocols reports on the **Enterprise Overview** page, if the user also has access to All Groups.
  - **Interfaces** page reports for the interfaces that are accessible to the user.
  - Existing reports on the **Custom Reporting**, **Flow Forensics**, and **Analysis** pages.
  - Menus that an administrator has assigned to the User role.
 The Role and Permission Group settings determine whether the user can also run existing reports, create reports, and manage reports. To create reports, a user must have access to All Groups.
- **None**  
Has no access to a data source. The user who has this product privilege cannot log in to the NFA console or drill down from a NetOps Portal view to the NFA console. By default, all users have this product privilege setting for all data sources.

For more information about user accounts and product privileges, see the [DX NetOps Network Flow Analysis documentation](#).

### **NOTE**

The same user account can have different privileges for different data sources.

## **CA Unified Communications Monitor Product Privileges**

The following list summarizes the product privileges that are applicable to the CA Unified Communications Monitor management console:

- **Administrator**  
Gives access to all functions, including administrative tasks, such as creating and editing:
  - Locations
  - Media devices
  - Thresholds
  - Call Watch definitions
  - Incident responses
  - Roles
  - User accounts
- **User**  
Gives access to report pages, and the ability to perform basic functions that the administrator selects. The User permission does not give access to administrative functions.

## Manage Product Access

Grant access to product features and data as you create each user account. Use the following method to verify and modify the role rights for a specific user.

### Follow these steps:

1. Log in as a user with the required administrative role rights.
2. Select **Administration**, **User Settings**, and click **Users**.  
The Manage Users page opens.
3. Select the user account that you want to edit.

#### NOTE

: The rights and privileges that are assigned to the predefined administrator account, 'admin', cannot be modified. This user account must have administrator access to all registered data sources.

The Create New User wizard opens.

4. Click **Product Privileges**.

#### NOTE

All registered data sources appear on the Product Privileges page.

5. Click the values that are shown in the Product Privileges column to enable drop-down lists.
6. Select one of the following product privileges from the drop-down lists:
  - **Administrator**  
Performs all functions, including creating and editing groups, menus, dashboards, roles, and user accounts.
  - **Power User**  
Creates menus and dashboards, and edits and creates roles.
  - **User**  
Views menus and dashboards that are designated by an administrator or power user.
  - **None**  
Has no access to a data source. This setting prevents the user from following a drilldown path from a view to the data source's user interface. By default, all users have this product privilege setting for all data sources.
7. Click **Save**.  
The changes to product privileges are saved to the selected user account.

## Manage User Accounts

You can manage user accounts by adding, cloning, editing, proxying existing, disabling, and deleting user accounts. You can also manage the user accounts assigned to a specific role.

You can manage user accounts using NetOps Portal and using the `users` endpoint for the NetOps Portal REST web service. This article describes how to manager user accounts using NetOps Portal.

For more information about how to manage user accounts using the `users` endpoint, see [Users Web Service](#).

In this article:

- [View User Account Settings](#)
- [Create and Configure User Account Workflow](#)
- [Add a User Account](#)
- [Edit a User Account](#)
- [Assign a Permission Group to a User Account](#)
- [Prevent a User Account from Accessing NetOps Portal](#)
- [Delete a User Account](#)

## View User Account Settings

You can view high-level settings for a user account, such as the role assigned to the user account, on the **Manage Users** page. In a multi-tenant environment, the global administrator sees a list of user accounts that are not associated with a tenant. Tenant administrators only see user accounts for their tenant.

To view the **Manage Users** page, hover over **Administration**, **User Settings**, and then click **Users**.

## Create and Configure User Account Workflow

Use the following process to create and configure a user account:

1. Confirm that the appropriate groups exist, or [create them if necessary](#). User account parameters include the groups that the user can view and manage.
  2. Confirm that the appropriate user account roles exist, or [create them if necessary](#).
  3. [Add a user account](#), which includes:
    - Assigning a role to it.  
Administrators assign *roles* to user accounts and control user account access to NetOps Portal functionality and dashboards. Based on the user's job functions, the role grants administrative access to NetOps Portal configuration by way of *role rights*. NetOps Portal provides predefined roles, with different role rights. Users with the Administer Users and Administer Roles role rights can create additional roles and assign them to user accounts.  
For more information about the predefined roles that are available, see [Role Rights](#).
    - Assigning permission groups to it.  
Administrators organize performance data and allocate operator access to that data using the predefined groups, or system groups. However, a more secure and better managed system is based on custom groups that are assigned to user accounts as permissions. Permission groups comprise the scope of the managed items (the data) that each user can view and monitor. Administrators can create custom groups of managed items—such as applications, servers, networks, routers, and interfaces—to reflect each user's area of responsibility. When the Administrator assigns a group to a user account as permissions, it is referred to as a permission group.
- NOTE**  
By default, new user accounts do not have access to groups. The **admin** and **user** predefined user accounts have access to all groups. For new user accounts, limit the groups that these users can see based on their responsibilities. The dashboards for new user accounts populate only after you assign one or more permission groups to the user account.
- Assigning group ownership to it. This allows the user to create and modify groups in a branch of the Groups tree.
- NOTE**  
You can assign group ownership only to user accounts that have the Administer Groups role right.
- Granting product privileges to registered data sources.  
The *product privilege* is a type of permission set associated with a user account that grants user access to functionality in selected data sources. It does not apply to DX NetOps Performance Management functionality. For more information about the product privileges that are available for registered data sources, see [Product Privilege](#).
- The following video examines how to create a user in NetOps Portal by assigning a role, defining access permissions, assigning groups, defining group administration rights, and defining product privileges. Then, once a user is configured, proxy the user, viewing NetOps Portal with their permissions to confirm the user is configured properly:
4. Test the user account by [temporarily proxying it](#).

## Example: Adding a Data Center Manager User Account

To understand user account parameters, consider the following example. A data center manager at your company is responsible for data centers, staff, and infrastructure in the Southwest region.

The Administrator completes the following steps in NetOps Portal:

1. Creates a group named Southwest.
2. Adds managed items to this group, including the routers, switches, applications, and servers that comprise the Southwest region.
3. Creates a custom role named Data Center Manager that includes the functionality and menus that the data center manager requires.

**NOTE**

The data center manager is not a network engineer. This user does not drill in to detailed data in the data sources. To manage the team, the data center manager requires the ability to create dashboards and assign them to the roles of the team members. The data center manager requires only the menus that contain high-level management and operations dashboards.

4. Adds the user account for the data center manager, selects the Data Center Manager role, and selects the permission groups containing the managed items within the Southwest group to which the data center manager requires monitoring access.

### **Add a User Account**

Add a user account for each user operating NetOps Portal. For security purposes, do not share user accounts.

To create user accounts quickly, clone an existing user. You use existing accounts, such as the **user** predefined user account, as templates for the new user accounts.

Administrators can also create user account templates that are based on job function. You can clone these templates to create individual accounts more easily. For security reasons and to prevent unintentional access to NetOps Portal, disable the **Enable user account** setting for templates. Instead, when creating user account by cloning templates, enable access to NetOps Portal as needed.

**Best Practice:** Check the permissions that are associated with each user account periodically to ensure that all items are being monitored. Each time that you register a new data source, new system groups are added to the Groups tree. Often, NetOps Portal operators do not monitor these new groups until you explicitly add them to user accounts.

**Prerequisites:** Before you create or edit user accounts, place managed items in custom groups and create the roles that you require. Typically, the predefined roles provide starting points for customization. Groups and roles are among the required user account parameters.

### **Follow these steps:**

1. On the **Manage Users** page, do one of the following:
  - To create a user account, click **New**.  
The **Create New User** dialog opens.
  - To create a user account based on an existing user account, select the user account from which you want to base the new user account, and then click **Clone**.  
The **Clone User** dialog opens.
2. Enter the following required information, make selections in the provided fields, and then advance the wizard:
  - **Name**
  - **Description**
  - **Preferred Language**  
Specify a language for the NetOps Portal user interface. NetOps Portal displays the selected language regardless of the language selected for the operating system or for the browser language.
  - **Email Address**
  - **Authentication Type**
  - **Password and Confirm Password**  
By default, user passwords must meet the following requirements:



**NOTE**

These requirements apply only to the passwords that NetOps Portal manages. Passwords managed by another provider are subject to that provider's requirements.

- Be different from the username
- Minimum length of 8 characters
- Maximum length of 30 characters
- Contain at least 3 of the following types of characters:
  - Special characters (ASCII only)
  - Uppercase letters (A-Z)
  - Lowercase letters (a-z)
  - Numbers (0-9)

For more information about how to view these settings using NetOps Portal, see [Configure the Password Security Settings](#).

– **Time Zone**

Specify the time zone that is applied to all dashboards viewed while the user is logged in to NetOps Portal. Typically, the time zone matches the locale of the computer that the operator uses to access NetOps Portal.  
**Default:** UTC (Coordinated Universal Time)

– **Role**

(Users with the Administer Users and Administer Roles role rights) Define the user account roles that you want to assign to this user account.

– **Account Status**

Defines the user's ability to access NetOps Portal.

**Options:**

- **Enabled:** Allows the user to access NetOps Portal.
- **Disabled:** Prevents the user from accessing NetOps Portal.

**Default:** Enabled

The **Access Permissions** step appears in the wizard.

3. Specify the access permissions settings, and then advance the wizard:

- By default, new user accounts do not have group assignment. Grant view and monitor access to managed items by assigning one or more groups or subgroups from the available list to the selected list (assigns the group to the user account as a permission). Map the user profile to the permission groups that makes sense for their initial view.

**IMPORTANT**

**Best Practice:** Do not assign the Collections group as part of a user's permission groups. Do not use this group for reporting.

- **Enable "My Custom Groups" Functionality** (Optional)

Allows the user to organize managed items into custom groups for troubleshooting and analysis.

**NOTE**

The custom groups that the user creates are only available to this user on the **My Custom Groups** page. They do not appear in the main Groups tree.

For more information about how users can access their custom groups, see [Manage Groups](#).

- **Default Group**

A default group is selected for the user automatically. When the user logs in, data from the default group appears in dashboards by default. The group selections on the **Access Permissions** dialog filter the Available Groups tree. This filtering prevents users from having administrative rights to a part of the tree from which they are prohibited.

**NOTE**

To prevent performance issues and to improve overall dashboard load times, do not use the All Groups as the default group.

The **Administer Groups** step appears in the wizard.

4. Move the desired group that you want this user to administer (user with the **Administer Groups** role right) from the available groups list to the selected groups list, and then advance the wizard.

**NOTE**

- Lock icons appear next to read-only groups, which users cannot administer.
- Administrators can administer all groups. This step is not required if this user account is an Administrator.
- Users with the **Administer Groups** role right can create groups *under* the selected group or subgroups. They can then modify or delete only those administered groups. Users cannot modify or delete groups that another user owns.  
For more information about this role right, see [Role Rights](#).
- Users with the **Administer Groups Owned by You** role right can manage a specific branch of the Groups tree.

The **Product Privileges** step appears in the wizard.

- For each data source, specify the access to functionality in the data source that you want to grant this user access, and then click **Save**:
  - **Administrator**  
Allow the user to perform all functions, including creating and editing groups, menus, dashboards, roles, and user accounts.
  - **Power User**  
Allow the user to create menus and dashboards and also edit and create roles.
  - **User**  
Allow the user to view menus and dashboards that are designated by an administrator or power user.
  - **None**  
Restrict user access to the data source. Prevent the user from following a drill-down path from a view in NetOps Portal to the data source user interface. By default, users have this product privilege setting for all data sources. The same user account can have different privileges for different data sources.

For more information, see [Product Privilege](#) and [Data Source Product Privileges](#).

The user account is added.

**Edit a User Account**

Modify a user account when its job responsibilities change, or when the permission groups or roles are created. Edit a user account to assign any new roles that you create.

**Follow these steps:**

- On the **Manage Users** page, select the user account that you want to change, and then click **Edit**. The **Edit User** dialog opens. The **Account Details** step appears in the wizard.

**NOTE**

You cannot modify the rights and privileges that are assigned to the **admin** predefined administrator user account. This user account must have administrator access to all registered data sources. If you select a group of accounts that includes this user account, you cannot modify any of the selected accounts.

- Edit the required information, make selections in the provided fields, and then advance the wizard.

**IMPORTANT**

Do not remove administrative role rights from the **admin** account. Administrative access to the console is required.

The **Access Permissions** step appears in the wizard.

- Specify the access permissions settings, and then advance the wizard.  
The **Administer Groups** step appears in the wizard.
- Move the desired group that you want the user to administer (user with the **Administer Groups** role right) from the available groups list to the selected groups list, and then advance the wizard.  
The **Product Privileges** step appears in the wizard.
- For each data source type, specify a product privileges, and then click **Save**.

The changes to the user account are saved.

### **Assign a Permission Group to a User Account**

Administrators can assign multiple permission groups to a user account when adding the user account. For example, you can assign the `North American Core Routers` and `North American Critical Applications` permission groups to the same user account.

To plan a strategy for creating a grouping and role structure, consult with a Broadcom technical representative. The best configuration meets your current requirements and is flexible enough to accommodate changes to your system.

### **Prevent a User Account from Accessing NetOps Portal**

User account status is "Enabled" or "Disabled". To prevent a user from accessing NetOps Portal, disable the account (change the status to "Disabled").

### **Delete a User Account**

Users with the Administer Users role right can delete user accounts. Users with the Administer Users role right *and* that have the Administrator DX NetOps Privilege can delete users that are assigned the Administrator role. When deleting a single user, you can assign their items (notifications, on-demand reports, (23.3.2 and higher) NetOps business reports, and scheduled reports) to another owner.

#### **NOTE**

Deleting a user account and assigning ownership of their scheduled reports to another user deletes the user-owned dashboards that are displayed in the **My Dashboards** menu. You can recreate these scheduled reports for the newly-assigned user.

**Prerequisite:** (If you plan to delete a single user account and assign ownership of their on-demand reports to another user) Ensure that the **Apply Changes** field for these reports is set to **For All Tenant Users**. For more information, see [Manage On-Demand Reports](#).

#### **Follow these steps:**

1. On the **Manage Users** page, select the user accounts that you want to delete, and then click **Delete**.  
The **Delete User** dialog opens.  
If you are deleting a single user account, the number of user-owned items that will be deleted with the user account are listed. If you are deleting more than one user account and you want to assign them to another user, click **No** to return to the **Manage Users** page, and delete each user account individually.
2. (If you are deleting a single user account) Choose one of the following options for the user-owned items from the **Assign Items to User** drop-down:
  - Assign the items to another user account by selecting the user. The user accounts listed are those that have sufficient role rights to manage notifications (the Create Notifications role right), to manage on-demand reports (the Create On-Demand Report Templates role right), and to manage scheduled reports (the Send Reports on a Schedule role right).
  - Delete user-owned items with the user account by selecting **No user selected**.  
The default is to delete the items with the user account (the **No user selected** option is selected).
3. Confirm the deletion of the user account by clicking **Yes**.  
The user account is deleted.
4. (If you deleted a single user account and you assigned ownership of the items to another user account) Complete the following:
  - Ensure that the newly-assigned user has access to the groups to which the items reference.  
For more information:

- About how to update notifications, see [Configure Notifications](#).
  - About how to update on-demand reports, see [Manage On-Demand Reports](#).
  - About how to update scheduled reports, see [Manage Scheduled Reports](#).
  - About how to update user account access permissions, see [Manage User Accounts](#).
- Recreate the scheduled reports for the newly-assigned user.  
For more information, see [Manage Scheduled Reports](#).

## Enable or Disable Users

If a user is disabled, for example, after multiple failed login attempts, you can enable them.

You can enable or disable a user's account status settings using the Single Sign-On Configuration (SsoConfig) tool or using NetOps Portal. This article details how to enable a disabled user using the SsoConfig.

For more information about how to enable or disable a user using NetOps Portal, see [Manage User Accounts](#).

### Follow these steps:

1. Log in to the server where NetOps Portal or a supported data source is installed (as root or with the 'sudo' command).
2. Launch the tool by running the `./SsoConfig` command in the `<installation_directory>/PerformanceCenter` directory.
  - **installation\_directory**  
The default installation directory for NetOps Portal.  
**Default:** `/opt/CA`

You are prompted to select an option. The available options correspond to data sources that are running on the local server.

Use the following commands as needed while you are selecting settings:

  - `q` (quit)
  - `b` (go back to the previous menu)
  - `u` (update, which overwrites the existing value for the property)
  - `r` (reset)
3. At the **SSO Configuration** prompt, enter the number for the **Dx NetOps** option to configure NetOps Portal security settings.
4. At the **SSO Configuration/DX NetOps** prompt, enter the number for the **Enable or Disable a user account** option.
5. At the **Enter user name to enable or disable** prompt, enter the username of the user that you want to enable or disable account status settings. For example, **admin**.  
If the user is enabled, the following message appears:  
`User account <username> is currently enabled.`  
If the user is disabled, the following message appears:  
`User account <username> is currently disabled.`
6. At the prompt, do one of the following steps:
  - (If the user is enabled) Disable the user by entering the number for the **Disable user <username>** option.  
The following message appears:  
`User account <username> has been disabled.`
  - (If the user is disabled) Enable the user by entering the number for the **Enable user <username>** option.  
The following message appears:  
`User account <username> has been enabled.`
7. Enter **q**.

SsoConfig closes. The user is enabled or disabled.

## Proxy Users and Tenants

Allow administrators to see and interact with NetOps Portal as another user or tenant by proxying that user or tenant.

Proxying a user or tenant is useful for the following tasks:

- Verifying permission groups and role rights.
- Configuring and testing menus and dashboards for a user.
- Configuring a tenant environment.

Proxying a user proxies the menus, dashboards, permission groups, and user settings. Proxying a tenant shows the users, items, and settings of that tenant.

### NOTE

Proxying a user or tenant does not proxy the role rights and product privileges to other data sources. The proxied user account only persists within NetOps Portal. If you follow a drill-down path to a data source, the user account defaults to the original user account settings.

The following video examines how to create and proxy a user in NetOps Portal:

In this article:

- [Proxy a Tenant](#)
- [Proxy a User](#)

### Proxy a Tenant

**Follow these steps:**

1. Hover over **Administration**, **Group Settings**, and then click **Tenants**.  
The **Manage Tenants** page appears.
2. Select the tenant that you want to proxy, and then click **Administer**.  
The **Manage Users** page appears.
3. Select the user that you want to proxy, and then click **Proxy**.

A banner at the top of NetOps Portal displays the tenant (the **Administering Tenant** indicator) and the user that you are proxying (the **Proxy User** indicator).

To stop proxying the tenant and user, click the **X** next to the **Administering Tenant** or **Proxy User** indicator.

### Proxy a User

#### TIP

To proxy a user account from another tenant, first proxy that tenant.

**Follow these steps:**

1. Hover over **Administration**, **User Settings**, and then click **Users**.  
The **Manage Users** page appears.
2. Select the user account that you want to proxy, and then click **Proxy**.

A banner at the top of NetOps Portal displays the user that you are proxying (the **Proxy User** indicator).

To stop proxying the user, click the **X** next to the **Proxy User** indicator.

## Multi-tenancy

You can create separate monitoring environments that you administer from a single instance of NetOps Portal by adding custom tenants.

With multi-tenancy, you can monitor discrete customer environments separately and securely. A tenant represents a customer environment that a managed service provider (MSP) administers. Each tenant environment is independent and effectively functions as a separate instance of NetOps Portal. Each instance can contain multiple users and roles that are not shared among tenants.

### **Managed Items and Tenant Definition**

By default, managed items and their data are associated with the Default Tenant. The basic tenant definition contains parameters to identify the MSP customer. Each tenant is associated with one or more IP domains. All items in the associated IP domains are associated with that tenant. The global administrator or the tenant administrator sets up the monitoring environment for the tenant.

The following configuration items are scoped to the tenant:

- [SNMP profiles](#)
- [User accounts](#)
- [Roles](#)
- [Custom groups](#)
- [Custom dashboards](#)
- [Custom menus](#)
- [Discovery profiles](#)
- [Threshold profiles](#)

### **Administrator Roles for Multi-tenancy**

When multi-tenancy is deployed, NetOps Portal supports the following roles:

- **Global Administrator**  
The Default Tenant administrator, usually representing an MSP. NetOps Portal settings and data are not shared among tenants, but the Default Tenant Administrator can access and modify all the settings. This user must have the predefined Administrator role.
- **Tenant Administrator**  
A limited administrator who is associated with a single tenant. This operator cannot access shared infrastructure or configuration belonging to the host (usually, the MSP). Tenant user accounts can include one or more of these administrator accounts. When you create a tenant, the user interface prompts you to create a tenant administrator.

## **Multi-tenancy Deployment Considerations**

Consider the following factors when you create tenants or IP domain definitions:

- The size, scope, and organization of your monitoring environment
- The data sources that you plan to install and register
- Data source support for the multi-tenancy features

These factors work together to determine your strategy. For example, some data sources do not detect IP domains that are created within tenants.

#### **IMPORTANT**

Do not change the tenant or IP domain definitions after data collection has started. Data sources collect and aggregate data that retains a database association with the *original* IP domain or tenant.

### **Domain Monitoring Considerations**

The IP domains feature supports environments where multiple enterprise systems must be monitored separately. For example, a managed services provider wants to monitor the systems and networks of different customers separately. The

MSP administrator creates a tenant in NetOps Portal for each customer enterprise. The data and the configuration for each tenant are hidden from all other tenant users.

However, in other situations, you can deploy multiple IP domains in NetOps Portal without multi-tenancy. In other words, some deployment models consist of *multiple IP domains within the Default Tenant*.

The IP domain lets you control data collection parameters. Use custom IP domains to determine which collection devices monitor the managed items in your infrastructure. Each collection device, such as a data collector or CA Unified Communications Monitor Collector, operates within a single IP domain.

The following list provides some examples of environments where you can deploy multiple IP domains within the Default Tenant:

- A deployment that includes a CA Application Delivery Analysis data source.  
CA Application Delivery Analysis monitors IP domains without a concept of tenants. The IP domains that you create within custom tenants are not detected. When the data aggregator or DX NetOps Network Flow Analysis monitors items in those domains, they appear as duplicates in CA Application Delivery Analysis. The duplicate data is not aggregated.
- A large deployment that requires load balancing.  
For example, your enterprise includes ten routers with many interfaces, IP SLA testing, and QoS policies in place. Such a deployment would have a polling load similar to an environment with hundreds of servers being monitored for CPU and memory statistics only.  
To monitor the busy routers, you can create an IP domain and can deploy a powerful system for the data collector within that domain. And you can monitor the servers in another IP domain, using a less powerful system for the data collector. By running discoveries in the appropriate IP domains, you can determine the devices that each data collector is polling.
- A method to minimize the potential network impact of bulk statistics collection.  
For example, you can deploy a data collector close to the devices that it is monitoring. Data collectors can process a massive amount of bulk statistics and can reduce them to a much smaller set of monitored metrics, which are sent to the data aggregator. As a result, less data passes across the network between the two components.
- Isolation of potentially sensitive SNMP traffic to a specific area, such as a DMZ.  
For example, security policies do not let SNMP traffic travel across the router that limits an area of network. One option is to deploy a data collector behind the router. A path to return the processed metrics back to the data aggregator must be open.  
The metric data traveling between the components is not encrypted. However, it is packaged and compressed in a way that makes it less "sniffable." As a result, the data is more secure than raw SNMP flows. To accomplish this setup, create an IP domain for the DMZ and deploy a data collector within that IP domain.

## **Deployment Process**

The global administrator who is associated with the default Tenant space performs the initial steps to create a multi-tenant environment.

Use the following process to set up a multi-tenant deployment:

1. Collect data about MSP customer virtual and physical systems.
2. Make a list of IP domains and SNMP versions, communities, or passwords for each MSP customer.
3. Create tenants. The tenant definition consists of a few simple parameters to identify the associated customer. The tenant definition also includes tenant administrator and user accounts.
4. Set the scope to a tenant to administer tenant configuration while logged in as a global administrator.
5. Create at least one IP domain to represent customer networks.
6. Create at least one SNMP profile to enable SNMP polling of devices supporting customer infrastructure.
7. Exit tenant administration. Repeat the previous steps for each tenant.



If data sources are already registered and are collecting data, wait a few minutes. NetOps Portal creates system groups based on items that are discovered during monitoring. These groups are useful for creating custom groups that you can then allocate to users as permissions.

When system groups are available, take the following steps:

1. Set the scope to a tenant to administer tenant configuration, or log in as the tenant administrator.
2. Create any custom groups that are required to represent the customer networks and systems.
3. Edit the default tenant user account to add permission groups.  
Consider the likely role of this user, and the managed items that this user manages.
4. Create any other custom roles, user accounts, SNMP profiles, dashboards, and menus that are required for this customer.

Work with the IT staff of each customer to designate a user to act as the tenant administrator. The tenant administrator can complete the tenant configuration by creating custom groups and additional user accounts, if desired.

### **Data Source Support for Multi-tenancy**

DX NetOps Performance Management fully support multi-tenant monitoring. Integrated environments that include the following data sources offer limited multi-tenancy support:

#### **DX NetOps Network Flow Analysis**

##### **Supported Features**

Full support for multi-tenancy.

Assign a tenant and IP domain to each Harvester and router. Tenant assignment determines the configuration items that are available (such as tenant SNMP profiles).

A few limitations are noted in [Multi-tenancy and DX NetOps Network Flow Analysis](#).

#### **CA Application Delivery Analysis**

##### **Supported Features**

- IP Domains
- Custom Tenant configuration

CA Application Delivery Analysis does not segregate data within the data source interface.

CA Application Delivery Analysis monitors IP domains without a concept of tenants. As a result, NetOps Portal receives all items from CA Application Delivery Analysis in the Default Tenant. However, CA Application Delivery Analysis does support IP domains. NetOps Portal can thus associate these items with tenants according to their IP domain.

##### **NOTE**

Some managed items are duplicated between the Default Tenant and custom tenants.

#### **DX NetOps Spectrum**

##### **Supported Features**

DX NetOps Spectrum supports multi-tenancy. However, tenancy is only visible in DX NetOps Performance Management, and not in OneClick.

OneClick receives IP domains from NetOps Portal. Models in IP domains are synchronized them with NetOps Portal (and thus, associated with custom tenants).



Tenants are not visible in DX NetOps Spectrum OneClick. However, tenants in NetOps Portal have associated DX NetOps Spectrum devices based on IP domain. The global administrator can also make these items available for monitoring by tenant users by placing them into the Global Tenant Items group.

## **CA Unified Communications Monitor**

### **Supported Features**

CA Unified Communications Monitor supports multi-tenancy. Locations are automatically associated with IP domains by their subnets.

## **DX Application Performance Management**

### **Supported Features**

DX Application Performance Management does not support multi-tenancy.

All items from these data sources are associated with the Default Tenant and Default IP Domain. Grant tenant access to these data sources by adding items from them to service provider groups.

### **Roles for Multi-tenancy**

When multi-tenancy is deployed, the following roles are available in DX NetOps Performance Management:

- **Global administrator**  
This role is the Default Tenant administrator, usually representing an MSP. Product settings and data are not shared among tenants, but the Default Tenant Administrator can access and modify all settings. This user must have the predefined Administrator role.
- **Tenant Administrator**  
This role is a limited administrator associated with a single tenant. This operator cannot access shared infrastructure or configuration belonging to the host (usually, the MSP). Tenant user accounts can include one or more of these administrator accounts.

When you create a tenant, NetOps Portal prompts you to create a tenant administrator and a tenant user account. Operators who use these accounts can only perform monitoring or administrative tasks within this tenant. They cannot access the managed items and parameters associated with other tenants.

## **Configure a Tenant Environment**

Configuration tasks in a multi-tenant environment are split between the global administrator and the tenant administrator. The global administrator creates the tenant and configures monitoring profiles. The tenant administrator configures discovery and reporting.

Use the following process to configure a tenant environment:

1. [Verify the prerequisites.](#)
2. [Set up the tenant.](#)
3. [Discover devices.](#)
4. [Configure monitoring collections.](#)

### **NOTE**

You can also configure tenants by way of NetOps Portal and Data Administrator REST web services. For more information, see [Automate Tenant Configuration with REST Web Services](#).

## **Verify the Prerequisites**

Before configuring a tenant environment, work closely with the tenant customer to plan the tenant deployment. Collect basic information about the customer environment, such as IP domains and SNMP profiles. Knowledge of physical and virtual system topology is useful for creating a custom grouping structure to represent the customer environment.

Select a user to act as the tenant administrator. The tenant administrator can then complete the tenant configuration by creating custom groups, roles, users, SNMP profiles, menus, and dashboards.

Gather the following information:

- The list of IP domains for the tenant.  
Each IP domain requires a dedicated data collector.
- The list of SNMP communities for the tenant networks.
- The desired username of the tenant administrator.  
For example, select a representative of the customer site that is monitored.
- A dedicated data collector or a multi-tenant data collector with available capacity.
- The configured monitoring profiles.  
For more information, see [Configure Monitoring Profiles](#).

## **Set Up the Tenant**

As the global administrator, complete the following tasks to set up the tenant.

### **Create a Tenant**

The global administrator creates the tenant definition and tenant administrator. The tenant administrator can only see data and configuration for a single tenant. Data from other tenants is not accessible to a tenant administrator.

For more information, see [Manage Tenants](#).

The tenant administrator can complete the following procedures. Alternatively, the global administrator can administer the tenant environment to complete configuration.

### **Create an IP Domain**

Each tenant requires an associated IP domain.

#### **Follow these steps:**

1. Hover over **Administration, Configuration Settings**, and then click **IP Domains**.  
The **Manage IP Domains** page appears.
2. Click **New**.  
The **IP Domains Administration** dialog opens.
3. Specify a domain name and description.
4. (Optional) To assign a primary and secondary DNS address for the domain in DX NetOps Network Flow Analysis, select **DNS Settings**, and specify the required values.

#### **NOTE**

Only DX NetOps Network Flow Analysis uses the DNS settings.

5. Click **Save**.

The IP domain is created.

### **Assign a Tenant and IP Domain to a Data Collector**

For a tenant environment with a dedicated data collector, assign, or bind, the tenant and the IP domain to the data collector.

**NOTE**

For a multi-tenant data collector environment, configure the association through the Data Aggregator REST Web Services.

For more information, see [Configure Tenant-Agnostic Data Collectors](#).

**Follow these steps:**

1. Hover over **Administration, Monitored Items Management**, and then click **Data Collectors**.  
The **Data Collectors** page appears.
2. Select the data collector to which you want to assign the IP address, and then click **Assign**.  
The **Assign Data Collector** dialog opens.
3. Select the tenant and IP domain, and then click **Save**.

The IP domain and tenant are assigned to the data collector.

**Discover Devices**

As the tenant administrator, complete the following tasks to discover devices. Discovery is the process that DX NetOps Performance Management uses to build an inventory of devices in your network.

For more information about how to discover devices, see [Discovery](#).

**Configure Monitoring Collections**

In a tenant environment, the tenant administrator defines monitoring by assigning device collections to monitoring profiles. Monitoring profiles control how often to poll devices and which information to collect. Assigned metric families determine which metrics the system collects. Device collections are system groups that group devices for monitoring. Associating a collection with a monitoring profile causes DX NetOps Performance Management to monitor the devices according to the parameters in that profile. The global administrator defines monitoring profiles for all tenants.

**Organize Devices into a Collection**

For more information about device collections, see [Configure Monitoring in a New Environment](#).

**Assign Monitoring Profiles to a Device Collection**

Reporting, dashboards, and threshold monitoring function as normal in a tenant environment.

For more information, see [Configure Reporting in a New Environment](#).

**NOTE**

You can also assign device collections to monitoring profiles.

For more information, see [Configure Monitoring Profiles](#).

**Follow these steps:**

1. From the **Collections** page, select a device collection, and click the **Monitoring Profiles** tab.
2. Click **Manage**.
3. Select monitoring profiles for the **Assigned Monitoring Profiles** list, and click **Save**.  
The monitoring behavior that is defined in the assigned monitoring profiles is applied to the device collection.

The monitoring profiles are assigned to the device collection.

**Manage Tenants**

You can manage your tenants using NetOps Portal.

*Tenants* provide you with separate monitoring environments that you administer from a single user interface (UI). Tenants are not required for all NetOps Portal deployments. With multi-tenancy, managed service providers (MSPs) can monitor discrete customer networks and systems from a single instance of NetOps Portal.

#### NOTE

To edit the SNMP profiles, IP domains, or other configuration for a tenant, [administer the tenant](#).

Global administrators (users with the Administrator role) manage tenants.

In this article:

- [Create a Tenant](#)
- [Clone a Tenant](#)
- [Edit a Tenant](#)
- [Disable a Tenant](#)
- [Enable a Disabled Tenant](#)
- [Delete a Tenant](#)

### Create a Tenant

During tenant creation, you can also create a tenant administrator and a tenant user. The tenant administrator can only see and access data and configuration for a single tenant/MSP customer.

#### Follow these steps:

1. Hover over **Administration, Group Settings**, and then select **Tenants**.  
The **Manage Tenants** page appears. The tenants are listed.
2. Click **New**.  
The **Add Tenant** dialog opens.
3. Complete the following fields:
  - **Name**  
Defines the name for the tenant.
  - **Account ID**  
The ID that identifies the tenant and often corresponds to the MSP account number. This is also known as the tenant ID.
  - **Description**  
The description for the tenant.
  - **Status**  
Select whether the tenant is enabled or disabled.  
**Options:**
    - **Enabled:** Specifies that the tenant is enabled.
    - **Disabled:** Specifies that the tenant is disabled.**Default:** Enabled
  - **Theme**  
Specifies the theme that controls the appearance of NetOps Portal pages for the tenant. This theme is applied to all user accounts that are associated with this tenant.  
**NOTE**  
Applying a custom UI theme to a tenant changes the appearance (the PDF output) of exported reports. If you do not deploy multi-tenancy, the theme applies to the default tenant.  
**Options:** CA-Blue and CA-White  
**Default:** CA-Blue
  - **Language**  
Defines the language (locale) for the tenant. This is also known as the tenant culture.

**Options:**

- English (US)
- French
- Japanese

**Default:** English (US)

- Specify the credentials for the tenant administrator (admin) by completing the following fields in the **Default Administrator** section.

- **Administrator**

This user has administrative privileges in the tenant.

**Default:** TenantName\_admin

- **Password**

The user account's password.

**IMPORTANT**

This password is a one-time password. It expires immediately. On first login, when prompted to change this password, this user must change it by clicking the change password link.

For more information about this user account, see [Manage User Accounts](#).

- **Confirm password**

Enter the password again.

- Specify the credentials for the default tenant user (user) by completing the following fields in the **Default User** section.

- **User**

This user can access tenant-specific dashboards, but cannot access any administrative functions (the **Administration** menu).

- **Password**

The user account's password.

**IMPORTANT**

This password is a one-time password. It expires immediately. On first login, when prompted to change this password, this user must change it by clicking the change password link.

For more information about this user account, see [Manage User Accounts](#).

- **Confirm password**

Enter the password again.

- Click **Save**.

The new tenant definition is created, and the tenant administrator and a tenant user are created.

**Next Step:** [Configure monitoring in the tenant by administering the tenant](#).

**Clone a Tenant**

To use status, UI theme, and language of an existing tenant as a template for a new tenant, from the **Manage Tenants** page, select the existing tenant, and then click **Clone**.

**NOTE**

Tenant-specific information, such as users, groups, and SNMP profiles, is not copied. Configure monitoring in a tenant clone as if the clone is a new tenant.

**Edit a Tenant**

To modify any of the parameters that are defined during tenant creation, from the **Manage Tenants** page, select the tenant, and then click **Edit**. To change monitoring configuration in the tenant, administer the tenant.

If you changed the **Status** of the tenant, see [the "Disable a Tenant" section](#) or [the "Enable a Tenant" section](#).

To remove a tenant and the associated data from NetOps Portal, see [the "Delete a Tenant" section](#).

## **Disable a Tenant**

Stop active monitoring of the infrastructure of a tenant by disabling the tenant.

After you disable the tenant, the following results occur *after* the next synchronization with the data aggregator:

- The data aggregator stops all data collectors that are associated with the disabled tenant. The data collectors' **Status** is "Not Collecting Data". For any new data collector installations for a disabled tenant, the **Status** is "Not Collecting Data".

### **NOTE**

If you reenable the tenant, [restart each data collector](#).

- Polling is stopped for any devices and components that are being monitored for the tenant.
- Historical data on the devices and components for the tenant remains accessible.
- You cannot run discovery profiles that are associated with the tenant. The discovery profiles have the "Tenant Disabled" state (the **Status** is "Disabled").
- If a discovery profile is invalidated while a discovery is running, the discovery is aborted.
- An audit event is generated on the data aggregator device for the disabled tenant.

### **Follow these steps:**

1. From the **Manage Tenants** page, select the tenant, and then click **Edit**.
2. Change the tenant **Status** to "Disabled".
3. Save your changes.

## **Enable a Disabled Tenant**

Restart active monitoring of the infrastructure of a tenant by enabling the disabled tenant.

After you enable a disabled tenant, the following results occur *after* the next synchronization with the data aggregator:

- Discovery profiles that are associated with the enabled tenant are validated. The discovery profiles display their current state.
- An audit event is generated on the data aggregator device for the tenant.

You can restart polling by [restarting the data collectors](#) that are associated with the tenant. After the restart, the following results occur:

- Polling is restarted for any devices and components that are being monitored on for the tenant.
- Discovery profiles that are associated with the enabled tenant can be run. If the discovery profile is schedule, the system runs discovery at the next scheduled time.

### **Follow these steps:**

1. From the **Manage Tenants** page, select the tenant, and then click **Edit**.
2. Change the tenant **Status** to "Enabled".
3. Save your changes.

## **Delete a Tenant**

Remove a tenant and the associated data from NetOps Portal by deleting the tenant.

### **IMPORTANT**

You cannot recover information in deleted tenants. Instead, to stop monitoring in a tenant, disable the tenant.

Deleting a tenant deletes the following definitions from NetOps Portal:

- Data sources
- SNMP profiles
- IP domains
- Threshold profiles
- User accounts
- Roles
- Groups
- Custom dashboards
- Custom menus

After you delete the tenant, the following results occur *after* the next synchronization with the data aggregator:

- Polling on the deleted devices and device components is stopped.
- Historical data on the deleted devices and device components is no longer accessible.
- An audit event is generated on the data aggregator device for the deleted tenant.

#### NOTE

If you delete a tenant when an associated data collector is down, the data collector immediately stops when it restarts. An error message is logged in the `<DC_installation_directory>/apache-karaf/shutdown.log` file.

- ***DC\_installation\_directory***  
The default installation directory for the data collector.  
**Default:** `/opt/IMDataCollector`

## Administer Tenants

After you add a tenant, set up the required monitoring parameters and user access. To set up a tenant environment, log in as a Tenant Administrator that is associated with that tenant. Administrators can access NetOps Portal from the perspective of the tenant.

When you set the tenant scope to a selected tenant, only the configuration items available to that tenant appear. You can then administer the tenant, creating the required IP domains, user accounts, and more. They will only be available to users with permission to see the items that belong to that tenant.

### Modify Tenant Monitoring Parameters as a Tenant Administrator

Administrators and Tenant Administrators can modify the monitoring parameters of a tenant. Custom definitions that you create while administering a tenant are specific to that tenant, and are not shared among tenants.

Tenant Administrators can modify the IP domain, SNMP profile, user, role, and group definitions for a tenant. Administrators (the administrator for the Default Tenant) must set the tenant scope to the selected tenant to gain access to these definitions.

#### NOTE

Administrators can create Tenant Administrator user accounts for each tenant.

When the tenant scope has been set, the procedures for administering a tenant are identical to the procedures to perform in a single-tenant environment.

#### Follow these steps:

1. Log in as a Tenant Administrator associated with this tenant.

#### NOTE

You can also set the tenant scope to access tenant configuration while logged in as an Administrator.

The **Administering Tenant** indicator appears to show that you are administering the selected tenant environment. You can now see and modify only definitions that are associated with this tenant.

2. Click the **Administration** tab, and then select an item to modify:
  - IP Domains
  - SNMP Profiles
  - Groups
  - Menus
  - Roles
  - Users
3. Follow the procedures specific to the selected item.
4. Save your changes.  
The modifications are only visible to Administrators and IT Operators whose user accounts were created within this tenant environment.

### **Set up a Tenant Environment as an Administrator**

Administrators can set up the environment for a tenant that you have already created. For example, you can add custom IP domains, user accounts, or groups to the tenant. Set the scope to the tenant to access NetOps Portal from the perspective of the tenant.

#### **Follow these steps:**

1. Log in as an Administrator.
2. Hover over **Administration**, **Group Settings**, and then click **Tenants**.  
The **Manage Tenants** page appears.
3. Select the tenant that you want to administer, and then click **Administer**.  
The configuration that is associated with the selected tenant appears. You can create the IP domains, SNMP profiles, roles, users, menus, and groups that are required to represent and monitor this tenant environment. Use the menus under the **Administration** tab to configure the tenant.
4. (Optional) Change the tenant scope to another tenant by clicking the [change] link.  
The **Manage Tenants** page appears, and you can select another tenant.

Administrators and Tenant Administrators can modify the monitoring parameters of a tenant. Custom definitions that you create while administering a tenant are specific to that tenant, and are not shared among tenants.

To modify the IP domain, SNMP profile, user, role, and group definitions for a tenant, the Tenant Administrator simply logs in. Administrators (the administrator for the Default Tenant) must set the tenant scope to the selected tenant to gain access to these definitions.

#### **NOTE**

Administrators can create Tenant Administrator user accounts for each tenant.

When the tenant scope has been set, the procedures for administering a tenant are identical to the procedures to perform in a single-tenant environment.

#### **Follow these steps:**

1. Log in as a Tenant Administrator associated with this tenant.  
You can set the tenant scope to access tenant configuration while logged in as the Administrator.  
The **Administering Tenant** indicator appears to show that you are administering the selected tenant environment. You can now see and modify only definitions that are associated with this tenant.
2. Click the **Administration** tab, and select an item to modify:



- IP Domains
  - SNMP Profiles
  - Groups
  - Menus
  - Roles
  - Users
3. Follow the procedures specific to the selected item.
  4. Save your changes.  
The modifications are only visible to Administrators and IT Operators whose user accounts were created within this tenant environment.

## Automate Tenant Configuration with REST Web Services

As an MSP administrator, you use a third-party or proprietary tool to manage tenants. You want to automate the configuration of your tenant information for monitoring and reporting purposes. This scenario shows how you can configure tenant information using a combination of NetOps Portal and Data Aggregator REST web services. For automated tenant configuration, you write an application or script that leverages these web services.

Use the following process to configure tenant information using REST web services.

You can use a REST client editor or an HTTP tool that sends requests and gets responses. This scenario refers to the REST client editor.

### TIP

You can also configure a tenant environment by way of NetOps Portal.

For more information, see [Configure a Tenant Environment](#).

Use the following process to configure a tenant environment using the `tenants` REST web service:

1. [Considerations Before You Begin](#)
2. [Create a Tenant](#)
3. [Create an IP Domain](#)
4. [Enable SNMP Polling of Devices](#)
5. [Assign the Data Collector to the Tenant](#)
6. [Create a Discovery Profile and Run a Discovery](#)
7. [Review the Discovered Devices and Instances](#)

### **Considerations Before You Begin**

Before you automate tenant configuration using the REST web services, consider the following information:

- You cannot interchange the identifiers from NetOps Portal REST web services and Data Aggregator REST web services.
- After an upgrade, verify that the version of the web services you were using has changed and is still compatible. Review the [Release Notes](#). If the REST web services have changed, update your application or script.

### **Create a Tenant**

As an MSP administrator, you can create tenants with specific parameters. The basic tenant definition contains a few parameters to identify your customer and let other operators access managed items and configuration for the customer. An administrator account is a required component of the tenant definition so that the customer can perform some tenant setup.

**Follow these steps:**

1. Launch a REST client.
2. Enter the following URL for the `tenants` endpoint for the NetOps Portal RESTful web services API, selecting **POST** for the **HTTP Method**:  
`http://<PC_host>:8181/pc/center/webservice/tenants/`
3. Provide a valid username and password for a user account that has administrator access to NetOps Portal.
4. Select **application/xml** as the Body Content-type in the Body settings.
5. Enter the following information within the **Body** text section:
  - **tenantName**  
Provides a name for the tenant you are creating.  
**Required:** Yes
  - **tenantDesc**  
Describes the tenant. Specifying a description is optional.  
**Required:** Yes
  - **accountIdentifier**  
Identifies this tenant. This value usually corresponds to the MSP account number. If you supply a value as input, it must be unique across all defined tenants.  
**Default:** null  
**Required:** No
  - **status**  
Defines the status of the tenant, and whether tenant user accounts that are associated with this tenant can use NetOps Portal.  
**Options:**
    - **Enabled:** The tenant is enabled and tenant user accounts that are associated with this tenant can use NetOps Portal.
    - **Disabled:** The tenant is disabled and tenant user accounts that are associated with this tenant cannot use NetOps Portal.**Default:** Enabled  
**Required:** No
  - **removable**  
States whether the item can be deleted (removed from the database).  
**Values:** true or false  
**Default:** true  
**Required:** No
  - **theme**  
Specifies the theme that controls the appearance of NetOps Portal pages for this tenant. This theme is applied to all user accounts that are associated with this tenant.  
**Options:** CA-Blue and CA-White  
**Default:** CA-Blue  
**Required:** No
  - **defaultCulture**  
Specifies the language (locale) for this tenant. Supply a language identifier from the following list:
    - `en_US` (English, United States)
    - `ja_JP` (Japanese)
    - `fr_FR` (French, France)**Default:** `en_US`  
**Required:** No

**Example:**

```

<tenant>
  <tenantName>John Doe</tenantName>
  <tenantDesc>John Doe Corporation tenant</tenantDesc>
  <accountIdentifier>JD1234</accountIdentifier>
  <status>Enabled</status>
  <removable>false</removable>
  <theme>CA-Blue</theme>
  <defaultCulture>en-US</defaultCulture>
</tenant>

```

## 6. Issue the web service call.

The tenant is created.

### Create an IP Domain

A tenant definition must contain at least one IP domain that identifies the IP addresses of managed items in the tenant environment. Create an IP domain and add it to the tenant.

#### Follow these steps:

1. Launch a REST client.
2. Enter the following URL for the `domains` endpoint for the NetOps Portal RESTful web services API, selecting **POST** for the **HTTP Method**:

```
http://<PC_host>:8181/pc/center/webservice/domains/
```

3. Provide a valid username and password for a user account that has administrator access to NetOps Portal.
4. Select **application/xml** as the Body Content-type in the Body settings.
5. Enter the following information within the **Body** text section:

- **groupId**

(Optional) Defines an internal (database) identifier for the group definition that is associated with a domain. You provide this number. To specify a valid `groupId`, issue a POST request to the `groups` endpoint for the NetOps Portal REST service:

```
http://<PC_host>:8181/pc/center/webservice/groups
```

- **itemDesc**

Describes the IP domain. The `itemDesc` tag is required, but an actual description is optional.

- **itemName**

Defines a name for the IP domain.

- **TenantID**

Defines an internal (database) identifier for the tenant definition. The tenant ID was assigned when you created the tenant and the tenant information was synchronized with the data aggregator.

**TIP**

Determine the tenant ID by entering the following URL for the `tenants` endpoint for the NetOps Portal RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<PC_host>:8181/pc/center/webservice/tenants/tenantName/tenant_name
```

- **primaryDNSAddress**

The IP address of the primary name server for this domain. This tag is required, but an actual IP address is optional.

- **primaryDNSPort**

The port number that the primary name server uses. This tag is required, but an actual port number is optional.

- **secondaryDNSAddress**

The IP address of the secondary name server for this domain. This tag is required, but an actual IP address is optional.

- **secondaryDNSPort**

The port number that the secondary name server uses. This tag is required, but an actual port number is optional.

- **isDnsProxyEnabled**

(Optional) Indicates whether the proxy address is enabled for this IP domain.

**Default:** false

- **dnsProxyAddress**

(Optional) Indicates the IP address of the DNS proxy server.

**Default:** null

**For example:**

```
<domain>
  <groupId>7</groupId>
  <itemDesc>{Description }</itemDesc>
  <itemName>{IP Domain Item Name}</itemName>
  <tenantID>8</tenantID>
  <primaryDNSAddress>0.0.0.0</primaryDNSAddress>
  <primaryDNSPort>0.0.0.0</primaryDNSPort>
  <secondaryDNSAddress>0.0.0.0</secondaryDNSAddress>
  <secondaryDNSPort>0.0.0.0</secondaryDNSPort>
  <isDnsProxyEnabled>false</isDnsProxyEnabled>
</domain>
```

## 6. Issue the web service call.

The IP domain is created and added to the tenant.

## Enable SNMP Polling of Devices

Enable SNMP polling of devices by providing SNMP credentials in an SNMPv1/SNMPv2c or SNMPv3 profile. During inventory discovery, the data collector uses SNMP profiles to determine what credentials it uses when accessing a device. NetOps Portal maintains a ranked list of profiles. During inventory discovery, each profile is tried for device access. The profile with the highest rank that can access a device is used.

For this scenario, you create an SNMPv3 profile.

### **NOTE**

Community strings and credentials are encrypted when they are stored in NetOps Portal and when they are sent to the data aggregator and the data collector. NetOps Portal REST web services use the HTTP protocol instead of the HTTPS protocol. Therefore, the community strings and passwords are sent over the Internet in clear text, and, can be compromised. Use this method only on the NetOps Portal host.

### **Follow these steps:**

1. Launch a REST client.
2. Get the SNMP profile by entering the following URL for the `profiles` endpoint for the NetOps Portal RESTful web services API, selecting **GET** for the **HTTP Method**:  
`http://<PC_host>:8181/pc/center/webService/profiles/`
3. Select **application/xml** as the Body Content-type in the Body settings.
4. Provide a valid username and password for a user account that has administrator access to NetOps Portal.
5. Enter the following information within the **Body** text section:
  - **name**  
Defines a name for the SNMP profile.

**NOTE**

Profile names must be unique, cannot be duplicated across SNMP versions, and are not case-sensitive.

- **port**  
(Optional) Identifies the port that is used to make SNMP connections to devices associated with this profile.  
**Default:** 161
- **userName**  
(Optional) Identifies the user for the profile, whose secret keys were used potentially to authenticate and encrypt the SNMPv3 packets. The user name is a character string.  
**Default:** The encrypted value of the name
- **context**  
(Optional) Specifies a collection of management information that is accessible by an SNMP entity. The context name is necessary for providing end-to-end identification and for retrieving data from an SNMPv3 agent. The context name is an octet string.  
**Default:** an empty string

**NOTE**

The data aggregator does not use Context Name on the SNMPv3 profiles to communicate with the device.

- **version**  
Specifies the version of SNMP that the profile uses. Specify Version3 for this scenario.  
**Default:** Version1or2C
- **securityLevel**  
(Optional) Specifies the security level to use.  
**Options:**
  - **NoAuthNoPriv:** Specifies no authentication and no privacy.
  - **AuthNoPriv:** Specifies authentication and no privacy.
  - **AuthAndPriv:** Specifies authentication and privacy.**Default:** NoAuthNoPriv
- **authProtocol**  
(Optional) Specifies the authentication protocol to use when contacting devices associated with this profile, and for authenticating SNMPv3 packets:  
**Options:**
  - None (Do not attempt authentication)
  - MD5 (Message Digest 5)
  - SHA (Secure Hash Algorithm)**Default:** None
- **authPassword**  
Specifies the password for authentication using SNMPv3 and the selected authentication protocol. The password must contain a minimum of eight characters. This value is required for the MD5 and the SHA authentication protocol values.
- **privProtocol**  
(Optional) Specifies the encryption protocol to use for data flows sent to any devices or servers that are associated with this profile.  
**Options:**
  - None (Does not encrypt communications. Use only with the NoPriv options.)
  - DES
  - AES (128-bit encryption)
  - Triple DES**Default:** None

**NOTE**

The privacy protocol option is only enabled when authentication is enabled for this profile.

- **privPassword**  
Defines the password that is used when exchanging encryption keys. This value is required for DES, AES, and Triple DES privacy protocols.
- **rank**  
(Optional) Specifies the rank of the profile in the global list of SNMP profiles.  
**Default:** 0
- **enabled**  
(Optional) Indicates whether the information in this profile is used when not explicitly assigned to a device. Select true.  
**Default:** true
- **tenantID**  
The internal (database) identifier for the tenant definition.  
**Default:** The ID of the Default Tenant

**NOTE**

This number must be the same number that you chose when you created the IP domain.

**Example: No authentication and no privacy**

```
<SnmpProfile>
  <name>Tokyo</name>
  <port>161</port>
  <userName>myuser</userName>
  <context></context>
  <version>Version3</version>
  <securityLevel>NoAuthNoPriv</securityLevel>
  <authProtocol>None</authProtocol>
  <authPassword>None</authPassword>
  <privProtocol>None</privProtocol>
  <privPassword>None</privPassword>
  <rank>4</rank>
  <enabled>true</enabled>
  <tenantID>7</tenantID>
</SnmpProfile>
```

**Example: Authentication and no privacy**

```
<SnmpProfile>
  <name>Brasil</name>
  <port>161</port>
  <userName>myuser</userName>
  <context></context>
  <version>Version3</version>
  <securityLevel>AuthNoPriv</securityLevel>
  <authProtocol>MD5</authProtocol>
  <authPassword>test</authPassword>
  <privProtocol>None</privProtocol>
  <privPassword>None</privPassword>
  <rank>3</rank>
  <enabled>true</enabled>
  <tenantID>7</tenantID>
```

```
</SnmpProfile>
```

### Example: Authentication and privacy

```
<SnmpProfile>
  <name>Boston</name>
  <port>161</port>
  <userName>myuser</userName>
  <context></context>
  <version>Version3</version>
  <securityLevel>AuthAndPriv</securityLevel>
  <authProtocol>MD5</authProtocol>
  <authPassword>test</authPassword>
  <privProtocol>TripleDES</privProtocol>
  <privPassword>test</privPassword>
  <rank>1</rank>
  <enabled>true</enabled>
  <tenantID>7</tenantID>
</SnmpProfile>
```

6. Create the profile by issuing a POST request to the `profiles` endpoint for the NetOps Portal REST service:

```
http://<PC_host>:8181/pc/center/webservice/profiles/saveProfile/{true|false}
```

– **{true|false}**

Specifies a Boolean value for the `rankTiesAscendingByDate` parameter. **True** indicates that the profile you are adding is the last in rank order (as determined by the creation date of the SNMP profile).

The XML returns true when the operation succeeds.

The SNMP profile is automatically synchronized with the data aggregator and is available for inventory discovery to use it. SNMP credentials are provided.

### Assign the Data Collector to the Tenant

Each data collector installation is assigned to a tenant and IP domain. Assign a data collector to the tenant.

#### Prerequisites:

- The data collector is installed.
- A discovery has not run against this data collector.

#### Follow these steps:

1. Launch a REST client.
2. Determine the tenant ID is by entering the following URL for the `tenants` endpoint for the NetOps Portal RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/tenants
```

3. Make a note of the tenant ID of the tenant to which you want to assign to the data collector.
4. Determine the IP domain ID of the IP domain that you want to assign to the data collector by entering the following URL for the `ipdomains` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/tenant/tenantID/ipdomains
```

– **tenantID**

The ID of the tenant that you want to assign to the data collector instance.

**NOTE**

You determined this ID in Step 2.

- Determine the data collector ID by entering the following URL for the `dcms` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/dcms
```

- Make a note of the ID of the data collector installation to which you want to assign your tenant and IP domain.
- Enter the following URL for the `genericWS` endpoint for the Data Aggregator RESTful web services API (the *generic* REST web service), selecting **POST** for the **HTTP Method**:

```
http://<DA_host>:8581/genericWS/dcm/collectorID
```

- **collectorID**

The ID of the data collector installation that you want to assign to your tenant and IP domain.

**NOTE**

You determined this ID in Step 5.

- Enter the following information within the **Body** text section:

```
<DCM>
  <DCM>
    <IP_DOMAIN_ID>IPdomainID</IP_DOMAIN_ID>
  </DCM>
</DCM>
```

- **IPdomainID**

The IP domain ID for the tenant to which you want to assign the data collector.

**NOTE**

You determined this ID in step 4.

The data collector is assigned to the tenant.

### **Create a Discovery Profile and Run a Discovery**

*Discovery profiles* specify how discovery operates in your data aggregator environment. Within a discovery profile, specify the IP addresses, IP address ranges, or host names you want to discover devices for.

You create a discovery profile and run a discovery in one operation when using the REST web services. When discovery is run, devices are discovered based on the discovery profile you create.

For more information about discovery profiles, see [Create Discovery Profiles](#).

#### **Follow these steps:**

- Launch a REST client.
- Select application/xml as the Body Content-type in the **Body** settings.
- Find and make a note of the default tenant ID by entering the following URL for the `tenants` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:  

```
http://<DA_host>:8581/rest/tenants
```
- Retrieve the XSD schema XML (which defines the structure of the discovery profile) for the data aggregator entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:  

```
http://<DA_host>:8581/rest/discoveryprofiles/XSD/getlist.xsd
```
- Include IP addresses and host names so that devices in your network are discovered. Do one or more of the following actions:
  - Enter individual IP addresses for which you want to discover devices in the **IP Address List** field.
  - Enter the host names for which you want to discover devices in the Host List field.



**NOTE**

These fields accept comma-delimited values. Single quotes, double quotes, backward slashes, forward slashes, and ampersands are not permitted.

## 6. Set the following required attributes:

– **RunStatus**

Specifies whether to run the discovery. For this scenario, set the attribute to START.

**NOTE**

To rerun discovery later, update (PUT) the run status to START.

– **Name**

Specifies a descriptive name for the discovery profile. This field cannot contain single quotes, double quotes, backward slashes, forward slashes, and ampersands.

– **IPDomainID**

An internal ID that is used to identify IP domains. The data aggregator generates this number.

**TIP**

Determine the IP domain ID by entering the following URL for the `tenant` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/tenant/<tenantID>/ipdomains
```

## 7. (Optional) Limit the discovery to specific SNMP profiles using the following attributes:

– **SNMPProfileIDList**

Specifies a ranked list of SNMP Profile ItemIDs for use during discovery. This list is used only if the `UseListOfSnmpProfiles` attribute equals true.

– **UseListOfSnmpProfiles**

A boolean indicating whether a specific list of SNMP profiles is used.

**Options:**

- **true:** Use a specific list of SNMP profiles.
- **false:** Use the global list of SNMP profiles.

## 8. (Optional) Change the priority in which you want the discovered device to be named, using the following attribute:

– **DeviceNameRankingList**

Specifies the priority in which you want discovered devices to be named.

9. (Optional) Set the `IcmpDiscoveryEnabled` attribute to true to enable ICMP discovery.**Example: Discovery profile XML that includes an optional list of specific SNMP profiles.**

```
<?xml version="1.0" encoding="UTF-8"?>
<DiscoveryProfile version="1.0.0">
  <ActivationStatus>true</ActivationStatus>
  <SNMPProfileIDList><SNMPProfileID>478</SNMPProfileID>
    <SNMPProfileID>2239</SNMPProfileID>
  </SNMPProfileIDList>
  <UseListOfSnmpProfiles>true</UseListOfSnmpProfiles>
  <IPRangesList>
    <IPRanges>10.0.64.202-10.0.64.206</IPRanges>
  </IPRangesList>
  <HostNamesList>
    <HostNames>ahost</HostNames>
  </HostNamesList>
  <IPListList>
    <IPList>10.10.10.10</IPList>
  </IPListList>
  <RunStatus>START</RunStatus>
```

```

<Item version="1.0.0">
  <Name>BigRange_TestProfileViaAPI</Name>
</Item>
<IPDomainMember version="1.0.0">
  <IPDomainID>285</IPDomainID>
</IPDomainMember>
<DeviceNameRankingList>
  <DeviceNameRanking>{http://im.ca.com/inventory}
    ManageableDevice.SystemName</DeviceNameRanking>
  <DeviceNameRanking>{http://im.ca.com/inventory}
    Device.HostName</DeviceNameRanking>
  <DeviceNameRanking>{http://im.ca.com/inventory}
    Device.PrimaryIPAddress</DeviceNameRanking>
</DeviceNameRankingList>
<IcmpDiscoveryEnabled>true</IcmpDiscoveryEnabled>
</DiscoveryProfile>

```

10. Save and run the new discovery profile by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **POST** for the **HTTP Method**:

`http://<DA_host>:8581/rest/tenant/<tenantID>/discoveryprofiles`

– **tenantID**

The internal database identifier for the tenant definition. The tenant ID was assigned when you created the tenant and the tenant information was synchronized with the data aggregator.

**TIP**

Determine the tenant ID by entering the following URL for the `tenants` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

`http://<DA_host>:8581/rest/tenants`

The discovery profile is created and discovery runs.

The discovered devices are automatically added to the appropriate out-of-the-box device collection or another user-created device collection.

The discovery profile is created and discovery is run.

## **Review the Discovered Devices and Instances**

After you create a discovery profile and run a discovery, verify your results. View the summary of the number of new pingable and manageable devices that were discovered. The summary also includes details about the discovered devices, including IP address, type, vendor association, SNMP profiles, and model type.

### **Follow these steps:**

1. Launch a REST client.
2. View a list of discovery profiles by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

`http://<DA_host>:8581/rest/discoveryprofiles`

A list of discovery profiles and their instance IDs is returned.

3. Find the discovery instance ID for the discovery profile that you created previously.
4. View the discovery instance details by entering the following URL for the `discoveryinstances` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

`http://<DA_host>:8581/rest/discoveryinstances/<instance_ID>`

The instance XML returns information about new and existing devices and SNMP profiles that were tested.

**Example:**

```

<?xml version="1.0" encoding="UTF-8"?>
  <DiscoveryInstance version="1.0.0">
    <ID>236</ID>
    <IPSweepTotalSuccess>5</IPSweepTotalSuccess>
    <CompletionTime>Thu Apr 12 12:36:35 CDT 2012</CompletionTime>
    <ExistingFoundDevicesList>
      <ExistingFoundDevices>241</ExistingFoundDevices>
      <ExistingFoundDevices>239</ExistingFoundDevices>
      <ExistingFoundDevices>240</ExistingFoundDevices>
      <ExistingFoundDevices>238</ExistingFoundDevices>
    </ExistingFoundDevicesList>
    <IPSweepCompletionTime>Thu Apr 12 12:36:35 CDT 2012</
IPSweepCompletionTime>
      <ExistingFoundManageableDevicesList>
        <ExistingFoundManageableDevices>241</ExistingFoundManageableDevices>
        <ExistingFoundManageableDevices>239</ExistingFoundManageableDevices>
        <ExistingFoundManageableDevices>240</ExistingFoundManageableDevices>
        <ExistingFoundManageableDevices>238</ExistingFoundManageableDevices>
      </ExistingFoundManageableDevicesList>
      <StartTime>Thu Apr 12 12:36:31 CDT 2012</StartTime>
      <NewlyCreatedDevicesList>
        <NewlyCreatedDevices>383</NewlyCreatedDevices>
      </NewlyCreatedDevicesList>
      <TestedCommProfilesList>
        <TestedCommProfiles>199</TestedCommProfiles>
        <TestedCommProfiles>198</TestedCommProfiles>
      </TestedCommProfilesList>
      <IPSweepStartTime>Thu Apr 12 12:36:31 CDT 2012</IPSweepStartTime>
      <ProgressPercentage>100</ProgressPercentage>
      <IPSweepTotal>7</IPSweepTotal>
      <ProfileID>235</ProfileID>
      <NewlyCreatedManageableDevicesList>
        <NewlyCreatedManageableDevices>383</NewlyCreatedManageableDevices>
      </NewlyCreatedManageableDevicesList>
      <PingResponseDeviceCount>5</PingResponseDeviceCount>
      <CompletionStatus>SUCCESS</CompletionStatus>
      <Item version="1.0.0">
        <CreateTime>Thu Apr 12 12:36:31 CDT 2012</CreateTime>
      </Item>
    </DiscoveryInstance>

```

5. (Optional) Review information about specific devices (new or existing) by entering the following URL for the devices endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

`http://<DA_host>:8581/rest/devices/<device_ID>`

– **device\_ID**

Specifies the ID of the referenced datapoint, such as devices. Use the ID numbers that were returned when you viewed the discovery instance details.

6. (Optional) Review information about the SNMP profiles that were tested during the discovery by entering the following URL for the `profiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP**

**Method:**

```
http://<DA_host>:8581/rest/profiles/<profile_ID>
```

– ***profile\_ID***

Specifies the ID of the SNMP profile. The `<TestedCommProfilesList>` section in the XML displays the SNMP profile IDs.

The tenant information is configured using REST web services.

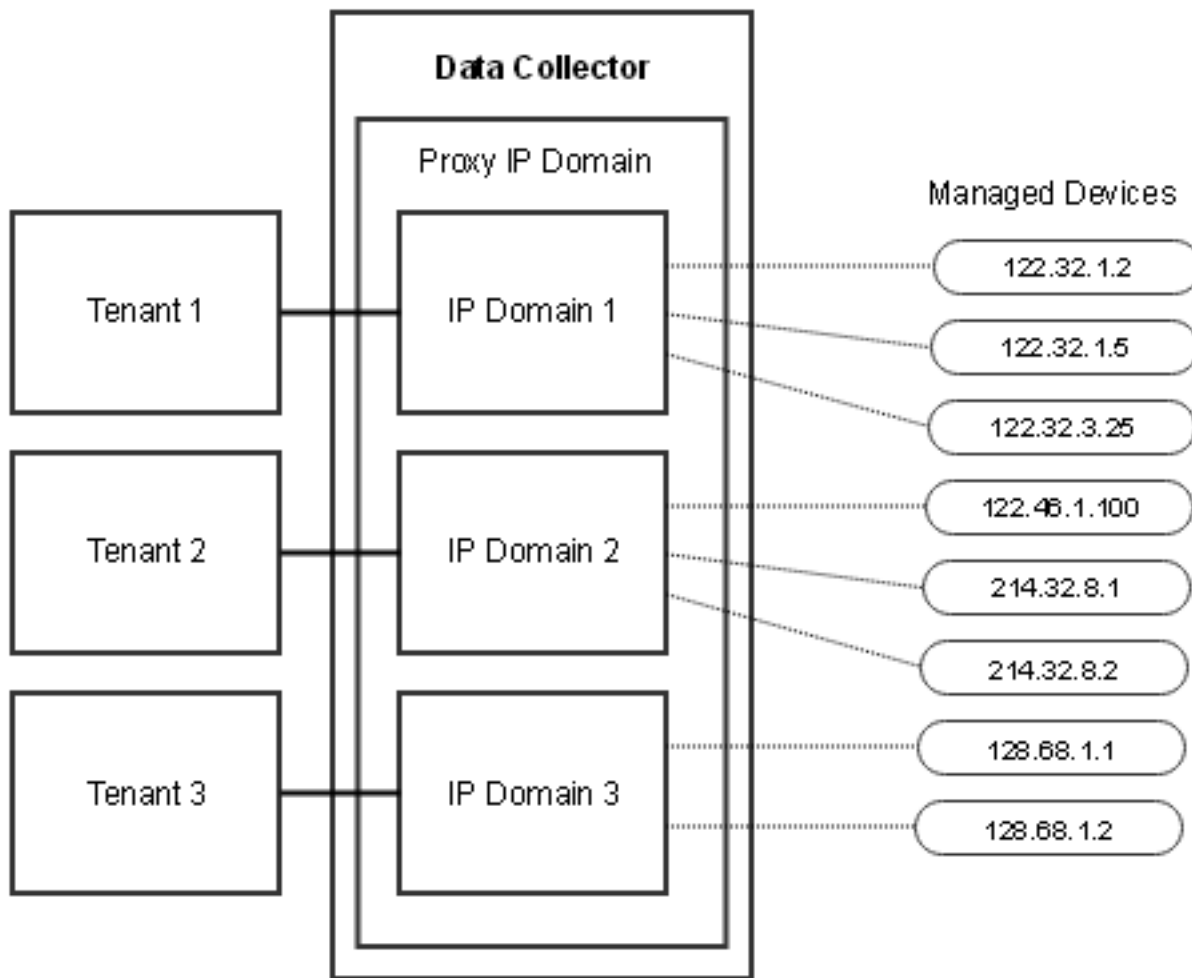
**TIP**

For automated tenant configuration, write an application or script that leverages the web services in this scenario.

## Configure Tenant-Agnostic Data Collectors

In a standard tenant deployment, each tenant has a dedicated data collector. For multiple tenants that reside in the same IP routing space, you can configure NetOps Portal to use fewer data collectors.

The following diagram illustrates the basic architecture of a tenant-agnostic data collector:

**Figure 135: Multi-tenant Data Collector Architecture**

The data collector is associated with the proxy IP domain. This domain represents the IP domain of the managed environment, meaning that IP addresses must be unique across all tenants in the environment.

**TIP**

The IP proxy domain is the only new object to create this feature. You must create this object through REST.

Multiple IP domains are associated with the data collector within the proxy IP domain. Each IP domain is associated with a single tenant. The data collector sends requests to the managed devices defined in each tenant.

Use the following process to configure tenant-agnostic data collectors:

1. [Review the Limitations](#)
2. [Configure the Data Collector Through REST](#)
3. [Get Item IDs](#)
4. [Configure the Data Collector](#)
5. [Add Tenants to the Proxy IP Domain](#)
6. [Configure the Tenant](#)

## Review the Limitations

A tenant-agnostic data collector configuration has the following limitations:

- IP addresses in tenant IP domains cannot overlap. IP addresses must be unique across all tenants.

### IMPORTANT

The owner of the environment is responsible for ensuring IP address uniqueness. Without unique IP addresses, tenants can monitor and manage devices that belong to other tenants. To ensure IP address uniqueness, enforce limited IP ranges in discovery profiles.

- Devices cannot be migrated between tenants.
- Existing tenant deployments cannot be migrated to tenant-agnostic data collectors.
- Tenant-agnostic data collectors cannot host DX NetOps Mediation Manager (DX NetOps MM) or DX NetOps Virtual Network Assurance (VNA). For each tenant, a separate data collector is required for DX NetOps MM and VNA.

## Configure the Data Collector Though REST

To configure a tenant-agnostic data collector, use REST to create the required items and associations. The following process uses NetOps Portal and Data Aggregator REST.

### IMPORTANT

If you script these steps, account for synchronization time between steps. Synchronization can take up to 5 minutes.

## Get Item IDs

Several item IDs are required to configure the data collector for multiple tenants. Get the item IDs before you begin configuration. Use the GET method with your preferred REST client for the following item IDs:

- Data Collector ID

`http://<da_host>:<port>/rest/dcms`

- **da\_host**

Specifies the data aggregator host name.

- **port**

Specifies the data aggregator required port number.

**Default:** 8581

For more information about the data aggregator server ports that should be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).

- Pseudo Tenant ID

`http://<da_host>:<port>/rest/pseudotenants`

- **da\_host**

Specifies the data aggregator host name.

- **port**

Specifies the data aggregator required port number.

**Default:** 8581

### NOTE

The pseudo tenant proxy (Pseudo Tenant Proxy) is predefined. The pseudo tenant works with the proxy IP domain to support multiple tenants and IP domains on a single data collector.

## Configure the Data Collector

Configure the data collector to handle multiple tenants. Complete this procedure for each tenant-agnostic data collector.

For the bold IDs, use the IDs from the previous step.

**Follow these steps:**

1. Create a proxy IP domain:
  - **Method:** POST
  - **Endpoint:** `http://<da_host>:<port>/rest/tenant/{pseudo_Tenant_ID}/proxyipdomains/`
    - **da\_host**  
Specifies the data aggregator host name.
    - **port**  
Specifies the data aggregator required port number.  
**Default:** 8581
  - **Body:**

```
<ProxyIPDomain version="0.0.0">
  <IPDomain version="0.0.0">
    <Name>Proxy IP Domain</Name>
    <Description>Description for Proxy IP Domain</Description>
    <DNSProxyEnabled>false</DNSProxyEnabled>
    <DNS1Address>0.0.0.0</DNS1Address>
    <DNS2Address>0.0.0.0</DNS2Address>
    <DNS1Port>0</DNS1Port>
    <DNS2Port>0</DNS2Port>
  </IPDomain>
</ProxyIPDomain>
```
2. Get and note the proxy IP domain ID:
  - **Method:** GET
  - **Endpoint:** `http://<da_host>:<port>/rest/proxyipdomains`
    - **da\_host**  
Specifies the data aggregator host name.
    - **port**  
Specifies the data aggregator required port number.  
**Default:** 8581
3. Associate the proxy IP domain with the data collector:
  - **Method:** PUT
  - **Endpoint:** `http://<da_host>:<port>/genericWS/dcm/{DataCollectorID}`
    - **da\_host**  
Specifies the data aggregator host name.
    - **port**  
Specifies the data aggregator required port number.  
**Default:** 8581
  - **Body:**

```
<DCM>
  <IP_DOMAIN_ID>{proxyDomainId}</IP_DOMAIN_ID>
</DCM>
```

**Add Tenants to the Proxy IP Domain**

Create tenants and assign those tenants to the data collector. Complete this procedure once for each tenant.

For more information about the tenant and IP domain XML, see [Automate Tenant Configuration with REST Web Services](#).

**TIP**

You can create the required tenants and the IP domains in those tenants, through NetOps Portal. For more information see [Manage Tenants](#) and [IP Domains](#). If you create these items in the UI, start this procedure from step 4.

**Follow these steps:**

## 1. Create the tenant:

- **Method:** POST
- **Endpoint:** `http://<PC_host>:<port>/pc/center/webservice/tenants/`
  - ***PC\_host***  
Specifies the NetOps Portal host name.
  - ***port***  
Specifies the NetOps Portal required port number.  
**Default:** 8181
- **Body:**

```
<tenant>
  <tenantName>tenant_name</tenantName>
  <tenantDesc>Tenant description</tenantDesc>
  <accountIdentifier>AccountNumber</accountIdentifier>
  <status>Enabled</status>
  <removable>false</removable>
  <theme>CA-Blue</theme>
  <defaultCulture>en-US</defaultCulture>
</tenant>
```

## 2. Get and note the tenant ID:

- **Method:** GET
- **Endpoint:** `http://<PC_host>:<port>/pc/center/webservice/tenants/tenantName/tenant_name`
  - ***PC\_host***  
Specifies the NetOps Portal host name.
  - ***port***  
Specifies the NetOps Portal required port number.  
**Default:** 8181

## 3. Create the tenant IP domain:

- **Method:** POST
- **Endpoint:** `http://<PC_host>:<port>/pc/center/webservice/domains/`
  - ***PC\_host***  
Specifies the NetOps Portal host name.
  - ***port***  
Specifies the NetOps Portal required port number.  
**Default:** 8181
- **Body:**

```
<domain>
  <itemDesc>IP domain description</itemDesc>
  <itemName>IP domain name</itemName>
  <TenantID>{tenantID}</TenantID>
  <primaryDNSAddress>0.0.0.0</primaryDNSAddress>
  <primaryDNSPort>0.0.0.0</primaryDNSPort>
  <secondaryDNSAddress>0.0.0.0</secondaryDNSAddress>
```



```

    <secondaryDNSPort>0.0.0.0</secondaryDNSPort>
    <isDnsProxyEnabled>>false</isDnsProxyEnabled>
  </domain>

```

4. Get the tenant IP domain ID:

- **Method:** GET
  - **Endpoint:** `http://<da_host>:<port>/rest/ipdomains`
    - ***da\_host***  
Specifies the data aggregator host name.
    - ***port***  
Specifies the data aggregator required port number.  
**Default:** 8581
- For more information about the data aggregator server ports that must be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).

5. Associate the proxy IP domain with the tenant IP domain:

- **Method:** PUT
  - **Endpoint:** `http://<da_host>:<port>/rest/ipdomains/{tenantIPDomainID}`
    - ***da\_host***  
Specifies the data aggregator host name.
    - ***port***  
Specifies the data aggregator required port number.  
**Default:** 8581
  - **Body:**

```

<IPDomain version="0.0.0">
  <ProxyIPDomain>{proxyIpDomainID}</ProxyIPDomain>
</IPDomain>

```
- For more information about the data aggregator server ports that must be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).

## Configure the Tenant

Use the standard methods to [configure monitoring in the tenant space](#).

## Multi-tenancy and Application Delivery Analysis

CA Application Delivery Analysis (ADA) does not offer a native multi-tenancy feature, but it supports the multi-tenancy features in NetOps Portal. ADA does support IP domains and data segregation per user account, so tenant users only see the data within their associated tenant.

For more information, see [assign-ca-application-delivery-analysis-items-to-an-ip-domain](#).

## Other Deployment Considerations

Take care when selecting user account product privileges. The product privilege to a data source enables a user to drill down from a view back to the source of the data.

Assuming that you have carefully segregated data from different customer environments into separate tenants, you probably want to prevent users from returning to ADA. In ADA, tenant separation is not applied, and all data is available for viewing in reports. You set User product privileges using the Add or Edit User Account wizard.

**IMPORTANT**

Assign the User product privilege to any user who does not require access to all data in the ADA data source.

## Multi-tenancy and Network Flow Analysis

The DX NetOps Network Flow Analysis (NFA) data source does not offer a native multi-tenancy feature. However, the harvesters and routers can associate monitored items with the appropriate tenant by means of the IP domain.

Interfaces and CVIs inherit their initial tenant-domain setting from the parent router and harvester. The setting is inherited when the parent harvester is added, and the router and interfaces first become active. If the harvester is not associated with a custom domain, the routers and interfaces are assigned to the Default Domain as they become active.

You can edit the settings for interfaces and CVIs to associate them with any tenant and domain at any time. The setting does not have to match the parent router or harvester.

Changing this setting can affect which operators have access to the interface's data. The setting does not affect which SNMP profiles are used for polling. The router tenant determines the set of SNMP profiles for polling.

### Follow these steps:

1. Open the **Active Interfaces** page:
  - a. Select **Administration** from the NFA console menu.  
The **Administration** page opens.
  - b. Select **Interfaces: Physical & Virtual** from the **Administration** menu.  
The **Active Interfaces** page opens.
2. Select the check box next to one or more interfaces that you want to associate with a tenant and domain.
  - To search for parent routers, interfaces, or CVIs, enter all or part of a router IP address, a router or interface name, or an interface description in the **Search** field, and then click **Search**. Expand the router details.
  - To navigate to an interface or CUI manually, go to the page that contains the parent router and click the arrow next to the router name. The router details appear, showing the interfaces and CVIs.
3. Click **Edit**.  
The editing dialog opens. The **Domain** selection list is included in the dialog only if multiple domains exist.
4. Select a tenant/domain option from the **Domain** list, and then click **Save**.  
The dialog closes. The changes are shown on the **Active Interfaces** page.

**NOTE**

You can also change the tenant-domain setting for harvesters and routers.

### Interface Domain Changes

When an administrator allocates the interfaces from a single router to multiple MSP customer domains, interfaces are often switched to different IP domains. NFA detects interface domain changes. These changes are not applied to data being sent to the data aggregator, however, because the data aggregator only identifies interfaces at the device level.

As a result, when interfaces are switched to a different domain after monitoring and data collection have begun, the global administrator sees two different interface items for each affected interface. The two interfaces continue to be monitored in separate IP domains, and their data is not aggregated. In this situation, individual tenant users do not see any problem. Device-level statistics for the routers or switches that contain these interfaces are not affected. Only the global administrator sees two interfaces to represent a single monitored interface.

## Multi-tenancy and CA Unified Communications Monitor

CA Unified Communications Monitor supports the multi-tenancy features of NetOps Portal. You can perform all tenant and IP domain configuration in NetOps Portal. The collector then associates monitored items with the appropriate tenant by means of the IP domain.

In the CA Unified Communications Monitor management console, you can instruct collectors to associate the items that they discover with custom domains in CAPC. The act of creating a single custom domain in NetOps Portal enables domain associations for Locations, voice gateways, and call servers in any registered data sources.

Items appear with domain designations as soon as they are discovered from call traffic. Items discovered previously do not receive retroactive associations.

Locations are automatically associated with IP domains by the subnets that they contain. To preserve the flow of data collection and the appropriate association of data with IP domains, take care when moving Locations to new IP domains. Follow the procedure provided in the CA Unified Communications Monitor Online Help to change IP domain assignments.

**NOTE**

Managed items that are not associated with a custom IP domain by a data source are associated with the Default Domain. This assignment is transparent to users who are not deploying custom IP domains.

Instruct collectors to associate items with custom IP domains.

**Follow these steps:**

1. In the CA Unified Communications Monitor management console, click **Administration, Data Collection, Collectors**.
2. Edit each collector to select its domain for the IP Domain parameter.
3. Reload the collectors to send them the domain information.

Domains are populated with managed items after the next product synchronization.

## Multi-tenancy and DX NetOps Spectrum

When you register the DX NetOps Spectrum (Spectrum) data source, database synchronization occurs.

Spectrum retrieves a list of IP domains. DX NetOps Performance Management sends all IP domain definitions to Spectrum, regardless of their association with individual tenants. OneClick displays these NetOps Portal IP Domain models in the same area as Spectrum Global Collections.

The IP Domain models have the same names as the IP domain definitions. Use these IP domains to determine which models are synchronized with DX NetOps Performance Management. To include a device model in DX NetOps Performance Management monitoring and make it available in NetOps Portal dashboards, add it to an IP domain in OneClick. If you attempt to add a model to multiple IP domains, you see an error message. Spectrum devices can be members of only a single IPDomain model type. Add only device models that Spectrum should synchronize with NetOps Portal.

When Spectrum synchronizes the device models, they are associated with the corresponding IP domain in NetOps Portal. The NetOps Portal IP domain might belong to the Default Tenant or to any custom tenant. Device interfaces are automatically added to the IP domain with which their device is associated.

For more information about multi-tenancy when DX NetOps Performance Management is integrated with Spectrum, see [Integrate with DX NetOps Spectrum for Fault Management](#).

## NetOps Portal Administration

To maximize the value of reporting and to make DX NetOps Performance Management fit the needs of your business, customize NetOps Portal.

The articles in this section provide information about how to customize NetOps Portal.

### Back Up NetOps Portal

Back up NetOps Portal before an upgrade. To avoid losing valuable data, back up and archive NetOps Portal regularly.

You can create back ups of NetOps Portal configuration and event manager data.

Use the following process to back up NetOps Portal:

1. [Back Up the NetOps Portal Database](#)
2. (If you configured Single Sign-On (SSO)) [Back Up the Single Sign-On Configuration Files](#)
3. [Back Up the Custom UI Settings](#)
4. [Back Up the Reports](#)
5. (If you have script notification actions) [Back Up the Script Notification Actions](#)
6. (If NetOps Portal is HTTPS-enabled) [Back Up the HTTPS Files and Directories](#)

### **Back Up the NetOps Portal Database**

Create a backup archive of the current database anytime that you plan to reinstall or upgrade NetOps Portal.

#### **Follow these steps:**

1. Log in to the NetOps Portal host as the root user or as the sudo user.
2. Estimate the size of the backup by issuing the following command:

```
du -hs <installation_directory>/data
```

#### **Example:**

```
du -hs /opt/CA/MySQL/data
```

#### **– *installation\_directory***

Specifies the default installation directory for MySQL.

**Default:** /opt/CA/MySQL

3. Stop the NetOps Portal services by issuing the following commands in the order listed:

Stop the NetOps Portal Console service by issuing the following command:

```
systemctl stop caperfcenter_console
```

Stop the device manager by issuing the following command:

```
systemctl stop caperfcenter_devicemanager
```

Stop the event manager by issuing the following command:

```
systemctl stop caperfcenter_eventmanager
```

Stop the SSO service by issuing the following command:

```
systemctl stop caperfcenter_sso
```

4. Change to a directory where you want to save the database archive by issuing the following command:

```
cd /<backupDir>
```

#### **Example:**

```
cd /opt/CA/backup
```

#### **– *backupDir***

The location for the back up file.

5. Create a MySQL dump (backup files) of the following items:

#### **– The NetOps Portal configuration data (the `netqosportal` database)**

Issue the following command:

#### **NOTE**

Leave out the optional password syntax from the following command to be prompted for the password.

```
<installation_directory>/bin/mysqldump --routines -u root -p<password>  
netqosportal > <backupdir>/netqosportal.sql
```

#### **Example:**

```
/opt/CA/MySQL/bin/mysqldump --routines -u root -pnetqos  
netqosportal > /opt/CA/backup/netqosportal.sql
```

- **installation\_directory**  
Specifies the default installation directory for MySQL.  
**Default:** /opt/CA/MySQL

- **backupDir**  
The location for the back up file.

**NOTE**

During the installation, the `Select a Location for the MySQL Data Directory` prompt appears. The default directory is `<installation_directory>/data`.

- **installation\_directory**  
Specifies the default installation directory for MySQL.  
**Default:** /opt/CA/MySQL

– **The event manager data (the em database)**

Issue the following command:

```
<installation_directory>/bin/mysqldump -u root -p<password> em > <backupDir>/em.sql
```

**Example:**

```
/opt/CA/MySQL/bin/mysqldump -u root -pnetqos em > /opt/CA/backup/em.sql
```

- **installation\_directory**  
Specifies the default installation directory for MySQL.  
**Default:** /opt/CA/MySQL

- **backupDir**  
The location for the backup file.

6. Save space by compressing the backup files by issuing the following commands:

```
tar czvf netqosportal.tgz netqosportal.sql
tar czvf em.tgz em.sql
```

7. Remove the uncompressed MySQL dump files by issuing the following commands:

```
rm netqosportal.sql
rm em.sql
```

8. Start the NetOps Portal services on the system as follows:

- Start the SSO service by issuing the following command:  
`systemctl start caperfcenter_sso`
- Start the event manager by issuing the following command:  
`systemctl start caperfcenter_eventmanager`
- Wait one minute, then start the device manager by issuing the following commands  
`systemctl start caperfcenter_devicemanager`
- Wait one minute, then start the NetOps Portal Console service by issuing the following command:  
`systemctl start caperfcenter_console`

The NetOps Portal database is backed up.

### **Back Up the Single Sign-On Configuration Files**

If you have configured SSO, back up the configuration settings (the SSO configuration files). You can back them up automatically using `rsync` or another preferred method, such as a script. Back up the following directories:

- `<installation_directory>/sso/webapps/sso/configuration`
- `<installation_directory>/sso/etc`
- `<installation_directory>/sso/conf`
- `<installation_directory>/PC/etc`
- `<installation_directory>/PC/conf`

**NOTE**

***installation\_directory*** is the installation directory for NetOps Portal.

**Default:** `/opt/CA/PerformanceCenter`

**Back Up the Custom UI Settings****Complete the following steps:**

1. If your NetOps Portal deployment includes custom OpenAPI applications, they are located in the `<installation_directory>/PC/webapps/pc/apps` directory. Back up this directory.
2. If NetOps Portal uses custom UI themes that includes custom logos, back up the `customLogo` file in the following `<installation_directory>/images` directories:
  - `<installation_directory>/PC/webapps/pc/css/CA-Blue/images`
  - `<installation_directory>/PC/webapps/pc/css/CA-White/images`

**NOTE**

***installation\_directory*** is the default installation directory for NetOps Portal.

**Default:** `/opt/CA/PerformanceCenter`

The custom UI settings are backed up.

**Back Up the Reports**

NetOps Portal temporarily saves full PDF reports for download in the `<installation_directory>/DM/repository` directory. Back up this directory.

- ***installation\_directory***  
Specifies the default installation directory for NetOps Portal.  
**Default:** `/opt/CA/PerformanceCenter`

**Back Up the Script Notification Actions**

If you have script notification actions, back up the `<installation_directory>/NotificationScripts` directory.

- ***installation\_directory***  
Specifies the default installation directory for NetOps Portal.  
**Default:** `/opt/CA/PerformanceCenter`

**Back Up the HTTPS Files and Directories**

If you have enabled HTTPS for NetOps Portal, back up the following Java truststore file and the configuration files with the website and port settings:

- `<installation_directory>/PC/start.d/http.ini`
- `<installation_directory>/sso/start.d/http.ini`
- `<installation_directory>/PC/start.d/ssl.ini`
- `<installation_directory>/sso/start.d/ssl.ini`
- `<installation_directory>/jetty/etc/keystore`
- The Java truststore file directory

**NOTE**

- **installation\_directory** is the default installation directory for NetOps Portal.  
**Default:** /opt/CA/PerformanceCenter
- **truststore** is the Java truststore file directory.  
**Default:** /opt/CA/jre/lib/security/cacerts

## Manage the Themes for NetOps Portal

You can create and deploy custom user interface (UI) themes for NetOps Portal to use.

The UI theme primarily controls the appearance, colors, and logos that NetOps Portal uses in browser windows (the UI) and when generating PDF files. DX NetOps Performance Management includes the CA-Blue and CA-White themes, but you can define more site-specific UI themes for NetOps Portal to include. The CA-Blue theme is the default theme for NetOps Portal, and displays the menus at the top of the screen. The CA-White theme displays the menus on the left side of the screen.

**IMPORTANT**

- Carefully consider whether you require custom UI themes. Only experienced administrators should create and deploy them.
- Custom UI themes require that you reconfigure each time that you upgrade or install DX NetOps Performance Management.
- Custom UI theme directories are not preserved after an uninstall of DX NetOps Performance Management.

Use the following process to create and deploy a custom UI theme:

1. [Create a Custom UI Theme](#)
2. [Deploy the Theme](#)
3. [Apply the Theme to a Specific Tenant](#)

You can also [change the UI coloring, icon coloring, or the PDF font and coloring](#) and [upload and apply a unique logo to a theme, which applies it to PDF output and to PDF report headers](#).

### Create a Custom UI Theme

You can change the UI coloring, icon coloring, or the PDF font and coloring while creating a theme.

**Follow these steps:**

1. Go to the following location by issuing the following command:  
`<installation-directory>/PC/webapps/pc/css`
  - **installation\_directory**  
The installation directory for NetOps Portal.  
**Defaults:** /opt/CA/PerformanceCenter
2. Package one of the following standard UI theme directories using the following command:
  - **CA-White theme directory:** /opt/CA/PerformanceCenter/PC/webapps/pc/css/CA-White
  - **CA-Blue theme directory:** /opt/CA/PerformanceCenter/PC/webapps/pc/css/CA-Blue

```
cd /opt/CA/PerformanceCenter/PC/webapps/pc/css
zip -r /tmp/<ThemeName>.zip <ThemeName>
```

**Example:**

```
cd /opt/CA/PerformanceCenter/PC/webapps/pc/css
zip -r /tmp/CA-White.zip CA-White
```

- **ThemeName**  
The name of the standard theme.

**Example:** CA-White

3. Copy the ZIP file (for example, CA-White.zip ) to your local computer, and then unzip it.
4. Rename the extracted UI theme directory using the XXXXXXXX-Theme naming convention, where XXXXXXXX is an identifier that you define, for example, My-Custom-Theme . Exclude spaces or special characters in the name.
5. Do any of the following tasks:
  - **Change the coloring for the UI theme.** Complete the following steps:
    - a. Edit the `<ThemeDirectory>/includes/capc_theme.css` file, locate the page background color comment, and replace the colors that are specified.
      - **ThemeDirectory**  
The local directory for your custom theme.  
**Example:** My-Custom-Theme
      - TIP**  
This CSS file is minified. You can make this file more legible for editing by auto formatting it using pretty print, or by using an editor that does it for you.

For some colors, such as primary accent colors, you can specify them in the following ways:

      - **By absolute color values.** For example, the primary accent color in the CA-Blue theme is the #376180 absolute color value.
      - **By reference.** For example, the primary accent color in the CA-Blue theme is the `var(--mnrl-palette-slate-600)` reference.
    - Example:**  
To replace the background color with the color pink, replace the color #f7f7f7 with pink .
    - b. Save your changes.
      - TIP**  
You can reminify the CSS file by using a website such as <https://minifier.org> .
    - c. Copy and replace this file in your working directory.
  - **Change the icon coloring for a theme.** Go to the `<ThemeDirectory>/images` directory, and then edit the relevant files in a graphics editor.
    - **ThemeDirectory**  
The local directory for your custom theme.  
**Example:** My-Custom-Theme
    - Example:**  
To change the color of the icons to white, edit the `capc_icons24.svg` file.
  - **Change the PDF font and font coloring for a theme.** Edit the `<ThemeDirectory>/xsl/pdf/defaultPageContent.xml` file.
    - **ThemeDirectory**  
The local directory for your custom theme.  
**Example:** My-Custom-Theme

The custom UI theme is created.
6. Zip the modified UI theme directory that is on your local computer.

**File requirements:**

**NOTE**

- The maximum file size is 100 MB.
- The file must have a ZIP file extension.
- A top-level theme directory must exist in the ZIP file.
- The name of the theme directory must end with -Theme .

## **Deploy a Custom UI Theme**

Deploy the custom UI theme so that NetOps Portal can use it.



**NOTE**

UI theme deployment can hang when you use the Chrome browser.  
For more information about this known issue, see [Known Issues](#).

**Follow these steps:**

1. On the **Theme Settings** page, in the **Theme Deployment/Deploy a Theme** section, complete the following fields:
  - **Theme**  
To deploy the custom UI theme, click **Browse**, select the custom UI theme zip file that you want to deploy, and then click **Open**.
  - **Replace existing themes**  
To replace an existing deployed custom UI theme, select **Replace existing themes**.
2. Click **Add**.  
The custom UI theme is deployed.

**Apply the Theme to a Specific Tenant**

[Apply the theme to a specific tenant.](#)

**Update a UI Theme**

You can change the UI coloring, icon coloring, or the PDF font and coloring by updating the UI theme.

**Follow these steps:**

1. Go to the following location by issuing the following command:  
`<installation-directory>/PC/webapps/pc/css`
  - **installation\_directory**  
The installation directory for NetOps Portal.  
**Defaults:** /opt/CA/PerformanceCenter
2. Package the custom UI theme directory that you want to update using the following command:  
`cd /opt/CA/PerformanceCenter/PC/webapps/pc/css`  
`zip -r /tmp/<ThemeName>.zip <ThemeName>`

**Example:**

```
cd /opt/CA/PerformanceCenter/PC/webapps/pc/css
zip -r /tmp/My-Custom-Theme.zip My-Custom-Theme
```

– **ThemeName**

The name of the custom theme.

**Example:** My-Custom-Theme

3. Copy the ZIP file (for example, My-Custom-Theme.zip ) to your local computer, and then unzip it.
4. Do any of the following tasks:
  - To change the coloring for the UI theme, complete the following steps:
    - a. Open the `<ThemeDirectory>/includes/capc_theme.css` file.
      - **ThemeDirectory**  
The local directory for your custom theme.  
**Example:** My-Custom-Theme
    - b. Edit the file. For example, to change the background color to pink, locate the comment `page background color` , and then replace the color `#f7f7f7` with `pink` .  
**TIP**  
The CSS file is minified. You can make this file more legible for editing by auto formatting it using pretty print, or by using an editor that does it for you.
    - c. Save your changes.

**TIP**

You can minify the file by using a website such as <https://minifier.org>.

d. Copy and replace this file in your working directory.

- To change the icon coloring for a theme, go to the `<ThemeDirectory>/images` directory, and then edit the relevant files in a graphics editor.

- **ThemeDirectory**

The local directory for your custom theme.

**Example:** `My-Custom-Theme`

**Example:**

To change all icons to white, edit the `capc_icons24.svg` file.

- To change the PDF font and font coloring for a theme, go to the `<ThemeDirectory>/xsl/pdf` directory and edit the `defaultPageContent.xml` file.

- **ThemeDirectory**

The local directory for your custom theme.

**Example:** `My-Custom-Theme`

The custom UI theme is updated.

5. Zip the modified custom UI theme directory that is on your local computer.

**File requirements:**

**NOTE**

- The maximum file size is 100 MB.
- The file must have a ZIP file extension.
- A top-level theme directory must exist in the ZIP file.
- The name of the theme directory must end with `-Theme`.

## **Apply a Custom Logo to a UI Theme**

You can upload and apply a custom logo to a theme, which applies it to PDF output and to PDF report headers.

For more information about scheduled reports, see [Manage Scheduled Reports](#).

**Follow these steps:**

1. Log in as a user with the Administrator role.
2. Hover over **Administration**, **Group Settings**, and then click **Themes**.  
The **Theme Settings** page opens.
3. In the **Theme Settings/Choose a Different Image File for Logos in PDF files (Per Theme)** section, complete the following fields:
  - **Logo image file**  
To upload a unique logo, click **Browse**, select the image file for your logo, and then click **Open**.
  - **Apply to theme**  
If you have chosen to upload a unique logo, to apply the logo to a UI theme, select the theme from the drop-down list. If you do not deploy multi-tenancy, leave **All Themes** selected.
4. Click one of the following:
  - **Reset Images**  
To restore the default theme settings, click **Reset Images**.
  - **Upload Image**  
If you have chosen to upload a unique logo, click **Upload Image**.
5. **Next Step:** [Edit the tenant to select the theme that you modified](#).

The logo is applied to PDF output.

## Update the NetOps Portal IP Address and Hostname

When you change the IP address or hostname for the server on which NetOps Portal is installed, update the NetOps Portal IP address and hostname.

Update the NetOps Portal IP address or hostname if the following changes occur to the server on which the NetOps Portal is installed:

- The server is assigned a new IP address or hostname.
- The server has moved to a new subnet.

Use the following process to update the NetOps Portal IP address and hostname:

1. [Stop the NetOps Portal services.](#)
2. [Update the hosts file.](#)
3. [Start the NetOps Portal services.](#)
4. [Update the event manager data source IP address or hostname.](#)
5. [Validate the NetOps Portal data source IP address or hostname.](#)
6. [Review the other considerations.](#)

### **Stop the NetOps Portal Services**

[Stop the NetOps Portal services.](#)

### **Update the Hosts File**

As needed, open the `/etc/hosts` file on the NetOps Portal server, and then edit the NetOps Portal IP address or hostname with the new NetOps Portal IP address or hostname.

### **Start the NetOps Portal Services**

[Start the NetOps Portal services.](#)

### **Update the Event Manager Data Source IP Address or Hostname**

**Follow these steps:**

1. In NetOps Portal, [update the event manager data source with the new NetOps Portal IP address or hostname \(the Host Name field\).](#)
2. On the event manager data source, set the values for the `LastEvent` and `ConsumerID` attributes in the `data_sources` table to 0 by issuing the following command from the `/opt/CA/MySQL/bin` directory, and then enter the password when prompted:

```
./mysql -unetqos -p -e "update em.data_sources set LastEvent=0,ConsumerID=0;"
```

The `em` event manager MySQL database includes the `data_sources` and `general` tables.

The event manager data source IP address or hostname is updated.

The `netqosportal` MySQL database includes the `data_sources2` and `performance_center_properties` tables. The data sources known to the event manager gather new details for the hosts from which they gather information. Prior NetOps Portal updates to the entries in the `data_sources2` and `performance_center_properties` tables drive the updated values.

During the global and data source synchronization cycles, NetOps Portal updates the NetOps Portal `Host` and `ConsoleHost` attributes in the `data_sources2` table with the new IP address.

**NOTE**

NetOps Portal stores its IP address in these attributes, and sets its initial entry using the default priority 0 (zero) value. NetOps Portal sets these values during the initial installation.

Initial synchronization updates the default priority 0 (zero) values for the `NpcWebSiteHost` and `NpcWebServiceHost` attributes in the `performance_center_properties` table to the NetOps Portal hostname if it can determine it, otherwise it updates the values to the first IPv4 non-localhost IP address that the system provides to the synchronization code. Initial synchronization also updates the NetOps Portal device item name with the new IP address.

During the event manager inventory synchronization cycle, NetOps Portal uses the highest priority value for the `NpcWebSiteHost` attribute to set the value for the `ConsoleHost` attribute in the `data_sources2` table, and uses the highest priority value for the `NpcWebServiceHost` attribute to set the value for the `Host` attribute in the `data_sources2` table for the NetOps Portal data source. This update triggers NetOps Portal to update the higher priority (0 or 1) values for those attributes in the `data_sources` table.

**NOTE**

NetOps Portal does not synchronize the priority 2 value with other data sources.

If the entries in the `data_sources2` or `data_sources` tables use a hostname, and the DNS and/or the `/etc/hosts` file on the NetOps Portal server has the new NetOps Portal IP address, the hostname resolves to the correct IP address. No change is required.

**NOTE**

By default, the **Web Service Host** and **Web Site Host** properties are set with the IP address, but they could have been overridden using the Single Sign-On Configuration (SsoConfig) tool, where the **Remote Value** option was used to set the IP address or hostname for the Event Manager service with a different value.

1. Check the values for the following properties:
  - **Web Service Host**  
For more information about this property, see [Configure the Web Service Security Settings Using the SSO Configuration Tool](#).
  - **Web Site Host**  
For more information about this property, see [Configure the Basic Security Settings Using the SSO Configuration Tool](#).
2. If either of the **Remote Value** or **Local Override** options show an IP address set for these properties, follow the prompts to update the value to use the new NetOps Portal IP address.

### **Validate the NetOps Portal Data Source IP Address or Hostname**

#### **Follow these steps:**

1. Ensure that the event manager data source is set with an IP address by issuing the following query from the `/opt/CA/MySQL/bin` directory, and then enter the password when prompted:

```
./mysql -unetqos -p -e "select <ConsoleName>,ConsoleHost,Host from data_sources2"
```

- **ConsoleName**

The console name.

**Example:** `%EventManager%`;

#### **Sample output when set with an IP address:**

```
[root@lvntest002828 bin]# ./mysql -unetqos -p -e "select ConsoleName,ConsoleHost,Host
from netqosportal.data_sources2 where ConsoleName like '%EventManager%';" Enter password:
enteredPasswordIsNotShownWhileTyped
+-----+-----+-----+ |
ConsoleName | ConsoleHost | Host |
+-----+-----+-----+ |
EventManager@1.1.1.1 | 1.1.1.1 | 1.1.1.1 |
```

+-----+-----+-----+

2. If the output shows that the event manager data source is set with an IP address, update the event manager data source.

For more information about how to update the event manager data source, see [Migrate NetOps Portal](#).

The event manager updates the general table. During the next inventory synchronization cycle for the event manager data source, the event manager re-registers with the other data sources, and then sends the other data sources a new event manager URL value.

The data sources now have the URL with the new IP address for use when requesting events from the event manager.

The NetOps Portal data source IP address or hostname is validated.

### Other Considerations

If the NetOps Portal server is configured with dual network interface controller (NIC) cards, ensure that the server is using the new IP address. Often in these situations, only one IP address is the right one to use for successful NetOps Portal operation. NetOps Portal cannot guarantee that the IP address it should use is the one that it will select.

You can override the IP address that NetOps Portal uses (the priority value 0 (zero) for the `NpcWebSiteHost` and `NpcWebServiceHost` attributes in the `performance_center_properties` table) using `SsoConfig`.

For more information:

- About how NetOps Portal determines the IP address of its host, see [the How does Performance Center determine the IP address of its host KB](#).
- About how to override the IP address (the **Web Site Host** property), see [Configure the Basic Security Settings Using the SSO Configuration Tool](#).

## Move the NetOps Portal Database to a Separate Node

Complete these steps to use a separate MySQL database.

### NOTE

In the following steps, *original node* refers to the node which has the single node deployment of NetOps Portal. After you complete these steps, this node hosts only the NetOps Portal services.

*Database node* refers to the new node that will host the externalized database.

### Follow these steps:

1. Upgrade to the latest DX NetOps Performance Management release by running the NetOps Portal installation on the original node.

For more information, see [Upgrading](#).

2. Run the NetOps Portal installation on the database node. Select **Advanced** as the installation type, and then install the Database feature.

For more information, see [Install NetOps Portal](#).

3. Stop the NetOps Portal services on the original node by issuing the following commands in the order listed:

Stop the NetOps Portal Console service:

```
systemctl stop caperfcenter_console
```

Stop the device manager:

```
systemctl stop caperfcenter_devicemanager
```

Stop the event manager:

```
systemctl stop caperfcenter_eventmanager
```

Stop the SSO service:

```
systemctl stop caperfcenter_sso
```

4. Dump the NetOps Portal `netqosportal` and `em` databases on the original node by issuing the following commands:
 

```
<installation_directory>/MySQL/bin/mysqldump --routines -u root -p netqos --databases netqosportal > <backupdir>/netqosportal.sql
<installation_directory>/MySQL/bin/mysqldump --routines -u root -p netqos --databases em > <backupdir>/em.sql
```

**Example:**

```
/opt/CA/MySQL/bin/mysqldump --routines -u root -p netqos --databases netqosportal > backupdir/netqosportal.sql
/opt/CA/MySQL/bin/mysqldump --routines -u root -p netqos --databases em > backupdir/em.sql
```

– **installation\_directory**

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

5. Transfer the dumped database SQL files to a backup directory, such as `backupdir`, on the database node. Import the dumped files into the database by issuing the following commands:

```
mysql -unetqos -p netqos -e 'source <backupdir>/netqosportal.sql'
```

```
mysql -unetqos -pnetqos -e 'source <backupdir>/em.sql'
```

6. Remove the database from the registry and from the `<installation_directory>/PerformanceCenter/CAPC_Database_SeedFile.xml` database seed file by running the `<installation_directory>/PerformanceCenter/Tools/bin/uninstall_database_component.sh` script on the original node.

– **installation\_directory**

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

7. Re-run the NetOps Portal installation on the original node.

For more information, see [Install NetOps Portal](#).

When prompted, provide the hostname and port (3306) of the externalized database on the database node.

8. Verify that you can now access NetOps Portal through port 8181 on the original node.

9. Uninstall MySQL from the original node by running the following script:

```
<installation_directory>/PerformanceCenter/Uninstall_MySql
```

– **installation\_directory**

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

**NOTE**

This procedure does not remove the stored data. To remove the stored data, issue the following command:

```
rm -rf /opt/CA/MySQL
```

## Modify Maximum Memory Usage for NetOps Portal

Modify the maximum memory usage for the NetOps Portal daemons to let the daemons run effectively.

You can configure the maximum amount of memory for a NetOps Portal deployment of your scalability requirements.

Configure memory allocation during or after the installation.

**Follow these steps:**

1. On the server where NetOps Portal is installed, issue the following command:

```
more /proc/meminfo
```

The total memory usage of the server is displayed.

2. Make a note of the total memory.
3. For each NetOps Portal daemon, modify the maximum memory as follows:
  - a. Edit the `<installation_directory>/PerformanceCenter/<service_subdirectory>/conf/wrapper.conf` file.
    - **installation\_directory**  
The default installation directory for NetOps Portal.  
**Default:** `/opt/CA`
    - **service\_subdirectory**  
The NetOps Portal service that you want configure memory allocation.  
**Options:**
      - PC (Console daemon)
      - DM (Device Manager daemon)
      - EM (Event Manager daemon)
  - b. Search for the `wrapper.java.maxmemory` parameter.
  - c. Adjust the setting, and then save the file:  
Use the following scale values:
    - **Small:** 3072 MB
    - **Units:** MB
4. [Stop and restart each daemon](#) (except the SSO service).

The maximum amount of memory is configured for a deployment of your scalability requirements.

## Restore NetOps Portal

Restore an existing backup of NetOps Portal.

Use the following process to restore NetOps Portal:

1. Do one of the following:
  - After a reinstallation, migration, or a disaster recovery, [restore the database](#).
  - If an error occurs during an upgrade of NetOps Portal, [recover from an upgrade failure](#).
2. (If your NetOps Portal deployment uses Single Sign-On (SSO)) [Restore the Single Sign-On Settings](#)
3. (If your NetOps Portal deployment includes custom OpenAPI applications) [Restore the Custom Settings](#)
4. [Restore Reports](#)
5. (If you have script notification actions) [Restore the script notification actions](#)
6. (If the new NetOps Portal hostname is different from the original NetOps Portal host) [Update the NetOps Portal disaster recovery script](#)
7. [Run the NetOps Portal disaster recovery script](#)
8. [Start NetOps Portal](#)
9. [Resynchronize the event manager database](#)

### Restore the Database

**Prerequisite:** You have reinstalled NetOps Portal on the same machine, on a migration machine, or on a disaster recovery machine.

Restoring the database from a backup preserves data continuity and enables most historical reporting after a failure occurs.

**IMPORTANT**

Restore the database *only* if a failure occurs. Before attempting to upgrade again, clean up after a failed installation, then complete this procedure.

For more information about how to clean up after a failed installation, see [Uninstall NetOps Portal](#).

**Follow these steps:**

1. Log in to the server as 'root', or use the 'sudo' account that you configured for the installation.
2. Stop the NetOps Portal services by issuing the following commands in the order listed:  
 Stop the event manager:  

```
systemctl stop caperfcenter_eventmanager
```

 Stop the device manager:  

```
systemctl stop caperfcenter_devicemanager
```

 Stop the SSO service:  

```
systemctl stop caperfcenter_sso
```

 Stop the NetOps Portal Console service:  

```
systemctl stop caperfcenter_console
```
3. Change to the directory where you saved the backup archive by issuing the following command:  

```
cd <backupDir>
```

 – **backupDir**  
 The location for the back up file.
4. Uncompress the database backup archives for NetOps Portal and the event manager by issuing the following commands:  

```
tar zxvf netqosportal.tgz
```

```
tar zxvf em.tgz
```
5. Import the uncompressed NetOps Portal backup file by issuing the following command:  

```
mysql netqosportal -u root -p<password> -e 'source $backupDir/netqosportal.sql'
```
6. Import the uncompressed Event Manager backup file by issuing the following command:  

```
mysql em -u root -p<password> -e 'source <backupDir>/em.sql'
```

 – **backupDir**  
 The location for the back up file.
7. Start the NetOps Portal services on the new system as follows by completing the following steps:
  - a. Start the SSO service by issuing the following command:  

```
systemctl start caperfcenter_sso
```
  - b. Start the event manager by issuing the following command:  

```
systemctl start caperfcenter_eventmanager
```
  - c. Wait one minute, then start the device manager by issuing the following commands:  

```
systemctl start caperfcenter_devicemanager
```
  - d. Wait one minute, then start the NetOps Portal Console service by issuing the following command:  

```
systemctl start caperfcenter_console
```
8. Delete the uncompressed archive files to save space by issuing the following command:  

```
rm netqosportal.sql
```

```
rm em.sql
```
9. Log in to NetOps Portal as an Administrator.



10. Verify that your configuration data appears in the Administration pages.

The NetOps Portal database is restored.

### **Recover from an Upgrade Failure**

Complete this procedure before attempting to upgrade again after a failed upgrade. This procedure restores the NetOps Portal database and verifies the previous database schema version.

#### **Follow these steps:**

1. Log in to the server as 'root', or use the 'sudo' account that you configured for the installation.

2. Stop the NetOps Portal services by issuing the following commands in the following order:

a. Stop the event manager:

```
systemctl stop caperfcenter_eventmanager
```

b. Stop the device manager:

```
systemctl stop caperfcenter_devicemanager
```

c. Stop the SSO service:

```
systemctl stop caperfcenter_sso
```

d. Stop the NetOps Portal Console service:

```
systemctl stop caperfcenter_console
```

3. Change to the directory where you saved the backup archive by issuing the following command:

```
cd <backupDir>
```

– **backupDir**

The location for the back up file.

4. Uncompress the database backup archives for NetOps Portal and the event manager by issuing the following commands:

```
tar zxvf netqosportal.tgz
```

```
tar zxvf em.tgz
```

5. Import the uncompressed NetOps Portal backup file by issuing the following command:

#### **NOTE**

Leaving out the optional password syntax from the following command prompts you for the password.

```
mysql netqosportal -u root -p<password> -e 'source <backupDir>/netqosportal.sql'
```

– **backupDir**

The location for the back up file.

6. Import the uncompressed event manager backup file by issuing the following command:

```
mysql em -u root -p<password> -e 'source <backupDir>/em.sql'
```

– **backupDir**

The location for the back up file.

7. Verify the database version by issuing the following command:

```
mysql -P3306 -D netqosportal -u root -p<password>
```

```
mysql> select InstallDate, version, dbschemaversion from revision_info order by  
InstallDate asc;
```

The output lists installation dates and versions of the NetOps Portal and database schema.

The latest database version must match the previous installed DX NetOps Performance Management version.

8. Upgrade the database schema to the installed DX NetOps Performance Management version by running the `npcshell.sh` database utility:

a. Change to the following installation directory by issuing the following command:

```
cd <installation_directory>/PerformanceCenter/Tools/bin
```

- **installation\_directory**  
The location for the back up file.  
**Default:** /opt/CA

b. Run the `npcshell.sh` database utility by issuing the following command:

```
./npcshell.sh upgradedb
```

9. Import the database translation files by issuing the following commands:

```
<installation_directory>/jre/bin/java -jar <installation_directory>/
PerformanceCenter/SQL/seedlu/bin/seedlu.jar -resfile "<installation_directory>/
PerformanceCenter/SQL/messages_en_US.properties" -ctrlfile "<installation_directory>/
PerformanceCenter/SQL/control.sdlctrl" -connection "jdbc:mysql://localhost:3306/
netqosportal?useUnicode=true&characterEncoding=UTF-8" -user netqos -pwd password -
lang en-US<installation_directory>/jre/bin/java -jar
<installation_directory>/PerformanceCenter/SQL/seedlu/bin/seedlu.jar -resfile
"<installation_directory>/PerformanceCenter/SQL/messages_fr_FR.properties" -ctrlfile
"<installation_directory>/PerformanceCenter/SQL/control.sdlctrl" -connection
"jdbc:mysql://localhost:3306/netqosportal?useUnicode=true&characterEncoding=UTF-8" -
user netqos -pwd password -lang fr-FR<installation_directory>/jre/bin/java -jar
<installation_directory>/PerformanceCenter/SQL/seedlu/bin/seedlu.jar -resfile
"<installation_directory>/PerformanceCenter/SQL/messages_ja_JP.properties" -ctrlfile
"<installation_directory>/PerformanceCenter/SQL/control.sdlctrl" -connection
"jdbc:mysql://localhost:3306/netqosportal?useUnicode=true&characterEncoding=UTF-8" -
user netqos -pwd password -lang ja-JP
```

#### Example:

```
/opt/CA/jre/bin/java -jar /opt/CA/PerformanceCenter/SQL/seedlu/bin/seedlu.jar -
resfile "/opt/CA/PerformanceCenter/SQL/messages_en_US.properties" -ctrlfile "/opt/
CA/PerformanceCenter/SQL/control.sdlctrl" -connection "jdbc:mysql://localhost:3306/
netqosportal?useUnicode=true&characterEncoding=UTF-8" -user netqos -pwd password -
lang en-US
/opt/CA/jre/bin/java -jar /opt/CA/PerformanceCenter/SQL/seedlu/bin/seedlu.jar -
resfile "/opt/CA/PerformanceCenter/SQL/messages_fr_FR.properties" -ctrlfile "/opt/
CA/PerformanceCenter/SQL/control.sdlctrl" -connection "jdbc:mysql://localhost:3306/
netqosportal?useUnicode=true&characterEncoding=UTF-8" -user netqos -pwd password -
lang fr-FR
/opt/CA/jre/bin/java -jar /opt/CA/PerformanceCenter/SQL/seedlu/bin/seedlu.jar -
resfile "/opt/CA/PerformanceCenter/SQL/messages_ja_JP.properties" -ctrlfile "/opt/
CA/PerformanceCenter/SQL/control.sdlctrl" -connection "jdbc:mysql://localhost:3306/
netqosportal?useUnicode=true&characterEncoding=UTF-8" -user netqos -pwd password -
lang ja-JP
```

Replace the `password` variable with the password.

- **installation\_directory**  
The installation directory for NetOps Portal.  
**Default:** /opt/CA

10. Update the information that NetOps Portal uses for dashboards and views by issuing the following commands:

- **Infrastructure management administration pages:**  
`./npcshell.sh dbmigrate -package com.ca.im.plugin.pc -path ../../SQL/plugins/pc/`

– **Event-related views:**

```
./npcshell.sh dbmigrate -package com.ca.im.plugin.em -path ../../SQL/plugins/
eventmanager/
```

– **Data aggregator administration pages and views:**

```
./npcshell.sh dbmigrate -package com.ca.im.plugin.da -path ../../SQL/plugins/
polaris/
```

The database schema is upgraded.

11. Verify that the database version matches the installed DX NetOps Performance Management version by issuing the following command:

```
mysql -P3306 -D netqosportal -u root -p<password>
mysql> select InstallDate, version, dbschemaversion from revision_info order by
InstallDate asc;
```

12. Start the NetOps Portal services by issuing the following commands in the order listed:

a. Start the SSO service:

```
systemctl start caperfcenter_sso
```

b. Start the event manager:

```
systemctl start caperfcenter_eventmanager
```

c. Wait one minute, and then start the device manager:

```
systemctl start caperfcenter_devicemanager
```

d. Wait one minute, and then start the NetOps Portal Console service:

```
systemctl start caperfcenter_console
```

13. Delete the uncompressed archive files to save space by issuing the following commands:

```
rm netqosportal.sql
rm em.sql
```

14. Log in to NetOps Portal as an administrator.

15. Verify that you can see inventory and report data.

You have recovered from an upgrade failure.

## **Restore the Single Sign-On Settings**

If your NetOps Portal deployment uses SSO, restore the SSO configuration settings.

### **Follow these steps:**

1. Restore the following files:

- <installation\_directory>/PerformanceCenter/sso/start.ini
- <installation\_directory>/PerformanceCenter/PC/start.ini

**NOTE**

**installation\_directory** is the installation directory for NetOps Portal.

**Default:** /opt/CA

2. Restore the follow directories:

- <installation\_directory>/PerformanceCenter/sso/webapps/sso/configuration
- <installation\_directory>/PerformanceCenter/sso/etc
- <installation\_directory>/PerformanceCenter/sso/conf
- <installation\_directory>/PerformanceCenter/PC/etc
- <installation\_directory>/PerformanceCenter/PC/conf

**NOTE**

**installation\_directory** is the installation directory for NetOps Portal.

**Default:** /opt/CA

3. (If you have enabled HTTPS for NetOps Portal) Restore the following files:

- <installation\_directory>/PerformanceCenter/sso/start.d/ssl.ini
- <installation\_directory>/PerformanceCenter/PC/start.d/ssl.ini

**NOTE**

**installation\_directory** is the installation directory for NetOps Portal.

**Default:** /opt/CA

The SSO settings are restored.

### **Restore the Custom Settings**

If your deployment includes custom OpenAPI applications, restore the custom settings.

**Follow these steps:**

1. Restore the <installation\_directory>/PerformanceCenter/PC/webapps/pc/apps directory.

- **installation\_directory**

The installation directory for NetOps Portal.

**Default:** /opt/CA

2. (If you use custom logos for your themes) Restore the following files:

- <installation\_directory>/PerformanceCenter/PC/webapps/pc/css/CA-Blue/images/customLogo
- <installation\_directory>/PerformanceCenter/PC/webapps/pc/css/CA-White/images/customLogo

**NOTE**

**installation\_directory** is the installation directory for NetOps Portal.

**Default:** /opt/CA

The custom settings are restored.

### **Restore Reports**

Full PDF reports are temporarily saved for download. Restore the <installation\_directory>/PerformanceCenter/DM/repository directory.

- **installation\_directory**

The installation directory for NetOps Portal.

**Default:** /opt/CA

### **Restore the Script Notification Actions**

If you have script notification actions, restore the <installation\_directory>/PerformanceCenter/NotificationScripts directory.

- **installation\_directory**

The installation directory for NetOps Portal.

**Default:** /opt/CA

### **Update the NetOps Portal Disaster Recovery Script**

If the new NetOps Portal hostname is different from the original NetOps Portal host, update the data aggregator and the event manager data sources.

**Follow these steps:**

1. Open the `<installation_directory>/PerformanceCenter/Tools/bin/update_pc_da_database_references.sh` NetOps Portal disaster recovery script.
  - **installation\_directory**  
The installation directory for NetOps Portal.  
**Default:** /opt/CA
2. On the NetOps Portal host in the recovery system, update the bold sections of the script to match your system:

```
...
#####
# UPDATE THE FOLLOWING PC/DA VARIABLES TO REFLECT NEW ENVIRONMENT
#####
NEW_PC_IP_ADDRESS="<Recovery/New PC IP Address>"
NEW_PC_HOSTNAME="<Recovery/New PC Hostname>"
NEW_PC_EVENT_PRODUCER_PORT=8181
NEW_PC_EVENT_PRODUCER_PROTOCOL="http" # change to "https" if using SSL
NEW_DA_IP_ADDRESS="<Recovery/New DA IP Address>"
NEW_DA_HOSTNAME="<Recovery/New DA Hostname>"
NEW_DA_PORT_NUMBER=8581
...
```

The data aggregator and event manager data sources are updated.

### **Run the NetOps Portal Disaster Recovery Script**

Run the `<installation_directory>/PerformanceCenter/Tools/bin/update_pc_da_database_references.sh` disaster recovery script.

- **installation\_directory**  
The installation directory for NetOps Portal.  
**Default:** /opt/CA

### **Start NetOps Portal**

**Follow these steps:**

1. Start the SSO service by issuing the following command:  
`systemctl start caperfcenter_sso`
2. Start the event manager by issuing the following command:  
`systemctl start caperfcenter_eventmanager`
3. Wait one minute, then start the device manager by issuing the following commands:  
`systemctl start caperfcenter_devicemanager`

- Wait one minute, then start the NetOps Portal Console service by issuing the following command:

```
systemctl start caperfcenter_console
```

NetOps Portal is restarted.

### **Resynchronize the Event Manager Database**

The backup and restore procedures include steps to back up the event manager database. To prevent problems from occurring when the event manager attempts to synchronize with NetOps Portal, resynchronize the event manager. The synchronization can fail because NetOps Portal has outdated event manager information. The newly-installed event manager includes the new information. If this problem occurs, resynchronize these two databases.

#### **NOTE**

This procedure does not include a step to restore the event manager database. Notifications are not preserved. You must recreate them. Otherwise, the event manager runs normally.

#### **Follow these steps:**

- [Delete the event manager data source.](#)
- Remove the properties that are related to NetOps Portal from the `em.general` database table by issuing the following command:  

```
DELETE from em.general where Attribute LIKE 'NPC.%';
```
- Restart the event manager by issuing the following command:  

```
systemctl restart caperfcenter_eventmanager
```
- [Register the event manager data source.](#)

The event manager database is resynchronized.

## **Configure the Email Server**

You can allow users to send reports by email, either manually or on a schedule, or to archive reports to the report repository by configuring an email server. The email server can send dashboard data as scheduled reports, and can archive them.

### **Configure an Email Server**

Configure an email server so that users can send reports by email, either manually or on a schedule. NetOps Portal sends reports on a schedule or as needed. Select a server to which the NetOps Portal server has network access.

#### **NOTE**

You can generate full reports without an email server.

#### **Follow these steps:**

- Log in as a user with administrative role rights.
  - Hover over **Administration, Configuration Settings**, and then click **Email Server**.  
The **Manage Email Server Settings** page opens.
  - Select the **Enable Email** check box.  
The page refreshes to highlight the required fields.
  - Complete the following fields as necessary:
    - **Server Type**  
Defines how the email server sends and saves dashboard data as scheduled reports.
- Options:**

- **Use SMTP Mail Server Only:** The email server sends dashboard data as reports from scheduled reports attached to emails. You can schedule to send these reports attached to emails, or you can download the reports and send them on your own.

For more information about how to set up schedules, see [Manage Scheduled Reports](#).

- **Use Repository Service Only:** NetOps Portal saves dashboard data as scheduled reports (enables archive reports). When the scheduled report runs, NetOps Portal saves the results to disk. You can view the results as archived reports.

For more information about how to view a list of archived reports for a selected scheduled report, see [Manage Scheduled Reports](#).

#### **IMPORTANT**

Choosing this option can affect the disk space of your report repository. Closely monitor the disk space over time.

For more information, see [View System Status](#).

- **Use Both SMTP Mail Server and Repository Service:** The email server sends dashboard data as reports from scheduled reports attached to emails *and* NetOps Portal saves the data as scheduled reports (enables archive reports). When the scheduled report runs, NetOps Portal saves the results to disk. You can view the results as archived reports.

#### **IMPORTANT**

Choosing this option can affect the disk space of your report repository. Closely monitor the disk space over time.

For more information, see [View System Status](#).

**Default:** Use SMTP Mail Server Only

#### – **SMTP Server Address**

The IP address or hostname of the server to use to send reports by email.

#### – **SMTP Server Port**

The port on the email server that is used to send messages.

**Default:** Port 25

For more information about the ports that DX NetOps Performance Management requires, see [Review Installation Requirements and Considerations](#).

#### – **Email Reply Address**

The email address from which NetOps Portal sends reports. An administrator monitors this address for responses to email messages that NetOps Portal sends.

5. If you want to send email using a secure connection from NetOps Portal, enable SSL encryption.

#### **NOTE**

To enable SSL encryption, you must enable HTTPS for NetOps Portal.

For more information, see [Set up HTTPS](#).

If the email server supports the StartTLS email protocol command, NetOps Portal automatically tries to send all emails securely.

6. (Optional) Enable Simple Mail Transfer Protocol (SMTP) authentication by completing the following steps:

#### **NOTE**

SMTP authentication is disabled by default.

- a. Select **Enable Authentication**.
- b. Enter the username for SMTP authentication in the **Username** field.
- c. Enter the authentication password in the **Password** field.
- d. Enter the authentication password again in the **Confirm Password** field.

7. Click **Save**.

The email server is configured.

## Configure the Email Server as a Trusted Connection

Configure the email server as a trusted connection by importing the email server certificate.

### NOTE

Store the certificate and private key files, such as \*.pem, \*.cer, \*.crt, \*.key files, that are referenced in configuration files during this process in a secure location. If the certificate and private key files are temporary files that are not referenced in configuration files after this process is complete, move or delete them.

### Follow these steps:

1. Import the email server certificate into the Java trusted certificate keystore by issuing the following command:

```
keytool -importcert -keystore <installation_directory>/jre/lib/security/cacerts -
storepass cacertspassword -alias alias_name -file filename.cer
```

- **installation\_directory**

Specify the Java truststore installation directory.

**Default:** /opt/CA

- **cacertspassword**

Specify the password for the Certificate Authority (CA) keystore.

### NOTE

The default password for the CA keystore is **changeit**.

- **alias\_name**

Specify an alias that can be used to refer to the keystore entry that is created for the email server certificate.

- **filename.cer**

Specify the file to which the certificate is exported. Use a full pathname that does not place the file in the current directory.

**Example:** /tmp/capcCert.cer

2. Import the root and intermediate certificates into the Java trusted certificate keystore for each certificate by issuing the following command:

```
keytool -importcert -keystore <installation_directory>/jre/lib/security/cacerts -
storepass <cacertspassword> -alias alias_name -file <filename>.cer
```

- **installation\_directory**

Specify the Java truststore installation directory.

**Default:** /opt/CA

- **cacertspassword**

Specify the password for the CA keystore.

### NOTE

The default password for the CA keystore is **changeit**.

- **alias\_name**

Specify an alias that can be used to refer to the keystore entry that is created for the root or intermediate certificate.

- **filename.cer**

Specify the file to which the certificate is exported. Use a full pathname that does not place the file in the current directory.

**Example:** /tmp/capcCert.cer

3. Confirm that you trust the certificate.

4. Verify that the imported keystore is available by issuing the following command:

```
keytool -list -keystore <installation_directory>/jre/lib/security/cacerts
```

- **installation\_directory**

Specify the Java truststore installation directory.

**Default:** /opt/CA

The email server is now a trusted connection.



## Configure Users for Internal Communications

You can configure which user NetOps Portal uses for internal communications.

You can configure users for internal communications within NetOps Portal among the DX NetOps Performance Management components (NetOps Portal, Event Manager, Device Manager, and the Data Aggregator). Out-of-the-box, these components make REST calls to each other as the predefined admin user with the user ID of 1. You can assign a dedicated user for these internal communications.

You can configure which user NetOps Portal uses for internal communications:

- [Configure a user for internal communications](#)

Configure this user in the following cases:

- A username is logged in various logs; You can have the logs distinguish between the user for *internal communications* and an actual NetOps Portal user account.
- You have security requirements that all users be in the Lightweight Directory Access Protocol (LDAP) system, and the predefined admin user cannot use an LDAP password. You can configure NetOps Portal to use a different user for *internal communications*, and then disable the admin user.

For more information about this user, see [Manage Roles and User Accounts](#).

- [Configure a user for user administration](#)

Configure this user if you have security requirements that require a dedicated user that performs *user administration*.

### Configure the User for Internal Communications

**Prerequisite:** The user that you want to configure for *internal communications* exists, and is assigned to the Administrator role. Administrators perform all administrative tasks in DX NetOps.

For more information about this role, see [Manage Roles and User Accounts](#).

#### Follow these steps:

1. Retrieve the details of the user account that you want to configure for *internal communications* by issuing a GET request to the users endpoint for the NetOps Portal REST web service:

```
GET http://<PC_host>:8181/pc/center/webservice/users/userName/<user>
```

#### Example:

The following request retrieves the details for the bob user account.

```
GET http://<PC_host>:8181/pc/center/webservice/users/userName/bob
```

#### – **user**

The name of the user account that you want to configure for *internal communications*.

#### Example output:

```
<user>
  <userId>456</userId>
  <permissionId>1</permissionId>
  <defaultGroupId>1</defaultGroupId>
  <name>bob</name>
```

.....

2. Launch the SSO Configuration tool (SsoConfig) by running the `./SsoConfig` command in the `<installation_directory>/PerformanceCenter` directory.

#### – **installation\_directory**

The default installation directory for NetOps Portal.

**Default:** `/opt/CA`

The **SSO Configuration** line displays.

You are prompted to select an option. The available options correspond to the data sources running on the local server.

Use the following commands as needed while you are selecting settings:

- q (quit)
- b (go back to the previous menu)
- u (update, which overwrites the existing value for the property)
- r (reset)

3. Enter the number for the **DX NetOps** option.

The **SSO Configuration/DX NetOps** line displays.

4. Enter the number for the **Performance Center Local Password Authentication** option.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication** line displays.

A listing of the properties and their values display.

5. Enter the number for the option to override the settings on this NetOps Portal instance (local override):

- **Remote Value**

Select this option if you are enabling LDAP/LDAPS authentication for the data sources that are registered to this instance of NetOps Portal, such as DX NetOps Network Flow Analysis and CA Application Delivery Analysis. This includes the Event Manager service, which embeds the NetOps Portal URL. NetOps Portal uses these settings only if a corresponding Local Override value is not present.

- **Local Override**

Select this option if you are enabling LDAP/LDAPS authentication only for NetOps Portal and the data aggregator. These settings take precedence over the **Remote Value** settings and the default settings.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication/Local Override** line appears.

A listing of the properties and their values display.

You are prompted to select a property to configure. Configure these property values, even if they match the default values.

6. Set the name of the user account (username) that you want to configure for *internal communications* by completing the following steps:

- a. Enter the number for the **Set name for internal REST user credentials (default: admin):** option.
- b. When prompted, enter **u** to overwrite the existing value with a new value.
- c. Enter a value, for example, bob.

**IMPORTANT**

Configure this property value, even if it matches the default value.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication/Local Override** line appears.

You are prompted to select a property to configure.

7. Set the ID of the user account that you want to configure for *internal communications* by completing the following steps:

- a. Enter the number for the **Set ID for internal REST user credentials (default: 1): (Local Override)** option.
- b. When prompted, enter **u** to overwrite the existing value with a new value.
- c. Enter a value, for example, 456.

**IMPORTANT**

Configure this property value, even if it matches the default value.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication/Local Override** line appears. You are prompted to select a property to configure.

8. Enable NetOps Portal to use this user account for *internal communications* by completing the following steps:

- a. Enter the number for the **Enable specific user ID for internal REST user credentials (default: Enabled):** option.
- b. When prompted, enter **u** to overwrite the existing value with a new value.
- c. Enter the number for the **Enabled** option.

**IMPORTANT**

Configure this property value, even if it matches the default value.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication/Local Override** line appears.

You are prompted to select a property to configure.

9. [Verify the configuration.](#)

The user account for *internal communications* is configured.

### **Configure the User for User Administration**

**Prerequisite:** The user that you want to configure for *user administration* exists, and is assigned to the User Administrator role. User administrators can perform all user administrative tasks, with some limitations.

For more information about this role, see [Manage Roles and User Accounts](#).

**Follow these steps:**

1. Retrieve the details of the user account that you want to configure for *user administration* by issuing a GET request to the `users` endpoint for the NetOps Portal REST web service:

```
GET http://<PC_host>:8181/pc/center/webservice/users/userName/<user>
```

**Example:**

The following request retrieves the details for the alice user account.

```
GET http://<PC_host>:8181/pc/center/webservice/users/userName/alice
```

– **user**

The name of the user account that you want to configure for *user administration*.

**Example output:**

```
<user>
  <userId>789</userId>
  <permissionId>1</permissionId>
  <defaultGroupId>1</defaultGroupId>
  <name>alice</name>
```

.....

2. Launch SsoConfig.

The **SSO Configuration** line displays.

You are prompted to select an option. The available options correspond to the data sources running on the local server.

Use the following commands as needed while you are selecting settings:

- q (quit)
- b (go back to the previous menu)
- u (update, which overwrites the existing value for the property)
- r (reset)

3. Enter the number for the **DX NetOps** option.

The **SSO Configuration/DX NetOps** line displays.

4. Enter the number for the **Performance Center Local Password Authentication** option.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication** line displays.

A listing of the properties and their values display.

5. Enter the number for the option to override the settings on this NetOps Portal instance (local override):

– **Remote Value**

Select this option if you are enabling LDAP/LDAPS authentication for the data sources that are registered to this instance of NetOps Portal, such as DX NetOps Network Flow Analysis and CA Application Delivery Analysis. This

includes the Event Manager service, which embeds the NetOps Portal URL. NetOps Portal uses these settings only if a corresponding **Local Override** value is not present.

– **Local Override**

Select this option if you are enabling LDAP/LDAPS authentication only for NetOps Portal and the data aggregator. These settings take precedence over the **Remote Value** settings and the default settings.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication/Local Override** line appears.

A listing of the properties and their values display.

You are prompted to select a property to configure. Configure these property values, even if they match the default values.

6. Set the ID of the user account that you want to configure for *user administration* by completing the following steps:
  - a. Enter the number for the **Set ID for internal user administration credentials (default: 1):** option.
  - b. When prompted, enter **u** to overwrite the existing value with a new value.
  - c. Enter a value. For example, 789.

**Default: 1**

**IMPORTANT**

Configure this property value, even if it matches the default value.

The **SSO Configuration/DX NetOps/Performance Center Local Password Authentication/Local Override** line appears.

A listing of the properties and their values display.

7. Enter **q**.  
The configuration tool closes.
8. [Verify the configuration](#).

The user account for *user administration* is configured.

### **Verify the Configuration**

You can verify which users NetOps Portal uses for internal communications by retrieving the details of the user account for *internal communications* and *user administration* by issuing a GET request to the `admin` endpoint for the NetOps Portal REST web service:

```
GET http://<PC_host>:8381/sso/webservices/admin/whoami/testRestUsers
```

NetOps Portal responds with the configured IDs for *internal communications* and *user administration* and verifies that these users can log in with their assigned roles.

## **Configure the Proxy Server on NetOps Portal**

You can configure the proxy server on NetOps Portal to access the Internet.

NetOps Portal requires access to the Internet to load geographic maps for the NetOps Portal map web service. If you cannot open the firewall to allow the map web service access to the map server, you can configure the proxy server on NetOps Portal to access the Internet.

### **Follow these steps:**

1. Open the `<installation_directory>/PC/webapps/pc/WEB-INF/cfg/portal.console.properties` file.
  - **installation\_directory**  
The default installation directory for NetOps Portal.  
**Default:** `/opt/CA/PerformanceCenter`

This file displays a section for configuring the proxy server. By default, a proxy server is not configured:

```
#
```

```
# define proxy server for accessing internet (like maptiler api server) in PC server.
# Leave proxy.host empty if no proxy server is configured
#
proxy.protocol=https
proxy.host=
proxy.port=443
# authentication type supported: Basic/Digest
proxy.auth_type=
proxy.realm=
proxy.user=
proxy.password=
proxy.passwordEncrypted=true
```

## 2. Configure the file as follows:

- **proxy.auth\_type**  
The authentication type depends on the proxy settings.  
**Values:** Basic or Digest
- **proxy.realm**  
Complete this parameter if the proxy server requires a realm. Otherwise, leave this parameter empty.
- **proxy.passwordEncrypted**  
Defines whether to use an encrypted password.  
**Options:**
  - **true:** The proxy server uses an encrypted password.
  - **false:** The proxy server does not use an encrypted password.

**Default:** true

### NOTE

You can generate an encrypted password by issuing the following command:

```
<installation_directory>/Tools/bin/doEncryption.sh
$./doEncryption.sh <plain_password>
```

### Example:

```
/opt/CA/PerformanceCenter/Tools/bin/doEncryption.sh
$./doEncryption.sh mypassword
```

- **installation\_directory**  
The default installation directory for NetOps Portal.  
**Default:** /opt/CA/PerformanceCenter
- **plain\_password**  
Defines the password in plain text.

### Sample configuration:

```
proxy.protocol=https
proxy.host=yourhost.net
proxy.port=443
proxy.auth_type=Basic
proxy.realm=
proxy.user=proxy
proxy.password=mypassword
proxy.passwordEncrypted=true
```

## 3. Save your changes.

- Restart the NetOps Portal Console service by issuing the following command:

```
systemctl restart caperfcenter_console
```

The proxy server on NetOps Portal is configured to access the Internet.

## Configure a Reverse Proxy

You can configure a reverse proxy and use it to connect to NetOps Portal.

Ensure that the reverse proxy meets the following requirements based on your configuration:

- (If you have configured NetOps Portal and SSO externally to use different ports, for example, 8181 /8381 or 8182 /8382 ) [Configure a single reverse proxy using unique ports.](#)
- (If you have configured NetOps Portal and SSO externally to use a single port, for example 8181 or 443 ) [Configure a single reverse proxy using a single port.](#)
- (If you require use of more than one reverse proxy) [Configure multiple reverse proxies.](#)

In all cases, notification emails that include a link to the item (or context page) in NetOps Portal, or scheduled emails that include PDFs with links to reports to download, use the `Web Site Scheme /Web Site Host /Web Site Port` property values for the URLs. For a multiple reverse proxy configuration, the links redirect the user to the reverse proxy configured in the `Web Site Host` property value. The `ODataQuery` page uses the `Sso Scheme /Web Site Host /Sso Port` property values when sending users to the SSO login page.

Users can directly access NetOps Portal using the NetOps Portal host name (the alias, the fully-qualified host name (FQHN), or the IP).

For more information:

- About email notification actions, see [Configure Notifications.](#)
- About how to directly access NetOps Portal, see [Install NetOps Portal.](#)

### Configure a Single Reverse Proxy using Unique Ports

If you have configured NetOps Portal and SSO externally to use different ports, ensure the following:

- The reverse proxy and NetOps Portal are using the same protocol (http or https).
- The NetOps Portal `Web Site Port` property matches the reverse proxy port for NetOps Portal redirection. For more information about how to set this property, see [Configure the Basic Security Settings Using the SSO Configuration Tool.](#)
- The SSO port (the `Port` property) matches the reverse proxy port for SSO redirection. For more information about how to set this property, see [Configure the Port and Website for HTTPS.](#)
- You have completed *one or both* of the following options:
  - You have configured the reverse proxy for the X-Forwarded-Host (XFH) header. Example of configuring Nginx for the XFH header:
 

```
proxy_set_header X-Forwarded-Host $host;
```

 where Nginx replaces `$host` with the host name (the alias, FQHN, or IP) that the client uses to access the reverse proxy.
  - You have configured the reverse proxy for the Host header. Example of configuring Nginx for the Host header:
 

```
proxy_set_header Host $host:$port;
```

 where Nginx replaces `$host` with the host name (the alias, FQHN, or IP) and replaces `$port` with the port that the client uses to access the reverse proxy.
- You have set the following using the Single Sign-On Configuration tool (SsoConfig):
  - The `Web Site Scheme` property value to the reverse proxy protocol.
  - The `Web Site Host` property value to either the host name alias, the FQHN, or the IP of the reverse proxy.
  - The `Web Site Port` property value to the reverse proxy port.

For more information about how to set these properties, see [Configure the Basic Security Settings Using the SSO Configuration Tool](#).

### **Configure a Single Reverse Proxy using a Single Port**

If you have configured NetOps Portal and SSO externally to use a port, ensure the following:

- The reverse proxy and NetOps Portal are using the same protocol (http or https).
- The NetOps Portal `Web Site Port` property matches the reverse proxy port for NetOps Portal redirection. For more information about how to set this property, see [Configure the Basic Security Settings Using the SSO Configuration Tool](#).
- The NetOps Portal `Web Site Host` property contains one of the reverse proxy host (the alias or FQHN) values. For more information about how to set this property, see [Configure the Basic Security Settings Using the SSO Configuration Tool](#).
- The SSO port (the `Port` property) matches the reverse proxy port for NetOps Portal redirection. For more information about how to set this property, see [Configure the Port and Website for HTTPS](#).
- The internal SSO service is running on a different port, where the reverse proxy is configured to redirect requests to SSO.
- You have completed *one or both* of the following options:
  - You have configured the reverse proxy for the XFH header. Example of configuring Nginx for the XFH header:
 

```
proxy_set_header X-Forwarded-Host $host;
```

 where Nginx replaces `$host` with the host name (the alias, FQHN, or IP) that the client uses to access the reverse proxy.
  - You have configured the reverse proxy for the Host header. Example of configuring Nginx for the Host header:
 

```
proxy_set_header Host $host:$port;
```

 where Nginx replaces `$host` with the host name (the alias, FQHN, or IP) and replaces `$port` with the port that the client uses to access the reverse proxy.

If you have configured the reverse proxy for *only* the Host header, NetOps Portal cannot determine if the host name is the reverse proxy host name (the alias, FQHN, or the IP) or the NetOps Portal host name (the alias, FQHN, or the IP), and when the request host name does not match the `Web Site Host` property value, it uses the value set for the SSO scheme/port for the data source. Clients accessing the reverse proxy must use the host name that is set for the `Web Site Host` property value.

For more information, about how to set the `Web Site Host` property value, see [Configure the Basic Security Settings Using the SSO Configuration Tool](#).

### **Configure Multiple Reverse Proxies**

If you require use of more than one reverse proxy, ensure that they are configured to use the same protocol/ports. If you have configured multiple reverse proxies, ensure that each sets the XFH header.

## **Configure Java Options for NetOps Portal Services**

You can set additional Java options for NetOps Portal services.

Complete this procedure for each NetOps Portal service for which you want to configure Java options. NetOps Portal preserves these additional Java options on upgrade.

### **Follow these steps:**

1. Edit the `<installation_directory>/<service_subdirectory>/conf/wrapper-user.conf` file.
  - ***installation\_directory***  
The installation directory for NetOps Portal.

**Default:** `/opt/CA/PerformanceCenter`

– **service\_subdirectory**

The NetOps Portal service for which you want to configure additional Java option.

**Options:**

- PC: The NetOps Portal Console service Java options.
- DM: The Device Manager service Java options.
- EM: The Event Manager service Java options.
- sso: The SSO service Java options.

2. Set one option per line using the following format, and then save the file:

```
wrapper.java.additional.<number>=-D<Java_option>=<value>
wrapper.java.additional.<number>=-D<Java_option>=<value>
```

**Example:**

The following example sets an additional Java option (setting 31) for Java to prefer IPv6 when available:

```
wrapper.java.additional.31=-Djava.net.preferIPv6Addresses=true
```

– **number**

Specifies the number for the additional Java option.

**NOTE**

Settings 1 through 30 are reserved for NetOps Portal.

**Values:** 30 and higher

– **Java\_option**

Specifies the Java option that you want to configure.

**Example:** `java.net.preferIPv6Addresses`

– **value**

Specifies the value for the additional Java option that you want to configure.

**Example:** `true`

The additional Java option for the NetOps Portal service is configured.

## NetOps Portal Scripts

The DX NetOps Performance Management installation includes the following scripts for NetOps Portal:

- [doEncryption.sh](#)
- [parseSyncTimes.pl](#)
- [RemoteEngineer/re.sh](#)
- [reset\\_mysql\\_password.sh](#)
- [reset\\_pc\\_db\\_password.sh](#)
- [uninstall\\_database\\_component.sh](#)
- [update\\_alias\\_name.sh](#)
- [update\\_pc\\_da\\_database\\_references.sh](#)

### doEncryption.sh

The `doEncryption.sh` script updates the MySQL data repository admin password.

### parseSyncTimes.pl

The `parseSyncTimes.pl` script parses the `DMSERVICE.log` file to find how long each sync phase is taking.

### RemoteEngineer/re.sh



The `re.sh` script (the CA Remote Engineer (CARE) tool) gathers the DX NetOps Performance Management configuration and database into a compressed file for Broadcom Support triage purposes.

For more information about when to run this script, see [Unable to Resolve Issue](#).

### **reset\_mysql\_password.sh**

The `reset_mysql_password.sh` script updates the MySQL password for root and netqos accounts.

### **reset\_pc\_db\_password.sh**

The `reset_pc_db_password.sh` script updates the NetOps Portal Console, Device Manager, Event Manager, and SSO service property files with new MySQL password.

### **uninstall\_database\_component.sh**

The `uninstall_database_component.sh` script moves the NetOps Portal database to a separate node.

For more information about when to run this script, see [Move the NetOps Portal Database to a Separate Node](#).

### **update\_alias\_name.sh**

The `update_alias_name.sh` script sets/updates the alias names for one or more monitored devices, interfaces, and components simultaneously.

For more information about when to run this script:

- See [Manage IP Domains](#).
- See [Set Alias Names Using a Script](#).

### **update\_pc\_da\_database\_references.sh**

The `update_pc_da_database_references.sh` script is a NetOps Portal migration/disaster recovery script that resets the MySQL password.

For more information about when to run this script:

- See [Install a Disaster Recovery System](#).
- See [Restore NetOps Portal](#).

## **Data Aggregator Administration**

The articles in this section relate to administration tasks for data aggregators and data collectors.

### **Automate Device Inventory Synchronization**

You can automate the maintenance of the data aggregator device inventory through automated synchronization with a CMDB source.

You can use a configuration management database to maintain an inventory of all devices within your environment. This information is kept in one place, where it is then provisioned out to various monitoring tools. You want to build an automated means of keeping data aggregator device inventory in synchronization with your CMDB. To do so, install NetOps Portal and the data aggregator *before* you write the integration script.

You can manually execute web service calls using either a REST client editor or an HTTP tool that sends requests and gets responses to perform these procedures. This article refers to the REST client editor.

**NOTE**

For automated inventory control, you would write an application or script that leverages the web services that are described in this article.

This process requires making calls to two different web servers:

- The NetOps Portal web server is used to create SNMP profiles.
- The Data Aggregator web server is used to create discovery profiles and review discovery results.

Use the following process to discover devices and monitor components in your network using REST web services:

1. [Provide SNMP Profile Credentials](#). Use NetOps Portal REST web services.
2. [Create a Discovery Profile and Run a Discovery](#). Use the Data Aggregator REST web services.
3. [Review the Discovered Devices and Instances](#). Use the Data Aggregator REST web services.

**Provide SNMP Profile Credentials**

*SNMP profiles* are definitions that contain the information necessary to enable secure queries of device MIBs using SNMP. These definitions provide SNMP parameters from NetOps Portal to the Data Aggregator data source, as needed, while ensuring data security.

For this scenario, you will create an SMPv3 profile.

**NOTE**

The methods that return SNMP profile definitions do not include community strings, user names, or passwords. The save method takes a community string, user name, or passwords that are not encrypted. We recommend that you only invoke this method on the computer where NetOps Portal is installed.

**Follow these steps:**

1. Access the NetOps Portal web server.
2. Launch a REST client that is configured to access the NetOps Portal server. Log in using administrator credentials.  
For more information about how to configure your REST client, see the REST client documentation.
3. Set the Content-type to application/xml. Enter the **Body** text and modify the attributes as needed:
  - **Profile Name**  
Defines a name for the SNMP profile. Profile names must be unique, cannot be duplicated across SNMP versions, and are not case-sensitive.  
**Required:** Yes
  - **SNMP Version**  
Specifies the version of SNMP that the profile uses. Enter SNMPv3 for this scenario.  
**Required:** Yes
  - **Port**  
Identifies the port that the data aggregator uses to make SNMP connections to devices associated with this profile.  
**Default:** 161  
**Required:** Yes
  - **User Name**  
Identifies the user for the profile, whose secret keys were used potentially to authenticate and encrypt the SNMPv3 packets. The User Name is a character string.  
**Required:** Yes
  - **Context Name**  
Specifies a collection of management information that is accessible by an SNMP entity. The Context Name is necessary for providing end-to-end identification and for retrieving data from an SNMPv3 agent. The Context Name is an octet string.

**NOTE**

The data aggregator does not use Context Name on the SNMPv3 profiles to communicate with the device.

**Required:** Yes

– **Security Level**

Specifies the security level to use. Enter one of the following values:

- NoAuthNoPriv
- AuthNoPriv
- AuthAndPriv

**Required:** Yes

– **Authentication Protocol**

Specifies the authentication protocol to use when contacting devices associated with this profile. Enter one of the following algorithms for authenticating SNMPv3 packets:

- None (Do not attempt authentication.)
- MD5 (Message Digest 5)
- SHA (Secure Hash Algorithm)

**Required:** Yes

– **Authentication Password**

Specifies the password for authentication using SNMPv3 and the selected authentication protocol. The password must contain a minimum of eight characters.

**Required:** Yes

– **Verify Authentication Password**

Confirms the authentication password.

**Required:** Yes

– **Privacy Protocol**

Specifies the encryption protocol to use for data flows sent to any devices or servers that are associated with this profile.

**Options:**

- None (Does not encrypt communications. Use only with NoPriv option.)
- DES
- AES (128-bit encryption)
- Triple DES

**NOTE**

The privacy protocol option is only enabled when authentication is enabled for this profile.

**Required:** No

– **Privacy Password**

(SNMPv3 only) Defines the password that is used when exchanging encryption keys.

**Required:** Yes

– **Verify Privacy Password**

(SNMPv3 only) Confirms the password that is used when exchanging encryption keys.

**Required:** Yes

– **Rank**

Specifies the rank of the profile in the global list of SNMP profiles.

**Required:** Yes

– **Enabled**

For NetOps Portal, this indicates whether the information in this profile is used when not explicitly assigned to a device. For the data aggregator, this value must be set to **true**.

**Required:** Yes

– **TenantID**

Specifies the tenant ID.

**Default:** 8

**Required:** Yes

**Example: No authentication and no privacy**

```
<SnmpProfile>
  <name>Tokyo</name>
  <port>161</port>
  <userName>myuser</userName>
  <context></context>
  <version>Version3</version>
  <securityLevel>NoAuthNoPriv</securityLevel>
  <authProtocol>None</authProtocol>
  <authPassword>None</authPassword>
  <privProtocol>None</privProtocol>
  <privPassword>None</privPassword>
  <rank>4</rank>
  <enabled>true</enabled>
  <tenantID>8</tenantID>
</SnmpProfile>
```

**Example: Authentication and no privacy**

```
<SnmpProfile>
  <name>Brasil</name>
  <port>161</port>
  <userName>myuser</userName>
  <context></context>
  <version>Version3</version>
  <securityLevel>AuthNoPriv</securityLevel>
  <authProtocol>MD5</authProtocol>
  <authPassword>test</authPassword>
  <privProtocol>None</privProtocol>
  <privPassword>None</privPassword>
  <rank>3</rank>
  <enabled>true</enabled>
  <tenantID>8</tenantID>
</SnmpProfile>
```

**Example: Authentication and privacy**

```
<SnmpProfile>
  <name>Boston</name>
  <port>161</port>
  <userName>myuser</userName>
  <context></context>
  <version>Version3</version>
  <securityLevel>AuthAndPriv</securityLevel>
  <authProtocol>MD5</authProtocol>
```

```

<authPassword>test</authPassword>
<privProtocol>TripleDES</privProtocol>
<privPassword>test</privPassword>
<rank>1</rank>
<enabled>true</enabled>
<tenantID>8</tenantID>
</SnmpProfile>

```

4. Create the profile by entering the following URL for the `profiles` endpoint for the NetOps Portal RESTful web services API in the **URL** field, selecting **POST** as the **HTTP Method**:

```
http://<PC_host>:8181/pc/center/webservice/profiles/saveProfile/{true|false}
```

- **PC\_host**

Specifies the NetOps Portal host name. (8181 is the required port.)

- **{true|false}**

Specifies a Boolean value for the `rankTiesAscendingByDate` parameter. **True** indicates that the profile you are adding will be the last in rank order (as determined by the creation date of the SNMP profile).

The XML returns true when the operation succeeds.

The SNMP profile is automatically synchronized with the data aggregator and is available for inventory discovery to use it.

### **Create a Discovery Profile and Run a Discovery**

*Discovery profiles* specify how discovery operates in the data aggregator environment. Within a discovery profile, you can specify the IP addresses, IP address ranges, and host names you want to discover devices for.

#### **NOTE**

This scenario is for an enterprise environment and discovers devices using the SNMP protocol. However, MSPs may also want to limit discovery to specific SNMP profiles. A full list of attributes is available by accessing the following URL:

```
http://<DA_host>:<port>/rest/discoveryprofiles/documentation
```

You create a discovery profile and run a discovery in one operation when using the REST web services. When discovery is run, devices are discovered based on the discovery profile you create.

#### **Follow these steps:**

1. Access the Data Aggregator web server.
2. Launch a REST client.
3. Set the Content-type to application/xml.
4. Find and make a note of the default tenant ID by entering the following URL for the `tenants` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:<port>/rest/tenants
```

- **DA\_host:port**

Specifies the data aggregator host name and the port number.

**Default port:** 8581

5. Enter the Body text for the discovery profile in a REST client and modify the attributes as needed.

#### **NOTE**

You can view the XSD schema which defines the structure of the discovery profile XML by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:<port>/rest/discoveryprofiles/XSD/getlist.xsd
```

6. Include IP addresses and host names so that the end-user switches in your network are discovered. Do one or more of the following actions:

- Enter individual IP addresses for which you want to discover devices in the **IP Address List** field.
- Enter the host names for which you want to discover devices in the **Host List** field.

**NOTE**

These fields accept comma-delimited values. Single quotes, double quotes, backward slashes, forward slashes, and ampersands are not permitted.

## 7. Set the following attributes:

- **Name**  
Specifies a descriptive name for the discovery profile. This field cannot contain single quotes, double quotes, backward slashes, forward slashes, and ampersands.
- **RunStatus**  
Specifies whether to run the discovery. For this scenario, set the attribute to START.

**NOTE**

You can rerun discovery later by updating (PUT) the run status to START.

**Example:**

```
<?xml version="1.0" encoding="UTF-8"?>
  <DiscoveryProfile version="1.0.0">
    <ActivationStatus>true</ActivationStatus>
    <IPRangesList>
      <IPRanges>10.0.64.202-10.0.64.206</IPRanges>
    </IPRangesList>
    <HostNamesList>
      <HostNames>ahost</HostNames>
    </HostNamesList>
    <IPListList>
      <IPList>10.10.10.10</IPList>
    </IPListList>
    <RunStatus>READY</RunStatus>
    <Item version="1.0.0">
      <Name>BigRange_TestProfileViaAPI</Name>
    </Item>
    <IPDomainMember version="1.0.0">
      <IPDomainID>285</IPDomainID>
    </IPDomainMember>
  </DiscoveryProfile>
```

8. Save and run the new discovery profile by entering the following URL for the `tenant` endpoint for the Data Aggregator RESTful web services API, selecting **POST** for the **HTTP Method**:

```
http://<DA_host>:<port>/rest/tenant/<default_tenant_ID>/discoveryprofiles
```

– **tenant/default\_tenant\_ID**

Specifies the ID number for the default tenant workspace. Enterprise customers typically only use the default tenant workspace. For the POST operation, enter the default tenant ID when there are no other tenants.

The discovery profile is created and discovery runs.

All discovered devices are automatically added to the appropriate out-of-the-box device collection or another user-created device collection.

9. Verify that the discovery profile displays in the discovery profiles list, and then note the ID of the new discovery profile by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method** (the discovery profile is used to review the discovery results):

```
http://<DA_host>:<port>/rest/discoveryprofiles
```

## Review the Discovered Devices and Instances

Verify that your initial discovery was successful by reviewing the discovered devices and instances.

### Follow these steps:

1. Access the Data Aggregator web server.
2. View a list of discovery profiles by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:<port>/rest/discoveryprofiles
```

– **DA\_host:port**

Specifies the data aggregator host name and the port number.

**Default port:** 8581

A list of discovery profiles and their instance IDs is returned.

3. Find the discovery instance ID for the discovery profile that you created, for example:

```
<DiscoveryInstanceIDList relatesURL="relatesto/
instances"rootURL="discoveryinstances">
  <ID>236</ID>
</DiscoveryInstanceIDList>
```

4. View the discovery instance details by entering the following URL for the `discoveryinstances` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:<port>/rest/discoveryinstances/<instance_ID>
```

– **instance\_ID**

Specifies the ID of the referenced discovery instance.

The instance XML returns information about new and existing devices and SNMP profiles that were tested.

### Example:

```
<?xml version="1.0" encoding="UTF-8"?>
  <DiscoveryInstance version="1.0.0">
    <ID>236</ID>
    <IPSweepTotalSuccess>5</IPSweepTotalSuccess>
    <CompletionTime>Thu Apr 12 12:36:35 CDT 2012</CompletionTime>
    <ExistingFoundDevicesList>
      <ExistingFoundDevices>241</ExistingFoundDevices>
      <ExistingFoundDevices>239</ExistingFoundDevices>
      <ExistingFoundDevices>240</ExistingFoundDevices>
      <ExistingFoundDevices>238</ExistingFoundDevices>
    </ExistingFoundDevicesList>
    <IPSweepCompletionTime>Thu Apr 12 12:36:35 CDT 2012</IPSweepCompletionTime>
    <ExistingFoundManageableDevicesList>
      <ExistingFoundManageableDevices>241</ExistingFoundManageableDevices>
      <ExistingFoundManageableDevices>239</ExistingFoundManageableDevices>
      <ExistingFoundManageableDevices>240</ExistingFoundManageableDevices>
      <ExistingFoundManageableDevices>238</ExistingFoundManageableDevices>
    </ExistingFoundManageableDevicesList>
    <StartTime>Thu Apr 12 12:36:31 CDT 2012</StartTime>
    <NewlyCreatedDevicesList>
      <NewlyCreatedDevices>383</NewlyCreatedDevices>
    </NewlyCreatedDevicesList>
    <TestedCommProfilesList>
```

```

<TestedCommProfiles>199</TestedCommProfiles>
<TestedCommProfiles>198</TestedCommProfiles>
</TestedCommProfilesList>
<IPSweepStartTime>Thu Apr 12 12:36:31 CDT 2012</IPSweepStartTime>
<ProgressPercentage>100</ProgressPercentage>
<IPSweepTotal>7</IPSweepTotal>
<ProfileID>235</ProfileID>
<NewlyCreatedManageableDevicesList>
<NewlyCreatedManageableDevices>383</NewlyCreatedManageableDevices>
</NewlyCreatedManageableDevicesList>
<PingResponseDeviceCount>5</PingResponseDeviceCount>
<CompletionStatus>SUCCESS</CompletionStatus>
<Item version="1.0.0">
  <CreateTime>Thu Apr 12 12:36:31 CDT 2012</CreateTime>
</Item>
</DiscoveryInstance>

```

5. (Optional) Review information about specific devices (new or existing) by entering the following URL for the `devices` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

`http://<DA_host>:<port>/rest/devices/<endpoint_ID>`

– **endpoint\_ID**

Specifies the ID of the referenced endpoint, such as devices.

**NOTE**

You found this ID in a previous step.

6. (Optional) Review information about the SNMP profiles that were tested during discovery by entering the following URL for the `profiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

`http://<DA_host>:<port>/rest/profiles/<profile_ID>`

– **profile\_ID**

Specifies the ID of the referenced SNMP profile.

The discovered devices and instances appear.

## Back Up the Data Aggregator

To avoid losing your settings and custom certifications due to an unexpected failure, back up the data aggregator.

Back up the data aggregator during specific events, such as before an upgrade or before migrating the data aggregator to another (new) host.

You do not need to stop the data repository, the data collector, or the data aggregator services before backing up the data aggregator. Backups are stored in the location that you specify, which can be on the data aggregator system or on a different backup host system.

**Prerequisite:** You have root or sudo privileges.

Use the following process to back up the data aggregator:

1. [Create a Backup Directory](#)
2. (If you configured Single Sign-On (SSO)) [Back Up the Single Sign-On Configuration Files](#)
3. [Back Up the Files](#)
4. (If the data aggregator is HTTPS-enabled) [Back Up the HTTPS Files and Directories](#)



5. (If you have secured communication between the data aggregator and the data collectors) [Back up the ActiveMQ Configuration File](#)
6. [Package the Files](#)

## **Create a Backup Directory**

### **Follow these steps:**

1. From a command prompt, create a backup directory in a secure location on the same or different backup host system by issuing the following command:

```
mkdir <DA_Backup>
```

#### **Example:**

```
mkdir /tmp/DA_Backup
```

#### **– DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

2. Create subdirectories within the <DA\_Backup> directory by issuing the following commands:

```
mkdir <DA_Backup>/deploy_backup
mkdir <DA_Backup>/cert_backup
mkdir <DA_Backup>/CustomDeviceType_backup
mkdir <DA_Backup>/etc_backup
mkdir <DA_Backup>/activemq_backup
mkdir <DA_Backup>/data
```

#### **Example:**

```
mkdir /tmp/DA_Backup/deploy_backup
mkdir /tmp/DA_Backup/cert_backup
mkdir /tmp/DA_Backup/CustomDeviceType_backup
mkdir /tmp/DA_Backup/etc_backup
mkdir /tmp/DA_Backup/activemq_backup
mkdir /tmp/DA_Backup/data
```

#### **– DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

## **Back Up the Files**

### **Follow these steps:**

1. Back up the deploy directory by issuing the following commands:

```
cp -r <installation_directory>/apache-karaf/deploy <DA_Backup>/deploy_backup
rm -f <DA_Backup>/deploy_backup/deploy/local-jms-broker.xml
```

#### **Example:**

```
cp -r /opt/IMDataAggregator/apache-karaf/deploy /tmp/DA_Backup/deploy_backup
rm -f /tmp/DA_Backup/deploy_backup/deploy/local-jms-broker.xml
```

#### **– installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

#### **– DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

2. Back up the certifications in the `certifications` directory by issuing the following command based on your environment:

- Non-fault-tolerant environment:

```
cp -r <installation_directory>/data/certifications <DA_Backup>/cert_backup
```

**Example:**

```
cp -r /opt/IMDataAggregator/data/certifications /tmp/DA_Backup/cert_backup
```

- Fault-tolerant environment:

```
cp -r <DASharedRepo>/certifications <DA_Backup>/cert_backup
```

**Example:**

```
cp -r /DASharedRepo/certifications /tmp/DA_Backup/cert_backup
```

#### NOTE

- ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** /opt/IMDataAggregator
- ***DA\_Backup***  
Specifies the directory path and name of the backup directory for the data aggregator.  
**Example:** /tmp/DA\_Backup
- ***DASharedRepo***  
Specifies the shared data directory that you set up when you installed the fault-tolerant data aggregators.  
**Example:** /DASharedRepo

3. Back up the custom device subtype XML file, `DeviceTypes.xml`, in the `devicetypes` directory by issuing the following command based on your environment:

- Non-fault-tolerant environment:

```
cp <installation_directory>/data/custom/devicetypes/DeviceTypes.xml <DA_Backup>/CustomDeviceType_backup/DeviceTypes.xml
```

**Example:**

```
cp /opt/IMDataAggregator/data/custom/devicetypes/DeviceTypes.xml /tmp/DA_Backup/CustomDeviceType_backup/DeviceTypes.xml
```

- Fault-tolerant environment:

```
cp <DASharedRepo>/custom/devicetypes/DeviceTypes.xml <DA_Backup>/CustomDeviceType_backup/DeviceTypes.xml
```

**Example:**

```
cp /DASharedRepo/custom/devicetypes/DeviceTypes.xml /tmp/DA_Backup/CustomDeviceType_backup/DeviceTypes.xml
```

#### NOTE

- ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** /opt/IMDataAggregator
- ***DA\_Backup***  
Specifies the directory path and name of the backup directory for the data aggregator.  
**Example:** /tmp/DA\_Backup
- ***DASharedRepo***  
Specifies the shared data directory that you set up when you installed the fault-tolerant data aggregators.  
**Example:** /DASharedRepo

4. Back up the `etc` directory files and the HTTPS configuration by issuing the following command:

```
cp -r <installation_directory>/apache-karaf/etc <DA_Backup>/etc_backup
```

**Example:**

```
cp -r /opt/IMDataAggregator/apache-karaf/etc /tmp/DA_Backup/etc_backup
```

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

5. Back up the data/log directory by issuing the following command:

```
cp -r <installation_directory>/apache-karaf/data/log <DA_Backup>/data/
```

**Example:**

```
cp -r /opt/IMDataAggregator/apache-karaf/data/log /tmp/DA_Backup/data/
```

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

6. Back up the ActiveMQ configuration file, activemq.xml , in the conf directory by issuing the following command:

```
cp <installation_directory>/IMDataAggregator/broker/<apache-activemq-*>/conf/activemq.xml <DA_Backup>/activemq_backup/activemq.xml
```

**Example:**

```
cp /opt/IMDataAggregator/broker/apache-activemq-5.18.3/conf/activemq.xml /tmp/DA_Backup/activemq_backup/activemq.xml
```

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt

– **apache-activemq-\***

The installation directory for Apache ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

**Back Up the HTTPS Files and Directories**

If you have enabled HTTPS for the data aggregator, back up the Java truststore file directory by issuing the following command:

```
cp <installation_directory>/jre/lib/security/cacerts <DA_Backup>/
```

**Example:**

```
cp /opt/IMDataAggregator/jre/lib/security/cacerts /tmp/DA_Backup/
```

• **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

• **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

### **Back up the Client Key and Broker Key for the Data Aggregator**

If you have secured Apache ActiveMQ communication between the data aggregator and the data collectors, back up the `broker.ks` and `client.ts` files by issuing the following commands:

```
cp <installation_directory>/broker/<apache-activemq-*>/conf/broker.ks <DA_Backup>/
cp <installation_directory>/broker/<apache-activemq-*>/conf/client.ts <DA_Backup>/
```

#### **Example:**

```
cp /opt/IMDataAggregator/broker/apache-activemq-5.18.3/conf/broker.ks /tmp/DA_Backup/
cp /opt/IMDataAggregator/broker/apache-activemq-5.18.3/conf/client.ts /tmp/DA_Backup/
```

- ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** /opt/IMDataAggregator
- ***apache-activemq-\****  
The installation directory for Apache ActiveMQ.  
**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`
- ***DA\_Backup***  
Specifies the directory path and name of the backup directory for the data aggregator.  
**Example:** /tmp/DA\_Backup

For more information about how to secure communication, see [Authenticate and Encrypt ActiveMQ Communication](#).

### **Package the Files**

Package the files by issuing the following command:

```
tar czf DA.tgz <DA_Backup>
```

#### **Example:**

```
tar czf DA.tgz /tmp/DA_Backup
```

- ***DA\_Backup***  
Specifies the directory path and name of the backup directory for the data aggregator.  
**Example:** /tmp/DA\_Backup

The data aggregator is backed up.

## **Choose Another Host in a Cluster When Selected Host Fails**

If you have installed the data repository in a cluster, you can point the database connections to another host in the cluster.

If more than one host in the data repository cluster fails, the data repository and the data aggregator shut down automatically. The data repository cluster is only capable of losing one host.

If a single host in the cluster that is *not* specified when you installed the data aggregator disconnects from the network (for example, because a firewall was put in place, or the Ethernet cable was removed), the data aggregator shuts down. If you set up the automatic recovery of the data aggregator process when you installed the data aggregator, the data aggregator restarts automatically.

If you stop a single host in the cluster that is *not* specified during the data aggregator installation by selecting the **Kill Vertica Process on Host** option on the **Advanced** menu of the Vertica Administration Tools utility (`adminTools`), the data aggregator continues to function.

In either case, after the host that is offline becomes available, complete the following:

1. Return that host to the cluster.
2. On the main menu of the Vertica Administration Tools utility, select the **Restart Vertica on Host** option, and then follow the prompts.

#### Follow these steps:

1. Open the `<installation_directory>/apache-karaf/etc/dbconnection.cfg` file on the data aggregator host.
  - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
2. Modify the following line in the file. Modify the line to reference a hostname or IP address of one of the data repository cluster hosts that is still up and running:
 

```
dbUrl=jdbc:vertica://<database server hostname>:<database server port>/databasename?
use35CopyFormat=true&BinaryDataTransfer=false
```

#### Example:

If `host2` is up and running in the cluster and you choose database connections to point to `host2`, your updated `dbUrl` entry could look like the following line:

```
dbUrl=jdbc:vertica://host2:5433/mydatabasename?
use35CopyFormat=true&BinaryDataTransfer=false
```

- **database server hostname:database server port**  
The hostname or IP address of the data repository and the data repository port number that you entered when you installed the data aggregator.  
**Default port number:** 5433
3. Save the file.
  4. Do one of the following steps:
    - Start the Data Aggregator service by issuing the following command:
 

```
systemctl start dadaemon
```
    - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary:
 

```
<installation_directory>/scripts/dadaemon activate
```
  5. Ensure that the data aggregator is not still running by issuing the following command:
 

```
Ps - ef | grep java | grep - v grep
```

The data aggregator processes are not returned when the data aggregator is not running.  
Database connections point to the specified host in the cluster going forward.

Another host in the cluster is chosen.

**Next step:** [Restart the data aggregator.](#)

## Update the Data Aggregator IP Address and Hostname

You can update the data aggregator IP address and hostname.

Update the data aggregator IP address and hostname if the following changes occur to the server on which the data aggregator is installed:

- The server is assigned a new IP address or hostname.
- The server has moved to a new subnet.

Use the following process to update the data aggregator IP address and hostname:

1. [Verify the prerequisites.](#)
2. [Stop the services.](#)
3. [Update the IP address or hostname.](#)
4. [Start the services.](#)
5. (If you are running in a (disaster) recovery system) [Update and run the disaster recovery script.](#)

### **Verify the Prerequisites**

Before updating the data aggregator IP address and hostname, ensure that you have completed the following prerequisite steps:

- (In a fault-tolerant environment) You have ensured that you are using the correct "Ready" and "Maintenance" status changes, and not the non-fault-tolerant data aggregator stop/start commands.
- You have updated the hostname for the server in the `/etc/hosts` file and in the DNS in the network to translate to the new IP for the existing hostname. This ensures that the data collector and the data aggregator can communicate.

### **Stop the Services**

For more information, see [Stop the Data Aggregator and ActiveMQ services.](#)

### **Update the IP Address or Hostname**

Follow these steps:

1. On the data collector hosts that communicate with the data aggregator, complete the following:
  - a. Open the `<DC_installation_directory>/apache-karaf/etc/com.ca.im.dm.core.collector.cfg` configuration file.  
**Example:**  
`/opt/IMDataCollector/apache-karaf/etc/com.ca.im.dm.core.collector.cfg`
    - ***DC\_installation\_directory***  
 The installation directory of the data collector.  
**Default:** `/opt/IMDataCollector`
  - b. If the `collector-manager-da-hostname` entry is set using the IP addresses for the data aggregator host instead of the hostnames, update the IP addresses with the new data aggregator host IP address. If it is set using the hostname, and the hostname remains the same, no change is required.
2. (For non-fault-tolerant environments only) In NetOps Portal, [update the data aggregator data source with the new data aggregator IP address or hostname \(the Host Name field\).](#)
3. On the data collector hosts that communicate with the data aggregator, complete the following:
  - a. Open the `<DC_installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml` ActiveMQ services file.  
**Example:**  
`/opt/IMDataCollector/broker/apache-activemq-5.18.3/conf/activemq.xml`
    - ***DC\_installation\_directory***  
 The installation directory of the data collector.  
**Default:** `/opt/IMDataCollector`
    - ***apache-activemq-\****

The installation directory for Apache ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

- b. If the `networkConnector` entry is set using the IP addresses for the data aggregator host instead of the hostnames, update the IP addresses in the following entries with the new data aggregator IP address, and then save your changes:
  - **host**  
The data aggregator with which the data collector should communicate. Typically, this entry is set using the hostname instead of the IP address. There are `networkConnector` entries for each of the following ports: 61616, 61618, 61620, and 61622. Ensure that the data collector and the data aggregator can communicate by updating the IP address with the new data aggregator IP address.
  - (In fault-tolerant environments) List the data aggregator hosts in each `networkConnector` entry. Ensure that the correct IP addresses are set for both.
- c. Open the `/opt/DCM.cfg` file.
- d. Update the value for *one* of the following properties with the new data aggregator IP address or hostname based on which property includes the existing IP address or hostname, and then save your changes:
  - `DA_HOST=`  
The data aggregator IP address.
  - (For fault-tolerant data aggregators) `DA_STANDBY_HOST=`  
The fault-tolerant data aggregator IP address.

The IP address or hostname is updated.

### **Start the Services**

For more information, see [Start the Data Aggregator and ActiveMQ services](#).

The data aggregator recalculates the IP Address and hostname.

### **Update and Run the Data Repository Disaster Recovery Script**

#### **IMPORTANT**

Complete this step *only* if you are running in a (disaster) recovery system.

If the new data aggregator IP address or hostname is different from the original data aggregator IP address or hostname, update the data repository disaster recovery script, and then run it.

For more information, see [Install a Disaster Recovery System](#).

The data repository is updated with the new data aggregator IP address.

## **Configure the Data Collectors When the Data Aggregator IP Address Changes**

Complete *one* of the following:

- If the data collector uses the hostname to communicate with the data aggregator, [restart the data collector](#).
- If the data collector uses the IP address to communicate with the data aggregator, update the data collectors to point to the new IP by completing the following procedure.

#### **Follow these steps:**

1. Log in to the data collector host.
2. Stop the Data Collector service by issuing the following command:  
`systemctl stop dcmd`
3. Open in the `<installation_directory>/apache-karaf/etc/com.ca.im.dm.core.collector.cfg` file.  
– ***DC\_installation\_directory***

The default installation directory for the data collector.

**Default:** /opt/IMDataCollector

4. Locate and edit the IP address in the following line, and then save the file:

```
...
collector-manager-da-hostname=
DA_host_IP
```

...

5. Locate and edit the `<DC_installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml` file.

**Example:**

/opt/IMDataCollector/broker/apache-activemq-5.18.3/conf/activemq.xml

– **DC\_installation\_directory**

The default installation directory for the data collector.

**Default:** /opt/IMDataCollector

– **apache-activemq-\***

The installation directory for Apache ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

6. Edit the IP address in the following lines:

```
...

<networkConnector name="da_manager" uri="static:(tcp://
DA_host_IP
:61616)" duplex="true" suppressDuplicateTopicSubscriptions="false"/>

<networkConnector name="da_manager-PRQ" uri="static:(tcp://
DA_host_IP
:61618)" duplex="true" suppressDuplicateTopicSubscriptions="false"/>

<networkConnector name="da_manager-IREP" uri="static:(tcp://
DA_host_IP
:61620)" duplex="true" suppressDuplicateTopicSubscriptions="false"/>

<networkConnector name="da_manager-blob" uri="static:(tcp://
DA_host_IP
:61622)" duplex="true" suppressDuplicateTopicSubscriptions="false"/>

...
```

In a fault-tolerant tolerant environment, the content of the `activemq.xml` file has the following format:

```
...

<networkConnector name="da_manager" uri="static:(failover:(tcp://
DA1_host_IP
:61616,tcp://
```



```

DA2_host_IP
:61616)?maxReconnectAttempts=3)" duplex="true"
  suppressDuplicateTopicSubscriptions="false"/>

<networkConnector name="da_manager-PRQ" uri="static:(failover:(tcp://
DA1_host_IP
:61618,tcp://
DA2_host_IP
:61618)?maxReconnectAttempts=3)" duplex="true"
  suppressDuplicateTopicSubscriptions="false"/>

<networkConnector name="da_manager-IREP" uri="static:(failover:(tcp://
DA1_host_IP
:61620,tcp://
DA2_host_IP
:61620)?maxReconnectAttempts=3)" duplex="true"
  suppressDuplicateTopicSubscriptions="false"/>

<networkConnector name="da_manager-blob" uri="static:(failover:(tcp://
DA1_host_IP
:61622,tcp://
DA2_host_IP
:61622)?maxReconnectAttempts=3)" duplex="true"
  suppressDuplicateTopicSubscriptions="false"/>

...

```

7. Start the Data Collector service by issuing the following command:
 

```
systemctl start dcmd
```
8. Verify that the correct address appears in the data collector list by completing the following steps:
  - a. Log in to NetOps Portal as an Administrator.
  - b. Hover over **Administration, Data Sources**, and then click a data aggregator data source.  
The **Monitored Devices** page appears.
  - c. Click **Data Collectors** from the **System Status** menu.  
The **Data Collectors** page appears.

The IP address of each data collector appears in the **Address** column. The **Polling Status** of each data collector is "Collecting Data".

The data collectors now point to the new data aggregator IP address.

## Update the Data Collector IP Address and Hostname

When you change the IP address or host name for the server on which the data collector is installed, update the data collector IP address and host name.

Update the data collector IP address or hostname if the following changes occur to the server on which the data collector is installed:

- The server is assigned a new IP address or hostname.
- The server has moved to a new subnet.

Use the following process to update the data collector IP address or hostname:

1. [Stop the services.](#)
2. [Update the Data Collector IP Address or Hostname.](#)
3. [Start the Data Collector Service.](#)  
The data collector updates the IP address and hostname at startup.
4. (If you are running in a (disaster) recovery system) [Update the disaster recovery script.](#)

### **Stop the Services**

Stop the Data Collector services.

For more information about how to stop the services, see [Restart the Data Collector](#).

### **Update the Data Collector IP Address or Hostname**

As needed, open the `/etc/hosts` file, and then edit the data collector IP address or hostname with the new IP address or hostname.

### **Start the Services**

Start the Data Collector service.

For more information about how to start this service, see [Restart the Data Collector](#).

The data collector sends the new information to the data aggregator to record in `http://<DA_HOST>:8581/rest/dcms`. The data collector updates the IP address and hostname at startup.

### **Update the Data Repository Disaster Recovery Script**

#### **IMPORTANT**

This step is required *only* if you are running in a disaster recovery system.

If the new data collector IP address or hostname is different from the original data collector IP address or hostname, update the data repository disaster recovery script.

For more information about how to update and run this script, see [Disaster Recovery](#).

The data repository is updated with the new data aggregator IP address.

## **Data Aggregator Configuration Changes During Network Disconnects to a Data Collector Host**

Occasionally, the connection between a Data Aggregator host and a Data Collector host breaks, such as, when a network disconnect occurs. If the Data Aggregator and Data Collector processes are running during a disconnect, you can make configuration changes to the Data Aggregator installation. In this case, polling continues on the Data Collector host according to the configuration that existed before the network disconnect. Once the connection between the Data Aggregator and the Data Collector hosts reestablishes, Data Collector downloads the new configuration and adjusts polling accordingly.

For example, you make one of the following configuration changes:

- Change the expression an SNMP vendor certification uses to calculate a value on a metric family.
- Change the metric family to poll a new operational metric.

When the connection between the Data Aggregator and the Data Collector hosts is broken, the changes cannot take effect. After reconnection, Data Collector begins polling the new SNMP MIB objects used in the new expression or in calculating the new operational metric.

## Assign Data Collectors to Tenants and IP Domains

You can manage the tenant or IP domain assignment for a data collector instance.

This procedure requires the Administer Data Sources role right and the Administrator product privilege for the data aggregator data source.

### NOTE

You can automate assigning data collectors to a tenant and IP domain by way of the data aggregator REST web services or use this API in your scripts for assigning data collectors.

For more information about this API, see [Manage Discovery Profiles using REST](#).

**Prerequisite:** The data collector is not already monitoring devices (the number of polled devices and components (**Polled Items**) is zero). The **Polled Items** column on the **Data Collectors** page lists the number of polled devices and components that are assigned to the data collector instance.

### Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Data Collectors**.  
The **Data Collectors** page appears.
2. Select the data collector instance that you want to assign to a tenant and IP domain, and then click **Assign**.  
The **Assign Data Collector** dialog opens.
3. Complete the following fields, and then click **Save**:
  - **Tenant**  
Select the tenant to which you want to assign this data collector instance from the drop-down list. Tenants represent customer environments that managed service providers administers. Each tenant environment is independent and effectively functions as a separate instance of NetOps Portal. Each instance can contain multiple users and roles that are not shared among tenants. The Default Tenant represents the tenant space for the managed service provider within the managed infrastructure. Assign the data collector to the Default Tenant if you are not deploying multi-tenancy. In a single-tenant environment, the Default Tenant is the space used for monitoring the entire infrastructure.  
All monitored devices and components that this data collector instance discovers are automatically associated with this tenant.  
For more information about multi-tenancy, see [Multi-tenancy](#).  
To use the Default Tenant, leave **Default Tenant** selected.
  - **IP Domain**  
Select the IP domain to which you want to assign this data collector instance. You can assign each data collector instance that does discovery requests to only *one* IP domain. IP domains are logical groupings that identify data from different devices and networks. You monitor IP addresses with associated interfaces or applications that belong to separate customer networks separately. When combined with appropriate permissions, you monitor IP domains from a single console, but you view data only for the domains that you monitor. The managed devices and components that this data collector instance discovers are automatically associated with this IP domain.  
For more information about IP domains, see [IP Domains](#).
  - **Use this data collector in standby mode**  
Specifies whether to use this data collector as a fault-tolerant data collector.  
For more information, see [Configure Data Collectors for Fault Tolerance](#).

The data collector is assigned to the tenant and IP domain.

## Configure Data Collectors for Fault Tolerance

You can configure data collectors for fault tolerance. The standby data collector allows data collection to continue in the event of a data collector failure, minimizing data gaps and maximizing network visibility.

To reduce data collector downtime, configure the data collectors for fault tolerance by assigning one or more data collectors to serve as standbys ("Standby" status) for one or more active data collectors ("Active" status). This creates a standby group of data collectors.

Standby groups are composed of one or more data collectors in standby mode and one or more data collectors that are actively polling. When an active data collector fails, the standby data collector swaps in to be the active data collector, and the recently failed data collector becomes the standby.

You can configure the standby groups as pairs where each active data collector has a standby or as groups where one or more data collectors are standing by for a specified set of active data collectors. Assigning a data collector to a standby group puts it in standby mode ("Standby" status).

If you have added a connection to DX NetOps Virtual Network Assurance (VNA) (VNA Gateway), and you have added this Gateway to a data collector that is part of a standby group, and the active data collector fails over to a standby data collector, VNA will reconnect to the data collector that becomes active.

For more information about connections to VNA, see [Manage Connections to Virtual Network Assurance](#).

#### Prerequisites:

- You have [reviewed the hardware requirements for fault-tolerant data collectors](#).
- The data collector is not already monitoring devices (the data collected has no polled items (the **Polled Items** column is zero (0))).
- The data collector is active ("Active" status). Standby data collectors ("Standby" status) can only be members of one standby group.
- The data collector has capacity set.
- The data collector is in the same IP domain as the standby data collector.
- The data collector is not a DX NetOps Mediation Manager (DX NetOps MM) data collector.

#### Follow these steps:

1. Hover over **Administration, Monitored Items Management**, and then click **Data Collectors**.

The **Data Collectors** page appears.

The **Configured** column shows the standby data collector host that is standing by for the host that makes up the standby group. The **Host Name/Standby For Hosts** column shows the standby groups. The following image shows a standby data collector:

## Data Collectors

<input type="checkbox"/>	Configured	Tenant	Host Name/Standby For Hosts <span>↑</span>	Address	ID
<input type="checkbox"/>	✓ Active	Branch	dcfailover-dc	10.35.209.92	dcfailover
<input type="checkbox"/>	✓ Active	Branch	dcfailover-dc2	10.35.209.91	dcfailover
<input type="checkbox"/>	✓ Active	Branch	dcfailover-dc3	10.35.209.108	dcfailover
<input type="checkbox"/>	✕ Standby	Branch	dcfailover-dc4 for dcfailover-dc3	10.35.209.97	dcfailover

For more information about these columns, see [View System Status](#).

2. Select the active data collector ("Active" status) that you want to put in standby mode, and then click **Assign**. The **Assign Data Collector** dialog opens, as shown in the following image:

## Assign Data Collector

### — Assign Data Collector to IP Domain —

Once you have assigned the data collector, you cannot make changes to the tenant/IP domain assignment if the collector has any polled items associated with it.

Data Collector:

dcfailover-dc4

Tenant:

Branch ▼

IP Domain:

Diet-Coke ▼

### — Assign Data Collector to Standby Group —

Data collectors that are in standby mode (standby data collectors) do not participate in discoveries or polling, but are used as failover by one or more active data collectors. If an active data collector fails, the standby data collector takes over processing, and the previously-active data collector changes to standby mode.



Use this data collector in standby mode

### — Data Collectors in this Standby Group —

Available

dcfailover-dc

dcfailover-dc2

Selected

dcfailover-dc3



Quick Filter ▼

Quick Filter ▼

Save

Cancel

3. In the **Assign Data Collector To Standby Group** section, complete the following steps, and then click **Save**:
  - a. For **Use this data collector in standby mode** specify whether to use this data collector as a standby data collector. If the data collector has no polled items, you can use standby data collectors to create standby groups. Data collectors in standby status swap in for any data collector that goes down within a configured group. When the standby data collector takes over, the host that was previously the standby data collector becomes the active data collector, and the previously active host becomes the standby data collector. Fault-tolerant data collectors can be in active or standby (running) status.
    - **Selected:** Use the data collector as a standby data collector, and populate the **Available** list in the **Data Collector in this Standby Group** section with active data collectors ("Active" status) that have the same capacity and that are within the selected IP domain.
    - **Cleared:** Do not use the data collector as a standby data collector.

**Default:** Cleared
  - b. In the **Data Collectors in This Standby Group** section, select the data collectors that you want to take over processing (put them in standby mode, or "Standby" status) if the active data collector ("Active" status) fails from the **Available** list, and move these data collectors to the **Selected** list using the arrow. You can modify the group of active data collectors for which this data collector will stand by.

The selected data collectors are configured as standby data collectors ("Standby" status) for the standby group. The **Data Collectors** page appears.

## Modify Maximum Memory Usage for the Data Aggregator and Data Collector

If you add more memory to the system, modify the maximum values to match the changes.

The default maximum memory for the data aggregator and the data collectors is 80% of the total system memory. Apache ActiveMQ uses 20% of total memory on both components. Both components reserve 2 GB of memory for the operating system. If you add more memory to the system, modify the maximum values to match the changes.

### Follow these steps:

1. Open a console and display the total memory usage by issuing the following command:

```
more /proc/meminfo
```

2. Make a note of this total memory.
3. Modify the maximum memory for the data aggregator by performing the following steps:

- a. Access the following file:

```
<installation_directory>/apache-karaf/bin/setenv
```

- **installation\_directory**  
The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

- b. Modify the following line:

```
IM_MAX_MEM=numberUnit
```

- **numberUnit**  
Specifies the maximum amount of memory. *number* is a positive integer. *Unit* is "G" for GB or "M" for MB. Reserve 2 GB for operating system operations and reserve 20 percent for AMQ. Use the following formula to determine the value:

```
total memory * 80% - 2 GB
```

### Example:

```
33554432 KB * 80% - 2G = 24 GB
```

```
IM_MAX_MEM=24G
```

- c. Save the file.
  - d. Do one of the following steps:
    - Start the Data Aggregator service:
 

```
systemctl start dadaemon
```
    - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:
 

```
<installation_directory>/scripts/dadaemon activate
```

      - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** /opt/IMDataAggregator
  - e. For the memory-setting change to persist during a data aggregator upgrade, modify the /etc/DA.cfg file, replacing the updated value for the da.memory property.  
**Example:** da.memory=24G
4. Modify the maximum memory for all data collector hosts by performing the following steps:
    - a. Access the setenv file.
    - b. Modify the following line, and then save the file:
 

```
IM_MAX_MEM=numberUnit
```

      - **numberUnit**  
Specifies the maximum amount of memory. *number* is a positive integer. *Unit* is “G” for GB or “M” for MB. Reserve 2 GB for operating system operations and reserve 20 percent for AMQ. Use the following formula to determine the value:
 

```
total memory * 80% - 2 GB
```
    - c. Restart the data collector hosts by issuing the following commands:
 

```
systemctl stop dcmd
systemctl start dcmd
```
    - d. For the memory-setting change to persist during a data collector upgrade, modify the /opt/DCM.cfg file, replacing the updated value for the IM\_MAX\_MEM property.

The maximum amount of memory is configured.

### Example:

The following example configures the maximum memory usage for the data aggregator where the total memory is 48 GB:

1. Open a console, and then issue the following command:

```
more /proc/meminfo
```

The following result appears:

```
MemTotal: 50331648KB
```

2. Calculate the maximum memory:  
Equation: total memory \* 80% - 2 GB = maximum memory  
Solution: 50331648 KB \* 80% - 2 GB ≈ 38 GB
3. Access the setenv file.
4. Modify the following line, and then save your changes:

```
IM_MAX_MEM=<numberUnit>
```

- **numberUnit**  
Specifies the maximum amount of memory. *number* is a positive integer. *Unit* is “G” for GB or “M” for MB. Reserve 2 GB for operating system operations and reserve 20 percent for AMQ. Use the following formula to determine the value:



```
total memory * 80% - 2 GB
```

**Example:** 38G

5. Restart the data aggregator.
6. Modify the `da.memory` value in the `/etc/DA.cfg` file, and then save your changes:

```
da.memory=38G
```

The maximum amount of memory is modified.

## Modify the External ActiveMQ Memory Limit

Fine tune ActiveMQ as needed.

The data aggregator installer calculates the memory that your system requires to accommodate the Apache ActiveMQ process. However, you can fine-tune ActiveMQ on the data aggregator by manually modifying the memory limit settings.

The following is a list of example circumstances for modifying the settings:

- When the system memory has changed.
- When the number of the data collectors has changed.
- To optimize the memory settings.
- When you have determined that ActiveMQ's performance is degraded, by monitoring either the JConsole or the DX NetOps Performance Management custom chart with ActiveMQ metrics.

### Follow these steps:

1. Calculate the amount of memory that the ActiveMQ broker requires for the data aggregator based on the following settings:
  - **Maximum java heap size**  
This value is set to 20% system memory by default.  
**Minimum:** 512M
  - **Initial minimum java heap size**  
Calculate this value to be 50% of maximum Java heap size.
  - **Memory limit for all messages**  
Calculate this value to be 50% of the maximum Java heap size.
  - **Memory limit per queue**  
Calculate this value based on the number of data collectors (data collector count).  
**Example:**  
The memory per queue (system memory for all messages)/5/(data collector count).
2. Calculate the amount of memory that the ActiveMQ broker requires for the data collector based on the following settings:
  - **Maximum java heap size (-Xmx)**  
This value is equal to 1GB, which is the recommended minimum.
  - **Initial java heap size (-Xms)**  
This value is equal to 50% of the maximum Java heap size. By default, the value is equal to 512M.
  - **Default Disk Cache Size**  
This value is equal to 50% of the maximum data collector memory. By default, this value is equal to 45 minutes, or 500K, of data when you use a 5-minute poll rate.
  - **Default Drop Rate**  
This value is equal to 10% of the default disk cache size. The checker captures the amount of data in the cache every 30 seconds, and drops 10% of the data when the maximum cache size is reached.
3. Log in to the data aggregator host as the root user or the sudo user.
4. Stop the ActiveMQ broker by issuing the following command:
 

```
service activemq stop
```

## 5. Modify the Java heap size for ActiveMQ:

## a. Access the following file:

```
<DC_installation_directory>/IMDataCollector/scripts/activemq
```

**Example:**

```
/opt/IMDataCollector/scripts/activemq
```

- **DC\_installation\_directory**

The installation directory of the data collector.

**Default:** /opt

## b. Modify the following line:

```
ACTIVEMQ_OPTS_MEMORY=" -Xms788M -Xmx2575M -Xmn394M -server -XX:SurvivorRatio=6 -
XX:+UseConcMarkSweepGC -XX:+UseParNewGC ...
```

Change the - Xms value to be the **Initial minimum java heap size**.

Change the - Xmx value to be the **Maximum Java heap size**.

## c. Save your changes.

## 6. Modify the ActiveMQ memory limit for the producer flow control:

## a. Access the following file:

```
<installation_directory>/IMDataCollector/broker/<apache-activemq-*>/conf/activemq.xml
```

**Example:**

```
/opt/IMDataCollector/broker/apache-activemq-5.18.3/conf/activemq.xml
```

- **DC\_installation\_directory**

The installation directory of the data collector.

**Default:** /opt

- **apache-activemq-\***

The installation directory of Apache ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

b. Locate the following line, and change the value to be the **Memory limit for all messages**:

```
<memoryUsage limit="value"/>
```

c. Locate the following line, and change the value to be the **Memory limit per queue**:

```
<policyEntry queue=">" producerFlowControl="true" memoryLimit="value"/>
```

## d. Save the changes.

## 7. To persist the memory setting change during a data aggregator upgrade, modify the /etc/DA.cfg file, replacing the updated value for the da.activemq.memory property.

**NOTE**

The static location of this file is the location of the data aggregator installation.

**Example:**

```
da.activemq.memory=value
```

## 8. Start the ActiveMQ broker by issuing the following command:

```
service activemq start
```

ActiveMQ starts.

ActiveMQ now uses the new memory limit settings.

## Configure Java Options for the Data Aggregator and the Data Collectors

You can set additional Java options for Apache Karaf and Apache ActiveMQ.

The data aggregator and the data collectors installation preserve these additional Java options on upgrade.

### Follow these steps:

1. Create a file named `_custom` in the `<installation_directory>/custom.d` directory.
  - **installation\_directory**  
The installation directory for the data aggregator and/or the data collectors.  
**Default:**
    - **The data aggregator**  
`/opt/IMDataAggregator`
    - **The data collectors**  
`/opt/IMDataCollector`
2. Inside this file, set the following variables with the Java options that you want to set, and then save your changes:
 

```
KARAF_CUSTOM_JAVA_OPTS="-D<Java_option>=<value> -D<Java_option>=<value>"
ACTIVEMQ_CUSTOM_JAVA_OPTS="-D<Java_option>=<value> -D<Java_option>=<value>"
```

### Examples:

```
KARAF_CUSTOM_JAVA_OPTS="-Dkaraf.optionA=true -Dkaraf.optionB=5"
ACTIVEMQ_CUSTOM_JAVA_OPTS="-Damq.optionA=true -Damq.optionB=5"
```

- **Java\_option**  
Specifies the Java option that you want to configure.  
**Examples:** `karaf.optionA`, `amq.optionA`
- **value**  
Specifies the value for the additional Java option that you want to configure.  
**Example:** `true`

The additional Java options for Apache Karaf and Apache ActiveMQ are configured.

## Configure the Data Aggregator Cleanup

Cleanups remove deleted items from the data aggregator. Removal of the deleted items helps keep NetOps Portal functioning well when there are a lot of deleted items in dynamic environments. Cleanups remove deleted items by pruning unnecessary data from Vertica. The removal of deleted items is less important or impactful in very static environments.

By default, cleanups run daily, but you can configure whether they run. You can also disable the cleanup for attribute instance tables or relationships.

To configure when cleanups occur (how often the data aggregator purges data), see [Schedule Data Purges](#).

### Follow these steps:

1. Go to the following location:
 

```
<installation_directory>/apache-karaf/etc
```

  - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
2. Create the `com.ca.im.dm.core.database.dao.impl.DeletedItemCleanupDAO.cfg` file in this location.
3. Specify the following parameters, and then save the file:
  - `cleanupEnabled`

Specifies whether the cleanup process runs.

**Default:** true

- `attributeCleanupEnabled`

Specifies whether to cleanup attribute instance tables.

**Default:** true

- `relationshipCleanupEnabled`

Specifies whether to cleanup relationships.

**Default:** true

**Example:**

```
cleanupEnabled=true
attributeCleanupEnabled=true
relationshipCleanupEnabled=true
```

## Monitor Data Aggregator System Health

You can monitor system usage and identify performance trends and degradation using the system health dashboards.

Detect abnormal changes by observing the value trend in the views.

To view these dashboards, click **System Health**. You must have the **View System Health Dashboards** role right and be assigned to the **System Health - Data Aggregator** group. By default, only the Administrator has this role right.

Use the following derived rollup metrics:

- **Average**  
The average value across the resolution window.
- **Maximum**  
The maximum value within the resolution window. Use this metric to visualize high peaks.
- **95th Percentile**  
95 percent of the samples within the resolution window have a lower value. Use this metric to smooth out spikes.

## Data Aggregator / Data Collector Health Dashboard

You can track the health of the data aggregator using the **Data Aggregator Health Dashboard**, and the health of the data collector using the **Data Collector Health Dashboard**.

To view this dashboard, hover over **System Health**, and then click **Data Aggregator Health/Data Collector Health**.

The following values indicate that the system is in good health:

- **Data Aggregator/Data Collector Heap Size**  
If the data collector heap size is too high, the data collector is polling too many items.
- **Data Aggregator/Data Collector Heap Percent**  
Less than 50%. The trend might increase slightly due to the storage of debug data in the memory logger.
- **Data Aggregator/Data Collector Throughput**  
Equal to 100%. If the data aggregator throughput does not equal 100%, it might be receiving too much data from the data collectors.
- **Data Aggregator ActiveMQ Broker**
- **Data Collector Pause Time**
- **Data Collector ActiveMQ Broker Enqueue Count**
- **Data Collector ActiveMQ Broker Dequeue Count**
- **Event List**

**NOTE**

The communication between the data aggregator and the data collectors is through Apache ActiveMQ. A high response time on the data aggregator indicates that ActiveMQ is overloaded.

For more information about ActiveMQ communication, see [Authenticate and Encrypt ActiveMQ Communication](#).

## Data Aggregator Polling Dashboard

Use the Data Aggregator Polling dashboard to determine the polling load on the data aggregator and on the data collectors.

Monitor the system health using following views:

- **Polled Item Count**  
To determine the threshold for the polled item count, use the [DX NetOps Performance Management Sizing Tool](#).
- **Calculated Metrics Per Second**  
This value is equal to the value that you identified when setting up your environment, unless there is a change in the metric families or items in your environment.
- **Device Polling Statistics**
- **Stopped Polls by Device**  
Use this view to identify problematic devices, which cause excessive polling, such as when the device is unavailable.
- **Count of Loaded Rows**
- **Data Aggregator ActiveMQ Broker**
- **Event List**

Consider the following when using this dashboard:

- If the data aggregator is polling too many items, redistribute polled items to other data collectors.
- Monitor the polling health and see whether the data aggregator is polling devices too quickly by observing the polling statistics.

**TIP**

If you have configured multiple data collectors, hover over the graph for accurate information for each data collector.

## Data Collector Polling Dashboard

Use the Data Collector Polling dashboard to determine the polling load on the data collectors.

Monitor the system health using following views:

- **Polled Item Count**  
To determine the threshold for the polled item count, use the [DX NetOps Performance Management Sizing Tool](#).
- **Calculated Metrics Per Second**  
This value is equal to the value that you identified when setting up your environment, unless there is a change in the metric families or items in your environment.
- **Data Collector ActiveMQ Broker Enqueue Count**
- **Data Collector ActiveMQ Broker Dequeue Count**
- **Event List**

Consider the following when using this dashboard:

- If the data aggregator is polling too many items, redistribute polled items to other data collectors.

**TIP**

If you have configured multiple data collectors, hover over the graph for accurate information for each data collector.

## Data Aggregator Queries Dashboard

You can monitor the quantity and timing of RIB and OpenAPI queries from the views on the **Data Aggregator Queries** dashboard. The RIB queries and OpenAPI queries views specify the number of corresponding queries for a specific time period.

Monitor the system health using following views:

- **RIB Query Processing Times**  
Less than 120 seconds. The value indicates that the system is in good health.
- **Number of RIB Queries**
- **OpenAPI Query Processing Times**
- **OpenAPI Queries**

## Data Aggregator General Processing Dashboard

The **Data Aggregator General Processing** dashboard monitors all data aggregator activities. To customize the views to fit the needs of your environment, specify thresholds on this dashboard.

The following values or conditions indicate that the system is in good health:

- **Rate Data Loading Times**  
Less than 30 seconds.
- **Roll-Up Calculation Times**  
Spikes in the Roll-Up Calculation Times after 00:30 UTC.  
Daily and weekly processing causes these normal spikes.
- **Baseline Calculation Times**  
Spikes in the Baseline Calculation Times after 00:30 UTC.  
Daily and weekly processing causes these normal spikes.
- **Reporting ETL Processing Times**  
Execution time does not exceed 15 minutes. Consistent processing time unless a significant number of new devices have been discovered.
- **Aggregated Components Calculation Times**
- **Data Repository Maintenance Times**  
Normal trend changes. The maintenance times vary based on the size of the schema.
- **Count of Loaded Rows**
- **Count of Roll-Up Processes**
- **Count of Baseline Processes**
- **Reporting ETL Count of Loaded Rows**  
The count is consistent with the number of items that the system is monitoring. This count tracks the number of dimension items, both device and component items, loaded into the reporting table.
- **Count of Aggregated Components Processes**
- **Data Repository Count of Maintenance Tasks**

## Data Aggregator Event Processing Dashboard

The **Data Aggregator Event Processing** dashboard has built-in thresholds, which can help you to identify issues with event processing.

You can monitor the health of event processing using the views on the **Data Aggregator Event Processing** dashboard.

The following values indicate that the system is in good health:

- **Event Queue**  
Less than two. An event queue size of greater than two indicates that the system is behind, or that the system is at risk.
- **Poll Cycle Percent**  
The threshold evaluation engine continues to run if the poll cycle percent is less than 50%. When you use a 1-minute poll rate, the evaluation has to complete in less than 30 seconds. When you use a 5-minute poll rate, the evaluation has to complete in less than 150 seconds.
- **Events Created and Cleared**  
The sum of created and cleared events during a 5-minute interval is less than 900.
- **Event Evaluations**  
Less than 150,000 during a 5-minute interval.
- **Event Calculation Times**
- **Event Producers by Metric Family**

#### **IMPORTANT**

Data aggregator event processing deteriorates if any of the following conditions occur:

- It is slow to process rules.
- You create and clear more than the recommended number of events.
- It is backing up.

## **Rebalance the Load on the Data Collectors**

Rebalance the workload on the data collectors by transferring the load from one overloaded data collector instance to other data collector instances within the same IP domain.

As a data collector monitors more devices, the capacity can be exceeded and it can become overloaded. You can transfer the workload from one overloaded data collector to other data collectors. You can rebalance the load on data collector in the following ways:

- Automatically. DX NetOps Performance Management automatically rebalances the load among the selected data collector instances.

#### **NOTE**

When the load on the data collectors is rebalanced, DX NetOps Mediation Manager and DX NetOps Virtual Network Assurance devices are not rebalanced. Only SNMP devices are rebalanced.

- Move selected devices from one data collector instance to another.

#### **IMPORTANT**

**Best Practice:** Rebalance or move many items during periods of maintenance or off-peak hours. These operations can negatively affect end-user performance.

## **View Polled Devices and Components**

You can view the number of devices and components that each data collector polls (polled items) from the **Monitored Devices** page. This page also shows the total number of devices that are assigned to each data collector instance, including devices that are not currently polled.

### **Follow these steps:**

1. Open NetOps Portal as an Administrator.
2. Select **Administration, Monitored Items Management, Data Collector**.  
The **Data Collectors** page appears.

3. Click **Monitored Devices** from the **Monitored Inventory** menu.  
The **Monitored Devices** page appears.

### **Automatically Rebalance the Load on Data Collector**

You can rebalance devices (polled items) between data collector instances within the same IP domain. Rebalancing polled items restarts the baseline average calculations for those items.

**Prerequisites:** The data collector instances that you want to rebalance are within the same IP domain.

#### **Follow these steps:**

1. On the **Data Collectors** page, select the data collector instances that you want to rebalance, and then click **Rebalance**.

A confirmation dialog displays the current device and polled item count for each selected data collector and the proposed resulting device and polled item counts.

#### **NOTE**

You can move devices only to data collector instances that can contact them.

2. Click **Yes**.

### **Move Selected Devices to a Specific Data Collector Instance**

You can move devices only to data collector instances that can contact them. Only data collector instances that are within the same IP domain are included for selection. Moving devices restarts the baseline average calculations for the those devices.

#### **Follow these steps:**

1. On the **Data Collectors** page, select the data collector instance from which you want to move selected devices.
2. In the **Devices** table, select the devices that you want to move to another data collector instance, and then click **Move Devices**.  
The **Move Devices to Selected Data Collector** dialog opens.
3. For **Target Data Collector**, select the data collector instance to which you want to move your selected devices from the drop-down list, and then click **Yes**.

## **Restore the Data Aggregator**

You can restore the data aggregator files that you have backed up. If the data repository remains intact, you can restore only the data aggregator component.

The data aggregator can remain running while you restore the backup. You can drop the files that are backed up in the desired directories while the data aggregator is running.

Back up and restore the data aggregator into the same DX NetOps Performance Management version.

#### **Prerequisites:**

- You have root or sudo privileges.
- [You have backed up the data aggregator.](#)

#### **Follow these steps:**

1. (Optional) If the data aggregator Apache Karaf service is not running, uninstall the existing data aggregator and then reinstall it.  
For more information, see [Uninstall the Data Aggregator](#).
2. Restore the `deploy` directory by issuing the following command:

```
cp -r <DA_Backup>/deploy_backup/deploy <installation_directory>/apache-karaf
```



**Example:**

```
cp -r /tmp/DA_Backup/deploy_backup/deploy /opt/IMDataAggregator/apache-karaf
```

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

If prompted, overwrite the existing files.

3. Restore the certifications in the `certifications` directory by issuing the following command based on your environment:

**Non-fault-tolerant environment:**

```
cp -r <DA_Backup>/cert_backup/certifications <installation_directory>/data
```

**Example:**

```
cp -r /tmp/DA_Backup/cert_backup/certifications /opt/IMDataAggregator/data
```

**Fault-tolerant environment:**

```
cp -r <DA_Backup>/cert_backup/certifications <DASharedRepo>
```

**Example:**

```
cp -r /tmp/DA_Backup/cert_backup/certifications /DASharedRepo
```

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **DASharedRepo**

Specifies the shared data directory that you set up when you installed the fault-tolerant data aggregators.

**Example:** /DASharedRepo

If prompted, overwrite the existing files.

4. Restore the custom device subtype XML file, `DeviceTypes.xml`, in the `devicetypes` directory by issuing the following command based on your environment:

**Non-fault-tolerant environment:**

```
cp <DA_Backup>/CustomDeviceType_backup/DeviceTypes.xml <installation_directory>/data/custom/devicetypes/DeviceTypes.xml
```

**Example:**

```
cp /tmp/DA_Backup/CustomDeviceType_backup/DeviceTypes.xml /opt/IMDataAggregator/data/custom/devicetypes/DeviceTypes.xml
```

**Fault-tolerant environment:**

```
cp <DA_Backup>/CustomDeviceType_backup/DeviceTypes.xml <DASharedRepo>/custom/devicetypes/DeviceTypes.xml
```

**Example:**

```
cp /tmp/DA_Backup/CustomDeviceType_backup/DeviceTypes.xml /DASharedRepo/custom/devicetypes/DeviceTypes.xml
```

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **DASharedRepo**

Specifies the shared data directory that you set up when you installed the fault-tolerant data aggregators.

**Example:** /DASharedRepo

If prompted, overwrite the existing files.

5. Restore the `etc` directory files and the HTTPS configuration by issuing the following command:

```
cp -r <DA_Backup>/etc_backup/etc <installation_directory>/apache-karaf
```

**Example:**

```
cp -r /tmp/DA_Backup/etc_backup/etc /opt/IMDataAggregator/apache-karaf
```

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

If prompted, overwrite the existing files.

6. Restore the ActiveMQ configuration file, `activemq.xml`, by issuing the following command:

```
cp <DA_Backup>/activemq_backup/activemq.xml <installation_directory>/broker/<apache-activemq-*>/conf/activemq.xml
```

**Example:**

```
cp /tmp/DA_Backup/activemq_backup/activemq.xml /opt/IMDataAggregator/broker/apache-activemq-5.18.3/conf/activemq.xml
```

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **apache-activemq-\***

The installation directory for Apache ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

7. If you have enabled HTTPS for the data aggregator, restore the Java truststore file directory by issuing the following command:

```
cp <DA_Backup>/cacerts <installation_directory>/jre/lib/security/cacerts
```

**Example:**

```
cp /tmp/DA_Backup/cacerts /opt/IMDataAggregator/jre/lib/security/cacerts
```

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

8. If you have secured Apache ActiveMQ communication between the data aggregator and the data collectors, restore the `broker.js` and `client.ts` files by issuing the following commands:

```
cp <DA_Backup>/broker.ks <installation_directory>/broker/<apache-activemq-*>/conf/
cp <DA_Backup>/client.ts <installation_directory>/broker/<apache-activemq-*>/conf/
```

**Example:**

```
cp /tmp/DA_Backup/broker.ks /opt/IMDataAggregator/broker/apache-activemq-5.18.3/conf/
cp /tmp/DA_Backup/client.ts /opt/IMDataAggregator/broker/apache-activemq-5.18.3/conf/
```

– **DA\_Backup**

Specifies the directory path and name of the backup directory for the data aggregator.

**Example:** /tmp/DA\_Backup

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **apache-activemq-\***

The installation directory for Apache ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

9. Wait for a few minutes for the data aggregator to synchronize automatically with NetOps Portal. When the connections between the data aggregator and the data collector hosts are established, the data collector hosts resume polling.

**NOTE**

To restore the data collector to a previous state, uninstall and reinstall it.

For more information, see [Uninstall the Data Collector](#).

The data aggregator is restored.

## View Data Aggregator Details

View the total number of manageable and pingable devices that the data aggregator monitors for all tenants.

Administrators can view individual device totals for each tenant, as well as the data aggregator version and build number. Tenant administrators can only view the total number of manageable and pingable devices that the data aggregator monitors for their tenant.

**Follow these steps:**

1. Log in to NetOps Portal as an Administrator.
2. Hover over **Administration**, **Data Sources**, and then click a data aggregator data source.
3. Click **Data Aggregator** from the **System Status** menu.  
The **Data Aggregator List** page opens.

## Configure the Failover Settings for Fault Tolerance

Configure the failover settings for fault tolerance.

By default, during failover, the data aggregator that is ready to take over if the *active* data aggregator goes offline (the data aggregator that is in "Ready" status) has 45 minutes to start. If it does not start within 45 minutes, Consul tries to start the other data aggregator host. This process repeats for each host every 45 minutes until a host start.

**NOTE**

If failover occurs, but the ActiveMQ process is still running, you can manually stop ActiveMQ.

For more information, see [Restart the ActiveMQ Broker](#).

**Follow these steps:**

1. Edit the <installation\_directory>/consul-ext/conf/config.json file.
  - **installation\_directory**  
The installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

2. Edit the following parameters:

– `startwait`

Defines the time for when the data aggregator in "Ready" status starts (**s** second, **m** minute, **h** hour).

**Best Practice:** Observe how much time passes between when you issue the command to start the Data Aggregator service and when the service is available. Adjust this parameter as appropriate.

For more information about how to start the Data Aggregator service, see [Restart the Data Aggregator](#).

**NOTE**

If the data aggregator always takes longer than 20 to 30 minutes to start, the hardware might be under-resourced, and DX NetOps Performance Management stops functioning. Check the sizing requirements for the data aggregator.

For more information, see the [DX NetOps Performance Management Sizing Tool](#).

**IMPORTANT**

To avoid data loss or system malfunction, set this parameter to 45 minutes or higher.

– `failwait`

Defines the failover wait time (**s** second, **m** minute, **h** hour). Failover happens only when the active data aggregator is unresponsive to the fault-tolerance heartbeat for longer than the set failover wait time. If you have limited network availability with periodical network outages or system thrashing that may last several minutes, then increase the failover wait time.

**Default:** 5 minutes

**IMPORTANT**

To avoid data corruption or data loss, set this parameter to 5 minutes or higher.

3. Save your changes.

## Install or Uninstall the Proxy Server

In a fault-tolerant environment, configure a proxy server (DAProxy) to send traffic from NetOps Portal to the active data aggregator.

Installing the proxy server installs Consul and Traefik.

Use the following process to install the proxy server:

1. [Verify the prerequisites](#).
2. (If you do not have root access to install and run the proxy server) [Configure the sudo user account for the proxy server](#).
3. [Install the proxy server](#).
4. (If you plan to use encrypted certificate files) [Install Nginx, and then configure the connection from NetOps Portal to the proxy server](#).

When you no longer need the proxy server, you can [uninstall the proxy server](#).

### Verify the Prerequisites

Before installing the proxy server, ensure that you have configured the failover settings for fault tolerance.

For more information, see [Configure the Failover Settings for Fault Tolerance](#).

### (Optional) Configure the Sudo User Account for the Proxy Server

If you do not have root access to install and run the proxy server, configure the sudo user account.

#### **Follow these steps:**

1. Locate the `/etc/sudoers` file on the proxy server host.

2. Add a command alias with the following permissions to the file:

```
Cmnd_Alias CA_DAPROXY = /tmp/installDAProxy.bin,/sbin/service daproxy *,/sbin/service
  consul *,/opt/CA/daproxy/Uninstall/Uninstall
## Allows the daproxy user to manage the proxy server
dasudouser_name ALL = CA_DAPROXY
```

This command alias details the commands that the sudo user must be able to run.

The sudo user account for the proxy server is configured.

## **Install the Proxy Server**

### **NOTE**

If a new version of Consul updates the configuration of the Consul service, in a fault-tolerant data aggregator environment, the data aggregator starts only *after* you upgrade the proxy server and the data aggregators.

### **Follow these steps:**

1. Log in to the proxy server host as the root or sudo user.
2. Copy the `installDAProxy.bin` file to the `/tmp` directory.
3. Change to the `/tmp` directory:
 

```
cd /tmp
```
4. Change permissions for the installation file:
 

```
chmod a+x installDAProxy.bin
```
5. Launch the installation using one of the following options:
  - If you have root access to install the proxy sever, launch the installation by issuing the following command:
 

```
./installDAProxy.bin -i console
```
  - If you have configured the sudo user account to install and run the proxy server, launch the installation by issuing the following command:
 

```
sudo ./installDAProxy.bin -i console
```

The installation is initiated.  
The **Choose Install Set** line appears.
6. Complete the following prompts:
  - **ENTER THE NUMBER FOR THE INSTALL SET, OR PRESS <ENTER> TO ACCEPT THE DEFAULT**  
Specifies the number for the option for a typical installation that includes both the proxy server and Consul.  
The **Specify the owner of the proxy services** line appears.
  - **Username (Default: consul)**  
Specifies the user name for the proxy server. The services run the proxy server processes using this user. This user must already exist on the system.  
**Default:** `consul`  
The **Choose Install Folder** line appears.
  - **ENTER AN ABSOLUTE PATH, OR PRESS <ENTER> TO ACCEPT THE DEFAULT**  
Specifies the location where you want the installer to create the proxy server directories.  
**Default:** `/opt/CA/daproxy`  
The **Specify the first DA host** line appears.
  - **First DA host**  
Specifies the hostname of the active data aggregator (the data aggregator that is in "Active" status).  
The **Choose second DA host** line appears.
  - **Second DA host**  
Specifies the hostname of the data aggregator that is ready to take over if the *active* data aggregator goes offline (the data aggregator that is in "Ready" status).

The **Specify the DA HTTP API port** line appears.

- **DA HTTP API port (Default: 8581)**

Specifies the port for HTTP communication between the data aggregator API and the fault-tolerant data aggregator proxy.

**Default:** 8581

The **Specify the Consul port** line appears.

- **Consul Port (Default: 8500)**

Specifies the port for communication with Consul. In fault-tolerant data aggregator environments, this port enables communication between the fault-tolerant proxy server and the data aggregators. This port must be open on the fault-tolerant proxy and on the data aggregators.

**Default:** 8500

(If multiple public IP addresses are configured) The **Specify the Consul host IP address** line appears.

- (If multiple public IP addresses are configured) **ENTER THE NUMBER FOR YOUR CHOICE, OR PRESS <ENTER> TO ACCEPT THE DEFAULT**

Specifies the number for the option to specify a different IP address.

(If you entered the number for the option to specify a different IP address) The **Specify the Consul host IP address** line appears.

- (If you entered the number for the option to specify a different IP address) **Consul IP address**

Specifies the bind address that the Consul agents use to communicate with each other. Specify an address that the other two hosts in the Consul cluster can reach. The Consul agents include the proxy host and both data aggregators in the cluster.

The proxy server is installed.

### **(Optional) Install Nginx and Configure the Connection from NetOps Portal to the Proxy Server**

If you plan to use encrypted certificate files, complete the following:

1. [Install Nginx.](#)
2. [Configure the connection from NetOps Portal to the proxy server using Nginx as the reverse proxy.](#)

### **Uninstall the Proxy Server**

If necessary, you can uninstall the proxy server.

#### **Follow these steps:**

1. Log in to the proxy server host as the root or sudo user.
2. Issue the following command:

```
<Proxy_installation_directory>/daproxymy/Uninstall/Uninstall
```

**Example:**

```
/opt/CA/daproxymy/Uninstall/Uninstall
```

- **proxy\_installation\_directory**

The installation directory of the proxy server.

**Default:** /opt/CA/daproxymy

The proxy server is uninstalled.

### **Install Nginx**

You can use Nginx instead of Traefik as your reverse proxy.

(21.2.5 and higher) This procedure is part of the process of configuring a fault-tolerant environment, which can include enabling the fault-tolerant data aggregators to use HTTPS.

For more information, see [Install Fault-Tolerant Data Aggregators](#).

**Follow these steps:**

1. Verify that you can install Nginx by issuing the following command:  

```
yum list nginx.x86_64
```
2. If the command returns an empty list (Nginx is not available), install the Extra Packages for Enterprise Linux (EPEL) repo (epel-release ) by issuing the following command:  

```
yum install -y epel-release
```
3. If Nginx is available or it is not and you have installed the epel-release repo, install Nginx by issuing the following command:  

```
yum install -y nginx.x86_64
```

Nginx is installed.

**Next step:** [Configure the connection from NetOps Portal to the proxy server using Nginx as the reverse proxy](#).

## Configure the Connection from NetOps Portal to the Proxy Server Using Nginx as Reverse Proxy

You can configure the connection from NetOps Portal to the proxy server, using Nginx as the reverse proxy instead of Traefik. The proxy server installation installs Traefik by default.

Configure the connection based on your requirements:

- [Configure the Connection as HTTP](#)
- (If you require that the data aggregator proxy server be configured for HTTPS and it will be using encrypted PEM certificate files) [Configure the Connection as HTTPS](#)

**Prerequisite:** Nginx is installed.

### Configure the Connection as HTTP

Use the following process to configure the connection as HTTP:

1. [Configure the Proxy Server](#)
2. [Restart the Proxy Service](#)
3. [Test the HTTP Connection to the Proxy Server](#)

### Configure the Proxy Server

**Follow these steps:**

1. Edit the `/etc/nginx/nginx.conf` file that Nginx creates by default.
2. Comment out the lines for the `server { ... }` section that is inside the `http { ... }` section using the `#` character.
3. Confirm that the following line exists inside the `http { ... }` section:  

```
include /etc/nginx/conf.d/*.conf;
```

 If it does not, add it after the `default_type` line.
4. Create the `/etc/nginx/conf.d` directory.
5. Edit the `/etc/nginx/conf.d/daproxy.conf` file.
6. Add the following lines to the file:  

```
upstream backend {
    ip_hash;
    server <DA1 FQHN or IP>:8581;
```

```

server <DA2 FQHN or IP>:8581;
}

server {
    listen 8581 default_server;
    listen [::]:8581 default_server;
    server_name <FQHN of daprox>;
    include /etc/nginx/default.d/*.conf;
    client_max_body_size 0;
    location / {
        proxy_pass http://backend;
    }
}

```

The proxy server is configured.

### **Restart the Proxy Service**

**Follow these steps:**

1. Disable the Traefik (daprox) service by issuing the following commands:

```

systemctl stop daprox
systemctl disable daprox

```

2. Enable then start the Nginx service by issuing the following commands:

```

systemctl enable nginx
systemctl start nginx

```

### **Test the HTTP Connection to the Proxy Server**

Test the connection to the proxy server using the following URL:

```
http://<daprox_hostname>:8581/
```

### **Configure the Connection as HTTPS**

Use the following process to configure the connection as HTTPS:

1. (If a keystore file already exists) [Back up the Existing Keystore File](#)
2. [Manage the Keystore File and Certificates](#)
3. [Configure the Proxy Server](#)
4. [Restart the Proxy Service](#)
5. [Test the HTTPS Connection to the Proxy Server](#)

### **Back up the Existing Keystore File**

If a keystore file already exists, back it up by issuing the following command:

```
cp <keystore> <keystore>.bak
```

#### **Example:**

```
cp /opt/CA/daprox/conf/cacerts /opt/CA/daprox/conf/cacerts.bak
```

- **keystore**  
Specifies the name of the existing keystore file.  
**Example:** /opt/CA/daprox/conf/cacerts



## Manage the Keystore File and Certificates

### Follow these steps:

1. Generate a private key and a public, self-signed certificate by issuing the following command:

```
keytool -genkeypair -ext SAN=dns:<fully_qualified_hostname>,dns:<hostname> -keystore <keystore> -storepass <store_password> -keyalg RSA -keysize 2048 -keypass <key_password> -alias <alias_name>
```

#### Example:

```
keytool -genkeypair -ext SAN=dns:myDaproxyHost.my.domain.net,dns:myDaproxyHost -keystore /opt/CA/daproxy/conf/cacerts -storepass changeit -keyalg RSA -keysize 2048 -keypass changeit -alias daproxykey
```

Specify the proxy host name with and without the domain within the SAN= argument.

Note your entries for the following variables:

- **fully\_qualified\_hostname**  
Specifies the fully qualified host name of the server. Enter the same value when you are prompted for your first and last name.
- **hostname**  
Specifies the host name of the server without the domain.
- **keystore**  
Specifies the name of the keystore file to create.  
**Example:** /opt/CA/daproxy/conf/cacerts
- **store\_password**  
Specifies a secure password for the keystore.
- **key\_password**  
Specifies a secure password for the private key.
- **alias\_name**  
Specifies an alias for the keystore entry created for the self-signed certificate.  
**Example:** daproxy

2. Proceed through the security prompt questions and confirm your responses.

#### IMPORTANT

If you will be creating a CA-signed certificate from this self-signed certificate, use the owner/issuer details (organization, country, province/state, etc.) that match the values that your organization requires.

#### Sample Output:

```
What is your first and last name?
[Unknown]: myDaproxyHost.my.domain.net
What is the name of your organizational unit?
[Unknown]: MyOrgUnit
What is the name of your organization?
[Unknown]: MyOrg
What is the name of your City or Locality?
[Unknown]: MyCity
What is the name of your State or Province?
[Unknown]: MyStateFullName
What is the two-letter country code for this unit?
[Unknown]: US
Is CN=myDaproxyHost.my.domain.net, OU=MyOrgUnit, O=MyOrg, L=MyCity, ST=MyStateFullName, C=US correct?
[no]: yes

Warning:
The JKS keystore uses a proprietary format. It is recommended to migrate to PKCS12 which is an industry
standard format using "keytool -importkeystore -srckeystore /opt/CA/daproxy/conf/cacerts -destkeystore /
opt/CA/daproxy/conf/cacerts -deststoretype pkcs12".
```

3. Verify that the key generated by issuing the following command:

```
keytool -list -keystore <keystore> -storepass <store_password> |grep daproxy
```

**Example:**

```
keytool -list -keystore /opt/CA/daproxy/conf/cacerts -storepass changeit |grep daproxy
```

**– keystore**

Specifies the name of the keystore file to create.

**Example:** /opt/CA/daproxy/conf/cacerts

**– store\_password**

Specifies a secure password for the keystore.

**Default:** changeit

**Sample output:**

```
daproxykey, Jun 24, 2020, PrivateKeyEntry,
```

**4. If a certificate file already exists, back it up by issuing the following command:**

```
cp <certificate_filename> <certificate_filename>.bak
```

**Example:**

```
cp /tmp/daproxyss.cer /tmp/daproxyss.cer.bak
```

**– certificate\_filename**

Specifies the name of the certificate file.

**Example:** /tmp/daproxyss.cer

**5. Export the self-signed certificate from the keystore by issuing the following command:**

```
keytool -exportcert -keystore <keystore> -storepass <store_password> -alias <alias_name> -file <filename>
```

**Example:**

```
keytool -exportcert -keystore /opt/CA/daproxy/conf/cacerts -storepass changeit -alias daproxkey -file /tmp/daproxyss.cer
```

**– keystore**

Specifies the name of the keystore file that you created previously.

**Example:** /opt/CA/daproxy/conf/cacerts

**– store\_password**

Specifies the password for the keystore that you created previously.

**Default:** changeit

**– alias\_name**

Specifies the alias for the keystore entry created previously for the self-signed certificate.

**Example:** daproxkey

**– filename**

Specifies the file to which the certificate is exported. Specify a full path that places the file outside the current directory.

**Example:** /tmp/daproxyss.cer

**Sample output:**

```
Certificate stored in file </tmp/daproxyss.cer>
```

Warning:

The JKS keystore uses a proprietary format. It is recommended to migrate to PKCS12 which is an industry standard format using

```
"keytool -importkeystore -srckeystore /opt/CA/daproxy/conf/cacerts -destkeystore /opt/CA/daproxy/conf/cacerts -deststoretype pkcs12".
```

```
ls -alrt /tmp
```

```
-rw-r--r--. 1 root root 945 Jun 24 13:56 daproxkeyss.cer
```

**6. Extract the private certificate and key as a .pem file by completing the following steps:**

- a. If a keystore file already exists (for example, cacerts.p12 ), back it up by issuing the following command:

```
cp <keystore>.p12 <keystore>.p12.bak
```

**Example:**

```
cp /opt/CA/daproxy/conf/cacerts.p12 /opt/CA/daproxy/conf/cacerts.p12.bak
```

- **keystore**

Specifies the name of the keystore file that you created previously.

**Example:** /opt/CA/daproxy/conf/cacerts

- b. Convert the keystore to unencrypted PEM (PKCS8) or PKCS12 file format by issuing the following command:

```
keytool -importkeystore -srckeystore <source_keystore> -srcstorepass store_password -
srckeypass <source_key_password> -srcalias <source_alias> -destalias <destination_alias> -
destkeystore <destination_keystore>.p12 -deststoretype PKCS12 -deststorepass <key_password> -
destkeypass <destination_key_password>
```

**Example:**

```
keytool -importkeystore -srckeystore /opt/CA/daproxy/conf/cacerts -srcstorepass changeit -srckeypass
changeit -srcalias daproxykey -destalias daproxykeyp12 -destkeystore /opt/CA/daproxy/conf/cacerts.p12
-deststoretype PKCS12 -deststorepass changeit -destkeypass changeit
```

Note your entries for the following variables:

- **source\_keystore**

Specifies the name of the keystore file that you created previously.

**Example:** /opt/CA/daproxy/conf/cacerts

- **store\_password**

Specifies the password for the keystore that you created previously.

- **source\_key\_password**

Specifies the password for the private key that you created previously.

- **source\_alias**

Specify the alias for the keystore entry created previously for the self-signed certificate.

**Example:** daproxy

- **destination\_alias**

Specifies an alias for the resulting keystore entry created for the self-signed certificate.

**Example:** daproxy

- **destination\_keystore**

Specifies the name of the resulting keystore file to create. Specify a full path.

**Example:** /opt/CA/daproxy/conf/cacerts

- **key\_password**

Specifies the password for the resulting keystore. Specify a secure password.

- **destination\_key\_password**

Specifies the password for the resulting private key. Specify a secure password.

**Sample output:**

```
Importing keystore /opt/CA/daproxy/conf/cacerts to /opt/CA/daproxy/conf/cacerts.p12...
```

- c. (Optional) If you have previously extracted the key to a file, remove or rename the file by issuing the following command:

```
rm -f /tmp/<key_name>.pem
```

**Example:**

```
rm -f /tmp/daproxyss_key.pem
```

- **key\_name**

Specifies the name of the key.

**Example:** /tmp/daproxyss\_key

- d. Extract the key as a .pem file by issuing the following command:

**NOTE**

The `openssl` command works on Linux only.

For more information about the openssl command line utility, see [the Red Hat documentation](#).

```
openssl pkcs12 -in <destination_keystore>.p12 -nodes -nocerts -out /tmp/<key_name>.pem
```

**Example:**

```
openssl pkcs12 -in /opt/CA/daproxy/conf/cacerts.p12 -nodes -nocerts -out /tmp/daproxyss_key.pem
```

- **destination\_keystore**

Specifies the name of the resulting keystore file to create. Specify a full path.

**Example:** /opt/CA/daproxy/conf/cacerts

- **key\_name**

Specifies the name of the key.

**Example:** /tmp/daproxyss\_key

**Output:**

```
Enter Import Password: <enter password value you passed in "-deststorepass" above>
```

```
MAC verified OK
```

- e. (Optional) If you have previously written the certificate as a .pem file, remove or rename the file by issuing the following command:

```
rm -f /tmp/<certificate_filename>.pem
```

**Example:**

```
rm -f /tmp/daproxyss.pem
```

- f. Write the proxy server certificate as a .pem file by issuing the following command:

```
openssl x509 -inform der -in /tmp/certificate_filename.cer -out /tmp/<certificate_filename>.pem
```

**Example:**

```
openssl x509 -inform der -in /tmp/daproxyss.cer -out /tmp/daproxyss.pem
```

- g. Verify the list of proxy server certificate files for the Proxy service (daproxy ) by issuing the following command:

```
ls -alrt /tmp/daproxyss*
```

**Sample output:**

```
-rw-r--r--. 1 root root 961 Jul 23 11:06 /tmp/daproxyss.cer
-rw-r--r--. 1 root root 1853 Jul 23 11:09 /tmp/daproxyss_key.pem
-rw-r--r--. 1 root root 1359 Jul 23 11:09 /tmp/daproxyss.pem
```

7. (Optional) Request a CA-signed certificate for the proxy server host using the self-signed certificate produced here. Copy and use your CA-signed certificate files, including any intermediate and root certificates, in the following steps.
8. Copy the proxy server certificate files for the Proxy service (daproxy ) and NetOps Portal by issuing the following commands:

**Example:**

```
// copy to daproxy config folder for use by Traefik on DAProxy...
cp /tmp/daproxyss.cer /opt/CA/daproxy/conf/daproxyss.cer
cp /tmp/daproxyss.pem /opt/CA/daproxy/conf/daproxyss.pem
cp /tmp/daproxyss_key.pem /opt/CA/daproxy/conf/daproxyss_key.pem
```

```
// move the certificate file(s) to a temp folder on PC...
cp /opt/CA/daproxy/conf/daproxyss.cer <myPcHost>/tmp/daproxyss.cer
```

**NOTE**

NetOps Portal uses Traefik as the HTTPS reverse proxy.

For more information about Traefik, see [the Traefik site](#).

## Configure the Proxy Server

### Follow these steps:

1. Edit the /etc/nginx/nginx.conf file that Nginx creates by default.

2. Comment out the lines for the `server { ... }` section that is inside the `http { ... }` section using the `#` character.

3. Confirm that the following line exists inside the `http { ... }` section:

```
include /etc/nginx/conf.d/*.conf;
```

If it does not, add it after the `default_type` line.

4. Create the `/etc/nginx/conf.d` directory.
5. Edit the `/etc/nginx/conf.d/daproxy.conf` file.
6. Add the following lines to the file:

```
upstream backend {
    ip_hash;
    server <DA1 FQHN or IP>:8582;
    server <DA2 FQHN or IP>:8582;
}

server {
    listen 8582 default_server;
    listen [::]:8582 default_server;
    server_name <FQHN of daproxy>;
    include /etc/nginx/default.d/*.conf;
    client_max_body_size 0;

    ssl_certificate "/opt/CA/daproxy/conf/daproxy.pem";
    ssl_certificate_key "/opt/CA/daproxy/conf/daproxyss_key.pem";
    ssl_session_cache shared:SSL:1m;
    ssl_session_timeout 5m;
    ssl_prefer_server_ciphers on;
    ssl_protocols TLSv1.2;
    ssl_ciphers HIGH:!aNULL:!MD5;

    location / {
        proxy_pass https://backend;
    }
}
```

The proxy server is configured.

## **Restart the Proxy Service**

Follow these steps:

1. Disable the Traefik (daproxy) service by issuing the following commands:

```
systemctl stop daproxy
systemctl disable daproxy
```

2. Enable then start the Nginx service by issuing the following commands:

```
systemctl enable nginx
systemctl start nginx
```

## **Test the HTTPS Connection to the Proxy Server**

Test the connection to the proxy server using the following URL:

```
https://<daproxy_hostname>:8582/
```

## Update the Proxy Server IP Address or Hostname

When you change the IP address or hostname for the proxy server (for the fault tolerance), update the proxy server's IP address or hostname in the data aggregator configuration files and in the proxy server.

Use the following process to update the IP address or hostname for the proxy server:

1. [Update it in the data aggregator configuration files.](#)
2. [Update it in the proxy server.](#)

### Update the IP Address or Hostname in the Data Aggregator Configuration Files

Complete the following steps on *both* fault-tolerant data aggregators.

**Follow these steps:**

1. Open the `<installation_directory>/consul/consul.cfg` file, and then if the `consul.primary_consul_server` entry is set with an IP address or hostname, ensure that it reflects the new IP address or hostname of the proxy server.
  - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
2. Open the `/etc/DA.cfg` file, and then if the `da.proxy.host` entry is set with an IP address or hostname, ensure that it reflects the new IP address or hostname of the proxy server.
3. Edit the `<installation_directory>/consul/conf/config.json` file.
  - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
4. Verify the following parameter:
  - `client_addr`  
Verify the IP address or hostname. The value for this parameter is normally set to the placeholder `0.0.0.0`. However, if it is set with an actual IP address or hostname, edit the IP address or hostname to reflect the new IP address or hostname of the proxy server.
  - `retry_join`  
Edit the IP address or hostname to reflect the new IP address or hostname of the proxy server.

The IP address or hostname of the proxy server is updated in the data aggregator configuration files.

#### **TIP**

If you have not yet updated the IP address or hostname of the proxy server in the configuration files for *both* fault-tolerant data aggregators, repeat this step for the other data aggregator's configuration files.

### Update IP Address or Hostname on the Proxy Server

Update the IP address or hostname for the proxy server in the following locations on the proxy server.

**Follow these steps:**

1. Open the following file based on your installation, and then edit the existing IP address or hostname in the `-bind=` entry with the new IP address or hostname of the proxy server:  
`/etc/systemd/system/consul.service`
2. Open the `<installation_directory>/consul/consul.cfg` file, and then edit the existing IP address or hostname in the `consul.bind_addr=` entry with the new IP address or hostname of the proxy server.
  - **installation\_directory**  
The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

3. Open the `<proxy_installation_directory>/conf/config.json` file, and then verify the IP address or hostname in the `client_addr` entry. This entry is normally set to the placeholder `0.0.0.0`. However, if it is set with an actual IP address or hostname, edit the IP address or hostname to reflect the new IP address or hostname of the proxy server.
  - **`proxy_installation_directory`**  
The location of the proxy server.

**Default:** /opt/CA/daproxy
4. In fault-tolerant environments, the data source must point to the proxy server. If the data source is configured to use the IP address instead of the hostname for the proxy server, edit the data source using NetOps Portal. For more information about how to edit the IP address of the data aggregator data source to use the hostname for the proxy server, and to test the connection to validate successful communication, see [Configure a Data Source](#).

## Data Aggregator Scripts

The DX NetOps Performance Management installation includes the following scripts for the data aggregator:

- [activemq](#)
- [cleanupComponents.sh](#)
- [dadaemon](#)
- [doEncryption.sh](#)
- [RemoteEngineer/re.sh](#)
- [reset\\_da\\_db\\_password.sh](#)
- [remove\\_not\\_present\\_items.sh](#)
- [rtt\\_rest\\_configuration\\_examples.py](#)
- [sslConfig.sh](#)
- [start-consul-ext.sh](#)

### **activemq**

The `activemq` script stops, starts, and restarts the ActiveMQ service/broker.

For more information about when to run this script, see [Restart the ActiveMQ Broker](#).

### **cleanupComponents.sh**

The `cleanupComponents.sh` script removes those items that do not pass the filter criteria for the monitoring profile filter.

### **dadaemon**

The `dadaemon` script stops, starts, restarts, activates, checks the status, and shuts down the data aggregator service.

For more information about when to run this script:

- See [Restart the Data Aggregator](#).
- See [Upgrade the Data Repository](#).
- See [Restore the Data Repository](#).
- See [Migrate the Data Repository](#).

### **doEncryption.sh**

The `doEncryption.sh` script encrypts the database password for config files.

For more information about when to run this script:

- See [Install Fault-Tolerant Data Aggregators](#).
- See [Install a Non-Fault-Tolerant Data Aggregator](#).

### **RemoteEngineer/re.sh**

The `re.sh` script (the CA Remote Engineer (CARE) tool) gathers the DX NetOps Performance Management configuration and database into a compressed file for Broadcom Support triage purposes.

For more information about when to run this script:

- See [Unable to Resolve Issue](#).
- See [Upgrade the Data Repository](#).

### **reset\_da\_db\_password.sh**

The `reset_da_db_password.sh` script updates the Vertica database administrator system account (the database administrator user) and/or the database administrator user password in the `<installation_directory>/apache-karaf/etc/dbconnection.cfg` file.

- ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`

### **remove\_not\_present\_items.sh**

The `remove_not_present_items.sh` script deletes components for a device that are marked not present from your network.

For more information about when to run this script, see [Delete Components That Are Not Present](#).

### **rtt\_rest\_configuration\_examples.py**

Run the `rtt_rest_configuration_examples.py` script to manage (create, delete, edit) the Round Trip Time (RTT) tests.

For more information about when to run this script, see [RTT Configuration Examples](#).

### **sslConfig.sh**

Run the `sslConfig.sh` script when enabling HTTPS for the data aggregator, and when reverting the data aggregator to using HTTP. This script sets up Subject Alternative Name (SAN) entries.

For more information about when to run this script:

- See [Configure the Connection from the Proxy Server to the Data Aggregators as HTTPS using a Self-Signed Certificate](#).
- See [Configure the Connection from the Proxy Server to the Data Aggregators as HTTPS using CA-Signed Certificates](#).
- See [Configure the Data Aggregator to Use HTTPS and ActiveMQ JMX to Use SSL Using the SSL Configuration Tool](#).
- See [Revert the Data Aggregator to Using HTTP](#).

### **start-consul-ext.sh**



The `start-consul-ext.sh` script restarts the `consul-ext` service on the data aggregator in a fault-tolerant environment. This script is related to HTTPS configuration.

For more information about when to run this script, see [Configure Consul as HTTPS](#).

## Data Repository Administration

The topics in this section provide information about how to configure the data repository.

### NOTE

The data repository uses Vertica database software. You cannot *customize* the data repository:

- Do not deploy custom projections or modify database parameters.
- The data repository *does not support* a Vertica k-safety value greater than 1.
- The data repository *does not support* the use of the Vertica Management Console tool. Management Console can modify the Vertica database configuration.

## Configure Data Retention Rates

The default data retention rates for the Data Repository help to conserve disk space and improve reporting, and can be customized.

Polled data is generated each poll cycle for all devices, and this data represents the most granular data available in the product. This raw, polled data is set to roll up at hourly, daily, and weekly aggregation levels. Because higher-level aggregated data requires less disk space, you can keep this data for a longer period than polled data. You can change how long the data repository retains the polled data, hourly rollup data, daily rollup data, and weekly rollup data. For example, you can change the polled data retention value to 30 days to conserve disk space. Find the balance that best suits your needs and environment.

By default, data is retained in the data repository for the following number of days:

- Polled data: 45 days
- Hourly rollup data: 90 days
- Daily rollup data: 365 days
- Weekly rollup data: 730 days

The minimum number of days that Data Repository can retain data for is as follows:

- Polled data: 2 days
- Hourly rollup data: 8 days

### NOTE

When the hourly retention rate is less than 32 days, the calendar heat chart views show daily data (only one sample per day). The charts appear sparse as a result.

- Daily rollup data: 31 days
- Weekly rollup data: 366 days

### Follow these steps:

1. Enter the following information in a web browser:

`http://hostname:port/rest/globalretentiondefinition`

#### – **hostname:port**

Use this parameter to specify the Data Aggregator host name and the port number.

**Default port:** 8581

2. Take note of the ID that is assigned to the `globalretentiondefinition`.
3. Look for the following elements:

- GtdRollupDataRetentionPeriod
- DailyRollupDataRetentionPeriod
- PolledDataRetentionPeriod
- HourlyRollupDataRetentionPeriod

This information helps you determine which types of data you want to modify the retention period for.

4. Look for the elements that has the HistoricConfigurationDataRetentionPeriod parameter. This information helps you specify the number of days to keep item repository attribute values.
5. Open a REST client editor or HTTP tool that sends requests and gets responses and set the Content-type to application/xml.
6. Enter the following criteria:
  - URL: `http://hostname:port/rest/globalretentiondefinition/ID`
    - **ID**  
A unique identification number that is assigned to the globalretentiondefinition parameter
  - HTTP method = PUT
  - Enter the retention periods that you want to change in the Body tab of the HTTP Request pane.  
For example:

```
<GlobalRetentionDefinition version="1.0.0">

<PolledDataRetentionPeriod>4</PolledDataRetentionPeriod>

</GlobalRetentionDefinition>
```

### IMPORTANT

Verify that there is no white space at the beginning of each of these lines, otherwise the PUT fails.

In this example, the polled data retention period has been changed to four days.

Results are returned in the **Body** tab of the HTTP Response pane.

For example:

```
<GlobalRetentionDefinitionList>
  <GlobalRetentionDefinition version="1.0.0">
    <ID>4</ID>
    <GtdRollupDataRetentionPeriod>730</GtdRollupDataRetentionPeriod>
    <HistoricConfigurationDataRetentionPeriod>7</
HistoricConfigurationDataRetentionPeriod>
    <DailyRollupDataRetentionPeriod>365</DailyRollupDataRetentionPeriod>
    <PolledDataRetentionPeriod>4</PolledDataRetentionPeriod>
    <HourlyRollupDataRetentionPeriod>90</HourlyRollupDataRetentionPeriod>
    <Item version="1.0.0">
      <CreateTime>Thu Dec 08 16:03:05 CST 2011</CreateTime>
      <Name>Global Retention Definition</Name>
    </Item>
  </GlobalRetentionDefinition>
</GlobalRetentionDefinitionList>
```

In this example, the polled data retention period has been changed to four days. The default retention periods for weekly rollup data, daily rollup data, and hourly rollup data remain.

## Back Up the Data Repository

Protect your data by backing up the data repository.

The first data repository backup is a full backup of all historical data. Subsequent backups are incremental and include database activity that occurred since the snapshot at the start of the previous backup. The data repository and data aggregator continue to run during a backup.

Perform full backups weekly. Perform incremental backups daily. Full backups occur only when the backup location is a new directory. The incremental snapshots store new files and hard links to unchanged files from the previous backup. Restoring to an incremental snapshot depends on the integrity of the files that are linked to in previous snapshots.

Backup processing can be resource-intensive, but DX NetOps Performance Management prioritizes the processing below other processing. To let backups to proceed more quickly, and to minimize the impact to other processing, perform backups during non-peak hours.

You can back up data repository to a remote host, or you can back it up to the same host. If you back up to the same host, save the backup to a different disk than the one that is used by the catalog and data directories.

For information about the size of the backup files, see [the DX NetOps Sizing Tool](#).

Use the following process to back up the data repository:

### IMPORTANT

Back up the data repository using this process. Do not back up the data repository by taking a virtual machine snapshot.

1. [Verify the Prerequisites](#)
2. (If you are backing up to a remote host) [Create a Remote Data Repository Backup Host](#)
3. [Configure the Data Repository Backup](#)

Also in this article:

- [Back Up Using the VBR Utility](#)
- [Recover Data from the iRep](#)

### Verify the Prerequisites

Before backing up the data repository, ensure that you have completed the following prerequisites:

- To ensure data integrity, you have backed up each node of the data repository to a dedicated backup host.
- You have verified the following information about the data repository host and the remote backup host:
  - The hosts are not connected to LDAP.
  - The hosts are not connected to Network Information Service (NIS), and they have the same Vertica Linux database administrator user.
  - Port 50000 is open on any firewalls so that the data repository host can access the custom rsync/ssh port 50000 on the backup host.

### NOTE

If you do not have a backup host, you can back up the data repository locally.

For more information, see [Configure the Data Repository Host for a Local Backup](#).

- You have enabled TCP forwarding by setting `AllowTcpForwarding = Yes` in the `/etc/ssh/sshd_config` file on the Vertica hosts, on both source and destination systems. Backing up the data repository and the `copycluster` command within Vertica using the `vbr` utility requires that TCP forwarding be enabled. This allows the utility to forward connections from database hosts to backup hosts. For more information, see [the Vertica documentation](#).

## Create a Remote Data Repository Backup Host

(If you are backing up to a remote host) Perform the following procedure for each data repository back up host.

### Follow these steps:

1. Log in to the backup host as the root user.
2. Create the Vertica Linux database administrator user on the remote backup host by issuing the following command:  

```
useradd <db_admin> -s /bin/bash
```

  - **db\_admin**  
The same Vertica Linux database administrator user that exists on the data repository hosts.
3. Set the Vertica Linux database administrator user password by issuing the following command:  

```
passwd db_admin
```
4. Create the Vertica directories on the remote backup host by issuing the following command:  

```
mkdir /opt/vertica/bin
mkdir /opt/vertica/oss
```
5. Change the owner of the Vertica directories by issuing the following command:  

```
chown -R db_admin /opt/vertica
```

  - **db\_admin**  
The same Vertica Linux database administrator user that exists on the data repository hosts.
6. [Set up passwordless ssh between the data repository hosts and the remote backup hosts.](#)
7. Log in to the remote backup host as the Vertica Linux database administrator user.
8. Copy the Vertica rsync and Python tools from the data repository host to the remote backup host by issuing the following commands:  

```
scp dradmin@drhost:/opt/vertica/bin/rsync /opt/vertica/bin
scp -r dradmin@drhost:/opt/vertica/oss /opt/vertica
```
9. Verify that the remote backup host has the following directories:
  - /opt/vertica/bin/rsync
  - /opt/vertica/oss/python3
10. Create the backup directory by issuing the following command:  

```
mkdir <backup_directory>
```

  - **backup\_directory**  
Specifies the directory where you want to save the backup files. Select a backup directory that is on a disk partition with a large amount of free space. If this directory is not writable by the database administrator user, give this user access to this directory.

The remote host is ready for the backup configuration file to be created and for the backup directory to be initialized.

A remote data repository backup host is created.

## Configure the Data Repository Backup

Configuring the data repository backup involves creating a configuration file for the backup. Vertica performs a full backup during the first backup into a new backup directory. Subsequent backups to the same directory are incremental backups, even if the snapshot name changes.

The node where you perform this procedure initiates the backup.

### Follow these steps:

1. Log in to the data repository host as the database administrator user.
2. Create the password file that is used to store the password for backups, for example /opt/vertica/config/password.txt.

**NOTE**

You can choose a different location for the password file.

```
[Passwords]
; Specified password for db admin account
dbPassword = DBpassword
; Specifies password for rsync user account - if different than DB admin
; serviceAccessPass = rsyncpwd
; Specifies password for the dest_dbuser Vertica account. Used only for restoring to
  alternate cluster.
; dest_dbPassword = DestinationPwd
```

3. As the file owner, or root, set the permissions for the `password.txt` file so that it is accessible only to the `dbadmin` user and the `verticadba` group or the user and group for the user running the backup by issuing the following command:

```
chmod 660 <passwordFile>
```

**Example:**

```
chmod 660 /home/dradmin/password.txt
```

**NOTE**

The password file name can vary.

4. Go to the Vertica `vbr` utility sample configuration files in the `/opt/vertica/share/vbr/example_configs` directory. The database administrator user requires write privileges for the directory.

**NOTE**

Vertica automatically installs sample configuration files in this directory.

For more information about these files, see [the Vertica documentation](#).

5. Copy, edit, and deploy a configuration file for backup. The `snapshotName` can be the same as previous backups. See the following examples from Vertica:

– **Example 1: Local backup** `backup_restore_full_local.ini`

Back up the data repository to a local area on the same machine. The backup can be a mount from an external shared drive or a local disk. You cannot use the same disk as `data/catalog`:

```
[Mapping]
; node_name = backup_host:backup_dir
; [] indicates backup to localhost
v_drdata_node0001 = []:/backups
v_drdata_node0002 = []:/backups
v_drdata_node0003 = []:/backups

[Misc]
; Backups with the same snapshotName form a time sequence limited by
  restorePointLimit.
; SnapshotName is used for naming archives in the backup directory, and for
  monitoring and troubleshooting.
; Valid values: a-z A-Z 0-9 - _
; The temp directory location on all database hosts.
; The directory must be readable and writeable by the dbadmin, and must implement
  POSIX style fcntl lockf locking.
; tempDir = /tmp/vbr
; Specifies the number of historical backups to retain in addition to the most
  recent backup.
```

```

; 1 current + n historical
snapshotName = backup_snapshot
restorePointLimit = 7

; Full path to the password configuration file
; Store this file in directory readable only by the dbadmin.
passwordFile = /opt/vertica/config/password.txt

; When enabled, Vertica confirms that the specified backup locations contain
; sufficient free space and inodes to allow a successful backup. If a backup
; location has insufficient resources, Vertica displays an error message explaining
; the shortage and
; cancels the backup. If Vertica cannot determine the amount of available space
; or number of inodes in the backupDir, it displays a warning and continues
; with the backup.
; enableFreeSpaceCheck = True

; When performing a backup, replication, or copycluster, specifies the maximum
; acceptable difference, in seconds, between the current epoch and the backup
; epoch.
; If the time between the current epoch and the backup epoch exceeds the value
; specified in this parameter, Vertica displays an error message.
; SnapshotEpochLagFailureThreshold = 3600

```

– **Example 2: External backup** backup\_restore\_full\_external.ini

Back up the data repository to a different machine. Replace the IP addresses with the IP address of the backup hosts:

```

[Mapping]
; node_name = backup_host:backup_dir
; In this "parallel backup" configuration, each node backs up to a distinct
; external host.
; To backup all database nodes to a single external host, use that single hostname/
; IP address in each entry below.
v_drdata_node0001 = 1.1.1.1:/backups
v_drdata_node0002 = 2.2.2.2:/backups
v_drdata_node0003 = 3.3.3.3:/backups

[Misc]
; Backups with the same snapshotName form a time sequence limited by
; restorePointLimit.
; SnapshotName is used for naming archives in the backup directory, and for
; monitoring and troubleshooting.
; Valid characters: a-z A-Z 0-9 - _
; The temp directory location on all database hosts.
; The directory must be readable and writeable by the dbadmin, and must implement
; POSIX style fcntl lockf locking.
; tempDir = /tmp/vbr

```

```

; Specifies the number of historical backups to retain in addition to the most
; recent backup.
; 1 current + n historical
snapshotName = backup_snapshot
restorePointLimit = 7

; Full path to the password configuration file
; Store this file in directory readable only by the dbadmin
; (no default)
passwordFile = /opt/vertica/config/password.txt

; When enabled, Vertica confirms that the specified backup locations contain
; sufficient free space and inodes to allow a successful backup. If a backup
; location has insufficient resources, Vertica displays an error message explaining
; the shortage and
; cancels the backup. If Vertica cannot determine the amount of available space
; or number of inodes in the backupDir, it displays a warning and continues
; with the backup.
; enableFreeSpaceCheck = True

; When performing a backup, replication, or copycluster, specifies the maximum
; acceptable difference, in seconds, between the current epoch and the backup
; epoch.
; If the time between the current epoch and the backup epoch exceeds the value
; specified in this parameter, Vertica displays an error message.
; SnapshotEpochLagFailureThreshold = 3600

```

6. (First Time Only) Initialize the backup directory before the first time that you run the backup by issuing the following command:

```

/opt/vertica/bin/vbr.py --task init --config-file
<configuration_directory_path_filename>

```

**Example:**

```

/opt/vertica/bin/vbr.py --task init --config-file /home/vertica/backup_restore_full_external.ini

```

– **configuration\_directory\_path\_filename**

Indicates the directory path and filename of the configuration file that you will reference when you restore. This file is located where you ran the backup utility.

Once initialized, multiple configuration files can use the directory if the files share the same backup directory location.

7. Back up the data repository by issuing the following command:

```

/opt/vertica/bin/vbr.py --task backup --config-file
<configuration_directory_path_filename>

```

**Example:**

```

/opt/vertica/bin/vbr.py --task backup --config-file /home/vertica/vert-db-
production.ini

```

– **configuration\_directory\_path\_filename**

Indicates the directory path and filename of the configuration file that you will reference when you restore. This file is located where you ran the backup utility.

8. If you are prompted about the authenticity of the host, answer **yes**.

The data repository starts the backup. This process can take a long time, especially for a full backup.

9. (Optional) If you do not want to retain the data repository password in clear text for future manual backups, complete the following steps:

**IMPORTANT**

The configuration file that is generated contains a clear text password. Automated backups require the password. This procedure prevents automated backups from this configuration file.

- a. Verify that the following line exists under the [Database] section:

```
dbPromptForPassword = True
```

- b. Remove the following line from the [Database] section:

```
dbPassword = password
```

10. Recommended next step: [Set up automatic backups of the data repository](#).

The data repository backup is configured.

### **Set Up Automatic Backups of the Data Repository**

To ensure regular backups of the data repository, create a cron job to schedule automatic backups. The first backup is a full backup and the following backups are incremental. Run a full backup weekly or biweekly. Vertica performs a full backup only when you use a new backup directory.

**TIP**

Run a full backup weekly. If disk space is limited, retain only two to three weeks of data. Delete the oldest backup file at the beginning of each week. Use the `vbr` utility to remove task to delete old backups. Vertica does not support removing backups through the file system.

**Follow these steps:**

1. Create a wrapper shell script that contains the following line by issuing the following command:

```
/opt/vertica/bin/vbr.py --task backup --config-file
<configuration_directory_path_filename>
```

**Example:**

```
/opt/vertica/bin/vbr.py --task backup --config-file /home/vertica/backup_restore_full_external.ini
```

– ***configuration\_directory\_path\_filename***

The directory path and filename of the configuration file that you will reference when you restore.

2. Save the contents to a new `backup_script.sh` script in a location of your choice, for example `/home/vertica/backup_script.sh`.
3. Change permissions for running the script by issuing the following command:

**NOTE**

`chmod 777` makes the file readable, writable, and executable by everyone. If you want only the script owner to run the file, use `chmod 700`. If you want only the root user to run the file, use `chmod 755`.

```
chmod 777 location_backup_script.sh/backup_script.sh
```

**Example:**

```
chmod 777 /home/vertica/backup_script.sh
```

4. As the database administrator user, open the crontab to define a cron job:

```
crontab -e
```

5. Add a cron job that runs the backup script.

**TIP**

Define the cron job to run the script daily at an off-peak time.

**Example:**

This following example cron job runs the backup script every day at 2:00 AM:

```
00 02 * * * /home/vertica/backup_script.sh >/tmp/backup.log 2>&1
```

The cron job runs a daily incremental data repository backup.



**TIP**

To disable scheduled data repository backups, comment the cron job:

```
# 00 02 * * * /home/vertica/backup_script.sh >/tmp/backup.log 2>&1
```

6. Add a script to copy the configuration file, and change the snapshot name in the configuration file. Also use a new backup directory in the configuration file to cause Vertica to perform a full backup.

**IMPORTANT**

Do not delete the previous configuration file. The original configuration file is required to remove a backup or restore from an older series of backups.

7. (Optional) Remove older backup sequences as required with the vbr utility with the remove task using the configuration that was used to create it by issuing the following command:

```
/opt/vertica/bin/vbr.py --task remove --archive=[<date>_<time>|"all"] --config-file
<configuration_directory_path_filename>
```

**Example:**

```
/opt/vertica/bin/vbr.py --task remove --archive=[<date>_<time>|"all"] --config-file </home/vertica/
backup_restore_full_external.ini
```

**IMPORTANT**

The `remove` command is destructive and removes the data and free space on the disk. Specify the archive to remove a single restore point, a comma separated list, or "all". To display the list of backups, issue the `--task listbackup` command.

**Back Up Using the VBR Utility**

You can back up and restore either the full database, or one or more schema and table objects of interest using the `vbr` utility.

**TIP**

You can also copy a cluster and list backups that you created previously using this utility.

For more information about how to use this utility, see [the Vertica documentation](#).

**Recover Data from the iRep**

In the event that you must recover your data, you can restore from the iRep. Export the iRep data from an existing system, and then import it into a different schema using the `/opt/CA/IMDataRepository_vertica<version>/caVerticaUtility.sh` script. You can then query the iRep data to determine how the original system was configured.

**TIP**

You can also export self-monitoring data or poll the data for debugging purposes using this script.

**Configure Passwordless SSH**

The Database Administrator user account (default: `dradmin`) requires passwordless SSH access between the Vertica clusters.

Use the following process to set up passwordless SSH:

1. Provide passwordless SSH access for the Database Administrator user account from each host in the primary cluster to each host in the recovery cluster.
2. Provide passwordless SSH access for the Database Administrator user account from each host in the target (recovery) cluster back to each host in the primary cluster.

Complete the following steps on each node, and then validate that you can reach the other nodes in the cluster without a password:

1. Verify that the Database Administrator user is set up with a passwordless SSH key by issuing the following command as that user:

```
ssh <other-hostname-in-the-cluster> ls
```

- **other-hostname-in-the-cluster**

The name of the remote host where the user requires passwordless SSH access.

2. If the Database Administrator user is set up with a passwordless SSH key, the user is not prompted for a password. If the user is prompted for a password, set up passwordless SSH by completing the following steps:

- a. Set up the Linux user account for the Database Administrator user with a passwordless SSH key by issuing the following commands:

```
ssh-keygen -N "" -t rsa -f ~/.ssh/id_rsa
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys2
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys
chmod 644 ~/.ssh/authorized_keys2
chmod 644 ~/.ssh/authorized_keys
```

Passwordless SSH is set up for the current node.

- b. Copy the SSH credentials to the other node in the cluster by issuing the following command:

```
ssh-copy-id -i dradmin@<other-hostname-in-the-cluster>
```

- **other-hostname-in-the-cluster**

The name of the remote host where the user requires passwordless SSH access.

- c. Verify that passwordless SSH is set up by issuing the following command:

```
ssh <other-hostname-in-the-cluster> ls
```

- **other-hostname-in-the-cluster**

The name of the remote host where the user requires passwordless SSH access.

## Configure the Data Repository Host for a Local Backup

For best data integrity, keep a dedicated backup host. If you do not have a dedicated backup host, you can back up the data repository locally.

For more information about how to back up the data repository, see [Back Up the Data Repository](#).

**Prerequisite:** Verify that the Database Administrator user is set up with a passwordless ssh key. For more information, see [Configure Passwordless SSH](#).

Create a backup directory for each host in the cluster by issuing the following command:

```
mkdir <backup_directory>
```

- **backup\_directory**

Specifies the directory where you want to save the backup files. Select a backup directory that is on a disk partition with a large amount of free space. If these directories are not writable by the database administrator user, give this user access to these directories.

The data repository host is ready for the backup to run.

## Update the Data Repository IP Address and Hostname

When you change the IP address or hostname for the server on which the data repository is installed, update the data repository IP address or hostname.

The following procedures are an overview of a basic IP address or hostname-change scenario when using a standard three-node database cluster.

Use the following process to update the data repository IP address or hostname:

1. [Verify the prerequisites.](#)
2. [Update the data repository IP addresses.](#)
3. [Update the data aggregator configuration file.](#)

### **Verify the Prerequisites**

Before updating the data repository IP address or hostname, verify that all nodes in the cluster are down. Use the Vertica Administration Tools utility, `admintools`. This utility is located in the `/opt/vertica/bin` directory.

For more information about how to verify that all nodes in the cluster are down, see [the Vertica documentation](#).

### **Update the Data Repository IP Addresses**

**Follow these steps:**

1. Log in to the database server as the database administrator (`dradmin`) user by issuing the following command:

```
su - dradmin
```

2. Create a mapping file that maps the old IP addresses to the new IP addresses, for example:

```
[dradmin@DR_Node tmp]$ more /tmp/mapping.out
192.0.2.254 198.51.100.255
192.0.2.255 198.51.100.256
192.0.2.256 198.51.100.257
```

For more information about how to create mapping files, see [the Vertica documentation](#).

3. Update the Vertica catalog with the new data repository IP addresses by issuing the following `admintools` command, and then enter `yes` when prompted:

```
admintools -t re_ip -f <mapping_file>
```

#### **– *mapping\_file***

The name of the mapping file.

#### **Example output:**

```
Parsing mapfile ...
New settings for Host 192.168.10.22 are:
  address: 192.168.122.1
  controlAddress: 192.168.122.1
  controlBroadcast: 192.168.122.255
```

```
The following databases would be affected by this tool: drdata
```

```
Checking DB status ...
```

```
Please check your new settings carefully, incorrect settings may cause database
damage!
```

```
Enter "yes" to write new settings or "no" to exit > yes
```

```
Backing up local admintools.conf ...
```

```
Writing new settings to local admintools.conf ...
```

```
Writing new settings to the catalogs of database drdata ...
```

```
The change was applied to all nodes.
```

```
Success. Change committed on a quorum of nodes.
You can start database drdata again.
```

```
Initiating admintools.conf distribution ...
Success. Local admintools.conf sent to all hosts in the cluster.
>>> Write new settings successfully.
```

The IP addresses are updated.

## Update the Data Aggregator Configuration File

Complete the following steps on each data aggregator host.

1. Open the `<installation_directory>/apache-karaf/etc/dbconnection.cfg` data repository configuration file.
  - ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
2. If the `dbHostNames` entry is set with the IP addresses for the previous data repository cluster nodes, update the IP addresses with the new IP addresses for the data repository cluster nodes, and then save your changes. Otherwise, no change is needed.
3. If the data aggregator is using the `/etc/hosts` file to resolve the data repository IP address, update this file with the new data repository cluster node IP addresses.

The data aggregator configuration file is updated.

## Restore the Data Repository

You can restore the data repository from an existing backup.

### IMPORTANT

To ensure that the data is fully compatible with schema changes, restore backups only from the current DX NetOps Performance Management version.

Usually, you restore the data repository to the same computer from where you backed it up. However, you *can* restore the data repository to a different computer.

### Prerequisites:

- The database administrator user is part of the sudoers file.
- (If you are restoring the data repository to a different computer) The computer to which you are restoring is configured in the same way that the computer from which you backed up the data repository is configured. In a cluster environment, each computer to which you restore is configured in the same way that each computer from where you backed up each data repository node is configured. The following configurations are the same: the IP address, the hostname, the catalog and data directories, the catalog and data directory permissions, the Vertica Linux database administrator user credentials, the database administrator user account credentials, and the database user account credentials.

### Follow these steps:

1. Stop the data collector hosts that are associated with the data aggregator by logging in to the computers where the data collector is installed as the root user or a sudo user with access to a limited set of commands. Open a command prompt, and then issue the following command:

```
systemctl stop dcmd
```

The data collector hosts stop.

2. Stop the data aggregator by doing the following steps:

- a. Log in to the computer where the data aggregator is installed as the root user or a sudo user with access to a limited set of commands.
- b. Do one of the following steps:

- Stop the Data Aggregator service by issuing the following command:

```
systemctl stop dadaemon
```

The data aggregator starts.

- (Fault-tolerant environment) If the local data aggregator is running, put the data aggregator in maintenance mode ("Maintenance" status) by issuing the following command:

```
<installation_directory>/scripts/dadaemon maintenance
```

- **installation\_directory**

The installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator status changes to "Ready" and is unavailable for failover.

The data aggregator stops.

3. Stop the data repository by doing the following steps:

- a. Log in to the database server you use for the data repository as the database administrator user, *not* as the root user.
- b. Open a command prompt, and then issue the following command:

```
/opt/vertica/bin/adminTools
```

The **Administration Tools** dialog opens.

- c. Select **(4) Stop Database**.
- d. Press the **Space** bar next to the database name, select **OK**, and then press the **Enter** key on your keyboard. You are prompted for the database password.
- e. Enter the database password, and then press the **Enter** key on your keyboard. The data repository stops.

**NOTE**

If the data repository does not stop, select **(2) Stop Vertica** on Host from the **(7) Advanced Tools Menu**.

- f. Select **Exit**, and then press the **Enter** key on your keyboard.

4. Prepare to restore the data repository backup by doing the following steps:

- a. Log in as the Linux user account for the database administrator user to the database server you use for the data repository.

When you set up automatic backups of the data repository, you configured the configuration file with a restore point of seven. You can restore the data repository to the most recent backup or to any of the previous seven incremental backups.

- b. Do *one* of the following steps:

- To restore the data repository to the most recent backup, issue the following command:

```
/opt/vertica/bin/vbr.py --task restore --config-file <configuration_directory_path_filename>
```

**Example:**

```
/opt/vertica/bin/vbr.py --task restore --config-file /home/vertica/vert-db-production.ini
```

- **configuration\_directory\_path\_filename**

Indicates the filename and directory path of the configuration file you created when you ran the backup configuration procedure. This file is located where you ran the `/opt/vertica/bin/vbr.py` backup utility.

- To restore the data repository to any of the previous seven incremental backups, issue the following command with the following options:

```
/opt/vertica/bin/vbr.py --task restore --config-  
file <configuration_directory_path_filename> --archive <archive_name>
```

**Example:**

```
/opt/vertica/bin/vbr.py --task restore --config-file myconfig.ini --archive  
20131020_170018
```

- **configuration\_directory\_path\_filename**

Indicates the filename and directory path of the specific configuration file you want to restore a specific archive from. You created this configuration file when you ran the backup configuration procedure. This file is located where you ran the `/opt/vertica/bin/vbr.py` backup utility.

- **archive\_name**

Indicates the name of the specific restore point that you want to restore to. Change to the backup directory that the configuration file for the restore point indicates. All of the restore points that are available are listed. Determine the archive name for the restore point that you want to restore to.

**NOTE**

In a cluster installation, you can run the `restore` task from any of the hosts that are participating in the cluster.

**TIP**

For a list of available restore points, issue the following command:

```
/opt/vertica/bin/vbr.py --task listbackup --config-file  
configuration_directory_path_filename
```

5. Restart the data repository by doing the following steps:

- a. Log in to the computer where the data repository is installed as the database administrator user, *not* as the root user.

- b. Open a command prompt, and then issue the following command:

```
/opt/vertica/bin/adminTools
```

The **Administration Tools** dialog opens.

- c. Select **(3) Start Database**.

- d. Press the **Space** bar on your keyboard next to the database name, select **OK**, and then press the **Enter** key on your keyboard.

You are prompted for the database password.

- e. Enter the database password, and then press the **Enter** key on your keyboard.

The data repository starts.

- f. Select **Exit**, and then press the **Enter** key on your keyboard.

6. Restart the data aggregator by doing the following steps:

- a. Log in to the data aggregator host as the root user or a sudo user.

- b. Do *one* of the following steps:

- Start the Data Aggregator service by issuing the following command:

```
systemctl start dadaemon
```

The data aggregator starts.

- (Fault-tolerant environment) If the local data aggregator is running, put the data aggregator into maintenance mode ("Maintenance" status) by issuing the following command:

```
<installation_directory>/scripts/dadaemon maintenance
```

- **installation\_directory**

The installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator changes to "Maintenance" status and is unavailable for failover.

7. Restart the data collectors. Log in to each data collector host as the root user or a sudo user, and start the Data Collector service by issuing the following command:

```
systemctl start dcmd
```

The data collectors restarts.

The data repository is restored.

## Add a Node to the Data Repository Cluster

In a cluster installation, you can increase system performance or replace a node by adding a node to the data repository cluster.

### NOTE

This process applies to expanding a node that was set up with a hostname. If you set up a node using localhost, the nodes table in Vertica has `node_address` or `node_ip` set to 127.0.0.1. To expand a node that was set up using localhost, contact Broadcom Support.

Use the following process to add a node to the data repository cluster:

### NOTE

If you are adding multiple nodes to the data repository cluster, prepare all hosts, including adding them to the cluster, before adding the nodes to the database.

For more details about this process, see [the Vertica documentation](#).

1. [Verify the prerequisites](#).
2. [Prepare the new host](#).
3. [Add the host to the cluster](#).
4. [Add the node to the database](#).
5. (If you are moving from one or two nodes to three or more nodes) [Create the buddy projection and mark the cluster k-safety of 1](#).
6. [Update the data repository configuration](#).
7. (If you are moving from one or two nodes to three or more nodes) [Segment the database](#).

### Verify the Prerequisites

Before you add the new node to the data repository cluster, verify the following prerequisite steps:

- You have stopped the data aggregator by issuing the following command:

```
systemctl stop dadaemon
```

### NOTE

Keep the data aggregator stopped until you complete this process.

- You have prepared the new host for the data repository installation.  
For more information see, [Prepare to Install the Data Repository](#).
- You have backed up the data repository.  
For more information, see [Back Up the Data Repository](#).

### Prepare the New Host

Prepare the new host using the option based on your access:

- [As a sudo user with passwordless Secure Shell \(SSH\) set up.](#)
- [As root or as a sudo user with root passwordless SSH set up.](#)

### **Prepare as a Sudo User With Passwordless SSH Set Up**

Before adding the host to the cluster, configure and validate the `drinstall.properties` file. Complete this procedure once per node.

#### **Follow these steps:**

1. Copy the `installDR.bin` file to `/tmp` on the new host.
2. On the new host, extract the installation files from the `BIN` file by issuing the following command:

```
sudo ./installDR.bin
```

3. Copy the `drinstall.properties` file from an active node.
4. Edit the copied file.

The following is an example of the contents of this file:

```
# Linux user created to serve as the database administrator
DbAdminLinuxUser=dradmin
# Home directory for database administrator Linux user
# Make sure the parent directory of the home directory exists
# before running dr_install.sh
DbAdminLinuxUserHome=/export/dradmin
# Location of Vertica's data directory
DbDataDir=/data
# Location of Vertica's catalog directory
DbCatalogDir=/catalog
# Comma-delimited list of hostnames for the Data Repository
# === Do NOT place any spaces in the list of host names ===
DbHostNames=<dbNode1Hostname>
# Database name
DbName=drdata
# Database password
# Comment out with # to provide password interactively
DbPwd=dbpass
# Override drive scheduler
# If set to false, drive type will be determined per drive and the scheduler will
# be set based on /sys/block/DRIVE/queue/rotational value
overrideDriveScheduler=false
# noop (best for SSDs) or deadline (best for HDs)
overrideDriveSchedulerType=noop
```

5. Modify the `DbHostNames` parameter with the name of the new host, and then save your changes.
6. Run the `dr_validate.sh` validation script by issuing the following command:

```
sudo ./dr_validate.sh -sp drinstall.properties
```

#### **NOTE**

The validation script might ask you to reboot.

The validate script validates the configuration of the new host. The host is prepared.



## **Prepare as Root or as a Sudo User with Root Passwordless SSH Set Up**

Before adding the host to the cluster, configure and validate the `drinstall.properties` file. Complete this procedure once for all new nodes that you are adding to the data repository cluster.

### **Follow these steps:**

1. On the first data repository node, complete the following:
  - a. Edit the `drinstall.properties` file.
  - b. Modify the following parameter:
    - **DbHostNames**  
The name of the current hosts. Append the new hostnames that you want to add to the data repository cluster.  
**Example:** `DbHostNames=<dbNode1Hostname>,<dbNode2Hostname>,<dbNode3Hostname>`
  - c. Save your changes.
2. Run the `dr_validate.sh` validation script by issuing the following command based on your access:
  - As root:
 

```
./dr_validate.sh -p drinstall.properties
```
  - As sudo user:
 

```
sudo ./dr_validate.sh -p drinstall.properties
```

### **NOTE**

The validation script might ask you to reboot.

The validate script configures root passwordless SSH to the new nodes if not already configured. It also validates the configuration of the new host. The host is prepared.

## **Add the Host to the Cluster**

Add the prepared host to the cluster.

### **Follow these steps:**

1. Log in to one of the existing data repository hosts.
2. Do one of the following tasks:
  - If you have root access, issue the following command:
 

```
/opt/vertica/sbin/update_vertica --add-hosts <hostname> -u <DbAdminLinuxUser> -l <DbAdminLinuxUserHome> -L <location>/resources/vlicense.dat --rpm <location>/resources/<vertica-release.rpm> -T -S default
```

    - **hostname**  
Specify the name of the new host that you want to add to the cluster.
    - **DbAdminLinuxUser**  
The database administrator user name for the cluster as specified in the properties file.
    - **DbAdminLinuxUserHome**  
The database administrator user home directory.
    - **location**  
The location where you extracted `installDR.bin` file.
    - **vertica-release.rpm**  
The current RPM file that exists in the extracted installation directory.
  - If you have passwordless SSH set up for root, issue the following command as the sudo user:
 

```
export SUDO_USER=root /opt/vertica/sbin/update_vertica --add-hosts <hostname> -u <DbAdminLinuxUser> -l <DbAdminLinuxUserHome> -L <location>/resources/vlicense.dat --rpm <location>/resources/<vertica-release.rpm> -T -S default
```
  - If you have passwordless SSH set up for the sudo user, issue the following command as the sudo user:

```
sudo /opt/vertica/sbin/update_vertica --add-hosts <hostname> -u <DbAdminLinuxUser>
-l <DbAdminLinuxUserHome> -L <location>/resources/vlicense.dat --rpm <location>/
resources/<vertica-release.rpm> -T -S default
```

The script installs Vertica on the new host and incorporates the host into the cluster.

#### NOTE

If the Database Administrator user does not already exist, the installation script creates the user. The script prompts you to assign a new password. If the database administrator user exists, but passwordless SSH is not set up, the script prompts for the password to set up.

3. Ensure that the Vertica database administrator can write to the `data` and `catalog` directories by issuing the following commands:

```
chown DbAdminLinuxUser.verticadba DbDataDir DbCatalogDir
chmod 755 DbDataDir DbCatalogDir
```

The new host is added to the cluster.

### Add the Node to the Database

#### Follow these steps:

1. Log in to any host in the cluster as the data repository administrator.
2. Open the Vertica `adminTools` utility by issuing the following command:  
`/opt/vertica/bin/adminTools`
3. Select **Advanced Menu**, and then press the **Enter** key on your keyboard.
4. Select **Cluster Management**, and then press the **Enter** key on your keyboard.
5. Select **Add Host(s)**, and then press the **Enter** key on your keyboard.
6. Select the database, and then press the **Enter** key on your keyboard.  
The console displays a list of unused hosts.
7. Select the new host, and then press the **Enter** key on your keyboard.
8. Confirm the selection, and then follow the instructions in the console. When asked for the directory to use for Database Designer outputs, enter `/tmp`, and then click **OK**.  
The data repository adds the node to the cluster.

#### NOTE

This process can take a long time.

The node is added to the database.

### (Optional) Create the Buddy Projection and Mark the Cluster K-safety of 1

#### NOTE

Complete this procedure *only* if you are moving from one or two nodes to three or more nodes.

Create the buddy projection for the `migration_status` table, and then make the expanded cluster k-safety of 1.

1. Log into `vsq` as the `dradmin` user. This is the Linux user that serves as the Vertica database administrator system account (the database administrator user).
2. Issue the following command:

```
select export_objects('/tmp/migration_status.sql', '<SCHEMA>.migration_status');
```

#### – SCHEMA

Defines the schema.

**Example:** `dauser`

3. Edit the `/tmp/migration_status.sql` file.
4. Remove the `CREATE TABLE` statement and everything above it.

5. Edit the `CREATE PROJECTION` statement as follows:
  - Change `migration_status_super` to `migration_status_super_b1`.
  - After the `ALL NODES` statement and before the semi-colon (;), add `OFFSET 1`.
  - Before the `SELECT MARK_DESIGN_SAFE` statement, add `SELECT START_REFRESH()` ; .
  - Change `MARK_DESIGN_KSAFE(0)` to `MARK_DESIGN_KSAFE(1)` .
6. Save your changes.
7. Issue the following command:
 

```
vsq1 -U dradmin -w <dradminPassword> -f /tmp/migration_status.sql
```

A prompt should confirm that the refresh background process has begun.

The buddy projection is created and the cluster is marked k-safety of 1.

### Update the Data Repository Configuration

On the data aggregator host, ensure that the `dbconnection.cfg` file points to the correct data repository hosts.

#### **Follow these steps:**

1. Edit the `<installation_directory>/apache-karaf/etc/dbconnection.cfg` file on the data aggregator host.

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

The following is an example of the contents of this file:

```
dbDriverClassName=com.vertica.jdbc.Driver
dbUrl=jdbc:vertica://{dbHost}:5433/drdata
dbUser=dauser
dbEncryptedPassword=QUVTOkMZrUfMytW85NBX+bj7T4w=
dbAdminUser=dradmin
dbEncryptedAdminPassword=QUVTOmwDpCHP5zlGhX7b2ce9I9M=
dataLoaderConnProps=&directBatchInsert=true&Label="Data Loader"
dbHostNames=<dbNode1Hostname>
```

2. Modify the `dbHostNames` parameter to reference the hostname or IP address of each data repository host, for example:

```
dbHostNames=<dbNode1Hostname>, <dbNode2Hostname>, <dbNode3Hostname>
```

3. Save your changes.
4. Do one of the following steps:
  - Start the Data Aggregator service by issuing the command:
 

```
systemctl start dadaemon
```
  - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the command:

```
<installation_directory>/scripts/dadaemon activate
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

The data repository configuration is updated.

## (Optional) Segment the Database

### NOTE

Segment the database *only* if you are moving from one or two nodes to three or more nodes.

For more information, see [Segment Database Tables](#).

## Data Repository Heartbeat Monitor Process

The heartbeat monitor process checks whether the data repository is running every 10 seconds.

If the heartbeat process fails to confirm that the database is running after 5 minutes, the data aggregator shuts down. An audit message is logged in the `<installation_directory>/apache-karaf/shutdown.log` file.

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

In a cluster environment, DX NetOps Performance Management checks the availability of the nodes in the cluster continuously every 10 seconds. If DX NetOps Performance Management cannot contact a node within 5 minutes, it generates and logs an event. DX NetOps Performance Management logs an audit message in the `shutdown.log` file.

If the data repository node that failed is the primary node, the data aggregator automatically switches to the next available node. DX NetOps Performance Management generates and logs an event.

### IMPORTANT

DX NetOps Performance Management interrupts certain administrative functions that occur during a high-availability failover, and then they fail. One poll cycle is lost. These functions do not resume after the data repository connects to another node in the cluster environment. The administrative functions that you perform after the data repository connects to another node in the cluster environment work as designed.

If more than one data repository node fails, the data aggregator shuts down.

The data aggregator shuts down automatically if it fails to connect to the data repository on start-up.

## Data Repository Audit Process

The license agreement states that the total data stored in Data Repository cannot exceed 64 TB. The audit process audits the database daily at 3:00 AM to calculate the total space that Data Aggregator data occupies.

To view the most recent result of an audit, access the following URL in your browser:

`http://da_host:8581/rest/datarepositorymaintenance/audit`

This URL returns XML. The "Current Size" tag displays the current size of Data Repository in bytes.

### WARNING

Review the audit results periodically. If you see a value greater than 64 TB, you are not in compliance with the license agreement. Contact CA Technical Support for further instructions.

## Run Data Repository Diagnostic Utilities

You can verify that the environment is ideal for the database by testing the performance of your hardware for the data repository. Run the following Vertica utilities *before* you upgrade the data repository:

Run these utilities at any time. They are available on each node in the `/opt/vertica/bin` directory.

### IMPORTANT

If the tests do not meet the recommendations, fix the issues before you continue the upgrade.

## **vcuperf**

The `vcuperf` utility measures the CPU processing speed of the host and compares the speed against benchmarks for common server CPUs. The utility measures how long the server requires to complete the test, and determines whether CPU throttling is enabled.

For more information about this utility, see the [Vertica documentation](#).

### **Follow these steps:**

1. Issue the following command on *each* data repository node:  

```
./vcuperf > /tmp/vcuperf.out
```
2. Verify that the performance meets the following requirements:
  - The CPU time is consistent with the benchmark values in the output.
  - The low load time and high load time are within 10 microseconds. If the difference is greater than 50 microseconds, CPU throttling might be enabled on your system. Disable CPU throttling.

### **Example:**

The following example shows the return from this utility:

```
$ /opt/vertica/bin/vcuperf
Compiled with: 4.1.2 20080704 (Red Hat 4.1.2-52)
Expected time on Core 2, 2.53GHz: ~9.5s
Expected time on Nehalem, 2.67GHz: ~9.0s
Expected time on Xeon 5670, 2.93GHz: ~8.0s
```

```
This machine's time:
CPU Time: 7.740000s
Real Time: 7.740000s
```

Some machines automatically throttle the CPU to save power.

This test can be done in <100 microseconds (60-70 on Xeon 5670, 2.93GHz).

Low load times much larger than 100-200us or much larger than the corresponding high load time

indicate low-load throttling, which can adversely affect small query / concurrent performance.

**This machine's high load time: 67 microseconds.**

**This machine's low load time: 64 microseconds.**

This test was performed on a system with 2.67-GHz processors, so the real time is acceptable. The difference between the high load time and low load time is within the expected tolerance.

## **vioperf**

The `vioperf` utility tests the performance of the disk input and output (I/O). The utility performs a series of reads and writes.

For more information about this utility, see the [Vertica documentation](#).

**Follow these steps:**

1. Issue the following commands on *each* data repository node:

```
./vioperf /data --duration=60sec > /tmp/vioperf.out
```

**/data** is the full path of the data directory.

```
./vioperf /catalog --duration=60sec > /tmp/vioperf.out
```

**/catalog** is the full path of the catalog directory.

2. Verify the Write and Read counter values at least 40 MB/s per core.

The recommended I/O is 40 MB/s per physical core on each node. For example, the recommended I/O rate for a node with 2 hyper-threaded six-core CPUs (12 physical cores) is 480 MB/s.

**IMPORTANT**

If the **thread count** column shows a value of 1, the utility cannot determine the number of cores. Add the following argument to the command to run the utility:

```
--thread-count=CORES
```

**Cores** defines the number of cores in the system as a fixed integer.

**Example:**

The following example shows the return from this utility for the data directory:

The minimum required I/O is 20 MB/s read and write per physical processor core on each node, in full duplex i.e. reading and writing at this rate simultaneously, concurrently on all nodes of the cluster. The recommended I/O is 40 MB/s per physical core on each node. For example, the I/O rate for a server node with 2 hyper-threaded six-core CPUs is 240 MB/s required minimum, 480 MB/s recommended.

Using direct io (buffer size=1048576, alignment=512) for directory "/drdata"

test	directory	counter value	counter value (10 sec avg)	counter value/core (10 sec avg)	counter name	thread count	%CPU	%IO Wait	elapsed time (s)	remaining time (s)
<b>Write</b>	/drdata	873	873	54.5625	MB/s	16	29	40	10	5
<b>Write</b>	/drdata	868	865	54.25	MB/s	16	28	30	15	0
ReWrite	/drdata	275+275	275+275	17.1875+17.1875	(MB-read+MB-write)/	16	13	21	10	5

ReWrite	/drdata				(MB-read+MB-write)/
s  242+242	178+178				15.125+15.125
11.125+11.125	16	7	17	15	
0					
<b>Read</b>	/drdata				MB/s
735	735				<b>45.9375</b>
	16	11	23	10	45.9375
					5
<b>Read</b>	/drdata				MB/s
786	786				<b>49.125</b>
	16	26	25	15	49.125
					0
SkipRead	/drdata				seeks/s
4511	4511				281.938
	16	14	19	10	281.938
					5
SkipRead	/drdata				seeks/s
4477	4407				279.812
	16	3	15	15	275.438
					0

This server has 16 cores. The Read and Write counter values indicate the I/O is greater than 40 MB/s per core.

### **vnetperf**

The `vnetperf` utility tests the network performance of the data repository hosts. The utility measures network latency and the throughput for the TCP and UDP protocols.

#### **IMPORTANT**

This utility causes a high network load and affects database performance. Do not run this utility while the database is running.

For more information about this utility, see the [Vertica documentation](#).

#### **Follow these steps:**

1. Log in as a user that has passwordless ssh between the nodes.
2. Issue the following command on *one* of the data repository nodes:

```
./vnetperf --hosts DAhost,DRhost1,DRhost2,DRhost3 > /tmp/vnetperf.out
```

Specify the hostname or IP address of the data aggregator host and each data repository host.

3. Verify that the network performance meets the following requirements:

- Round-trip time (RTT) latency of 200 microseconds or less
- Clock skew under 1 second
- Throughput of 800 MB/s or more

The utility runs a series of throttled tests. Verify the throughput for the highest speed test.

## **Segment Database Tables**

If you received the `unsegmented projections exist` table segmentation message while upgrading the data aggregator, segment the database tables on the data repository.

Table segmentation is a one-time task that is required for systems where the original DX NetOps Performance Management installation was DX NetOps Performance Management 2.3.2 or earlier. Segmenting the tables reduces the amount of disk space that is required for the database. Segmenting the tables also improves general query performance.

#### TIP

If you are unsure whether your database requires segmentation, attempt to identify tables that require segmentation using the `segment.py` script. This script is located in the `/opt/CA/IMDataRepository_vertica<version>/dr-health` directory. If segmentation is already complete or not required, the script does not return any tables.

Segmentation can take several hours to segment large tables in the database. During tests, migration of tables larger than 100 GB took over 10 hours. The segmentation time is not uniform to table size. Time depends on many factors including row count, column count, compression of the data, and system specifications.

No active monitoring of your infrastructure environment occurs when the data aggregator and data collectors are shut down.

The data aggregator and the data collectors can be running or shut down when you segment the database tables. When segmenting the tables in the database, if the data aggregator is running, at least 40 percent of the available disk space must remain free for query processing and other database activities. When the data aggregator is shut down, the total disk utilization during segmentation must not exceed 90 percent of available disk space. Tables that would cause the disk utilization to exceed these limits during segmentation are not segmented.

If you choose to segment the database tables while the data aggregator is running, the following restrictions apply:

#### NOTE

Segmentation is a resource-intensive process. Complete this process when the data aggregator and the data collectors are shut down.

- Do not perform any data aggregator-administrative functions, such as:
  - Modifying monitoring profiles
  - Associating collections to monitoring profiles
  - Increasing poll rates
  - Running new discoveries
- Minimize the report load.

Use the following process to segment the database tables:

1. [Identify the tables that require segmentation.](#)
2. [Back up the data repository.](#)
3. [Segment the tables with no data.](#)
4. [Determine the table segmentation time.](#)
5. [Segment the remaining database tables.](#)

### Identify the Tables that Require Segmentation

#### Follow these steps:

1. Log in to a host in the data repository cluster as the Vertica Linux database administrator user.
2. Run the `segment.py` script with the `--task tables` option:

```
./segment.py --task tables --pass <database_admin_user_password> [--name <database_name>] [--port <database_port>]
```

#### For example:

```
./segment.py --task tables --pass password --name mydatabase
```

- **database\_admin\_user\_password**

The Vertica Linux database administrator user password.



**Example:** password

– **database\_name**

The name of the database.

**Case-sensitive:** yes

**Default:** drdata

**NOTE**

This parameter is case-sensitive.

– **database\_port**

The Vertica port.

**Default:** 5433

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

A list of unsegmented table projections sorted from largest to smallest is returned. If tables require segmentation, continue this process.

## **Back up the Data Repository**

For information, see [Back Up the Data Repository](#).

**NOTE**

- After segmentation, the disk space for the data repository backup increases by the amount of data in the new segmented table projections. Verify that the data repository backup system has enough disk space available after segmentation is completed and before backups run.

For information about sizing guidelines, see [the DX NetOps Performance Management Sizing Tool](#).

- The data for the old unsegmented table projections is removed from data repository backup data one day after the time of the `restorePointLimit`.

To remove this data immediately:

- a. Change the `snapshotName` in the data repository backup configuration file.
- b. Do a full data repository backup.
- c. Archive the older data repository backup.
- d. Delete the data repository backup from the backup disk.

Use the pre-segmentation data repository backup only if you cannot use the backup that was created after segmentation completed. If you use the pre-segmentation data repository backup, segment the tables again.

## **Segment the Tables with No Data**

Tables with no data are segmented quickly and segmentation does not negatively affect performance. You can segment these tables while the data aggregator is running.

### **Follow these steps:**

1. Log in to any host in the data repository cluster as the Vertica Linux database administrator user.
2. Run the `segment.py` script with the `--task segment` and `--zerotables` options by issuing the following command:

```
./segment.py --task segment --zerotables --pass <database_admin_user_password> [--name <database_name>] [--port <database_port>]
```

**For example:**

```
./segment.py --task segment --zerotables --pass password --name mydatabase
```

– **database\_admin\_user\_password**

The Vertica Linux database administrator user password.

– **database\_name**

The name of the database.

**Default:** drdata

**NOTE**

This parameter is case-sensitive.

- **database\_port**  
The Vertica port.  
**Default:** 5433

The script segments the database tables with no data.

### **Determine the Table Segmentation Time**

To determine whether to segment the tables while the data aggregator is running or shut down, calculate the necessary time for segmentation.

#### **Follow these steps:**

1. While logged in to a host in the data repository cluster as the Vertica Linux database administrator user, retrieve a list of the tables that require segmentation by running the `segment.py` script with the `--task tables` option:

```
./segment.py --task tables --pass <database_admin_user_password> [--  
name <database_name>] [--port <database_port>]
```

**For example:**

```
./segment.py --task tables --pass password --name mydatabase
```

- **database\_admin\_user\_password**  
The Vertica Linux database administrator user password.  
**Example:** password
- **database\_name**  
The name of the database.  
**Case-sensitive:** yes  
**Default:** drdata
- **database\_port**  
The Vertica port.  
**Default:** 5433

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

The list sorts the tables from largest to smallest.

2. To ensure that scheduled data repository backups do not interfere with the segmentation process, [disable the automatic backup of the data repository](#).
3. Segment a table projection that is about 5 GB in size by running the `segment.py` script with the `--task segment` and `--table` options:

**NOTE**

You can run this command while the data aggregator is running, however run the command during a 2-3 hour maintenance window.

```
./segment.py --task segment --table <rate_table_name> --  
pass <database_admin_user_password> [--name <database_name>] [--port <database_port>]
```

**For example:**

```
./segment.py --task segment --table mytable --pass password --name mydatabase
```

- **rate\_table\_name**  
The 5-GB table that requires segmentation.
- **database\_admin\_user\_password**  
The Vertica Linux database administrator user password.

**Example:** password

– **database\_name**

The name of the database.

**Case-sensitive:** yes

**Default:** drdata

– **database\_port**

The Vertica port.

**Default:** 5433

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

4. [Reenable the automatic backup of the data repository.](#)

5. Use the segmentation time for the 5-GB table to determine how long segmentation might take to segment all of the tables that are less than 100 GB.

**NOTE**

The actual segmentation time for the database tables can vary based on the type and compression of the data in the tables. The values that are calculated here are rough estimates. When planning a scheduled maintenance window, add an extra hour of time for every 10 GB to 15 GB of database tables. For large databases, you might not be able to schedule a single maintenance window that is long enough to segment the entire database. In this case, you can segment the database tables over multiple maintenance windows.

## **Segment the Remaining Database Tables**

### **Follow these steps:**

1. While logged in to a host in the data repository cluster as the Vertica Linux database administrator user, if there are *more than ten* zero-length table projections that require segmentation, segment them by running the `segment.py` script with the `--task segment` and `--zerotables` options:

```
./segment.py --task segment --pass <database_admin_user_password> --zerotables [--name <database_name>] [--port <database_port>]
```

**For example:**

```
./segment.py --task segment --pass password --zerotables --name mydatabase --port 1122
```

– **database\_admin\_user\_password**

The Vertica Linux database administrator user password.

**Example:** password

– **database\_name**

The name of the database.

**Case-sensitive:** yes

**Default:** drdata

– **database\_port**

The Vertica port.

**Default:** 5433

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

2. (If there are table projections that are *greater than* 100 GB in size that require segmentation) Segment the table projections that are *less than* 100 GB by running the `segment.py` script with the `--task script` and `--lt100G` options:

```
./segment.py --task script --pass <database_admin_user_password> --lt100G [--name <database_name>] [--port <database_port>]
```

**For example:**

```
./segment.py --task script --pass password --lt100G --name mydatabase --port 1122
```

– **database\_admin\_user\_password**

The Vertica Linux database administrator user password.

**Example:** password

– **database\_name**

The name of the database.

**Case-sensitive:** yes

**Default:** drdata

– **database\_port**

The Vertica port.

**Default:** 5433

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

3. To ensure that scheduled data repository backups do not interfere with the segmentation process, [disable the automatic backup of the data repository](#).

4. Run the `segment-script.sh` script by issuing the following command:

```
nohup ./segment-script.sh
```

The script segments the unsegmented table projections that are *less than* 100 GB and sorts them from smallest to largest. The output is sent to `nohup.out`. If the shell is closed accidentally, the script continues to run.

5. Depending on your maintenance window size and the combined size of all of the tables under 100 GB, determine which tables can be segmented in the maintenance window. Modify the generated script by removing the tables that do not fit inside the maintenance window, and then segment those tables (by running the generated `segment-script.sh` script) during the maintenance window. If all of the tables under 100 GB could not be segmented in the maintenance window, regenerate the script, and run the `segment-script.sh` script during the next maintenance window until all of the tables have been segmented.

**IMPORTANT**

- When you run the script, any tables that cause disk utilization to exceed 90 percent display an error message. These tables are not segmented. To segment these tables, more available disk space is needed.
  - You are prompted for each table that can cause disk utilization to exceed 60 percent. Before segmenting these tables, [shut down \(stop\) the data aggregator](#).
  - The `segment-script.sh` script can take several hours to execute. Do not interrupt the script execution once it begins to avoid corruption of the database.
6. (If more segmentation is needed and you will do this in a future maintenance window) [Reenable the automatic backup of the data repository](#).
  7. Segment the table projections that are *over* 100 GB by running the `segment.py` script with the `--task script`:  

```
./segment.py --task script --pass <database_admin_user_password> [--  
name <database_name>] [--port <database_port>]
```

**For example:**

```
./segment.py --task script --pass password --name mydatabase --port 1122
```

– **database\_admin\_user\_password**

The Vertica Linux database administrator user password.

**Example:** password

– **database\_name**

The name of the database.

**Case-sensitive:** yes

**Default:** drdata

- **database\_port**  
The Vertica port.

**Default:** 5433

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

Those tables that might cause disk utilization to exceed 60 percent and 90 percent are indicated.

8. (If scheduled data repository backups are not already disabled) [Disable the automatic backup of the data repository](#).
9. Run the `segment-script.sh` script by issuing the following command:

```
nohup ./segment-script.sh
```

The script segments the unsegmented table projections and sorts them from smallest to largest.

10. Verify that all table projections have been segmented by running the `segment.py` script with the `--task tables`:  

```
./segment.py --task tables --pass <database_admin_user_password> [--  
name <database_name>] [--port <database_port>]
```

#### For example:

```
./segment.py --task tables --pass password --name mydatabase --port 1122
```

- **database\_admin\_user\_password**  
The Vertica Linux database administrator user password.

**Example:** password

- **database\_name**  
The name of the database.

**Case-sensitive:** yes

**Default:** drdata

- **database\_port**  
The Vertica port.

**Default:** 5433

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

The following message appears:

```
No tables found with unsegmented projections.
```

11. [Reenable the automatic backup of the data repository](#).
12. (If you segmented the database tables when data aggregator and data collector were shut down) [Restart these components](#).

## Move the Data Repository Data Directory

If necessary, you can move the data repository data directory from an existing location to another location on the same Vertica cluster.

This process involves the following steps:

1. Create the new location.
2. Move the data.
3. Drop the old location.

For more information, see the [Vertica documentation](#).

You might need to move the Data Repository data directory for the following reasons:

- You want to add new storage to an existing server.
- You want to move the database from an old mount point to a new mount point because of an updated server build template.

The following process causes some downtime for DX NetOps Performance Management while the move occurs.

The following factors impact the total amount of downtime:

- Size of the database
- Storage input and output speed

The following process covers only moving the data repository data directory.

For information about moving the data repository database from one server to another, see [Disaster Recovery](#).

#### Example:

- You are moving the data directory for a three-node cluster:
  - node0001
  - node0002
  - node0003
- The current data directory is: `/spare/dbdata/data`
- The new data directory is: `/opt/application/CA/drdata`
- The Vertica database name is `drdata`.
- The `dbadmin` user for Vertica is `dradmin`.
- The following end goals apply to this scenario:
 

```
/spare/dbdata/data/drdata/v_drdata_node0001_data >> /opt/application/CA/drdata/drdata/v_drdata_node0001_data
/spare/dbdata/data/drdata/v_drdata_node0002_data >> /opt/application/CA/drdata/drdata/v_drdata_node0002_data
/spare/dbdata/data/drdata/v_drdata_node0003_data >> /opt/application/CA/drdata/drdata/v_drdata_node0003_data
```

#### Follow these steps:

1. On each node, create the new data directory by issuing the following command:
 

```
mkdir -p /new_data_directory/dbname/v_dbname_node000x_data
```

  - **new\_data\_directory**  
Specify the new data directory.
  - **dbname**  
Specify the name of the database.
  - **x**  
Specify the existing node number.
2. Grant permissions to the new data directory by issuing the following command:
 

```
chown -R dradmin:verticadba /new_data_directory/dbname
```
3. For each parent in the path to `/new_data_directory/dbname`, grant permissions by issuing the following command:
 

```
chown dradmin:verticadba /parent
```
4. As `dradmin`, create the new storage location for each node in Vertica by issuing the following command:

```
create location '/new_data_directory/dbname/v_dbname_node000x_data' NODE
'v_dbname_node000x' USAGE 'DATA,TEMP' LABEL 'TO_DATA_TEMP';
```

– **TO\_DATA\_TEMP**

Specify a label for the new storage location. This label is required when you set the object policy.

5. Confirm the new storage locations by issuing the following command:

```
select * from storage_locations;
```

6. Set the object policy by issuing the following command:

```
select set_object_storage_policy('schema_name', 'TO_DATA_TEMP', true );
```

– **schema\_name**

Specify the name for the schema.

**Example:** dauser

To view a lists of the available schemas, issue the following command:

```
dradmin=> \dn
```

– **TO\_DATA\_TEMP**

Specify the same label used when you created the new storage location.

7. Trigger the object to move by issuing the following command:

```
select enforce_object_storage_policy('schema_name');
```

8. Confirm that the policy is applied by issuing the following command:

```
select * from storage_policies;
```

9. Retire the old storage locations for each node by issuing the following command:

```
select retire_location('/existing_data_directory/dbname/v_dbname_node000x_data',
'v_dbname_node000x', true);
```

– **existing\_data\_directory**

Specify the existing data location.

10. Confirm that the old storage locations are retired by issuing the following command:

```
select is_retired, location_path from storage_locations;
```

11. Drop the old storage location for each node:

```
select drop_location('/existing_data_directory/dbname/v_dbname_node000x_data',
'v_dbname_node000x');
```

12. Clear the storage policy by issuing the following command:

```
select clear_object_storage_policy( 'schema_name' );
```

13. Remove the label from the new storage location for each node by issuing the following command:

```
select alter_location_label('/new_data_directory/dbname/v_dbname_node000x_data',
'v_dbname_node000x', '');
```

## Data Repository Scripts

The DX NetOps Performance Management installation includes the following scripts for the data repository:

- [caVerticaUtility.sh](#)
- [dr\\_install.sh](#)
- [dr\\_validate.sh](#)
- [etlHealth.sh](#)
- [migrate\\_sdn\\_device\\_metrics.sh](#)
- [RemoteEngineer/re.sh](#)
- [reset\\_vertica\\_password.sh](#)
- [dr-health/segment.py](#)
- [update\\_da\\_dc\\_database\\_references.sh](#)
- [update\\_backup\\_password.sh](#)

### **caVerticaUtility.sh**

The `caVerticaUtility.sh` script exports the iRep data from an existing system that you can use as a backup or for Broadcom Support to use for triage purposes.

For more information about when to run this script, see [Back Up the Data Repository](#).

### **dr\_install.sh**

The `dr_install.sh` script installs or upgrades the Vertica database.

For more information about when to run this script:

- See [Install the Data Repository](#).
- See [Upgrade the Data Repository](#).

### **dr\_validate.sh**

The `dr_validate.sh` script verifies the OS settings and modifies the settings if necessary. This script runs the `etlHealth.sh` script.

For more information about when to run this script:

- See [Upgrade Fault-Tolerant Data Aggregators](#).
- See [Upgrade a Non-Fault-Tolerant Data Aggregator](#).
- See [Upgrade the Database](#).
- See [Install the Data Repository](#).
- See [Verify ETL Health](#).

### **etlHealth.sh**

The `etlHealth.sh` script verifies extract, transform, load (ETL) health. This script collects data. The `dr_validate.sh` validation script runs this script.

For more information about when to run this script:

- See [Upgrade the Database](#).
- See [Verify ETL Health](#).



**migrate\_sdn\_device\_metrics.sh**

The `migrate_sdn_device_metrics.sh` script migrates virtual disk usage data on existing SDN devices from the Virtual Disk metric family to the SDN Devices Metrics metric family, and unifies the data into a single metric family for VNA to send in the data aggregator.

For more information about when to run this script:

- See [Install the Data Repository](#).
- See [Migrate Virtual Disk Usage Data for SDN Devices](#).

**RemoteEngineer/re.sh**

The `re.sh` script (the CA Remote Engineer (CARE) tool) gathers the DX NetOps Performance Management configuration and database into a compressed file for Broadcom Support triage purposes.

For more information about when to run this script, see [Unable to Resolve Issue](#).

**reset\_vertica\_password.sh**

The `reset_vertica_password.sh` script updates the Vertica database administrator system account (the database administrator user) and/or the database administrator user password.

**dr-health/segment.py**

The `segment.py` script identifies tables that require segmentation.

For more information about when to run this script, see [Segment Database Tables](#).

**update\_da\_dc\_database\_references.sh**

When reinstalling or migrating the data collector, if the new data collector is going to run on a new IP address and hostname, use the `update_da_dc_database_references.sh` script to update the data aggregator database to reflect the new environment (the new host). Or when reinstalling or migrating the data aggregator, if the new data aggregator is going to run on a new IP address and hostname, use the `update_da_dc_database_references.sh` script to update the data aggregator database to reflect the new environment (the new host).

For more information about when to run this script:

- See [Reinstall or Migrate the Data Collectors](#).
- See [Reinstall or Migrate the Data Aggregator](#).
- See [Install a Disaster Recovery System](#).

**update\_backup\_password.sh**

The `update_backup_password.sh` script updates the configuration file that you reference when restoring or backing up the data repository (the `.ini` file) with the updated Vertica database administrator user password.

## Flow Administration

You can configure flow monitoring (from DX NetOps Network Flow Analysis (NFA)) in NetOps Portal.

You can do the following using NetOps Portal:

- [Manage Monitoring](#)
- [Configure the Application Settings](#)
- [Establish Secure Communication While Sending Emails for NFA Reports](#)
- [Configure SNMP Polling](#)
- [Configure Device-Level Flow Direction](#)
- [Configure the Settings for the Watchdog Services](#)
- [Manage a Harvester](#)
- [Edit an AS Name](#)
- [Manage an Application Mapping Rule](#)
- [Manage Port Priorities](#)

### **Manage Monitoring**

You can manage monitoring (enable and disable network flow processing for your DX NetOps Network Flow Analysis (NFA) interfaces, edit interfaces, and delete interfaces) using NetOps Portal.

#### **Follow these steps:**

1. Hover over **Administration, Data Sources**, and then click the NFA datasource.  
The page for the NFA data source opens. By default, the **Manage Monitoring** tab from the **Monitored Interfaces** menu is selected.
2. Change the settings as needed, and then click **Save**.

### **Configure the Application Settings**

You can configure many NFA settings from NetOps Portal.

#### **Follow these steps:**

1. From the page for the NFA data source, select **Application Settings** from the **Configuration Settings** menu.  
The **Application Settings** page appears and displays the current settings.
2. Change the settings as needed, and then click **Save**.

#### **NOTE**

The application settings that you leave blank default to the values that NetOps Portal receives from NFA.

### **Establish Secure Communication While Sending Emails for NFA Reports**

You can establish secure communication while sending emails for NFA reports using SMTP servers with TLS/SSL. NFA leverages the email server settings present in NetOps Portal. You can view the SMTP server details from the **Manage Email Server Settings** page. NetOps Portal email server settings always override the NFA Console email server settings.

#### **NOTE**

Ensure that you install the SMTP server certificate file on the NFA Console server to use SMTP servers with TLS/SSL.

For more information about how to configure the email server, see [Configure the Email Server](#).

#### **NOTE**

You cannot use the **SMTP Server** option on the **Application Settings** page (under **Administration, Data Sources, NFA data source, Configuration Settings**). The option is disabled. For easier navigation, you can click the link to the **Manage Email Server Settings** page from this location.

## Configure SNMP Polling

You can enable or disable SNMP polling for a router. This ability helps you decide whether you want to perform SNMP polling on a device and collect the related SNMP data. If you do not want to collect the SNMP data, you can disable the SNMP polling functionality. Similarly, to collect the data, enable the polling.

### Follow these steps:

1. Hover over **Administration, Data Sources**, and then click the NFA data source.
2. Verify that **Manage Monitoring** is selected under **Monitored Interfaces**.
3. Select the router for which you want to configure the SNMP polling.

#### NOTE

- You can also select multiple routers at the same time and enable or disable their SNMP polling.
- Use the **Quick Filter** to search for the required router if you have a long list.

4. Click **Edit**.

The **Edit Router(s)** dialog opens.

5. Update the following field, and then click **Save**:

#### – **SNMP Polling**

Specifies whether SNMP polling is enabled or disabled.

#### Options:

- **Enabled:** Indicates that you want to perform SNMP polling on devices.
- **Disabled:** Indicates that you do not want to perform SNMP polling on devices.

#### NOTE

The **Discover Profile**, **SNMP Refresh**, and **Test Profile** actions are disabled when you select the **Disabled** option.

6. (Optional) Add the **SNMP Polling** column to the table so that you can quickly view the SNMP polling status:
  - a. Hover over any existing column (for example, **Harvester Address**).
  - b. Click the gear icon.
  - c. Select **Column, SNMP Polling**.

The **SNMP Polling** column is added to the table along with the corresponding value (Disabled or Enabled) for each row.

The SNMP polling is configured successfully.

## Configure Device-Level Flow Direction

You can configure device-level flow direction so that you can eliminate duplication of data flow between the router and harvester. By eliminating the data duplication, you can view the traffic utilization, bandwidth information in the report.

### Follow these steps:

1. Hover over **Administration, Data Sources**, and then click the NFA data source.
2. Verify that **Manage Monitoring** is selected under **Monitored Interfaces**.
3. Select the router for which you want to configure the flow direction.

#### NOTE

- You can also select multiple routers at the same time and change their flow direction.
- Use the **Quick Filter** to search for the required router if you have a long list.

4. Click **Edit**.

The **Edit Router(s)** dialog opens.

5. Update the following field, and then click **Save**:

#### – **Flow Direction**

Defines the flow direction.

**Options:**

- **Either:** Indicates that the flow is observed in either Ingress or Egress direction.
- **Both:** Indicates that the flow is observed in both Ingress and Egress directions.

6. (Optional) Add the **Flow Direction** column to the table so that you can quickly view the flow direction value:

- Hover over any existing column (for example, **Harvester Address**).
- Click the gear icon.
- Select **Column, Flow Direction**.

The **Flow Direction** column is added to the table along with the corresponding value for each row.

The device-level flow direction is configured.

### **Configure the Settings for the Watchdog Services**

Use the Watchdog Services to monitor components. The Watchdog Services poll each server in your configuration once every hour to determine the status of all components. You can configure the Watchdog Services to use the SNMPv3 profile for communicating with the SNMP Agent installed on the NFA Console / Harvester Hosts. The Watchdog Service uses this configuration and retrieves the Health Status information about the System CPU, memory, or Disk and the status of the NFA Processes. You can also configure the polling frequency based on your requirement and you can establish thresholds, an email address for receiving messages, and other settings for the Watchdog Services to ensure that you are notified of issues with the components as soon as possible.

Before configuring the Watchdog service with SNMPv3, ensure that you configure the SNMPv3 agent for Console (Windows) and Harvester(s) (Windows/Linux) as follows:

- For Windows, to use SNMPv3, you must install the **third party** SNMPv3 agent, and then configure it for SNMPv3. For more information, see the respective agent documentation for configuration.
- For Linux, **Net-SNMP** agent is required and it is installed as a prerequisite for Network Flow Analysis. You can configure the Net-SNMP agent to enable SNMPv3. For more information see, [Configure SNMPv3 User on RHEL 8](#).

**NOTE**

In the **Watchdog Service** settings (user name, auth protocol, and auth password), you must use the configured SNMPv3 user details (password, protocol, and user name).

**Follow these steps:**

1. Hover over **Administration, Data Sources, Network Flow Analysis Data Source Display Name**.
2. Click **Watchdog Settings** under Configuration Settings, and then edit the settings as follows:
  - **SNMP Version (SNMP v1/v2c or SNMPv3)**  
Specifies the SNMP version to be used by the Watchdog service for polling the NFA services.
  - **(SNMPv3) User Name**  
Username for the SNMPv3 profile.
  - **(SNMPv3) Authentication Protocol**  
Protocol used for Authentication in the SNMPv3 profile.  
**Values:** None, MD5, SHA, SHA2\_256, SHA2\_512
  - **(SNMPv3) Authentication password**  
Password to be used in the Authentication process.
  - **(SNMPv3) Privacy Protocol**  
Use the Protocol for Communication in the SNMPv3 profile.  
**Values:** None, DES, AES, TripleDES, AES192, AES256
  - **(SNMPv3) Privacy password**  
Use the password for the communication process.
  - **(SNMPv1/SNMPv2C) Community String**

The SNMP string that the Watchdog Services use to verify the identity of components in a distributed deployment. The community string is used for gathering information from Harvesters. Use the same community name throughout the DX NetOps deployment:

- The **Watchdog Settings** page
- SNMP service on each Windows server
- The `snmpd.conf` file on each Linux server

**Default:** public

– **CPU Threshold**

The threshold for CPU utilization. You are notified by email when the CPU threshold on a server is exceeded on any server and an SNMP trap notification is generated, provided that the address and string are set.

**Default:** 80 percent CPU utilization

– **Disk Threshold**

The threshold for disk utilization. If the disk threshold on a server is exceeded, you are notified by email and an SNMP trap notification is generated, provided that the address and string are set.

**Default:** 80 percent disk utilization

– **Email Address**

The destination email address to use for email notifications when thresholds are exceeded. To notify multiple recipients, separate the addresses with commas. The **Email Address** setting has no default value.

**Default:** (none)

– **Memory Threshold**

Threshold for memory utilization. You are notified by email when the memory threshold on a server is exceeded, provided that the address and string are set.

**Default:** 80 percent memory utilization

– **SNMP Retries**

Number of times the program attempts to poll an SNMP device. A high number of SNMP Retries can affect performance, depending on your network configuration.

**Default:** 2

– **SNMP Timeout**

Number of seconds before an SNMP poll times out.

**Default:** 5

– **System Check Interval**

Number of minutes between Watchdog system checks.

**Default:** 60

– **Trap Community String**

SNMP string to use for sending traps to a third-party trap receiver. Use one of the community names that the trap receiver is configured to accept.

**Default:** public

– **Trap Destination**

IP address of the server that receives SNMP traps from the Watchdog Services. The traps are generated when thresholds for DX NetOps component performance are violated.

**Default:** (none)

3. Click **Save**.

The settings Watchdog Services are configured.

## **Manage a Harvester**

You can manage (view, add, edit, and delete) harvesters.

Follow these steps:

1. From the page for the NFA data source, select **Harvester** from the **Configuration Settings** menu.

The **Harvester** page appears and displays the current settings.

2. Do one of the following:

- To add a harvester, complete the following:
  - a. Click **Add**.  
The **Add Harvester** dialog opens.
  - b. Enter an **IP Address**, a **Domain**, and a **Description**, and then click **Save**.
- To edit a harvester, complete the following:
  - a. Select the row for the harvester that you want to edit, and then click **Edit**.  
The **Edit Harvester** dialog opens.
  - b. Enter an **IP Address**, a **Description**, and a **Domain**, and then click **Save**.
- To delete a harvester, complete the following:

**NOTE**

Deleting a harvester deletes the router and all its interfaces from the system and reduces the license count. You cannot undo this action.

- a. Select the row for the harvester that you want to delete, and then click **Delete**.  
The **Delete Harvester** dialog opens.
- b. Confirm the deletion by clicking **Yes**.

### **Edit an AS Name**

Interface reports that show data about AS traffic typically label the AS traffic by name and number. You can make the AS references in reports shorter or more descriptive by editing the AS Names as needed.

**NOTE**

AS names are domain-specific. The changes that you make affect reports about interfaces in the selected domain.

### **Follow these steps:**

1. From the page for the NFA data source, select **AS Names** from the **Define an Application** menu.  
The **AS Names** page appears and displays the AS Names associated with the selected domain.
2. Select the row for the AS number that you want to edit, and then click **Edit**.  
The **Edit AS Names** dialog opens.

**NOTE**

You can restore an AS name to its official (base) name by selecting the AS name, clicking **Reset**, and then clicking **Yes** in the **Reset** dialog that opens.

3. Enter the desired description in the **New Name** field, and then click **Save**.

**NOTE**

The **Old Name** value is the official (base) name for the AS. You cannot edit this value.

The AS Name is edited.

### **Manage an Application Mapping Rule**

Application mapping rules identify traffic in reports. You can manage the application mapping rules (type of Service (ToS), host, subnet, NBAR2) in NFA using NetOps Portal in the following ways:

- [Add an application mapping rule.](#)
- [Edit an application mapping rule.](#)
- [Delete an application mapping rule.](#)

## **Add an Application Mapping Rule**

Follow these steps:

1. From the page for the NFA data source, select **Application Definitions**, **Application Mapping** tab to view the application mapping rules.
2. Click **Add rule**.  
The **Add Rule** dialog opens.
3. Complete the following fields:
  - **Name**  
The name of the application mapping rule.
  - **Description**  
The description of the application mapping rule.
  - **Type of Rule**  
The type of application mapping rule.

**Options:**

- **ToS Mapping Rule:** Identifies traffic by its Type of Service (ToS) value.
- **Host Mapping Rule:** Identifies traffic based on its source host.

### **NOTE**

You can enter both IPv6 and IPv4 addresses in the **Host** field in **Host Mapping Rule**.

- **Subnet Mapping Rule:** Identifies traffic originating from a particular subnet mask.
  - **NBAR2 Classification Mapping Rule:** Identifies NBAR2 application traffic.
4. Add the application rule details.  
For more information about the rule details (settings) for each type of application rule, see the [DX NetOps Network Flow Analysis documentation](#).
  5. Test the entry for the destination port by clicking **Test**.  
A valid **Destination Port** value shows the following message at the top of the **Add Rule** dialog:  
`Traffic has been seen on this protocol before.`
  6. Save the application rule by clicking **Save**.

The application mapping rule is created.

## **Edit an Application Mapping Rule**

Follow these steps:

1. From the **Application Definitions Mapping** pane, select the row for the application mapping rule that you want to edit, and then click **Edit**.  
The **Edit Rule** dialog opens.
2. Change the rule details (settings) as needed, and then click **Save**.  
For more information about the rule details (settings) for each type of application rule, see the [DX NetOps Network Flow Analysis documentation](#).

## **Delete an Application Mapping Rule**

Follow these steps:

1. From the **Application Definitions Mapping** pane, select the row for the application mapping rule that you want to delete, and then click **Delete**.  
The **Delete** dialog opens.
2. Confirm the deletion by clicking **Yes**.

## Manage Reserved Seating

You can create Reserved Seating rules to help ensure that reports include the port and protocol combinations that interest you, regardless of traffic volume or rates. The rules create 'reserved seats' for the ports that are used by those protocols so the data is sure to be included in reports.

For example, during an application rollout you want to watch the traffic for a particular application, but the Top N Protocols reports for interfaces do not show the traffic for the application. The protocol that the application uses is not included in the Top N Protocol group; the group of protocols with the highest traffic volume or utilization rate. You create a Reserved Seating rule to collect data for the specific protocol and port that the application uses. The protocol now is included in the Top N Protocols reports.

### Follow these steps:

1. From the page for the NFA data source, select **Application Definitions, Reserved Settings** tab.
2. Select one of the following actions:
  - [Add a Reserved Settings Rule.](#)
  - [Edit a Reserved Settings Rule.](#)
  - [Delete a Reserved Settings Rule.](#)

## Add a Reserved Settings Rule

### Follow these steps:

1. Click **Add Rule**.
2. Specify ports as follows:
  - **Protocol:** Protocol of the data that is affected by the rule, either TCP or UDP
  - **Port:** Target port for the Reserved Seating rule. Enter the port number in the **Port** box, a value from 0 through 65535, expressed in Base 10 decimal format. If you do not enter a value, port 0 is assigned. The port and protocol combination must be unique; that is, it cannot match any other Reserved Seating rule. Data of the specified protocol type is reported for this port regardless of traffic rate or volume.
  - **Description:** (*Optional*) Identifying text for the Reserved Seating rule. The description appears in the list of **Reserved Seating** rules on the **Application Definitions** page.
3. Click **Save**.  
If you entered a valid port and protocol combination and you have not yet reached the maximum number of rules, the dialog closes. The new rule appears in the list of **Reserved Seating** rules.
4. Repeat this process for each Reserved Seating rule you want to add.

## Edit a Reserved Settings Rule

### Follow these steps:

1. Click the check box next to the rule you want to edit, then click **Edit**.  
The **Edit Reserved Seating** dialog opens.
2. Make the required changes, and then click **Save**.  
If the changes are successful, the **Edit Reserved Seating** dialog closes.

## Delete a Reserved Settings Rule

### Follow these steps:

1. Click the check box next to one or more rules you want to delete. To select the checkboxes for all the rules, click the check box in the heading row.
2. Click **Delete**.  
A confirmation message opens.



3. Click **Yes**.  
The list of **Reserved Seating** rules are updated to reflect your deletions.

### **Manage Port Priorities**

By default, NFA defines the server port and protocol as the lower number in the flow record.

**Source port:** 80

**Destination port:** 8000

In this case, NFA determines that the lower port is 80, and therefore http.

When the server port (TCP/UDP) is a high number, the port priority can be used.

For example:

**Server port:** 8888

**Client port:** 6000

By default, NFA uses the lower port (6000) as the server port. The Port Priority functionality tells it to use 8888 as the server port when it is present in the data.

Create port priority rules to ensure that the correct protocols are identified for each range of ports.

#### **Follow these steps:**

1. From the page for the NFA data source, select **Application Definitions, Port Priority** tab.  
The **Port Priority** page displays a list of the port priority rules.
2. Select one of the following actions:
  - [Add a Port Priority Rule](#).
  - [Edit a Port Priority Rule](#).
  - [Delete a Port Priority Rule](#).

### **Add a Port Priority Rule**

Follow these steps:

1. Click **Add Rule**.  
The **Add Port Priority** dialog opens.
2. Specify ports as follows:
  - **Protocol:** Protocol of the data that is affected by the rule, either **TCP** or **UDP**
  - **Start Port:** Target starting port for the Port Priority rule. Enter the port number in the **Start Port** box, a value from 0 through 65535, expressed in Base 10 decimal format. If you do not enter a value, port 0 is assigned.
  - **End Port:** Target ending port for the Port Priority rule. Enter the port number in the **End Port** box, a value from 0 through 65535, expressed in Base 10 decimal format. If you do not enter a value, port 0 is assigned.

**NOTE**  
The start port, end port, and protocol combination must be unique; that is, it cannot match any other **Port Priority** rule.

  - **Description:** (*Optional*) Identifying text for the Port Priority rule. The description appears in the list of **Port Priority** rules on the **Application Definitions** page.
3. Click **Save**.  
If you entered a valid start port, end port, and protocol combination and you have not yet reached the maximum number of rules, the dialog closes. The new rule appears in the list of Port Priority rules.
4. Repeat this process for each Port Priority rule you want to add.  
You can specify a maximum of 50 rules.

## Edit a Port Priority Rule

Follow these steps:

1. Click the check box next to the rule you want to edit, then click **Edit**.  
The **Edit Port Priority Rule** dialog opens.
2. Make the needed changes.
3. Click **Save**.  
If the changes are successful, the **Edit Port Priority** dialog closes.

## Delete a Port Priority Rule

Follow these steps:

1. Click the check box next to one or more rules you want to delete. To select the checkboxes for all the rules, click the check box in the heading row.
2. Click **Delete**.  
A confirmation message opens.
3. Click **Yes**.  
The list of **Port Priority** rules are updated to reflect your deletions.

## Data Extraction

You can export metric data from DX NetOps Performance Management using the following methods:

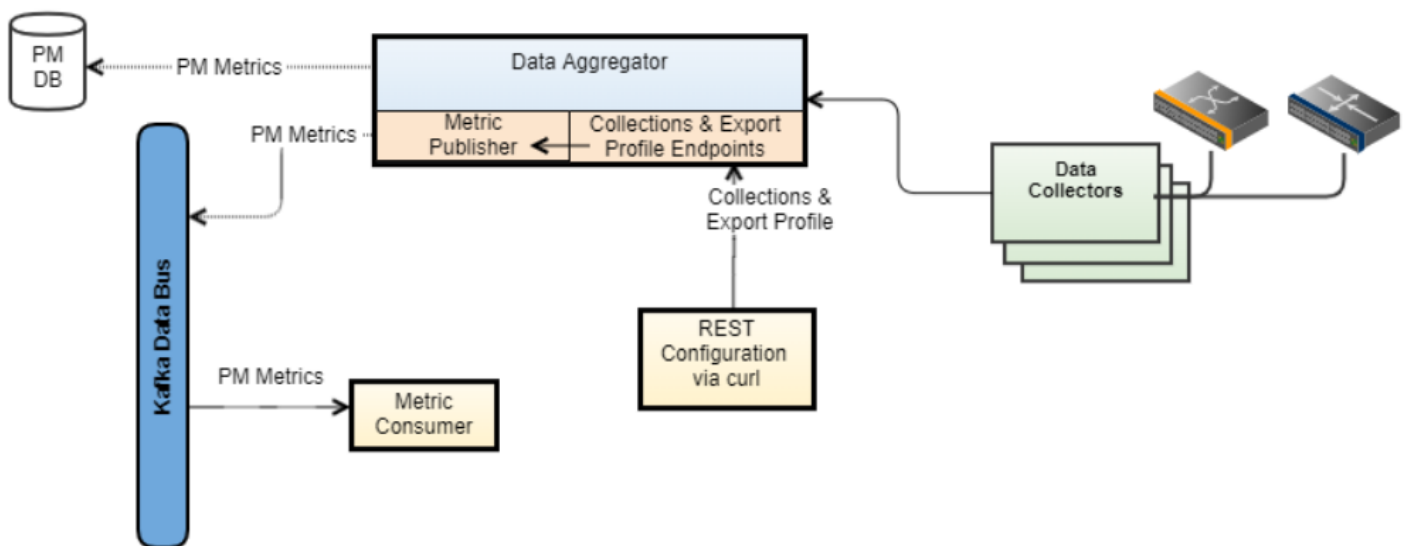
- (Preferred) [Configure Streaming Metric Export to a Kafka Cluster](#)
- [Configure Bulk Data Export as CSV Files](#)
- [Use Custom OpenAPI Query URLs](#)

This method is useful for when you want to export specific data, or when you want to run a one-time export.

## Configure Streaming Metric Export to a Kafka Cluster

Enable the communication between DX NetOps Performance Management and DX Operational Intelligence (metric export) by configuring the data aggregator to send, or export, collected performance metrics to Apache Kafka.

The following diagram shows streaming metric export to an Apache Kafka cluster:



The data collectors poll devices for desired metrics and forward these to the data aggregator. The data aggregator determines which devices the metric publisher should export metrics using device collections. The export profile defines which metrics the metric publisher in the data aggregator exports for those devices.

The metric publisher taps into the polled data that the data collectors send to the data aggregator in parallel with the normal polled data loading into the data repository. If the metric publisher encounters polled data that matches the configured export profile, it exports those performance metrics to the Kafka data bus for consumption by one or metric consumers.

#### **NOTE**

The metric publisher exports only non-null raw (as-polled) metrics.

Use the following process to configure streaming metric export to a Kafka cluster:

1. [Verify the Prerequisites](#)
2. [Enable and Configure the Kafka Export Producer on the Data Aggregator](#)

#### **IMPORTANT**

In a fault-tolerant data aggregator environment, complete this procedure on each data aggregator.

3. [Define the Devices and Metrics to Export](#)

#### **NOTE**

If you are setting up metric export for the OI Connector to use, the OI Connector already automatically defines the metrics to export, and you can skip this step.

### **Verify the Prerequisites**

Before you can configure streaming metric export, verify the following items:

- The data aggregator is installed.
- Kafka is installed on a server or cluster of servers that is accessible to the data aggregator, and the data aggregator has permission to access the broker ports on the Kafka servers (by default, port 9092).

#### **TIP**

If you do not already have a Kafka deployment, you can use NetOps Kafka, which is provided as a separate download.

For more information about how to install NetOps Kafka, see [Install NetOps Kafka](#).

- You have configured NetOps Portal to collect performance metrics for a given set of devices. You have created a monitoring profile, and applied the monitoring profile polling behavior to device collections.

For more information, see [Manage Monitoring Profiles](#).

### **Enable and Configure a Kafka Export Producer on the Data Aggregator**

You enable and configure a Kafka export producer on the data aggregator by configuring the Kafka producer config file. When configured, the data aggregator harvests and publishes DX NetOps Performance Management metrics to the Kafka data bus using the Kafka export producer. (For communication between DX NetOps Performance Management and DX Operational Intelligence) The OI Connector consumes the metrics from the Kafka data bus, and forwards these metrics to DX Operational Intelligence.

#### **IMPORTANT**

In a fault-tolerant data aggregator environment, complete this procedure on each data aggregator.

#### **NOTE**

The properties in the Kafka producer config file are prepended with `producer.` . Limit these properties to what is needed for SSL/TLS authentication.

For more information, see [the Kafka documentation](#).

For more information about Kafka configuration information, see [the "Kafka Configuration Information" section](#).

**Follow these steps:**

1. Edit the `<installation_directory>/apache-karaf/etc/kafkaexport.producer.cfg` file.

**Example:**

```
/opt/IMDataAggregator/apache-karaf/etc/kafkaexport.producer.cfg
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

2. Set the following properties in the file:

- `producer.bootstrap.servers`

Specifies the list of host/port pairs for the data aggregator to use to establish the initial connection to the Kafka cluster. Set the hostname (`changeme`) to the hostname of a Kafka broker, and the port (`9092`) to the broker port. For example, `producer.bootstrap.servers=kafkaservername:9092`.

**NOTE**

For Kafka clusters, you can provide a comma-separated list of `hostname:port` values.

**Default:** `producer.bootstrap.servers=changeme:9092`

- `feature.enabled`

Specifies whether a Kafka export producer on the data aggregator is enabled. Set this property to `on`.

**Options:**

- **off:** A Kafka export producer on the data aggregator is disabled. For example, `feature.enabled=off`.
- **on:** A Kafka export producer on the data aggregator is enabled. For example, `feature.enabled=on`.

**Default:** `feature.enabled=off`

- `topic`

Specifies the `metric-export` topic name in the metric publisher, which is the topic to which the metric published exports metrics.

**Default:** `topic=metric-export`

3. (If you have secured Kafka) Provide a keystore/truststore that includes a trusted certificate that can be used to access Kafka by adding the following properties to the file:

**Prerequisite:** The truststore and keystore are in place.

```
producer.security.protocol=SSL
producer.ssl.keystore.location=<path_to_keystore>
producer.ssl.keystore.password=<keystore_password>
producer.ssl.key.password=<kafka_broker_key_password>
producer.ssl.truststore.location=<path_to_truststore>
producer.ssl.truststore.password=<truststore_password>
```

**Example:**

```
producer.security.protocol=SSL
producer.ssl.keystore.location=/opt/IMDataAggregator/apache-karaf/etc/
kafka01.keystore.jks
producer.ssl.keystore.password=password
producer.ssl.key.password=password
producer.ssl.truststore.location=/opt/IMDataAggregator/apache-karaf/etc/
kafka.truststore.jks
producer.ssl.truststore.password=password
```

- **path\_to\_keystore**

The absolute path to the keystore file.

- **keystore\_password**

The password for the keystore file.

– ***kafka\_broker\_key\_password***

The password for the Kafka broker's key in the keystore. You can make this password different than the keystore password if you choose.

– ***path\_to\_truststore***

The location of the truststore file.

– ***truststore\_password***

The password to access the truststore.

4. Save your changes.

**IMPORTANT**

(Fault-tolerant data aggregator environments) If you have not yet enabled and configured a Kafka export producer on each of the data aggregators, repeat this procedure for each data aggregator.

The Kafka export producer is enabled and configured on the data aggregator.

## **Define the Devices and Metrics to Export**

The data aggregator determines which devices the metric publisher should export metrics using device collections. The export profile defines which metrics the metric publisher exports for those devices.

**NOTE**

- If you are setting up metric export for the OI Connector to use, the OI Connector already automatically defines the metrics that the metric publisher should export, and you can skip the following procedures.
- The following procedures send cURL commands using a username and prompt for a password. You can send credentials in the cURL commands using the alternate methods.  
For more information about these methods, see [NetOps Portal REST Web Services](#).

Use the following process to define the devices and metrics that the metric publisher should export:

1. [Create the Export Profile](#)
2. (If required) [Update the Export Profile](#)
3. [Identify the Tenants that Require Metric Export](#)
4. [Identify the Device Collections that Require Metric Export](#)
5. [Associate the Device Collections with the Export Profile](#)

## **Create the Export Profile**

**IMPORTANT**

Create only one export profile. If multiple export profiles exist, the data aggregator honors the profile with the lowest item ID, and ignores the remaining export profiles.

**Follow these steps:**

1. Issue the following cURL, with the metric families (and filtered metrics if desired) to be exported:

```
curl --user <NetOpsPortalUser> --location --request POST 'http://<DA_hostname>:8581/
rest/exportprofiles' \
--header 'Content-Type: application/xml' \
--data '      <ExportProfile version="1.1.0">
<!-- The list of metric families whose metrics are to be exported. If any metrics
      from a metric family are to be exported, the metric family must be listed here. -->

      <ExportedMetricFamilyList>
```

```

<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedReachabilityInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedCPUInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedPortInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedMemoryInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedAvailabilityInfo</
ExportedMetricFamily>
</ExportedMetricFamilyList>
<!-- This list filters the specific metrics to be exported if only a subset of
metrics should be exported for a given metric family. If any metrics are specified
here, then only those metrics are to be exported for the associated metric family.
If no metrics for a metric family listed in the ExportedMetricFamily attribute are
listed here then all metrics for that metric family are to be exported. -->
<ExportedMetricsList>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.Availability</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.UtilizationIn</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.UtilizationOut</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsIn</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsPerSecondIn</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsOut</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsPerSecondOut</
ExportedMetrics>
</ExportedMetricsList>
    <Item version="1.0.0">
        <Name><Export Profile Name></Name>
        <Description>Export Profile.</Description>
    </Item>
</ExportProfile>'

```

**Example:**

```

curl --user JaneDoe --location --request POST 'http://myDAHost:8581/rest/
exportprofiles' \
--header 'Content-Type: application/xml' \
--data '    <ExportProfile version="1.1.0">
<!-- The list of metric families whose metrics are to be exported. If any metrics
from a metric family are to be exported, the metric family must be listed here. -->

<ExportedMetricFamilyList>

```

```

<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedReachabilityInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedCPUInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedPortInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedMemoryInfo</
ExportedMetricFamily>
<ExportedMetricFamily>{http://im.ca.com/normalizer}NormalizedAvailabilityInfo</
ExportedMetricFamily>
</ExportedMetricFamilyList>
<!-- This list filters the specific metrics to be exported if only a subset of
metrics should be exported for a given metric family. If any metrics are specified
here, then only those metrics are to be exported for the associated metric family.
If no metrics for a metric family listed in the ExportedMetricFamily attribute are
listed here then all metrics for that metric family are to be exported. -->
<ExportedMetricsList>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.Availability</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.UtilizationIn</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.UtilizationOut</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsIn</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsPerSecondIn</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsOut</
ExportedMetrics>
<ExportedMetrics>{http://im.ca.com/normalizer}NormalizedPortInfo.BitsPerSecondOut</
ExportedMetrics>
</ExportedMetricsList>
    <Item version="1.0.0">
        <Name>Export Profile Name</Name>
        <Description>Export Profile.</Description>
    </Item>
</ExportProfile>'

```

- **NetOpsPortalUser**  
The NetOps Portal username.
  - **DA\_hostname**  
The IP address for the data aggregator.  
**Example:** myDAHost
  - **Export Profile Name**  
The name of the export profile.
2. When prompted, provide a password.
  3. Note the item ID of the export profile in the response to the request.

**Sample output:**

```
<success>{Export Profile Item ID}</success>
```

The export profile is created.

### **Update the Export Profile**

If required, you can modify the values in the existing export profile.

#### **Follow these steps:**

1. Issue the following cURL, with the data aggregator server hostname, the NetOps Portal username, the export profile ID that needs to be updated and any settings that need to be modified (example below modifies the profile name to New Export Profile Name):

```
curl --user <NetOpsPortalUser> --location --request PUT 'http://<DA_hostname>:8581/
rest/exportprofiles/<ExportProfileID>' \
--header 'Content-Type: application/xml' \
--data ' <ExportProfile version="1.1.0"> <Item version="1.0.0"> <Name><New Export
Profile Name></Name> </Item> </ExportProfile>'
```

#### **Example:**

```
curl --user JaneDoe --location --request PUT 'http://myDAHost:8581/rest/
exportprofiles/<ExportProfileID>' \
--header 'Content-Type: application/xml' \
--data ' <ExportProfile version="1.1.0"> <Item version="1.0.0"> <Name>New Export
Profile Name></Name> </Item> </ExportProfile>'
```

#### **– NetOpsPortalUser**

The NetOps Portal username.

#### **– DA\_hostname**

The IP address for the data aggregator.

**Example:** myDAHost

#### **– New Export Profile Name**

The updated name for the export profile.

2. When prompted, provide a password.

The export profile is updated.

### **Identify the Tenants that Require Metric Export**

#### **Follow these steps:**

1. List the details of the available tenants in the data aggregator by issuing the following cURL command:

```
curl --user <NetOpsPortalUser> --request GET 'http://<DA_hostname>:8581/rest/tenants' \
--header 'Accept: application/xml'
```

#### **Example:**

```
curl --user JaneDoe --request GET 'http://myDAHost:8581/rest/tenants' \
--header 'Accept: application/xml'
```

#### **– NetOpsPortalUser**

The NetOps Portal username.

#### **– DA\_hostname**

The IP address for the data aggregator.

**Example:** myDAHost



2. When prompted, provide a password.
3. Locate the tenants that you want to configure metric export by name, and then note the item IDs (the <ID> ) for each tenant.

### **Identify the Device Collections that Require Metric Export**

#### **Follow these steps:**

1. List the details of the available device collections in the data aggregator by issuing the following cURL command:

```
curl --user <NetOpsPortalUser> --request GET 'http://<DA_hostname>:8581/rest/monitoredgroups' \
--header 'Accept: application/xml'
```

#### **Example:**

```
curl --user JaneDoe --request GET 'http://myDAHost:8581/rest/monitoredgroups' \
--header 'Accept: application/xml'
```

- **NetOpsPortalUser**  
The NetOps Portal username.
- **DA\_hostname**  
The IP address for the data aggregator.

**Example:** myDAHost

2. When prompted, provide a password.
3. Locate the device collections that you want to configure metric export by name, and then note item IDs (the <ID> ) for each device collection.

### **Associate the Device Collections with the Export Profile**

Associate the device collections in the data aggregator with the export profile on a per-tenant basis.

#### **NOTE**

Complete this procedure for each tenant/collection pair that you want to associate with the export profile.

#### **Follow these steps:**

1. Issue the following cURL command:

```
curl --user <NetOpsPortalUser> --location --request PUT 'http://<DA_hostname>:8581/rest/tenant/<TenantItemID>/exportprofiles/<ExportProfileItemId>/relatesto/groups/<CollectionItemId>' \
--header 'Accept: application/xml' \
--header 'Content-Type: application/xml' \
--data ''
```

#### **Example:**

```
curl --user JaneDoe --location --request PUT 'http://myDAHost:8581/rest/tenant/132614/exportprofiles/1234567/relatesto/groups/56789' \
--header 'Accept: application/xml' \
--header 'Content-Type: application/xml' \
--data ''
```

- **NetOpsPortalUser**  
The NetOps Portal username.
- **DA\_hostname**  
The IP address for the data aggregator.

**Example:** myDAHost

– **TenantItemID**

The item ID for the tenant.

**Example:** 132614

– **ExportProfileItemID**

The item ID for the export profile with which you want to associate the device collection.

**Example:** 1234567

– **CollectionItemID**

The item ID for the device collection that you want to associate with the export profile.

**Example:** 56789

2. When prompted, provide a password.

The tenant/collection pair is associated with the export profile. The item ID for the export profile is recorded.

**NOTE**

If you have not yet issued the `cURL` command for each tenant/collection pair that you want to associate with the export profile, repeat this procedure for the other tenant/collection pairs.

### **Kafka Configuration Information**

The metric publisher exports metrics to the `metric-export` topic and heartbeat messages to the `metric-heartbeat` topic. It exports these metrics and messages using the default log retention, partition count, and replication settings that are configured in the Kafka broker.

Kafka export to a Kafka cluster requires that the `metric-export` and `metric-heartbeat` topics exist, and expects that Kafka will create them when the metric publisher exports metrics and messages to them if they do not exist. If the Kafka administrator has configured the Kafka broker to disallow automatic creation of topics when producers (such as the metric publisher) export metrics and messages to them, then the Kafka administrator must create these topics.

You can change the default log retention, partition count, and replication settings to suit your specific requirements using the Kafka utilities on the Kafka broker.

For more information, see [the Kafka documentation](#).

You can modify the `metric-export` topic name in the metric publisher by editing the name in the `topic` property in the `kafkaexport.producer.cfg` file.

**NOTE**

Long-term retention past an hour is not required for the `metric-heartbeat` topic.

## **Bulk Data Export**

You can export polled rate data from the data aggregator for analysis and custom reporting purposes.

As a systems administrator or architect, you can export this data as a continuous comma-separated value (CSV) export to your own reporting tool.

**NOTE**

You can bulk data export only in a non-fault-tolerant data aggregator environment.

**TIP**

You can also export metric data—for example, to export *specific* data, or to run a one-time export—using OpenAPI or by configuring streaming metric export to a Kafka cluster.

For more information, see [OpenAPI](#) and [Configure Streaming Metric Export to a Kafka Cluster](#).

You export data using the data aggregator. The data export is done at the frequency of the polling rate. When you start data export, the poll responses are written to CSV files. You can configure inclusive lists for metric families so that you only export the data you need.

The CSV files contain polled rate data and internal names, which are mapped to readable display names. If exported metric families include suppressed metrics, the CSV shows a null value for those metrics.

The size of the CSV output files that are generated can come close to or can exceed the following limits:

Environment	Limit per hour (uncompressed)
Large scale	50 GB
Medium scale	25 GB
Small scale	5 GB

#### TIP

To avoid exhausting the available disk space, regularly process the output CSV files. For example, copy the files to another system and remove the files from the data aggregator host.

For performance reasons, create a separate disk partition on the data aggregator system. Select this partition for the output of the CSV polled rate data.

Use the following process to export data from DX NetOps Performance Management:

1. [Configure the export file output options.](#)
2. (Optional) [Configure an inclusive list of metric families.](#)
3. (Optional) [Export extra columns for components.](#)
4. [Export the bulk data.](#)

If necessary, you can stop the export of bulk data. See [Stop bulk data export](#).

### Configure the Export File Output Options

Configure the output file options using the `streamexport.csvoutwriter.cfg` file. This configuration file is copied to the directory during the data aggregator installation.

#### Follow these steps:

1. Change to the `<installation_directory>/<apache-karaf-*>/etc` directory.
  - **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`
  - **apache-karaf-\***  
The installation directory for Apache Karaf.  
**Example:** (21.2.6 and higher) `apache-karaf-4.3.3`
2. Edit the `streamexport.csvoutwriter.cfg` file as needed:
  - **output.filenameExtension**  
Specifies the extension or suffix of the CSV files.
  - **output.csvFileDelimiter**  
Specifies the column delimiter that is used in the output CSV file. For example, if the bundle is started and this parameter is changed, a new file is written immediately using this new column delimiter.
  - **output.filenameLocationPath**  
Specifies the file path and the prefix of the output file name. (The output file also consists of the date and time.)

**IMPORTANT**

This file path must be on a different partition than the data aggregator installation. If DX NetOps Performance Management was installed with a sudo user, change the ownership of the directory to enable access to the sudo user.

The syntax of the complete file name is:

```
output.<filenameLocationPath>=<DC_host>_yyyy-MM-dd-
Thh-mm-sec-ms.output.<filenameextension>
```

**Example:**

The file is written on April 2, 2013 at 8:42:04 a.m. and 123 ms. The data collector hostname is server.abc.com, and the following parameters are configured:

```
output.filenameLocationPath=/opt/export_data/mydata
output.filenameExtension=.csv
```

The file name is:

```
mydata_server.abc.com_2013-04-02T08-42-04-123.csv
```

You can configure files with an absolute name, such as /myOutputDir/mydata . If the parent folders in the absolute path do not exist, the folders are created.

- **output.filesize**

Specifies the file size in bytes using a valid integer greater than 0. If the file size is exceeded, a new output file gets written.

If the value is -1, this parameter is ignored and *all* the data is written into a single file (infinite).

- **output.duration**

Specifies the number of minutes using a valid integer greater than 1. If a file is older than x minutes, then a new file is written.

If the value is 1, then this parameter is ignored. If the value is -1, then *all* the data is written into a single file (infinite).

The output.filesize and output.duration parameters affect each other. If output.filesize is exceeded or the file is older than output.duration , then a new file is written.

**Example:** streamexport.csvoutwriter.cfg

```
output.filesize=1000000000
output.filenameLocationPath=/opt/data_export/ratedata_
feature.enabled=on
output.duration=60
output.csvFileDelimiter=,
output.filenameExtension=.csv
```

### 3. Save your changes.

#### **(Optional) Configure an Inclusive List of Metric Families**

You can specify the metric families whose data you do want to export by configuring an inclusive list (whitelist). Data is collected only for the listed metric families.

The following files pertain to scoping by metric family:

- streamexport.allMetricFamilies.out

This file is auto-generated at the start-up of the data aggregator. The file includes the available metric families. The data aggregator periodically updates this file to include new metric families.

**Location:** <installation\_directory>/<apache-karaf-\*>/etc directory

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **apache-karaf-\***

The installation directory for Karaf.

**Example:** (21.2.6 and higher) apache-karaf-4.3.3

This file has the following format:

```
metricFamilyInternalName=metricFamilyDisplayNameinEnglish
```

**Example:**

```
#Metric Family Name List for Customer Reference
#Mon Jun 17 11:18:50 EDT 2013
normalizedmemoryinfo=Memory
normalizedavailabilityinfo=Availability
normalizedcpuinfo=CPU
normalizedportinfo=Interface
```

- streamexport.metricFamilyWhiteList.cfg

Specifies the metric families to export.

**Location:** <installation\_directory>/<apache-karaf-\*>/etc directory.

– **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

– **apache-karaf-\***

The installation directory for Karaf.

**Example:** (21.2.6 and higher) apache-karaf-4.3.3

This file has the following format:

```
whitelist.number=metricFamilyInternalName
```

**Example:**

```
#This is the whitelist file for the metric families.
#Mon Jun 10 17:25:46 EDT 2013
feature.enabled=on
whitelist.1 = NormalizedMemoryInfo
whitelist.2 = NormalizedPortInfo
```

**Follow these steps:**

1. On the data aggregator, using a text editor, open the following files:
  - streamexport.allMetricFamilies.out
  - streamexport.metricFamilyWhiteList.cfg
2. Copy one or more metric family internal names (metricFamilyInternalName ) from the streamexport.allMetricFamilies.out file, and the paste it into the streamexport.metricFamilyWhiteList.cfg file.

**Example:**

```
whitelist.1=<paste here>
whitelist.2=<paste here next metric family name>
Set feature.enabled=on
```

3. Save the changes to the streamexport.metricFamilyWhiteList.cfg file.

**(Optional) Export Extra Columns for Components**

You can include more information about component items by configuring the export to include the ifAlias and ifDescr columns.

**IMPORTANT**

If you configure this option *after* you start the data export, the output file might contain mixed data. Lines that are exported before you add the columns do not include the extra columns. Lines that are exported after the change include the columns. To avoid a mixed file, stop the data aggregator before you change the configuration.

**Follow these steps:**

1. Edit the `etc/streamexport.exportInfoResolver.cfg` file.
2. Set the `enableInterfaceOutput=true` parameter.
3. Save the file.

The data aggregator adds the `ifAlias` and `ifDescr` columns to the output file. This change occurs immediately.

**Export the Bulk Data**

You can export polled rate data using any REST client with a connection to the data aggregator server or an HTTP tool that can send requests and can get responses. This example uses a REST client.

**Follow these steps:**

1. Set the REST Content-type to `application/xml`.
2. Retrieve a list of the export profile IDs by issuing a GET request to the `dataexport` Data Aggregator REST web service:  
`http://da_host:8581/rest/dataexport/`
3. Take note of the ID of the export profile that you want to modify.

**NOTE**

By default, there is only one export profile.

4. Start bulk data export by setting the `Enabled` attribute to **true**. Issue a PUT request to the `dataexport` Data Aggregator REST web service, with the following body:

`http://da_host:8581/rest/dataexport/<id>`

**– id**

Specifies the ID of export profile.

**Body**

```
<DataExportInfo version="1.0.0">
<Enabled>true</Enabled>
</DataExportInfo>
```

**TIP**

You can review the other attributes that you can set at the following URL:

`http://da_host:8581/rest/dataexport/xsd/get.xsd`

5. Verify that your changes took effect by issuing a GET request to the following URL:

`http://da_host:8581/rest/dataexport/<id>`

The `dataexport` Data Aggregator REST web service starts automatically, and the temporary export file is created. When the export file is ready, the exporter automatically renames it to the previously configured file extension, such as `.csv`. You do not need to restart the services for a newly written file.

**Next step:** Copy the data to another system using the method required by that system.

**Stop Bulk Data Export**

Stop the export of polled rate data by setting the `Enabled` attribute to **false**. Issue a PUT request to the `dataexport` Data Aggregator REST web service, with the following body:

`http://da_host:8581/rest/dataexport/<id>`

- **id**  
Specifies the ID of export profile.

### Body

```
<DataExportInfo version="1.0.0">
<Enabled>false</Enabled>
</DataExportInfo>
```

## View System Status

You can view the overall status of your system, such as the status of the data aggregator, the data collectors, and the data repository.

You can view the overall status for the following from the **System Status** page:

### NOTE

To view the **System Status** page, log in as an Administrator, hover over **Administration**, **Data Sources**, and then click **System Status**.

- **Data Source Synchronization**  
Displays the overall system status of your data sources. If you are experiencing data source synchronization issues, determine the cause of the issue.  
For more information, see [Synchronize Data Sources](#).
- **Global Synchronization**  
Displays the global synchronization system status.
- **Data Aggregator**  
Displays the overall system status of the data aggregators.
- **Data Collector**  
View the overall system status of your data collectors. The data collector list shows the following information for the data collectors:
  - **Configured**  
Displays whether the data collector is configured as an active or standby (running) data collector host.  
**Values:** Active, Standby  
For more information about how to configure data collectors for fault tolerance, see [Configure Data Collectors for Fault Tolerance](#).
  - **Status**  
Displays the status of the data collector. On the **Data Collectors** page, this is the **Overall Status**.  
**Values:** Connected (Normal), Not Connected, Upgrade Pending
  - **IP Domain**  
Displays the tenant and IP domain to which each data collector installation is assigned.
  - **Version**  
Displays the data collector version.
  - **Polling Status**  
Displays the polling status for each data collector installation.  
**Values:** Collecting Data, Standby

The Administrator can see a list of data collector installations for all tenants. Tenant administrators can see only the data collector installations that are assigned to their tenant.
- **Data Repository**  
Displays data repository node status. NetOps Portal polls Vertica node status every 60 seconds by way of the data aggregator. If the data aggregator is down, the data aggregator overall system status shows as "Failed".

For more information about how to address common issues related to the data repository, such as startup issues or Vertica nodes in the cluster being down, see [Restart the Data Repository](#).

- **VNA Gateway**

Displays the overall system status of the DX NetOps Virtual Network Assurance (VNA) Gateways.

- **Scheduled Report Repository**

Displays the overall system status of your scheduled report repository.

**IMPORTANT**

If you have chosen to have NetOps Portal save dashboard data as scheduled reports (archive reports is enabled), closely monitor disk space over time. You can also have Broadcom Support relocate the report repository.

For more information about how to enable archive reports, see [Configure the Email Server](#).

- **Report Generation Services**

Displays the overall system status of your report generation services. You can also view the percentage of that last 12 hours that each operation status occurred.

**Operational statuses:**

- **Low Load:** 0-1 reports are waiting to run.
- **Medium Load:** 2-5 reports are waiting to run.
- **High Load:** 6 or more reports are waiting to run.

- **Alarm Status Service**

Displays the overall system status of your Alarm Status service.

- **Report Manager Service** (23.3.2 and higher)

If you have enabled the NetOps Report Manager Service, this section displays the overall system status of the NetOps Report Manager Service. From this section, you can view the NetOps Report Manager Service hostname, check to see if the NetOps Report Manager Service is up, view the name of Spectrum's MySQL "reporting" database (SRMDB), and check to see if SRMDB is up.

For more information about how to enable the NetOps Report Manager Service, see [Set up to Run NetOps Business Reports](#).

- **Event Manager Publication Service**

If you have configured the Event Manager service (`caperfcenter_eventmanager`) to publish events to a Kafka topic, this section displays the status of the publication. The **Latency (secs)** and **Latency Trend** metrics provide insight into whether the event publisher is falling behind, steady, or catching up. Latency is a measure of how long it is taking the event to make it to the Kafka topic. Latency trend defines whether the latency is steady, falling behind, or improving.

The following image shows this section:

**Figure 136: Event Manager Publication Service**

Event Manager Publication Service <span>✓ Normal</span>				
Quick Filter <span>▼</span>				
Source Name	Status	Description	Last Status Change	Last Polled On
Event Manager	<span>✓ Up</span>		Jan 3, 2023, 9:07:52 PM E...	Jan 6, 2023, 3:24:52 PM E...

- **Network Flow Analysis**

If you have integrated with DX NetOps Network Flow Analysis (NFA), this section displays the overall system status of your NFA components.

For more information about this integration, see [Integrate with DX NetOps Network Flow Analysis](#).

- **DX Operational Intelligence**



If you have integrated DX NetOps Performance Management and NFA with DX Operational Intelligence, this section displays the system status of the integration.

For more information about this integration, see [Integrate with DX Operational Intelligence](#).

- **Flow Services**

If Ingress is configured and NetOps Flow is deployed and configured with the Ingress information, this section displays NetOps Flow-related status.

For more information about how to deploy NetOps Flow, see [Deploy NetOps Flow](#).

## View Health Monitoring Information

A built-in mechanism monitors the health of the data aggregator and the data collector devices. Self-monitoring monitoring profiles determine the statistics that are discovered and polled for these devices. The discovered statistics are collected automatically.

### IMPORTANT

Do not change or stop this self-monitoring.

As an administrator, you want to monitor your data aggregator and data collector items so that you can manage their performances proactively and you can perform capacity planning.

In this scenario, you will view the monitoring profiles that are associated with data aggregator and data collector self-monitoring. You will also view the components that are being monitored on the data aggregator device. This information helps you to understand how the health of these items is managed.

You will also create dynamic trend views, where you can view changes that are occurring on data aggregator over time. This information is useful when you are troubleshooting performance issues and performing capacity planning.

In this article:

- [View Metric Families That Are Associated with Self-Monitoring](#)
- [View Monitored Components on the Data Aggregator](#)
- [Create a Group](#)
- [Create a Dashboard and View Relevant Information](#)

### View Metric Families That Are Associated with Self-Monitoring

You can view the metric families that are related to the self-monitoring of the data aggregator and the data collector. You can also view the metrics within these metric families, which are used for polling the devices. This information helps you to gain an understanding of what types of data is collected and at what rates the data is collected. You can then understand how the health of these devices is managed. You will see this information in the view that you will create later in this scenario.

#### Follow these steps:

1. Log in to NetOps Portal as an administrator.
2. Hover over **Administration, Data Sources**, and then click a data aggregator data source.  
The data aggregator admin pages open.
3. Click **Monitoring Profiles** from the **Monitoring Configuration** menu.  
The **Monitoring Profiles** page appears.
4. Select each DA Health (<Rate>) monitoring profile individually, and then view the corresponding metric families on the **Metric Families** tab.

#### NOTE

To view a description for each metric family, hover over it.

5. Make a note of the metric families.
6. Click **Metric Families** from the **Monitoring Configuration** menu.  
The **Metric Families** page appears.

- Find and select each metric family individually, and then view the specific metrics that are collected on the **Metrics** tab.

### **View Monitored Components on the Data Aggregator**

You can view which components are being monitored on the data aggregator. You can also view the polling status on the components. For example, for capacity planning purposes, you can see all of the data collector instances that are being monitored on the data aggregator instance.

#### **NOTE**

Only data collector instances for the tenant that you are administering are displayed. In an enterprise environment, you will likely only use the default tenant workspace, and therefore will see all of the data collector instances.

#### **Follow these steps:**

- Click **Monitored Devices** from the **Monitored Inventory** menu for a data aggregator data source.
- Expand the **All Data Aggregators** folder, and then select the data aggregator device.  
The **Polled Metric Families** tab shows the metric families that are associated with the data aggregator device. The **Components** table for a given metric family shows the polling status on the discovered components.

#### **NOTE**

When new components are discovered on the self-monitoring metric family, the components are monitored automatically. Performing an "Update Metric Family" operation on a self-monitoring metric family will result in no change in the state of monitoring.

### **Create a Group**

As an administrator, you can create a custom group to organize managed devices in NetOps Portal. In this scenario, you create a group for the data aggregator device. The Data Aggregator device becomes the context for the view you will create in a later step.

#### **Follow these steps:**

- Log in to NetOps Portal as a user with the Administrator role.
- Hover over **Administration**, **Group Settings**, and then click **Groups**.  
The **Manage Groups** dialog opens.
- To find a location for the new group, expand nodes in the Groups tree. Create the group under the All Groups node in the Groups tree, or within an existing custom or site group. You cannot add groups to system groups, which appear "locked" in the Groups tree.

#### **NOTE**

Create this group under a node where only you have access rights.

- Right-click the node, and select **Add Group**.  
The **Add Group** window opens. The **New** tab is selected by default.
- Enter "Data Aggregators" as the name of the group and optionally provide a description.
- Select **Custom** from the **Group Type** list, and then click **Save**.  
The new group appears in the **Groups** tree.
- Select the **Data Aggregators** group that you created.
- Click the **Items** tab in the right pane.
- Click **Add Item Type**.

The **Add Items** dialog opens.

The list of items refreshes to show items of the selected type that are available to add to the group. The available items depend on the item type, the data sources that are registered, and the items discovered.

#### **NOTE**

To see more pages of items, click the links below the list. Or use the Search field to search for an item in the list.

10. Select the `DataAggregator@ip_address` item, and then click **Add Items**.

– **ip\_address**

Indicates the IP address of the data aggregator device.

11. Click **Close**.

The **Add Items** dialog closes. The **Items** tab shows the `DataAggregator@ip_address` device that you added. Synchronization with data aggregator can take up to 5 minutes to begin.

### **Create a Dashboard and View Relevant Information**

Custom dashboards are useful for displaying data from a particular item or group of items. (*Item* can be a device, component, or interface.) In this scenario, you want to display self-monitoring data on the data aggregator item. In particular, you want to look at the Polled Item Count metric, the SNMP Poll Failure Count metric, and the SNMP Request Count metric. By looking at this information, you can view any changes that are occurring over time. Use this information when troubleshooting performance issues and when you are performing capacity planning.

#### **Follow these steps:**

1. From any dashboard, click the **More Options** icon, and then click **Add Dashboard**.  
The **Add Dashboard** page opens.
2. Complete the following fields:
  - a. Enter "Data Aggregator Health" in the **Menu Item** and **Dashboard Title** fields.
  - b. (Optional) Select a layout template for the dashboard.  
Each layout treats the page as a table with rows and columns for views. The layout buttons indicate the number of views in each column and row on the page.

#### **NOTE**

Select your layout before you add the Dynamic Trend View. If you change the layout after you add a view, the views will be removed.

3. Expand **Dynamic Views** in the left pane. Select and drag the Dynamic Trend View to the page layout, and drop the view where you want it to appear.
4. To create two more Dynamic Trend Views, select the view, and then click **Copy** twice.
5. Add the **Polled Item Count** metric to the first view. Add the **SNMP Poll Failure Count** metric to the second view. Add the **SNMP Request Count** metric to the third view. To add the metrics to the view, do the following steps:
  - a. Select the **Dynamic Trend View**, and then click **Edit**.
  - b. If it is not already selected, select **Device** from the **Context Type** drop-down box.
  - c. For the title, enter "Polled Item Count" for the first view, "SNMP Poll Failure Count" for the second view, and "SNMP Request Count" for the third view.
  - d. Keep the default, **Multitrend**, as the **View Type**.
  - e. Select the **DA Data Collector Polling Statistics** metric family for each view.

#### **NOTE**

It can take a few minutes for the available metric families to synchronize with NetOps Portal. If the self-monitor metric families are not immediately available, wait a few minutes and then try again.

- f. Select the **Polled Item Count - Average** metric value for the first dynamic trend view. Select the **SNMP Poll Failure Count - Total** metric value for the second view. Select the **SNMP Request Count - Total** metric value for the third view.
  - g. Select **Add Groups**, and then click **Add/Remove Groups**. Select the **Data Aggregators** group that you created previously from the list of available groups, and then move it to the list of selected groups. Click **OK**.
  - h. Click **Save**.  
Your report is generated automatically.
6. View any changes that are occurring over time.  
For example:

- Polled item count increases when you were *not* expecting it to - Did you unintentionally start monitoring more items?
- Polled item count remains the same, but you see spikes in the SNMP Request Count occasionally - Do you have change detection turned on in one of your monitoring profiles or perhaps scheduled discoveries?
- Polled item count remains the same, but you see changes in the SNMP Poll Failure Count - Perhaps there has been a network outage.
- Polled item count increases and the SNMP Poll Failure Count increases - Perhaps the data collector is overloaded and cannot respond to all of the requests.

The dashboard is created.

## Restart Performance Management Component Services

To restart DX NetOps Performance Management, restart the component services on the host servers.

For the procedures to restart the component services, see the topics in this section.

### Restart the Data Aggregator

Restart the Data Aggregator service when required.

During a planned shutdown, the data aggregator continues processing for five minutes before the service stops:

- Data that has been received continues processing and is sent to the data repository. Data that is not processed is saved, and loaded when the data aggregator restarts.
- Incomplete rollup processing resumes when the data aggregator restarts.

The data collectors queue poll data when the data aggregator is stopped. Data loading of queued data, threshold event processing, and rollup processing resume when the data aggregator restarts. If queued data exceeds the disk space limit on the data collector, the oldest data is discarded.

If the data aggregator stops ungracefully, polled data and threshold event processing information might be lost.

In this article:

- [Restart the Fault-Tolerant Data Aggregators](#)
- [Stop the Data Aggregator](#)
- [Start the Data Aggregator](#)

### Restart the Fault-Tolerant Data Aggregators

**Prerequisite:** You know which data aggregator is active ("Active" status) and which is ready to take over if the *active* data aggregator goes offline ("Ready" status).

For more information about how to check this, see [View System Status](#).

Use the following process to restart the fault-tolerant data aggregators:

1. Do the following steps for the data aggregator in "Ready" status:
  - a. [Put it into maintenance mode](#).
  - b. [Activate it](#).

This data aggregator is now available for failover.
2. Do the following steps for the data aggregator in "Active" status:
  - a. [Put it into maintenance mode](#).
  - b. [Activate it](#).

This data aggregator is now available for failover. The data aggregator that was ready to take over ("Ready" status) is now active.

## **Put a Data Aggregator into Maintenance Mode**

Prevent the data aggregator that is ready to take over ("Ready" status) from restarting by changing its status to "Maintenance", and making it unavailable for failover.

### **Follow these steps:**

1. Log in to the data aggregator host as the root user or a sudo user.
2. Issue the following command:
 

```
<installation_directory>/scripts/dadaemon maintenance
```

  - ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** /opt/IMDataAggregator

The status of the data aggregator changes to "Maintenance" and is unavailable for failover. Consul detects this change, switches (failover) to the data aggregator that is available for failover (the data aggregator that is in "Ready" status), and changes its status to "Active".

## **Activate a Data Aggregator**

Activate the data aggregator that is in "Maintenance" status so that it is available for failover and can take over if the *active* data aggregator goes offline.

### **Follow these steps:**

1. Log in to the data aggregator host that is in "Maintenance" status as the root user or sudo user.
2. Issue the following command:
 

```
<installation_directory>/scripts/dadaemon activate
```

  - ***installation\_directory***  
The installation directory of the data aggregator.  
**Default:** /opt/IMDataAggregator

The data aggregator is in "Active" status and is available for failover.

### **NOTE**

The data aggregator might take several minutes to start.

## **Stop the Data Aggregator**

Complete the following process:

1. Shut down the data aggregator in preparation for required tasks, such as upgrading the data aggregator.
2. Complete the required actions/tasks.
3. [Restart the data aggregator](#).

### **Follow these steps:**

1. Log in to the data aggregator host as the root user or the sudo user.
 

**NOTE**  
 If you installed the data aggregator as the sudo user, you set up a sudo command alias for the `systemctl dadaemon` command. Use the sudo commands.
2. Do one of the following steps based on your environment:
  - Stop the Data Aggregator service by issuing the following command:
 

```
systemctl stop dadaemon
```

 or, if non-root user:

```
sudo systemctl stop dadaemon
```

- (Fault-tolerant environment) If the local data aggregator is running, shut it down and prevent it from restarting until maintenance is complete by issuing the following command:

```
<installation_directory>/scripts/dadaemon maintenance
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator completes processing and the service stops.

## Start the Data Aggregator

After you complete any required actions/tasks, restart the data aggregator. In an environment with fault-tolerant data aggregators, activate each data aggregator.

### Follow these steps:

1. Log in to the data aggregator host as the root user or the sudo user.

**NOTE**

If you installed the data aggregator as the sudo user, you set up a sudo command alias for the `systemctl dadaemon` command. Use the sudo commands.

2. Do one of the following steps based on your environment:

- Start the Data Aggregator service by issuing the following command:

```
systemctl start dadaemon
```

or, if non-root user:

```
sudo systemctl start dadaemon
```

- (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:

```
<installation_directory>/scripts/dadaemon activate
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator starts and synchronizes with NetOps Portal automatically. If the data repository is unavailable, the data aggregator shuts down. The data collectors send any queued data to the data aggregator.

## Restart the Data Collector

Restart the Data Collector service when required.

For example, restart it when any of the following occur:

- The host loses power.
- You change the location of the host.
- You install an operating system patch.

### Follow these steps:

1. Log in to the data collector host as the root user or the sudo user.

**NOTE**

If you installed data collector as the sudo user, you set up a sudo command alias for the `dcmd` command. Use the sudo commands.

2. Stop the Data Collector, the ICMP daemon, and ActiveMQ services by issuing the following commands:

```
systemctl stop dcmd
```

```
systemctl stop icmpd
systemctl stop activemq
```

or, if non-root user:

```
sudo systemctl stop dcmd
sudo systemctl stop icmpd
sudo systemctl stop activemq
```

The data collector, the ICMP daemon, and ActiveMQ stop.

3. Complete any required tasks.

4. Start the Data Collector, the ICMP daemon, and ActiveMQ services by issuing the following commands:

```
systemctl start activemq
systemctl start icmpd
systemctl start dcmd
```

or, if non-root user:

```
sudo systemctl start activemq
sudo systemctl start icmpd
sudo systemctl start dcmd
```

The data collector, the ICMP daemon, and ActiveMQ restart and resynchronizes automatically. Scheduled polling and discovery resumes. If the data aggregator is not available, the data collector shuts down. If the ActiveMQ broker is unavailable, the data collector restarts that service.

## Restart the Data Repository

Restart the data repository when required.

Restart the data repository when required, for example when the following occurs:

- The data repository host loses power or locks up.
- The data repository host is relocated.
- You install an operating system patch.

### IMPORTANT

Attempting to restart a data repository node before another restarted node has completely rejoined the cluster can cause the entire cluster to go down. The cluster comes back up only after you manually start it.

Use the following process to restart the data repository:

1. [Stop the data aggregator](#). In fault-tolerant environments, put the data aggregator that is in "Ready" status into maintenance mode ("Maintenance" status), and then put the *active* data aggregator into maintenance mode ("Maintenance" status).
2. [Shut down the data repository](#).
3. [Restart the data repository](#).
4. [Restart the data aggregator](#).

## Shut Down the Data Repository

Shut down the data repository in preparation for any required tasks (for example, applying a security patch).

### IMPORTANT

If you try to restart a data repository node before another restarted node has completely rejoined the cluster, the whole cluster can go down. The cluster does not come back up until you manually start it.

Use the following process to shut down the data repository:

1. [Stop the database.](#)
2. [Stop Vertica.](#)

### **Stop the Database**

**Follow these steps:**

1. Log in to the database server as the database administrator (dradmin) user by issuing the following command:  

```
su - dradmin
```
2. Open the Vertica admintools utility from the `/opt/vertica/bin/adminTools` directory.
3. Select **(4) Stop Database**.
4. Select the database, select **OK**, and then press the **Enter/Return** key on your keyboard.
5. Provide the database password, and then press the **Enter/Return** key on your keyboard.  
 The data repository stops.

#### **TIP**

If the data repository does not stop, select **(2) Stop Vertica on Host** from **(7) Advanced Tools Menu**.

6. Select **(E) Exit**, and then press the **Enter/Return** key on your keyboard.

The database is stopped.

### **Stop Vertica**

Stop Vertica on one Node in the cluster.

**Follow these steps:**

1. With the Vertica admintools utility open, select **(7) Advanced Tools Menu**.
2. Select **(2) Stop Vertica on Host**.
3. Select the node to stop, and then press the **Enter/Return** key on your keyboard.
4. Provide the database password, and then press the **Enter/Return** key on your keyboard.  
 The data repository on the node stops.
5. Select **(M) Main Menu**, and then press the **Enter/Return** key on your keyboard.
6. Select **(E) Exit**, and then press the **Enter/Return** key on your keyboard.

Vertica is stopped. The data repository is shut down.

### **Restart the Data Repository**

After you complete required tasks, such as applying a security patch, restart the data repository.

#### **IMPORTANT**

To avoid the cluster from going down and having to manually restart it, ensure that other restarted nodes have completely rejoined the cluster before restarting a data repository node.

Use the following process to start the data repository:

1. [Restart the database.](#)
2. [Restart Vertica on a node that is down in the cluster.](#)

### **Restart the Database**

**Follow these steps:**

1. With the Vertica admintools utility open, select **(3) Start Database**.
2. Select the database, select **OK**, and then press the **Enter/Return** key on your keyboard.



You are prompted for the database password.

3. Provide the database password, and then press the **Enter/Return** key on your keyboard.  
The data repository starts.

**TIP**

If Vertica fails to start normally, try issuing the following command with the following options:

```
/opt/vertica/bin/admintools -t restart_node -d dbname --hosts host_ip_address --force
```

4. Select **(E) Exit**, and then press the **Enter/Return** key on your keyboard.

The database is started.

### **Restart Vertica on a Node that is Down in the Cluster**

**Follow these steps:**

1. With the Vertica admintools utility open, select **(5) Restart Vertica on Host**.
2. Select the database, select **OK**, and then press the **Enter/Return** key on your keyboard.
3. Select the node to start, and then press the **Enter/Return** key on your keyboard.
4. Provide the database password, and then press the **Enter/Return** key on your keyboard.  
The data repository on the node starts.
5. Select **(E) Exit**, and then press the **Enter/Return** key on your keyboard.

Vertica is restarted. The data repository is restarted.

## **Restart the ActiveMQ Broker**

Restart the ActiveMQ broker if required.

If the data aggregator or the data collector detect a problem with Apache ActiveMQ, they restart the ActiveMQ broker automatically. If the restart is unsuccessful, you can restart the ActiveMQ broker manually.

By default, the data aggregator and data collector determine the availability of the ActiveMQ broker every 30 seconds.

**Follow these steps:**

1. Open ActiveMQ, which is located in the following location:

```
<installation_directory>/broker/<apache-activemq-*>/bin
```

**Example:**

```
/opt/IMDataAggregator/broker/apache-activemq-5.18.3/bin
```

– **installation\_directory**

The installation directory of the data aggregator or data collector.

**Default:** (for the data aggregator) /opt/IMDataAggregator

**Default:** (for the data collector) /opt/IMDataCollector

– **apache-activemq-\***

The installation directory of Apache ActiveMQ.

**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2

2. Stop and start the ActiveMQ service by issuing the following commands:

```
systemctl stop activemq
systemctl start activemq
```

The ActiveMQ broker is restarted.

## Restart NetOps Portal

You can stop and restart NetOps Portal, such as when you are backing up your database.

To restart NetOps Portal successfully, stop and restart the relevant services in the following order:

1. [Stop NetOps Portal](#).
2. (If you are experiencing issues with NetOps Portal) [Restart the MySQL database](#).
3. [Start NetOps Portal](#).
4. [Verify the status of NetOps Portal services](#).

### Stop NetOps Portal

Issue the following commands:

1. Stop the NetOps Portal Console service:  

```
systemctl stop caperfcenter_console
```

  
or, if non-root user:  

```
sudo systemctl stop caperfcenter_console
```
2. Stop the device manager:  

```
systemctl stop caperfcenter_devicemanager
```

  
or, if non-root user:  

```
sudo systemctl stop caperfcenter_devicemanager
```
3. Stop the event manager:  

```
systemctl stop caperfcenter_eventmanager
```

  
or, if non-root user:  

```
sudo systemctl stop caperfcenter_eventmanager
```
4. Stop the SSO service:  

```
systemctl stop caperfcenter_sso
```

  
or, if non-root user:  

```
sudo systemctl stop caperfcenter_sso
```

### (Optional) Restart the MySQL Database

Usually, you do not need to restart MySQL database. If you are experiencing issues with NetOps Portal, you can restart the MySQL database.

**Prerequisite:** All NetOps Portal services are stopped.

Issue the following commands:

1. Stop MySQL:  

```
systemctl stop mysql
```

  
or, if non-root user:  

```
sudo systemctl stop mysql
```
2. Restart MySQL:  

```
systemctl start mysql
```

  
or, if non-root user:  

```
sudo systemctl start mysql
```

## **Start NetOps Portal**

Start NetOps Portal by restarting the component services.

### **Follow these steps:**

1. Start the SSO service:  
`systemctl start caperfcenter_sso`  
or, if non-root user:  
`sudo systemctl start caperfcenter_sso`
2. Start the event manager:  
`systemctl start caperfcenter_eventmanager`  
or, if non-root user:  
`sudo systemctl start caperfcenter_eventmanager`
3. Start the device manager:  
`systemctl start caperfcenter_devicemanager`  
or, if non-root user:  
`sudo systemctl start caperfcenter_devicemanager`
4. Wait one minute, then start the NetOps Portal Console service:  
`systemctl start caperfcenter_console`  
or, if non-root user:  
`sudo systemctl start caperfcenter_console`

NetOps Portal is started.

## **Verify the Status of NetOps Portal Services**

After you start NetOps Portal, you can verify the status of NetOps Portal services by issuing the following commands:

1. Verify the NetOps Portal Console service status:  
`systemctl status caperfcenter_console`  
or, if non-root user:  
`sudo systemctl status caperfcenter_console`
2. Verify the device manager status:  
`systemctl status caperfcenter_devicemanager`  
or, if non-root user:  
`sudo systemctl status caperfcenter_devicemanager`
3. Verify the event manager status:  
`systemctl status caperfcenter_eventmanager`  
or, if non-root user:  
`sudo systemctl status caperfcenter_eventmanager`
4. Verify the SSO service status:  
`systemctl status caperfcenter_sso`  
or, if non-root user:  
`sudo systemctl status caperfcenter_sso`

The following sample message indicates that the NetOps Portal console is running:

```
Performance Center Console is running: PID:12993, Wrapper:STARTED, Java:STARTED
```

## Syslog Integration

You can configure the following components to send Common Event Format (CEF) messages to Syslog:

- [NetOps Portal](#)
- [The data aggregator](#)

If you have configured and enabled a Syslog server such that it is logging the CEF messages that it receives from NetOps Portal and the data aggregator to the `/var/log/messages` file, [you can watch for these messages](#).

### Configure NetOps Portal to Send Messages to Syslog

You can configure NetOps Portal to send Common Event Format (CEF) messages to Syslog.

The following are examples of messages that NetOps Portal sends to Syslog:

- A user logs in or logs out by way of NetOps Portal or REST (when `syslogVerbose` is uncommented).
- A user changes (their own or another) user password.
- A user proxies to another user.
- A user creates, deletes, or edits (their own or another) user or role using NetOps Portal or REST.
- A user modifies the administration/security settings.
- A user modifies the administration/data sources.
- A NetOps Portal service (Console, SSO, Device Manager, Event Manager) has stopped or started.

For more information, see [Examples of Syslog Logging Output](#).

NetOps Portal Syslog supports standard log4j Syslog appender configuration.

For more information about log4j Syslog appenders, see [the Apache Log4j 2 documentation](#).

**Prerequisite:** You have configured and enabled a Syslog server that can accept CEF messages from NetOps Portal.

You can configure Syslog logging on the following NetOps Portal services:

- [The NetOps Portal Console service](#)
- [The SSO service](#)

### Configure Syslog Logging on the NetOps Portal Console Service

When configured, the NetOps Portal Console Service logging logs administrative and REST actions.

**Follow these steps:**

1. Edit the NetOps Portal version of the `log4j2.properties` file by issuing the following command:

```
vi /opt/CA/PerformanceCenter/PC/resources/log4j2.properties
```

2. Search for `Syslog` comments, and uncomment the lines per the instructions in the file.

**For example:**

```
logger.Syslog.appenderRef.Syslog.ref = Syslog
```

#### NOTE

To enable verbose logging, uncomment the `SyslogVerbose` line. For example, due to the amount of log entries it produces, user login/logout is part of verbose logging.

3. If your Syslog server is on a different host, to have log4j log events to this host, update the `appender.syslog.host` parameter.
4. Restart the NetOps Portal Console service:
 

```
systemctl stop caperfcenter_console
systemctl start caperfcenter_console
```

Syslog logging is configured on the NetOps Portal Console service.

## Configure Syslog Logging on the SSO Service

When configured, the SSO Service logging logs user authentication actions.

### Follow these steps:

1. Edit the SSO version of the `log4j2.properties` file by issuing the following command:

```
vi /opt/CA/PerformanceCenter/sso/resources/log4j2.properties
```

2. Search for `Syslog` comments, and uncomment the lines per the instructions in the file.

#### For example:

```
logger.Syslog.appenderRefs = Syslog
```

#### NOTE

To enable verbose logging, uncomment the `SyslogVerbose` line. For example, due to the amount of log entries it produces, user login/logout is part of verbose logging.

3. If your Syslog server is on a different host, to have `log4j` log events to this host, update the `appender.syslog.host` parameter.
4. Restart the SSO service:

```
systemctl stop caperfcenter_sso
systemctl start caperfcenter_sso
```

Syslog logging is configured on the SSO Service.

## Examples of Syslog Logging Output

The following are CEF message examples that NetOps Portal generates:

### Generated when a user is created:

```
Mar 25 10:39:09 myhostname.my.domain.net caperfcenter[162027] CEF:0|Broadcom|DX NetOps
Portal|22.2.7.1160|504|USER CREATED|1|cat=User Admin act=Create suid=1 suser=admin
cs1=8 cs1Label=Tenant ID cs2=Default Tenant cs2Label=Tenant Name msg=Created new
User( ID: 4, Name: testuser )
```

### Generated when a user is edited:

```
Mar 25 10:40:06 myhostname.my.domain.net caperfcenter[162027] CEF:0|Broadcom|DX NetOps
Portal|22.2.7.1160|505|USER MODIFIED|1|cat=User Admin act=Modify suid=1 suser=admin
cs1=8 cs1Label=Tenant ID cs2=Default Tenant cs2Label=Tenant Name msg=Updated User( ID:
2, Name: user )
```

### Generated when a user is deleted:

```
Mar 14 05:16:27 myhostname.my.domain.net caperfcenter[162027] CEF:0|Broadcom|DX NetOps
Portal|22.2.7.1160|506|USER DELETED|1|cat=User Admin act=Delete suid=1 suser=admin
cs1=8 cs1Label=Tenant ID cs2=Default Tenant cs2Label=Tenant Name msg=Deleted User( ID:
9, Name: sam )
```

### Generated when a role is edited:

```
Mar 23 15:21:29 myhostname.my.domain.net caperfcenter[162027] CEF:0|Broadcom|DX NetOps
Portal|22.2.7.1160|501|ROLE MODIFIED|1|cat=Role Admin act=Modify suid=1 suser=admin
cs1=8 cs1Label=Tenant ID cs2=Default Tenant cs2Label=Tenant Name msg=Updating Role( ID:
8, Name: IT Manager ) with Enabled: false
```

### Generated when NetOps Portal console service is started/stopped:

```
Mar 23 13:11:12 myhostname.my.domain.net caperfcenter_console: CEF:0|Broadcom|DX NetOps
Portal|22.2.7.1160|100|STARTING SERVICE|1|user=root
Mar 23 15:15:26 myhostname caperfcenter_console: CEF:0|Broadcom|DX NetOps Performance
Center Console|21.2.8|101|STOPPING SERVICE|1|user=root
```

If you have enabled verbose logging, the following are examples of the CEF messages that the SyslogVerbose logger generates:

#### Generated when a user logs in:

```
Mar 14 05:15:29 myhostname.my.domain.net caperfcenter[9476] CEF:0|Broadcom|DX NetOps
Portal|22.2.7.1160|303|USER LOGGED IN|1|cat=User Session act=Log In suser=User1
msg=User User1 Successfully Logged In
```

#### Generated when a user logs out:

```
Mar 24 21:26:58 myhostname.my.domain.net caperfcenter[2128] CEF:0|Broadcom|DX NetOps
Portal|2.0.0.0|302|USER LOGGED OUT|1|cat=User Session act=Log Out suid=1 suser=admin
cs1=8 cs1Label=Tenant ID cs2=Default Tenant cs2Label=Tenant Name msg=User Successfully
Logged Out
```

#### Generated when a user proxies as another user:

```
Mar 24 21:23:22 myhostname.my.domain.net caperfcenter[2128] CEF:0|Broadcom|DX
NetOps Portal|2.0.0.0|305|PROXY USER LOGGED IN|1|cat=User Session act=Log In suid=1
suser=admin cs1=8 cs1Label=Tenant ID cs2=Default Tenant cs2Label=Tenant Name
msg=Proxying User ( Name: User1, ID: 124 )
```

## Configure the Data Aggregator to Send Messages to Syslog

You can configure the data aggregator to send Common Event Format (CEF) messages to Syslog.

The following are examples of messages that the data aggregator sends to Syslog:

- REST User authentication failure/success (when `success` is uncommented).
- A service (dadaemon, activemq) has stopped or started.

For more information, see [Examples of Syslog Logging Output](#).

### Configure Syslog Logging on the Data Aggregator

**Prerequisite:** You have configured and enabled a Syslog server that can accept CEF messages from the data aggregator.

The data aggregator Syslog supports standard log4j Syslog appender configuration.

For more information, see [the Apache Log4j 2 documentation](#).

#### Follow these steps:

1. Edit the `org.ops4j.pax.logging.cfg` file by issuing the following command:

```
vi <installation_directory>/apache-karaf/etc/org.ops4j.pax.logging.cfg
```

#### Example:

```
vi /opt/IMDataAggregator/apache-karaf/etc/org.ops4j.pax.logging.cfg
```

#### – **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

2. Search for `Syslog.level` comments, and change the logger level from `OFF` to `INFO`.

**Change:**

```
log4j2.logger.Syslog.level = OFF
```

**To:**

```
log4j2.logger.Syslog.level = INFO
```

3. Search for `Syslog` comments, and uncomment the lines per the instructions in the file.

**For example:**

```
log4j2.logger.Syslog.appenderRefs = Syslog
```

4. If your Syslog server is on a different host, to have log4j log events to this host, update the `log4j2.appender.syslog.host` parameter.

Syslog logging is configured on the data aggregator.

### **Examples of Syslog Logging Output**

The following are CEF message examples that the data aggregator sends:

#### **Sent for each successful REST login:**

```
Apr 12 10:03:37 myhostname.my.domain.net dadaemon[31533] CEF:0|Broadcom|DX NetOps
Data Aggregator|21.2.10.374|300|AUTHENTICATION SUCCESSFUL|1|cat=Authentication
suid=1 suser=admin src=10.17.255.65 request=/consul/servicestatus msg=Successfully
authenticated user - Method: SSO token
```

#### **Sent for each failed REST login attempt:**

```
Apr 12 10:03:33 myhostname.my.domain.net dadaemon[31533] CEF:0|Broadcom|DX NetOps Data
Aggregator|21.2.10.374|400|AUTHENTICATION FAILURE|4|cat=Authentication src=10.17.255.64
request=/rest msg=Failed to authenticate user - Method: No auth header
```

## **Watch for Syslog Messages**

If you have configured and enabled a Syslog server such that it is logging the Common Event Format (CEF) messages that it receives from NetOps Portal and the data aggregator to the `/var/log/messages` file, you can watch for these messages.

For example, you can watch for CEF messages that NetOps Portal and the data aggregator send.

For more information about the configuration and examples of CEF messages:

- See [Configure NetOps Portal to Send Messages to Syslog](#).
- See [Configure the Data Aggregator to Send Messages to Syslog](#).

Watch for Syslog messages by issuing the following command:

```
tail -f /var/log/messages
```

#### **NOTE**

If you have configured your log4j Syslog appender to log events to localhost, you can view them in the `/var/log/messages` file.

## Logs

To investigate issues, review any log messages that occurred around the time of the issue. Each DX NetOps Performance Management component includes separate log files.

Support often requests log files or changes the log configuration. CARE collects the log files.

## Data Aggregator Logs

The data aggregator supports standard log4j configuration and logging levels. Configure the logs using the `<installation_directory>/apache-karaf/etc/org.ops4j.pax.logging.cfg` file. Data aggregator logs are available in the `<installation_directory>/apache-karaf/data/log` directory.

- **installation\_directory**  
The installation directory of the data aggregator.  
**Default:** `/opt/IMDataAggregator`

The following logs provide potentially useful information:

- `Exception.log`  
This log includes exceptions and warnings. When the system is operating normally, this log does not contain messages. Exceptions in this log require attention from Support.  
If an exception is repeated, recurring entries in the log do not include the stack trace. A recurrence count indicates how many times the exception has occurred.
- `karaf.log`  
This log is a catch-all for data aggregator information.
- `LongRunningAndFailedQueries.log`  
This log lists errors and information messages related to database queries.
  - Long running queries are requests that take longer than 60 seconds to complete.
  - Queries time out and fail after 110 seconds.
- `PollSummary.log`  
This log shows poll responses that the data aggregator receives from the data collectors. Use this log to verify that the data aggregator is receiving poll responses.
- `shutdown.log`  
This log includes audit messages when contact with the data repository is lost.
- `shutdown_details.log`  
This log includes heartbeat messages between the data aggregator and the data repository, as well as any data aggregator shutdowns for debugging purposes.

By default, the data aggregator is configured to save 1-10 backups with a max files size of 100 MB.

### NOTE

Undocumented logs in the `log` directory are for internal use only.

## NetOps Portal Logs

Use the most recent NetOps Portal log files to find errors that are associated with the database or data source synchronization.

NetOps Portal log filenames include the relevant date and time. NetOps Portal generates new log files automatically each day, and removes older log files automatically after 14 days to avoid consuming excessive disk space.



## Locate Log Files from Events

### Follow these steps:

1. In NetOps Portal, hover over **Performance, My Dashboard**, and then click **Events**.  
The **Events** dashboard opens.
2. Sort by **Status**.
3. To look at the related log file, note the event type and failure date and time. In the log directory, open the log file with the corresponding date in the filename.

## Log Files

NetOps Portal stores the following logs in subfolders that correspond to a service (or daemon). Locate the following log files in the `CA/PerformanceCenter/servicename/logs` directory:

Replace the *servicename* parameter with one of the following service names:

- **DM**

The Device Manager.

- `DMService.log`  
The device manager application log file. Output from the device manager service, primarily related to synchronization.
- `wrapper.log`  
Device manager service process logging.

- **EM**

The Event Manager.

- `EMService.log`  
Output from the event manager; includes details of events and alarms.
- `wrapper.log`  
Event manager service process logging.

- **PC**

NetOps Portal.

- `PCService.log`  
NetOps Portal-related logging; comprises user interface (UI) and view components.
- `wrapper.log`  
NetOps Portal Console service process logging.

- **SSO**

The NetOps Portal Single Sign-On (SSO) service.

- `SSOService.log`  
SSO logging, including HTTPS information when you have enabled HTTPS for NetOps Portal.
- `wrapper.log`  
SSO service process logging.

For issues with the SSO Configuration tool (SsoConfig), view the `/opt/CA/PerformanceCenter/sso/logs/application.log` application log.

- **MySQL**

- `hostname.err`  
The hostname is the name of the system. This file contains errors related to MySQL.  
NetOps Portal stores the MySQL error log, by default, in the `/opt/CA/MySQL/data` directory.

## SSO Audit Log

To support security auditing, Single Sign-On logs details about user login activity to a file. Each time a user attempts to log in to the NetOps Portal login page, SSO logs an entry in the audit log. If SSO redirects the user to a SAML2 server to log

in or re-authenticate, SSO only logs successful logins. Check the log to verify user activity. The log contains one line per login.

The following details are written to the log per login request:

- Time and date stamp when the user logged in
- Product code (for example, pc for NetOps Portal)
- Username
- Whether the Remember Me option was selected
- Single Sign-On version
- The remote host IP address

#### Follow these steps:

1. Log in to the server where a CA data source product is installed.
2. Open a command prompt, and cd to the following directory:

```
[InstallationDirectory]/PerformanceCenter/sso/logs
```

#### NOTE

The audit log is saved in the following location on Windows servers:

```
[InstallationDirectory]\Portal\SSO\logs.
```

3. Enter dir to see the contents of the directory.  
The filename of the log file is SingleSignOnAuditLogyyyy-mm-dd.log.
4. Enter the name of the audit file you want to view.  
The file opens in the local text editor application.

## Activate a Disaster Recovery System

If a large-scale disaster occurs, and the primary (source) system is unavailable, activate, or start, the recovery (target) system.

#### NOTE

Startup time for the recovery system takes the same time that is required to start the data aggregator.

**Prerequisite:** You have [installed a disaster recovery system](#).

Use the following process to activate the discovery recovery system:

1. [Start the Data Repository](#)
2. [Start the Data Aggregator](#)
3. [Start NetOps Portal](#)
4. [Start the Data Collectors](#)
5. (If the URL to access the data aggregator is a selected authorized URL) [Reconfigure URL Protection for the Data Aggregator](#)

### Start the Data Repository

#### Follow these steps:

1. Log in to the recovery database cluster as the database admin user.
2. Start the Vertica Administration Tools utility (adminTools ):  
`/opt/vertica/bin/adminTools`
3. Select option **3 (Start Database)**.
4. Press the **Space** bar on your keyboard next to the database name, select **OK**, and then press the **Enter** key on your keyboard.

You are prompted for the database password.

5. Enter the database password, and then press the **Enter** key on your keyboard.  
The data repository starts.
6. Select **Exit**, and then press the **Enter** key.
7. Run the `<data_repository_directory>/update_da_dc_database_references.sh` data repository disaster recovery script.
  - ***data\_repository\_directory***  
The installation directory for the data repository.  
**Default:** `/opt/CA/IMDataRepository_verticaVersion`

The data repository is started.

### **Start the Data Aggregator**

**Do one of the following steps:**

- Start the ActiveMQ and Data Aggregator services:
 

```
systemctl start activemq
systemctl start dadaemon
```
- (Fault-tolerant environments) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:
 

```
<installation_directory>/scripts/dadaemon activate
```

  - ***installation\_directory***  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`

The data aggregator is started. If the data repository is unavailable, the data aggregator shuts down.

### **Start NetOps Portal**

**Follow these steps:**

1. Restore the NetOps Portal backups.  
For more information, see [Restore NetOps Portal](#).
2. Run the `<installation_directory>/Tools/bin/update_pc_da_database_references.sh` NetOps Portal disaster recovery script.
  - ***installation\_directory***  
The installation directory for NetOps Portal.  
**Default:** `/opt/CA/PerformanceCenter`
3. Start the SSO service by issuing the following command:
 

```
systemctl start caperfcenter_sso
```
4. Wait one minute, then start the event manager and device manager by issuing the following commands:
 

```
systemctl start caperfcenter_eventmanager
systemctl start caperfcenter_devicemanager
```
5. Wait one minute, then start the NetOps Portal Console service by issuing the following command:
 

```
systemctl start caperfcenter_console
```

NetOps Portal is started.

### **Start the Data Collectors**

Start the Data Collector service by issuing the following command:

```
systemctl start dcmd
```

The data collector restarts. If the data aggregator is unavailable, the data collector shuts down.

**TIP**

For remote data collectors, update them to connect to the recovery data aggregator.

For more information, see [Configure Data Collector When the Data Aggregator IP Address Changes](#).

**Reconfigure URL Protection for the Data Aggregator**

If the URL to access the data aggregator is a selected authorized URL (it is listed in the **Select Authorized URLs** field on the **Security Settings** page in NetOps Portal), add the new data aggregator URL to the list.

1. Hover over **Administration**, **Data Sources**, and then click **Data Sources**. The **Data Sources** page appears.
2. Select the data aggregator data source, and then click **Edit**.  
The **Edit Data Source** page appears.
3. Confirm that the host name for the data aggregator data source is correct, and then click **Save**.
4. [Secure your system and require users to log in using only authorized URLs by enabling Single Sign-On \(SSO\) spoofing protection](#).

## Securing Performance Monitoring

---

You can configure NetOps Portal so that it is secure.

For the latest information about how to configure NetOps Portal so that it is secure, see the unified [securing section](#) in the [DX NetOps documentation](#).

The following is a list of topics that you might find helpful while securing NetOps Portal:

- [Single Sign-On](#)
- [Enable HTTPS](#)
- [Configure the DX NetOps Security Settings Using NetOps Portal](#)
- [Configure the DX NetOps Security Settings Using the SSO Configuration Tool](#)
- [Authenticate and Encrypt ActiveMQ Communication](#)
- [Change the ActiveMQ JMX Password](#)
- [Certificate Administration](#)

## Integrating

You can integrate DX NetOps Performance Management with other monitoring software to enrich the available data and provide more information about your infrastructure. Many integrations, such as CA Application Delivery Analysis and DX NetOps Network Flow Analysis, use the data source model. These integrations collect data and send it to NetOps Portal for visualization. NetOps Portal controls administrative functions.

For more information, see [Manage Data Sources](#).

Other integrations, such as DX NetOps Virtual Network Assurance and DX NetOps Mediation Manager, connect to the data collector, and inject the data into the data aggregator data source. With these integrations, data appears as native to the DX NetOps Performance Management environment.

**Table 1: Performance Monitoring Integrations**

Product Integrations with DX NetOps Performance Management	Integration Business Value
DX NetOps Spectrum	DX NetOps Spectrum is a services and infrastructure management system that monitors the state of managed elements including devices, applications, host systems, and connections. For more information, see <a href="#">Integrate with DX NetOps Spectrum for Fault Management</a> .
DX Operational Intelligence	Validate the impact of changes and solve problems faster by integrating with DX Operational Intelligence (OI). Validate application performance that is delivered over the network by applying network capabilities and service quality metrics. For more information, see <a href="#">Integrate with DX Operational Intelligence</a> .
DX NetOps Virtual Network Assurance	DX NetOps Virtual Network Assurance enables existing infrastructure management solutions to monitor software-defined networking (SDN) and network functions virtualization (NFV). For more information, see <a href="#">Integrate with Virtual Network Assurance</a> .
CA Application Delivery Analysis	CA Application Delivery Analysis provides end-to-end performance monitoring through dashboards and views. These dashboards and views show historically normal performance for users and metrics that cross acceptable performance thresholds. For more information, see <a href="#">Integrate with CA Application Delivery Analysis</a> .
DX NetOps Network Flow Analysis	DX NetOps Network Flow Analysis gives you an enterprise-wide view into the composition of traffic on every link and helps you detect threatening traffic patterns in the making. For more information, see <a href="#">Integrate with DX NetOps Network Flow Analysis</a> .
CA Unified Communications Monitor	CA Unified Communications Monitor passively monitors the performance of your unified communications systems to maintain and report on the quality of audio and video calls. For more information, see <a href="#">Integrate with CA Unified Communications Monitor</a> .

Product Integrations with DX NetOps Performance Management	Integration Business Value
DX Application Performance Management	DX Application Performance Management (APM) supports NetOps Portal, the data aggregator, and the data collector for instrumentation. It receives performance metrics about the component services and hosts. For more information, see <a href="#">Monitor Server Performance with DX Application Performance Management</a> .
DX NetOps Mediation Manager	DX NetOps Mediation Manager monitors the performance for non-SNMP based devices, such as mobile wireless, fiber-optic switch, radio access, and 3G or 4G voice data. DX NetOps Mediation Manager supports a wide range of protocols to access data, for example, SOAP, SSH, XML, SQL, JMS, SFTP, and HTTP. DX NetOps Mediation Manager is portable across all platforms. For more information, see <a href="#">Install DX NetOps Mediation Manager</a> .

## Integrate with DX NetOps Spectrum for Fault Management

You can integrate DX NetOps Performance Management with DX NetOps Spectrum (Spectrum) so that you can share models, Global Collections, and events between the two systems.

Spectrum contributes devices, interfaces, and groups to the NetOps Portal inventory, which DX NetOps Performance Management can monitor. DX NetOps Performance Management contributes infrastructure performance events to Spectrum, so you can see performance events and fault alarms side by side in Spectrum OneClick.

The following video shows the DX NetOps Spectrum/DX NetOps Performance Management integration process:

In this article:

- [Contribute Devices to NetOps Portal](#)
- [Drill Down from OneClick to DX NetOps Performance Management Data](#)
- [Model Classes and Device Subtypes](#)
- [Enable Synchronized Discovery](#)
- [Enable Control of Device Life Cycle States](#)
- [Interface Synchronization](#)
- [IP Domain Synchronization](#)
- [Group Synchronization](#)
- [Multi-Tenancy and Managing Devices](#)
- [Event Integration](#)
- [View Alarms in NetOps Portal](#)
- [Integration Architecture](#)
- [Synchronization with DX NetOps Spectrum](#)

### **Contribute Devices to NetOps Portal**

You determine which devices to contribute to the NetOps Portal inventory from Spectrum. NetOps Portal creates discovery profiles to monitor those devices. NetOps Portal automatically discovers interfaces that are associated with those devices. If the devices exist in NetOps Portal, the devices are reconciled to a single item.

#### **IMPORTANT**

Spectrum SNMP throttling does not apply to ongoing polling by the data aggregator. This feature protects critical devices from failing in case too many polling flows are configured. The throttling mechanism applies to any

monitoring or discovery activities. If you have configured throttling in Spectrum, apply the same setting in the data aggregator.

For more information, see [Poll Sensitive and Critical Devices Without a Performance Impact](#).

### **Drill Down from OneClick to DX NetOps Performance Management Data**

Access DX NetOps Performance Management data from Spectrum device and interface models, which provides rapid access to contextual information about device performance issues.

#### **TIP**

From OneClick, you can navigate directly to NetOps Portal by right-clicking a device or interface, and then clicking **NetOps Portal**. A window opens, and NetOps Portal loads the context page for the selected device or interface. Spectrum data sources require a username and password.

### **Model Classes and Device Subtypes**

Synchronization uses the Spectrum model class to determine the NetOps Portal device subtype. The following list shows the model class followed by the device subtype:

- Router = Router
- Switch-Router = Router
- Switch = Switch
- Workstation-Server = Server

All other model classes are listed as **Other** in NetOps Portal.

#### **IMPORTANT**

Devices that are synchronized from Spectrum (the **Synchronize device life cycle state from Spectrum** checkbox is selected for that data source) appears as pingable if any of the follow configuration issues occur:

- The SNMP profile in NetOps Portal does not have the correct contact information for the device.
- The SNMP profile is not assigned to the discovery profile.
- A local firewall is preventing communication from the data collector.

For more information about this checkbox, see [Configure a Data Source](#).

### **Enable Synchronized Discovery**

Synchronization reduces the administrative management that is required for device discovery. NetOps Portal adds the IP addresses of synchronized devices from Spectrum to a discovery profile. The discovery profile is specified for each IP domain. Run the discovery profile to identify the devices through SNMP.

For more information, see [Run Device Discovery](#).

You can enable synchronized discovery by ensuring that the **Synchronize component items that are not currently present on the monitored device** checkbox is selected for the data aggregator data source.

For more information about this checkbox, see [Configure a Data Source](#).

### **Enable Control of Device Life Cycle States**

You can enable Spectrum to control the life cycle state of devices in NetOps Portal (the **Synchronize device life cycle state from Spectrum** checkbox is selected for that data source). When enabled, changes in Spectrum trigger changes in NetOps Portal. If you change the state of a device in NetOps Portal, the state does not change again unless the state changes in Spectrum.

For more information about this checkbox, see [Configure a Data Source](#).



**IMPORTANT**

To avoid conflicts in device life cycle management, do not grant the Administer Life Cycle role to users in NetOps Portal.

For more information about this role, see [Role Rights](#).

When enabled, the state of the device in Spectrum overwrites the state of the device in NetOps Portal, and NetOps Portal uses the following behavior:

- Active devices in Spectrum are in "Active" state in NetOps Portal.
- Devices in Maintenance mode in Spectrum are in "Maintenance" state in NetOps Portal.
- Deleted devices in Spectrum are in "Retired" state in NetOps Portal.

For more information about how to view device states in NetOps Portal, see [Manage Device Life Cycles](#).

**Interface Synchronization**

For each device that Spectrum sends to DX NetOps Performance Management, the monitored interfaces from Spectrum are also synchronized. For each devices, DX NetOps Performance Management shows the following for interfaces:

- Those that are synchronized from Spectrum.
- Those that are discovered directly through SNMP.
- Those that other data sources, such as CA Application Delivery Analysis and DX NetOps Network Flow Analysis (NFA), monitor.

**NOTE**

DX NetOps Performance Management does not show interfaces that are filtered out at the monitoring profile-level. Right-clicking an interface model in OneClick does not show the drill-down option if that interface is filtered out of the corresponding monitoring profile in DX NetOps Performance Management.

For more information, see [Manage Monitoring Profiles](#).

**IP Domain Synchronization**

OneClick synchronizes DX NetOps Performance Management IP domains as NetOps Portal IP Domain models. The IP domain models appear in the same area as Global Collections in the OneClick Navigation panel. The IP domain models have the same names as the DX NetOps Performance Management IP domain definitions. If you add a device model to the IP domain in OneClick, the device is synchronized to that IP domain in DX NetOps Performance Management.

**NOTE**

OneClick synchronizes all IP domain definitions, regardless of tenant associations.

**Group Synchronization**

Spectrum Global Collections and landscapes are synchronized to DX NetOps Performance Management as groups in the NetOps Portal Groups tree. The groups are created under **Inventory/Data Sources** in the group tree.

Use these groups for the following tasks:

- Create reporting groups.
- Define site membership.
- Drive the content of other custom groups and collections.

**Multi-Tenancy and Managing Devices**

Spectrum devices that are synchronized to DX NetOps Performance Management belong to the tenant who owns the associated IP domain. For items that are shared among multiple tenants, add the items to the appropriate Service Provider groups in NetOps Portal.

For more information, see [Groups](#) and [Multi-tenancy](#).

## Event Integration

DX NetOps Performance Management performance events are converted to Spectrum alarm set or clear events that are asserted on models in each landscape. Polling for supported events begins when synchronization completes. Alarms that originate in DX NetOps Performance Management appear in Spectrum OneClick.

### NOTE

- Event processing cannot map events to the correct devices in Spectrum for those devices that are modeled as non-proxy on multiple landscapes.
- Spectrum processes performance events that are synchronized from DX NetOps Performance Management only for those device models that are active in Spectrum.

By default, Spectrum can handle events from the data aggregator and NFA data sources. You can configure OneClick to handle other events from the Event Manager.

The following image shows the alarm details for a DX NetOps Performance Management performance event in OneClick. The alarm includes the alarm type, time of occurrence, event ID, and source:

**Figure 137: Infrastructure Performance Events in DX NetOps Spectrum**

The alarm details tab shows information about the performance event.

Alarm Details | Incident Cause | Interfaces | Performance | Alarm History | Neighbors | Events | Path View

Threshold exceeded: VM Total Mem Util for nvminintegration.ca.com ( ) at Sysedge 5.0.2 & 2.0.2 AIM testing (vmmemutil[83.000000] >= 60.000000)  
Jul 6, 2011 8:12:17 AM EDT  
Performance Threshold Exceeded

Detail of Threshold Violation:  
1) Incident Start Time: Jun 30, 2011 9:45:03 AM EDT  
2) Event ID: 926928  
3) Event Source: NetVoyant  
3) Alert Message: Threshold exceeded: VM Total Mem Util for nvminintegration.ca.com ( ) at Sysedge 5.0.2 & 2.0.2 AIM testing (vmmemutil[83.000000] >= 60.000000)

A corresponding major Threshold Exceeded alarm will be generated.

Severity Major  
Impact 0  
Acknowledged [set](#)  
Clearable Yes  
Trouble Ticket ID [set](#)  
Assignment  
Landscape fcawinesx11 (0xffe00000)

Symptoms The monitored threshold has been exceeded.  
Probable Cause  
Actions Launch the "Performance View" to see incident details.

You can view the event in Performance Center for more details.

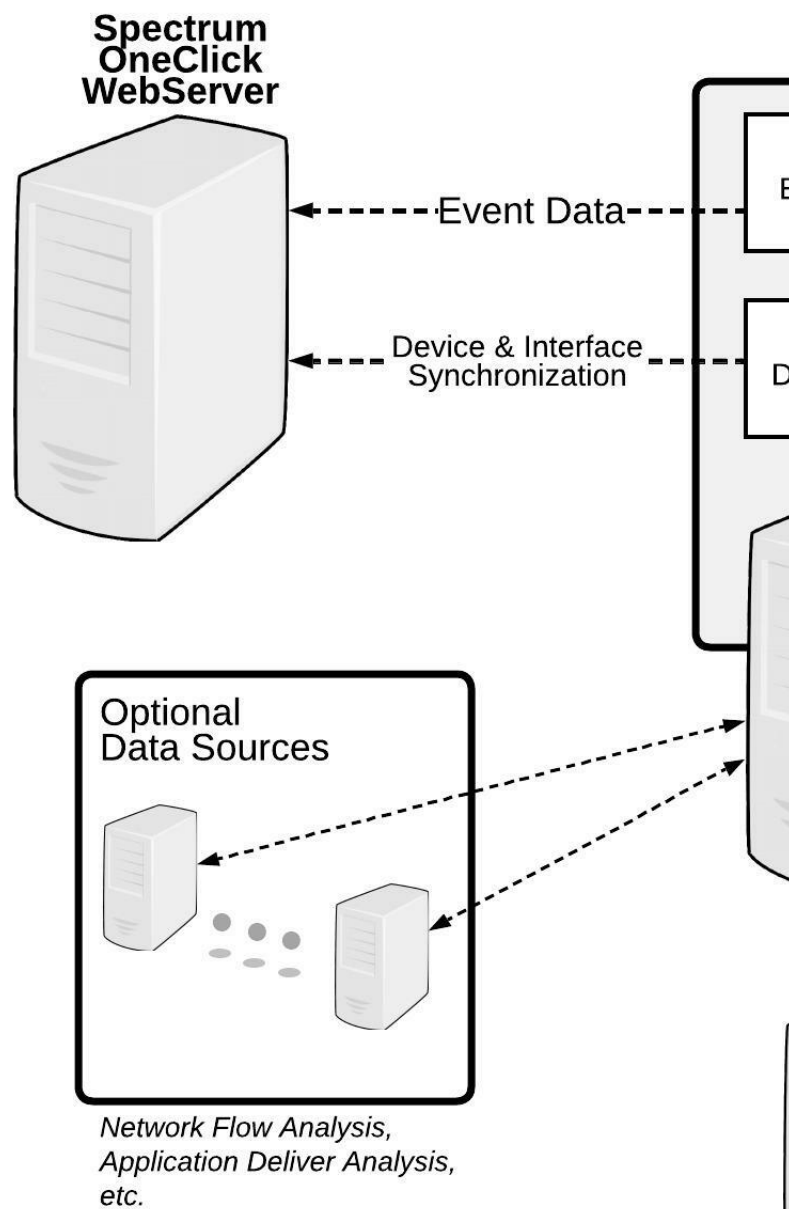
## View Alarms in NetOps Portal

You can view and manage Spectrum alarms using the Alarms views in NetOps Portal. An alarms view provides a prioritized list of Spectrum alarms, which helps you quickly focus on resolving the most impactful problems. The alarms view also provides visibility into other, potentially related, issues on the same device, or connected to a device.

For more information, see [Alarms View](#).

## Integration Architecture

The following image illustrates the architecture of the integration:

**Figure 138: Solution Architecture with Spectrum and DX NetOps Performance Management**

The OneClick server synchronizes devices with the device manager on the NetOps Portal host. The OneClick server polls the event manager for performance events:

- Synchronize One SpectroSERVER or a Distributed SpectroSERVER (DSS) with DX NetOps Performance Management by specifying the OneClick web server as a data source.
- Each landscape in the DSS is defined as a NetOps Portal group.
- Devices and interfaces in the DSS are synchronized with DX NetOps Performance Management and added to the appropriate landscape group.
- OneClick polls the event manager for events that are relevant to a specified landscape group. By default, this polling occurs every 60 seconds. Retrieved events are then translated to Spectrum events, which can generate or clear alarms.  
Spectrum only processes alarms for devices and interfaces that are synchronized to DX NetOps Performance Management.

### **Synchronization with DX NetOps Spectrum**

Consider the following guidelines:

- Whenever you restore the SpectroSERVER or OneClick database and you restore another data source, restart Tomcat.
- Enable data to synchronize between Spectrum and DX NetOps Performance Management.

## **Integrate with DX NetOps Spectrum**

Monitor DX NetOps Spectrum (Spectrum) devices and interfaces with NetOps Portal and process DX NetOps Performance Management events in Spectrum by integrating DX NetOps Performance Management with Spectrum.

NetOps Portal synchronizes the devices with Spectrum.

Use the following process to integrate with Spectrum:

1. [Verify the Prerequisites](#)
2. [Configure Spectrum as a Data Source](#)
3. [Enable Monitoring in NetOps Portal](#)
4. [Configure Event Polling in Spectrum](#)
5. [Send Traps to Spectrum](#)
6. [Review Multi-Tenancy and Spectrum](#)

### **Verify the Prerequisites**

Before integrating with Spectrum, ensure that you have completed the following prerequisite steps:

- [You have verified that the Spectrum and DX NetOps Performance Management versions are compatible.](#)
- To discover Spectrum devices with the data aggregator, you have verified the following:
  - [You have configured the Spectrum data source to contribute new data source inventory to the data aggregator \(the \*\*Contribute inventory to Data Aggregator\*\* checkbox is selected on the \*\*Edit Data Source\*\* dialog\).](#)
  - [You have enabled Spectrum to control the life cycle state of items in DX NetOps Performance Management \(the \*\*Synchronize device life cycle state from Spectrum\*\* checkbox is selected on the \*\*Edit Data Source\*\* dialog\).](#)
- [You have created IP domains in NetOps Portal for each IP routing space that Spectrum monitors.](#)
- For each IP domain, [you have installed a data collector and assigned it to that IP domain..](#)
- For event integration, you have verified that the Spectrum OneClick server can communicate to the NetOps Portal host on port 8281.

---

## Configure Spectrum as a Data Source

[Add, or register, Spectrum as a data source in NetOps Portal.](#)

### Enable Monitoring in NetOps Portal

Configure NetOps Portal to monitor Spectrum devices for performance data. Use the following process to enable monitoring:

1. [Create an SNMP Profile](#)
2. [Add the Spectrum Device as a Member to the IP Domain](#)
3. [Control Performance Metric Polling from NetOps Portal](#)
4. [Configure a Threshold Profile](#)

During the next synchronization, Spectrum sends the device IP addresses to the selected IP domains in DX NetOps Performance Management. DX NetOps Performance Management reconciles the synchronized devices with existing devices. If the IP address does not belong to a device that the data aggregator already monitors, the IP address is added to a discovery profile with the following properties:

- The name of the discovery profiles is the same as the name of the IP domain.
- By default, the discovery profile runs daily.
- To avoid the buildup of IP addresses, the IP address is removed from the discovery profile after the device is discovered.

#### TIP

To enable rediscovery, add the IP address of the Spectrum devices to another discovery profile in DX NetOps Performance Management.

For more information about discovery, see [Discovery](#).

#### NOTE

You can enable NetOps Portal device monitoring in Spectrum using *one* of the following options:

- Add devices to IP domains manually.
- Update IP domain membership dynamically.

For more information, see [the DX NetOps Spectrum documentation](#).

### Create an SNMP Profile

[Create an SNMP profile](#)

### Add the Spectrum Device as a Member to the IP Domain

Add the Spectrum device as a member to the IP domain by adding a device model to an IP domain. To define which Spectrum devices NetOps Portal discovers, add those devices to the DX NetOps Performance Management IP Domain global collections in OneClick. Integrating with Spectrum creates the Default Domain global collection. When synchronization occurs, NetOps Portal sends the IP domains to Spectrum.

#### TIP

You can update IP domain membership dynamically in Spectrum using search rules on the IP domain Global Collecton.

For more information, see [the DX NetOps Spectrum documentation](#).

## **Control Performance Metric Polling from NetOps Portal**

Use the following process to control performance metric polling from NetOps Portal:

1. Sort the Spectrum devices into collections.
2. [Configure a monitoring profile](#).

## **Configure a Threshold Profile**

[Configure a threshold profile](#) to have the data aggregator analyze configurable performance thresholds and generate events when performance metrics violate those thresholds. Spectrum processes these events, and then creates alarms on the appropriate models.

## **Configure Event Polling in Spectrum**

The SpectroSERVER polls the NetOps Portal host for threshold events on devices that are modeled in Spectrum. The alarms appear in Spectrum one to two minutes after the events occur in NetOps Portal.

For more information about how to configure (enable) event polling in Spectrum, see [the DX NetOps Spectrum documentation](#).

### **TIP**

[Customize event integration](#).

Spectrum starts polling events on the next poll cycle.

## **Send Traps to Spectrum**

Use the following process to send traps to Spectrum:

1. Create a matching SNMP profile in Spectrum and in DX NetOps Performance Management.
2. Map the traps from a device to specific Spectrum events using the MIB Tools application in OneClick.
3. Complete event customization using Event Configuration. In Event Configuration, define event processing rules, create the event message to display to users, and set other parameters.

For more information, see [the DX NetOps Spectrum documentation](#).

## **Review Multi-Tenancy and Spectrum**

[Review multi-tenancy and Spectrum](#).

# **Integrate with DX Operational Intelligence**

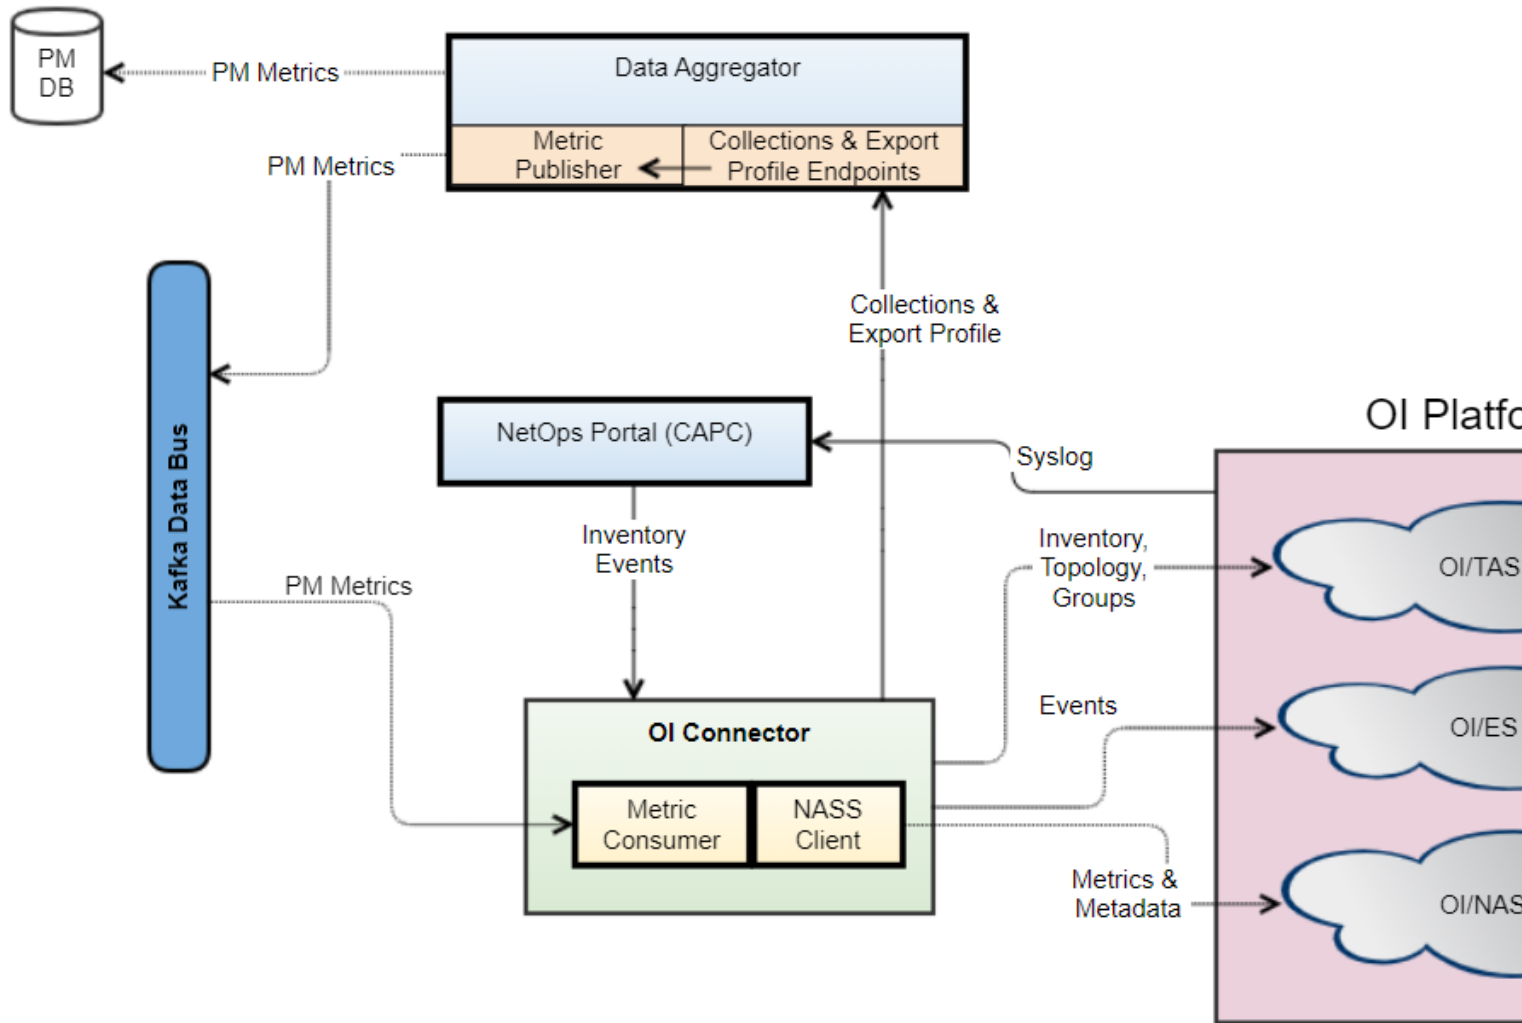
Integrate DX NetOps Performance Management with DX Operational Intelligence by way of the OI Connector. The OI Connector collects the metrics, events, inventory, and topology data and displays it in DX Operational Intelligence.

By integrating with DX Operational Intelligence, your enterprise can do the following:

- Validate application performance that is delivered over the network by applying network capabilities and service quality metrics.
- Solve problems faster by validating the impact of changes.

For more information about how to integrate with DX Operational Intelligence by way of the OI Connector, see [Install and Upgrade the OI Connector](#).

The following diagram explains the data flow from DX NetOps Performance Management to DX Operational Intelligence by way of the OI Connector:



### **DX NetOps Performance Management Metric Collection**

The OI Connector communicates with DX NetOps Performance Management and specifies the inventory and metrics that it should export by way of the Apache Kafka data bus. Kafka is a messaging protocol that is based on the publish/subscribe model.

### **OI Connector Service**

The OI Connector includes the OI Connector service (`caperfcenter_oiconnector`). This service sends inventory, events, group definitions, and DX NetOps Performance Management metrics to DX Operational Intelligence.

#### **NOTE**

You can manage (stop, start, restart) this service using the OI Connector wrapper service (`caperfcenter_oi`).

For more information, see [Configure the OI Connector](#).

## The Kafka Service

The OI Connector installation prompts for information about one or more installed Kafka brokers. DX NetOps Performance Management includes NetOps Kafka. With NetOps Kafka, you can install either a Kafka cluster of brokers or a single node broker. For the purposes of the OI Connector, a single Kafka broker is sufficient. You can install this broker co-located on the same node as the OI Connector.

For more information about how to install NetOps Kafka, see [Install NetOps Kafka](#).

## OI Connector (DX Operational Intelligence) Release Notes

The OI Connector includes the following enhancements and new features.

The following table includes the revision history for the previous OI Connector releases:

OI Connector Version	What's New
OI Connector 3.10.0	<ul style="list-style-type: none"> <li>This release includes third-party library upgrades.</li> <li>The installer failure issue using Java 11.0.20 has been resolved.</li> </ul>
OI Connector 3.4.0-3.9.0	<ul style="list-style-type: none"> <li>These releases include third-party library upgrades.</li> </ul>
OI Connector 3.3.0	<ul style="list-style-type: none"> <li>This release includes third-party library upgrades to Spring, jackson, Jetty, and Apache Commons.</li> </ul>
OI Connector 3.2.0	<ul style="list-style-type: none"> <li>This release includes third-party library upgrades to Spring and Jetty.</li> <li>The OI Connector installer now uses the kafka-installer merge module. This module tests to make sure Kafka is accessible. Kafka is no longer included with the OI Connector. You must download and install NetOps Kafka or provide your own Kafka instance.</li> <li>The OI Connector compiled code is now Java Runtime Environment 11 (JRE) 11.</li> <li>You can now create the topic that the OI Connector uses to send log events to NetOps Portal using the NetOps Kafka software development kit (SDK).</li> </ul>
OI Connector 3.1.0	<ul style="list-style-type: none"> <li>This release includes third-party and internal library upgrades to address Black Duck vulnerabilities.</li> <li>The OI Connector installer no longer provides a Kafka installation. As a prerequisite to installing NetOps Kafka, Kafka must be installed. You can install Kafka using the NetOps Kafka package downloaded from the Broadcom Support site, or using an Apache-downloaded package.</li> </ul>
OI Connector 3.0.0	<ul style="list-style-type: none"> <li>This release no longer includes the NFA/ADA integration. Starting in this release, the OI Connector does not harvest and send NFA and ADA metrics.</li> <li>The OI Connector now uses version 1.4.49 of the supportability jar.</li> <li>The OI Connector now attempts to ingest metrics into NASS using the Reactive NASS Client. If React support is not available on the OI server, the OI Connector now falls back to REST ingestion.</li> <li>This release includes Spring and other third-party library upgrades to address vulnerabilities.</li> <li>To prevent the loss of metrics on restart, the OI Connector now delays the consumption of metrics from the Kafka data bus until the first DX NetOps Performance Management inventory harvest has been completed. This can result in a longer delay between when the OI Connector is launched and metrics start flowing to DX Operational Intelligence.</li> <li>The heap requirements have been reduced. The OI Connector memory utilization can now successfully run the OI Connector with 16 GB heap.</li> </ul>
OI Connector 2.1.8	<ul style="list-style-type: none"> <li>This release includes third-party upgrades to Spring and Jackson to address vulnerabilities.</li> </ul>
OI Connector 2.1.7	<ul style="list-style-type: none"> <li>This release includes third-party upgrades to jackson and bc-fips.</li> <li>The OI Connector now requires Java Runtime Environment 11 (JRE) 11.</li> <li>The OI Connector now guards against non-person entity (NPE) when an interface has no parent global device id.</li> </ul>



OI Connector Version	What's New
OI Connector 2.1.6	<ul style="list-style-type: none"> <li>You can now receive events of type syslog (SysLog events) from DX Operational Intelligence using log analytics for Insights. For more information, see <a href="#">Install and Configure Log Analytics for Insights</a>.</li> <li>The OI Connector now uses log4j2.</li> </ul>
OI Connector 2.1.5	<ul style="list-style-type: none"> <li>This release includes third-party upgrades to log4j2 and RestEasy.</li> </ul>
OI Connector 2.1.4	<ul style="list-style-type: none"> <li>This release includes Apache Log4J vulnerability mitigation.</li> <li>This release includes defect fixes.</li> </ul>
OI Connector 2.1.3	<ul style="list-style-type: none"> <li>This release includes encrypted sensitive data in configuration files.</li> <li>The OI Connector now contributes health status to DX Operational Intelligence.</li> <li>You can now enable syslog monitoring by integrating syslog with NetOps Portal (Requires DX Operational Intelligence Log Analytics). For more information about log analytics, see Log Analytics in <a href="#">the DX Operational Intelligence documentation</a>.</li> <li>You can now download the OI Connector from DX Operational Intelligence.</li> </ul>
OI Connector 2.1.2	<ul style="list-style-type: none"> <li>This release includes an authenticated connection to NetOps Portal and the data aggregator APIs.</li> </ul>
OI Connector 2.1.1	<ul style="list-style-type: none"> <li>The OI Connector now includes a Kafka data bus for DX NetOps Performance Management metric harvesting. For more information, see <a href="#">Integrate with DX Operational Intelligence</a>.</li> <li>This release includes performance and scale enhancements.</li> </ul>

## Install and Upgrade the OI Connector

Integrate DX NetOps Performance Management with DX Operational Intelligence using the OI Connector.

Use the following process to install the OI Connector:

1. [Verify the Prerequisites](#)
2. [Install the OI Connector](#)
3. [Verify the Installation](#)
4. [Next steps](#)

To upgrade the OI Connector, see [Upgrade the OI Connector](#).

### Verify the Prerequisites

Before you install the OI Connector, ensure that you have met the following prerequisites:

- You know which OI Connector version is compatible with the DX NetOps Performance Management version that you have installed. This article includes information about how to install or upgrade to OI Connector 3.10.0. For more information:
  - About the OI Connector version that DX NetOps Performance Management includes, see [DX Operational Intelligence Interoperability](#).
  - About the latest version of the OI Connector, see [OI Connector \(DX Operational Intelligence\) Release Notes](#).
- You have installed Java Runtime Environment 11 (JRE) 11 on the node/device on which you are installing the OI Connector.

#### NOTE

(3.2.0 - 3.9.0) To avoid the installer from failing to launch with the following error, install the OI Connector using Java 11.0.19 or lower:

```
# ./OIConnectorInstaller.bin
Preparing to install
```

```

Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...
Launching installer...
Error: Could not find or load main class com.zerog.lax.LAX
Caused by: java.lang.ClassNotFoundException: com.zerog.lax.LAX that the connector cannot be
installed using Java 11.0.20 (you can install the OI Connector using a Java version below Java
11.0.20). Use an older Java 11 version to install the OI Connector.

```

After the OI Connector is installed, you can *run* the OI Connector with any version of Java 11, including 11.0.20.

- You have installed NetOps Kafka, with an adequate log retention period and sufficient disk space for log retention.

#### NOTE

During the installation of NetOps Kafka, you configure log retention. One or two(hours) is adequate as your retention period. The Kafka broker requires a maximum of 60 GB of disk space for each hour of log retention on a single-node NetOps Kafka cluster, based on the out-of-the-box values (metric families) in the `<OIConnector_installation_directory>/config.xml` configuration file.

- ***OIConnector\_installation\_directory***

The installation directory for the OI Connector.

**Default:** `/opt/CA`

For more information:

- About how to install NetOps Kafka, see [Install NetOps Kafka](#).
- About the configuration file, see [Configure the OI Connector](#).
- You have enabled the communication between DX NetOps Performance Management and DX Operational Intelligence (metric export) by [enabling and configuring a Kafka export producer on the data aggregator](#).

#### NOTE

(In fault-tolerant data aggregator environments) Enable and configure a Kafka export producer on each data aggregator.

When configured, the data aggregator harvests and publishes DX NetOps Performance Management metrics to the Kafka data bus using the Kafka export producer. The OI Connector consumes the metrics from the Kafka data bus, and forwards these metrics to DX Operational Intelligence.

- You have verified the sizing guidelines for the host that is running the OI Connector:

- **Disk:** 40 GB
- **RAM:**
  - **Small (100,000 polled items or less):** 12 GB
  - **Medium (500,000 polled items or less):** 16 GB
  - **Large (more than 500,000 polled items):** 28 GB

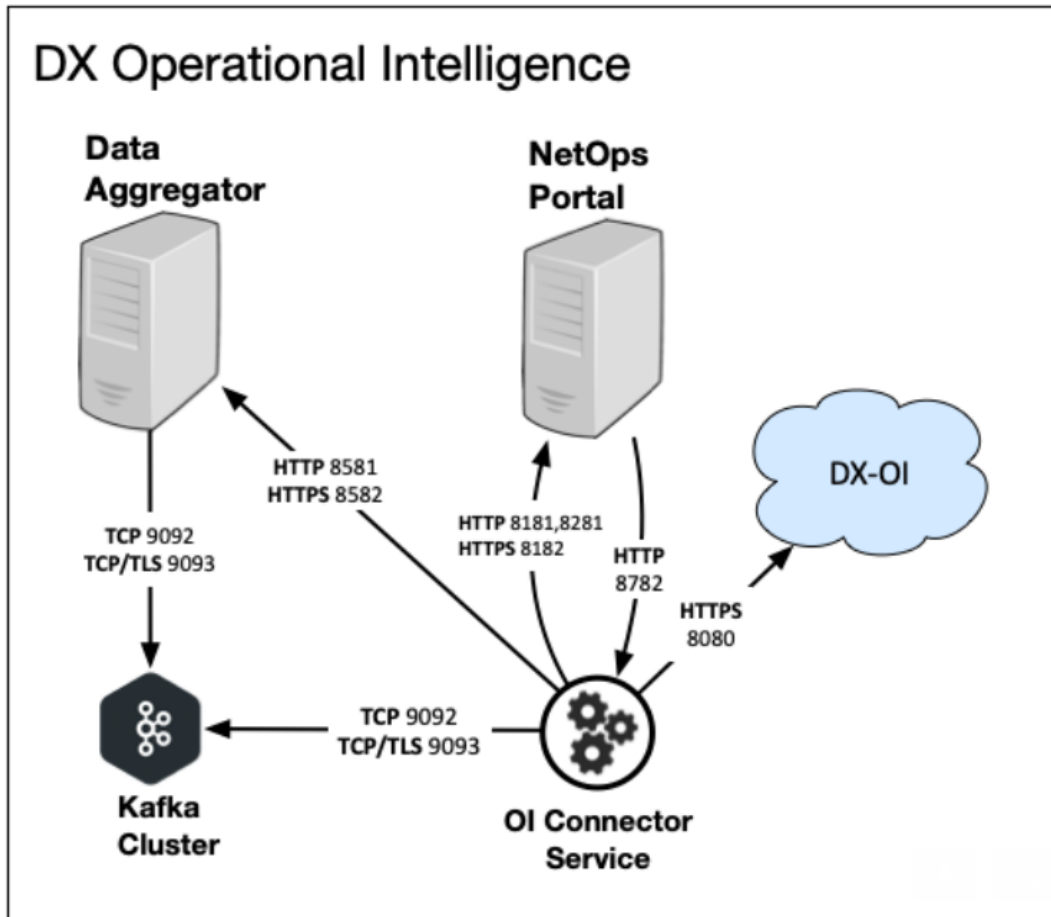
#### TIP

You can get the number of polled items from NetOps Portal (hover over **Inventory**, and then click **Interfaces**) and then look at the value for the total amount of inventory.

- **CPU:** 8 cores
- You have created an Administrator user account that is specific to the OI Connector, and you have set the password to be non-expiring (you have disabled password expiration for this user account).

For more information about how to disable password expiration for user accounts:

- Using NetOps Portal, see [Configure the Password Security Settings](#).
- Using the SSO Configuration tool (SsoConfig), see [Configure the Password Security Settings Using the SSO Configuration Tool](#).
- You have verified that you have configured your environment to allow communication between the various components of the integration per the ports shown in the following diagram:



This diagram shows the default port configurations. If you have modified the port configurations, use your modified values per your environment:

Function	Protocol and Port	Clients
Kafka	TCP 9092, TCP/TLS 9093 Enables communication between the data aggregator, the OI Connector, and Kafka.	Data aggregator, OI Connector
Data aggregator	HTTP 8581/HTTPS 8582 Enables REST communication from the OI Connector to the data aggregator for configuration of metric export.	OI Connector service
NetOps Portal	HTTP 8181/HTTPS 8182 Enables REST communication from the OI Connector to NetOps Portal.	OI Connector service
NetOps Portal	HTTP 8281 Enables communication from the OI Connector to the Event Manager for event harvesting.	OI Connector service

Function	Protocol and Port	Clients
OI Connector service	HTTP 8782 Enables the OI Connector to report its status to NetOps Portal. You can view this status on the <b>System Status</b> and <b>Manage Connector to DX Operational Intelligence</b> pages. <b>Tip:</b> For more information: <ul style="list-style-type: none"> <li>About the status on the <b>System Status</b> page, see <a href="#">View the Health of the System</a>.</li> <li>About the status on the <b>Manage Connector to DX Operational Intelligence</b> page, see <a href="#">Install and Configure Log Analytics for Insights</a>.</li> </ul>	NetOps Portal
DX Operational Intelligence	HTTPS 8080 Enables communication between the OI Connector service and DX Operational Intelligence.	OI Connector service

- You are using a dedicated OI Connector server. No other DX NetOps Performance Management component is sharing this server.

## Install the OI Connector

### Follow these steps:

- Download the OI Connector for NetOps Performance Statistics package (NetOps-OIConnector-`<latest_version>-Linux-RELEASE.tar.gz` ) from the [Broadcom Support site](#) to the node/device on which you are installing the OI Connector.  
**IMPORTANT**  
 Install the OI Connector on a separate node than the other DX NetOps Performance Management components.  
 The OI Connector installer is downloaded.
- Unpack the OI Connector installer by issuing the following command:  

```
tar xpvfz NetOps-OIConnector-<latest_version>-Linux-RELEASE.tar.gz
```

 The OI Connector installer is unpacked.
- Launch OI Connector installer by issuing the following command:  

```
./OIConnector.bin
```

 The OI Connector installer launches.
- At the prompt, press the **Enter** key on your keyboard to continue.  
**TIP**  
 If at any time during the installation you need to return to the previous step, enter `back` .
- When prompted to accept the terms of the license agreement, to proceed further, enter **Y** for Yes.
- At the **Choose Java Virtual Machine** prompt, select the Java Virtual Machine (JVM) for the OI Connector use. Enter the option based on your Java installation, and then press the **Enter** key on your keyboard to accept the option.  
 The following image shows an example of these options:

```

Choose Java Virtual Machine
-----

Please Choose a Java VM for Use by this Application

->1- /usr/bin/java
   2- /opt/CA/jre/bin/java

   3- Choose a Java VM already installed on this system

ENTER THE NUMBER FOR THE JAVA VM, OR PRESS <ENTER> TO ACCEPT THE
CURRENT SELECTION: 1

```

7. At the **Choose Install Folder** prompt, define the location where you want the installation to create the `OIConnector` and `kafka` directories. Press the **Enter** key on your keyboard to accept the default installation directory (`/opt/CA`) for the OI Connector installation, or enter the absolute path of the OI Connector installation directory, for example, `<OIConnector_installation_directory>`.
    - **`OIConnector_installation_directory`**  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`
  8. At the **OI Connector to NetOps Portal (CAPC) Configuration** prompt, enter the following information:
 

**NOTE**  
These are the default values. You can customize them.

    - **Protocol (http or https) (Default: http):**  
Accept the default, or if HTTPS is enabled for NetOps Portal, enter **https**.
    - **Hostname (Default: localhost):**  
Specify the hostname for NetOps Portal. For example, `capchost`.
    - **Port (Default: 8181):**  
Specify the NetOps Portal port number. For example, if HTTPS is enabled for NetOps Portal, 8182.  
For more information about the required ports for secured connections, see [Review Installation Requirements and Considerations](#).
    - **Tenant (Default: `_default_`):**  
For a single tenant environment, accept the default tenant (`_default_`), or if you want to use a specific tenant, then enter the specific tenant name. For an OI Connector configured for multiple tenants (multi-tenant environment), enter the specific tenant name.

**TIP**  
If you have not already configured the OI Connector for multiple tenants, you can complete this step after completing this OI Connector installation process.  
For more information about how to configure for multiple tenants, see [Configure the OI Connector](#).

    - **User Name (Default: admin):**  
Enter the credentials of the Administrator user account that is specific to the OI Connector.
- The following image shows these options:

```
=====
OI Connector to NetOps Portal (CAPC) Configuration
-----
```

```
Enter the protocol (http/https), hostname, port, tenant and REST user
your NetOps Portal server. This is the route over which the connector
communicate with NetOps Portal.
```

```
Protocol (http or https) (Default: http):
```

```
Hostname (Default: localhost):
```

```
Port (Default: 8181):
```

```
Tenant (Default: _default_):
```

```
User Name (Default: admin):
```

9. At the **NetOps Portal (CAPC) - Password** prompt, enter the credentials of the Administrator user account that is specific to the OI Connector.
10. At the **Link Back to NetOps Portal (CAPC) Configuration**, enter *one* of the following options, and then press the **Enter** key on your keyboard to continue:
  - **1 - No:** If the NetOps Portal URL is publicly available and you do not need to configure a public route.
  - **2 - Yes:** If the NetOps Portal URL is *not* publicly available and you need to configure a public route for those links back from DX Operational Intelligence to NetOps Portal. When prompted, enter the public route.

The following image shows these options:

```
=====
Link Back to NetOps Portal (CAPC) Configuration
=====
```

```
You have specified the following route to NetOps Portal in a previous step:
http:\\localhost:8181. If this URL is not publicly accessible, then you need
to configure a public route for links back to NetOps Portal by selecting
option 2 (Yes) below.
```

```
Do you need to configure a public route for links back from OI to NetOps
Portal?
```

```
->1- No
    2- Yes
```

```
ENTER THE NUMBER OF THE DESIRED CHOICE, OR PRESS <ENTER> TO ACCEPT THE
DEFAULT: █
```

11. At the **OI Gateway Connection** prompt, enter the following information:

- **Protocol (Default: https):**  
Specify the DX Operational Intelligence Gateway protocol for the metrics and topology (TAS) endpoint or NASS endpoint. The protocol is based on your implementation.  
**Options:** HTTP or HTTPS  
**TIP**  
The protocol is dependent on your OpenShift or Kubernetes implementation.
- **Hostname:**  
Specify the fully qualified DX Operational Intelligence Gateway hostname for the TAS endpoint or NASS endpoint.  
**TIP**  
You can get this information by issuing the command:
  - (OpenShift)
 

```
oc get routes -n<your_nameSpace> | egrep apmservices-gateway
```
  - (Kubernetes)
 

```
kubect1 get ingress -n<your_nameSpace> | grep apmservices-gateway
```
- **Port (Default: 80):**  
Specify the DX Operational Intelligence Gateway port number for the TAS endpoint or NASS endpoint, or press the **Enter** key on your keyboard to accept the default port. The port is based on your implementation.  
For example, for HTTP, 80, and for HTTPS, 443.  
**TIP**  
The port is either 80 or 443, depending on your OpenShift or Kubernetes implementation.
- **Tenant ID:**  
**Example:** AD07A694-748F-47EE-A0C2-343311E24A30 Specify the DX Operational Intelligence **tenant ID**.  
**TIP**  
You can get the tenant ID by completing the following:
  1. Obtain the Elastic ending by issuing the following command based on your implementation:
    - (OpenShift)

```
oc get routes -n<your namespace> | egrep jarvis-es
```

– (Kubernetes)

```
kubectl get ingress -n<your namespace> | egrep jarvis-es
```

2. Run a query against the elastic endpoint to list all tenants:

```
http(s)://elastic-endpoint/ao_dxi_tenants_1_1/_search?size=200&pretty
```

3. Locate your tenant name and extract the value of the associated `tenant_id` parameter.

– **Security Token:**

Specify the generated DX Operational Intelligence Gateway agent or tenant token.

**TIP**

You can get this information by completing the following based on your implementation:

- If you have installed *only* DX Operational Intelligence, log in to the Cluster Management console using the master admin account, and then generate a tenant token.  
For more information about this console and how to generate tokens, see the [DX Platform documentation](#).
- If you have installed DX Application Performance Management *and* DX Operational Intelligence, generate the token using one of the following methods:
  - In DX Application Performance Management, generate a security token.  
For more information, see the [DX Application Performance Management documentation](#).
  - In DX Operational Intelligence, from the Cluster Management console, generate a tenant token.

The following image shows these options:

12. At the **OIConnector Status Webservice** prompt, enter the following information:

**Port (Default: 8782):**

Enter the port for the communication from NetOps Portal to the OI Connector service so that it can query and report the DX Operational Intelligence overall system status on the **System Status** page. Accept the default port (8782) or enter the port on which to run the OI Connector webservice.

For more information about this status, see [View System Status](#).

The following image shows these options:

```
=====
OIConnector Status Webservice
=====

Enter the port for the OIConnector Status webservice.

Port (Default: 8782):
```

13. Set up the connection to NetOps Kafka by completing the following steps:

- a. At the **Number of Kafka Brokers** prompt, enter the following information, and then press the **Enter** key on your keyboard to continue:

**Number of Kafka Brokers (Default: 1):**

Enter the number of Kafka brokers for the cluster.

**Default:** 1

The following image shows this option:

- b. At the **Kafka Port** prompt, enter the following information, and then press the **Enter** key on your keyboard to continue:

**Kafka Port (Default: 9092):**

Enter the client port that the OI Connector will use to connect to the Kafka brokers.



**Default:** 9092

The following image shows this option:

- c. At the **Kafka Broker** prompt, enter the following information, and then press the **Enter** key on your keyboard to continue:

**Kafka Broker 1 :**

Specify the hostname for each Kafka broker in the cluster.

The following image shows this option:

- d. At the **Kafka SSL Configuration** prompt, enter the following information, and then press the **Enter** key on your keyboard to continue:

**Is Kafka Configured with SSL (Default: N):**

Specify whether Kafka is configured with SSL (secured).

**Default:** N

The following image shows this option:

- e. (If you chose that Kafka is configured with SSL) At the **Kafka Client Keystore Location** prompt, enter the following information, and then press the **Enter** key on your keyboard to continue:

**Kafka Client Keystore Location :**

Enter the location of the Kafka client keystore.

The following image shows this option:

- f. (If you chose that Kafka is configured with SSL) At the **Kafka Client Keystore Password** prompt, enter the following information, and then press the **Enter** key on your keyboard to continue:

**Please Enter the Password:**

Enter the password of the Kafka client keystore.

The following image shows this option:

The installer tests the connection to the cluster.

14. At the **Local http/https Proxy Configuration** prompt, enter *one* the following options, and then press the **Enter** key on your keyboard to continue:
- If your network requires communication through a local proxy, enter **2** for Yes, and then configure the local proxy (next step).
  - If your network *does not* require communication through a local proxy, enter **1** for No.
- The following image shows these options:

15. (If you have chosen to configure a local proxy) At the **Configure the Proxy** prompt, enter the following information:

- **Proxy Protocol:**

The communication protocol for the proxy.

**Options:** http or https

**Default:** http

- **Proxy Host:**

For example, 10.17.100.3 .

- **Proxy Port:**

For example, 9091 , 9093 . If you are securing the Kafka data bus, 9093 .

- **Proxy Username:**

Specify the username.

- **Proxy Password:**

Specify the password.

- **Non-proxy Hosts:**

Enter the details of the non-proxy hosts. Add localhost and DX NetOps Performance Management hostname to the non-proxy list.

The following image shows these options:

The OI Connector is installed, and the OI Connector service is started automatically.

## **Verify the Installation**

### **Follow these steps:**

1. Verify that the services are running by issuing the following commands:  

```
service caperfcenter_oiconnector status
```

  - **caperfcenter\_oiconnector**  
The core service of the OI Connector (the OI Connector service).
2. Verify that the DX Operational Intelligence overall system status is "Normal" (green) on the **System Status** page.
3. Verify that event queries are issued to DX NetOps Performance Management every minute from the OI Connector logs, and then tail the `OIConnector` log.

## **Next Steps**

After you have installed the OI Connector, [secure connections for the OI Connector](#).

## **Upgrade the OI Connector**

**Prerequisite:** (If you are upgrading from OI Connector 1.4 or lower) Delete the **OI Inventory Group** group that is located in the **OI Integration Group** group. The OI Connector recreates this group with the new structures.

1. Create a backup of your existing configuration files to a directory outside of the installation by issuing the following commands:

```
cp -R <OIConnector_installation_directory>/OIConnector/  
conf <OIConnector_installation_backup_directory>/connector-bkup
```

- **OIConnector\_installation\_directory**  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`
- **OIConnector\_installation\_backup\_directory**  
The backup directory for the OI Connector.

A backup of your existing configuration files are created. These files are available to you if you need to refer to them or if you need to re-apply the customizations to your configuration after upgrading to the new version of the OI Connector.

2. Uninstall the existing OI Connector by issuing the following command:

```
<OIConnector_installation_directory>/uninstall/Uninstall_caperfcenter_oiconnector
```

3. Install the new version of the OI Connector.

## **Secure Connections for the OI Connector**

Secure the communication between NetOps Portal and the OI Connector and between the OI Connector and DX Operational Intelligence. Secure access to DX Operational Intelligence over Hypertext Transfer Protocol Secure (HTTPS) by adding a certificates.

### **NOTE**

The following procedures use the `keytool` and `openssl` commands. The `keytool` command is part of Java.

For more information:

- About `keytool`, see [the Java documentation on the Oracle website](#).
- About `openssl`, see [the OpenSSL documentation](#).

#### NOTE

Store the certificate and private key files, such as \*.pem, \*.cer, \*.crt, \*.key files, that are referenced in configuration files during this process in a secure location. If the certificate and private key files are temporary files that are not referenced in configuration files after this process is complete, move or delete them.

#### Follow these steps:

- Export the certificates from the following:
  - (If DX Operational Intelligence Gateway is running in HTTPS (port 443)) Export the certificates from the DX Operational Intelligence Gateway. Go to the OpenShift or Kubernetes master server, and then issue the following command:
 

```
openssl s_client -connect apmgateway_Endpoint:<port> < /dev/null | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > tas.cer
```
  - (If NetOps Portal is running in HTTPS (port 443)) Export the certificates from the NetOps Portal server. Go to the NetOps Portal server, and then issue the following command:
 

```
openssl s_client -connect capc_hostname:<port> < /dev/null | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > capc.cer
```
- Import the certificates from the following into the OI Connector:
  - The DX Operational Intelligence Gateway certificate, by issuing the following command:
 

```
keytool -importcert -alias <certificate_alias> -file <filename> -keystore /<truststore>
```

**Example:**

```
keytool -importcert -alias tas -file tas.cer -keystore <truststore>
```
  - The NetOps Portal server certificate, by issuing the following command:
 

```
keytool -importcert -alias <certificate_alias> -file <filename> -keystore /<truststore>
```

**Example:**

```
keytool -importcert -alias tomcatssl -file capc.cer -keystore <truststore>
```
  - **certificate\_alias**  
Specify the alias of the certificate in the Java truststore.  
**Example:** tas
  - **filename**  
Specify the file for the certificate. Use a full pathname to the certificate.  
**Examples:**
    - **DX Operational Intelligence Gateway certificate:** /tmp/tas.cer
    - **NetOps Portal server certificate:** /tmp/capc.cer
  - **truststore**  
Specify the Java truststore location.  
**Default:** /opt/ca/OIConnector/lib/security/cacert
- Restart the OI Connector.  
For more information, see [Configure the OI Connector](#).

## Configure the OI Connector

Configure the OI Connector based on your requirement.

In this topic:

- [Check the OI Connector Version](#)
- [Configure the Java Heap Size for the OI Connector](#)
- [Manage the OI Connector Service](#)
- [Manage the Kafka Service](#)
- [Restrict the Data that the OI Connector Sends](#)
- [Configure the NetOps Portal-to-DX Operational Intelligence Tenant Mappings](#)
- [Configure the Disposition of Syslog Events](#)
- [Update the NetOps Portal Credentials](#)
- [Update the Local Proxy Configuration](#)
- [Configure Additional Metric Families](#)

### **Check the OI Connector Version**

If you are not sure which OI Connector version you have installed, you can check the version.

#### **Follow these steps:**

1. Open the `<OIConnector_installation_directory>/OIConnector/version.properties` file.
  - ***OIConnector\_installation\_directory***  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`
2. Look for the `oi.version` entry, for example `oi.version=2.1`.

### **Configure the Java Heap Size for the OI Connector**

The OI Connector expects a max scale of 100,000 polled items or less (small). You can configure the heap size based on the scale.

#### **Follow these steps:**

1. Open the `<OIConnector_installation_directory>/OIConnector/conf/wrapper.conf` file.
  - ***OIConnector\_installation\_directory***  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`
2. Adjust the heap settings (the `property wrapper.java.maxmemory` property):  
Use the following scale values:
  - **Small (100,000 polled items or less):** 8 GB
  - **Medium (500,000 polled items or less):** 12 GB
  - **Large (more than 500,000 polled items):** 24 GB

The Java heap size for the OI Connector is configured.

### **Manage the OI Connector Service**

You can manage (start, restart, stop) the OI Connector service (`caperfcenter_oiconnector`) using the OI Connector wrapper service (`caperfcenter_oi`). You start or restart the services after making changes to the OI Connector configuration.

You can do the following:

- **Start the service**  
`service caperfcenter_oi start`
- **Restart the service**

```
service caperfcenter_oi restart
```

- **Stop the service**

```
service caperfcenter_oi stop
```

## **Manage the Kafka Service**

If you chose to install the single Kafka broker while installing the OI Connector, you can manage (stop, start, restart, retrieve status) this broker and its required Kafka and Zookeeper services using the `kafka` service. This service is located in the `<OIConnector_installation_directory>/kafka` directory.

- **`OIConnector_installation_directory`**

The installation directory for the OI Connector.

**Default:** `/opt/CA`

You can do the following:

- **Retrieve the status of the service**

```
service kafka status
```

- **Start the service**

```
service kafka start
```

- **Stop the service**

```
service kafka stop
```

## **Restrict the Data that the OI Connector Sends**

The OI Connector installation defines an **OI Seed Group** for each NetOps Portal tenant that is mapped to a DX Operational Intelligence tenant. The OI Connector determines the set of inventory that it sends to DX Operational Intelligence using this group. For the Default Tenant, this group is located in the `All Groups\OI Integration\OI Seed Group` group path. For a specific NetOps Portal tenant, this group is located under the `Tenants\<TenantName>` node of the group tree.

By default, the **OI Seed Group** includes rules that populate it with the devices that NetOps Portal monitors within the tenant. You can restrict the inventory that the OI Connector sends by editing these rules.

The following image illustrates a group rule that collects router device types:

Rules				
<span>New Rule</span> <span>Edit Rule</span> <span>Remove Rule</span> <span>Run Rules</span> <span>Preview Results</span> <span>Cancel</span> <span>Save</span>				
Rule	Item Type	Condition	Operator	Value
Add Devices	Device	Device Item	is a member of	Routers

For more information about how to edit group rules, see [Manage Group Rules](#).

## **Configure the NetOps Portal-to-DX Operational Intelligence Tenant Mappings**

The initial OI Connector installation sets up a single NetOps Portal-to-DX Operational Intelligence tenant mapping. This mapping maps the Default Tenant within NetOps Portal to the DX Operational Intelligence tenant specified during the OI Connector installation. You can configure additional NetOps Portal tenant-to-DX Operational Intelligence tenant mappings.

The following tenant configuration files specify the mappings:

- `<OIConnector_installation_directory>/conf/tenants.properties`  
Defines the NetOps Portal-to-DX Operational Intelligence tenant mappings.

- ***OIConnector\_installation\_directory***  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`
- ***<OIConnector\_installation\_directory>/conf/tenants-token.properties***  
Defines the DX Operational Intelligence tenant-to-DX Operational Intelligence Gateway token mappings.
  - ***OIConnector\_installation\_directory***  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`

Add or change the mappings by modifying these files.

#### NOTE

NetOps Portal stores the information in the tenant configuration files in encrypted form for security. You can obtain the encrypted form of the DX Operational Intelligence Tenant ID (the `Cohort ID`) or of the DX Operational Intelligence Gateway token using the `encryptor.sh` utility.

#### Example: Obtain the encrypted form of the DX Operational Intelligence Tenant ID

Issue the following command:

```
<OIConnector_installation_directory>/bin/encryptor.sh <PlainText_OI_Tenant_ID>
```

- ***OIConnector\_installation\_directory***  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`

#### Example: Obtain the encrypted form of the DX Operational Intelligence Gateway token

Issue the following command:

```
/opt/CA/bin/encryptor.sh <PlainText_Gateway-Token>
```

#### Follow these steps:

1. Do the following:
  - a. Add or change the mapping in the `tenants.properties` tenant configuration file. Add a line for each of the NetOps Portal tenant-to-DX Operational Intelligence tenant mapping configurations with the following format:
 

```
<NetOps Portal_Tenant_Name>=<OI_Tenant_ID>
```

    - ***NetOps Portal\_Tenant\_Name***  
The tenant name as it is defined in NetOps Portal. Escape space, =, or : characters with a preceding \ character. For example, specify the tenant `My Tenant:X` as `My\ Tenant\:X`.
    - ***OI\_Tenant\_ID***  
The encrypted form of the DX Operational Intelligence Cohort ID.
  - b. Add or change the mapping in the `tenants-token.properties` tenant configuration file. Add a line for each of the DX Operational Intelligence tenants-to-DX Operational Intelligence Gateway token mapping configuration with the following format:
 

```
<OI_Tenant_ID>=<Gateway-Token>
```

    - ***OI\_Tenant\_ID***  
The encrypted form of the DX Operational Intelligence Cohort ID. Escape = characters with a preceding \ character. For example, specify the `XfGh3rft==` encrypted tenant ID as `XfGh3rft\=\=`.
    - ***Gateway-Token***  
The encrypted form of the DX Operational Intelligence Gateway token.
2. [Restart the OI Connector services.](#)

The NetOps Portal-to-DX Operational Intelligence tenant mappings are configured.

## Configure the Disposition of Syslog Events

**Prerequisite:** To receive events of type syslog (Syslog events) from DX Operational Intelligence, log analytics for Insights must be installed and configured. For more information, see [Install and Configure Log Analytics for Insights](#).

You configure Syslog pattern match rules in DX Operational Intelligence. These rules create notifications that the OI Connector polls. The OI Connector creates Syslog events from these notifications and sends them to the event manager in NetOps Portal. You monitor Syslog events for specified log patterns in NetOps Portal.

By default, triggered rules result in events of type syslog (Syslog events) and of subtype **Syslog Info**, which do not result in alarms. You can configure the OI Connector to generate Syslog events that generate alarms (Syslog events of subtype **Syslog Alarm**), or you can configure it to ignore the notifications it receives from DX Operational Intelligence. You can configure the disposition of Syslog events for each Syslog pattern match rule.

For more information:

- About log analytics for Insights, see [Insights](#).
- About the Syslog event type, see [Event Types](#).
- About how to define and configure log patterns, see Log Analytics in [the DX Operational Intelligence documentation](#).
- About how to monitor Syslog events, see [Use Log Analytics for Insights](#).

### Follow these steps:

1. Edit the `<OIConnector_installation_directory>/conf/syslog_rules.properties` file.

- **`OIConnector_installation_directory`**  
The installation directory for the OI Connector.  
**Default:** `/opt/CA`

2. Add a line for each rule as follows:

```
<rule_name>=ALARM|IGNORE
```

- **`rule_name`**  
The name of the pattern match rule defined in DX Operations Intelligence. Whitespace must be escaped with a backslash `\` character.

#### Example 1: Generate Alarms

The following line specifies that Syslog events generated by the Security Violation rule generate an alarm:

```
Security\ Violation=ALARM
```

#### Example 2: Ignore Notifications

The following line specifies that the OI Connector ignore notifications generated by the Minor Alert rule:

```
Minor\ Alert=IGNORE
```

3. [Restart the OI Connector](#).

The disposition of Syslog events is configured.

## Update the NetOps Portal Credentials

The OI Connector requires credentials for a (non-SAML) locally-defined user with Administrator role rights. NetOps Portal stores these credentials in encrypted form in the `<OIConnector_installation_directory>/config.xml` configuration file at the time that you install the OI Connector. If these credentials change, you must either re-install the OI Connector or update them. The following procedure explains how to update the credentials.

### Follow these steps:

1. Obtain the Base64-encoded form of the credentials. On Linux, use the following command:

```
echo -n "<username>:<password>" | base64 -w 0
```

- **`username`**

The username for the proxy server.

- **password**

The password for the proxy server.

2. Encrypt the Base64-encoded form of the credentials by issuing the following command:

```
<OIConnector_installation_directory>/bin/encryptor.sh <base64_encoded_credentials>
```

- **OIConnector\_installation\_directory**

The installation directory for the OI Connector.

**Default:** /opt/CA

3. Add the encrypted credentials into the `<encrypted_credentials>` section of the `config.xml` file:

```
<!-- Connection for NetOps Portal -->
<bean id="pcConnection"
    class="com.ca.im.oinet.connector.sources.PCDataConnection">
    <constructor-arg index="0" value="CAPC" />
    <constructor-arg index="1" value="http" />
    <constructor-arg index="2" value="localhost" />
    <constructor-arg index="3" value="8181" />
    <!-- Change this as needed to support login credentials -->
    <constructor-arg index="4" value="<encrypted_credentials>" />
</bean>
```

4. [Restart the OI Connector services.](#)

The NetOps Portal credentials are updated.

### Update the Local Proxy Configuration

NetOps Portal store the configuration of a local proxy in the `<OIConnector_installation_directory>/conf/java.conf` file at the time that you install the OI Connector. You can modify the configuration of the proxy server by modifying this file.

- **OIConnector\_installation\_directory**

The installation directory for the OI Connector.

**Default:** /opt/CA

The following example shows the proxy configuration options in the `java.conf` file:

```
#encoding=UTF-8

# Additional java parameters may be configured here to be used by all services

# HTTP Proxy
-Dhttp.proxyHost=myproxyhost -Dhttp.proxyPort=80 -Dhttp.proxyCredentials=iI+w3CwGdxZjnFpVpGoTtJeWSdf64xrVo4b/srfX4W4=

# Non-Proxyed Hosts
-Dhttp.nonProxyHosts=localhost
```

- **proxyHost**

The hostname or IP address of the proxy server.

- **proxyPort**



The port on which the proxy server runs.

- **proxyCredentials**

(Optional) Encrypted proxy credentials.

- **nonProxyHosts**

Pipe-delimited list of hosts (you can also use wildcards (\*)) that skip the proxy. Specify the localhost, NetOps Portal host, and the data aggregator hosts. Hosts not listed in this set will go through the proxy.

If the credentials required for the proxy server change, re-install the OI Connector or update the credentials in the `java.conf` file. The following procedure explains how to update the credentials.

#### Follow these steps:

1. Obtain the Base64-encoded form of the credentials by issuing the following command:

```
echo -n "username:password" | base64 -w 0
```

- **Username**

The username for the proxy server.

- **Password**

The password for the proxy server.

2. Encrypt the Base64-encoded form of the credentials by issuing the following command:

```
<OIConnector_installation_directory>/bin/encryptor.sh <base64_encoded_credentials>
```

3. Add the encrypted credentials into the `proxyCredentials` setting in the `java.conf` file.
4. [Restart the OI Connector services.](#)

The credentials are updated.

### Configure Additional Metric Families

You can define the metric families and specific metrics within those metric families that the OI Connector sends to DX Operational Intelligence by updating the values (adding metric families) in the `<OIConnector_installation_directory>/config.xml` configuration file.

- **OIConnector\_installation\_directory**

The installation directory for the OI Connector.

**Default:** `/opt/CA`

**Prerequisite:** You have ensured that the Kafka broker has sufficient disk space for the additional metric families. For more information, see [Install NetOps Kafka](#).

The OI Connector supports the following metric families out-of-the-box:

- `availability`
- `reachability`
- `port (interface)`
- `cpu`
- `memory`

#### NOTE

The OI Connector harvests all metrics within the metric families, except for the `port` metric family.

For more information about metric families, see [Manage Metric Families](#).

The following example shows the out-of-the-box values in the configuration file. It shows the five metric families, and the `port` metric family with seven metrics:

```
<!-- CAPM Metric Configuration -->
```

```

<bean id="pmMetricFamilyFilters"
class="com.ca.im.oinet.connector.sources.kafka.dao.PMMetricFamilyFilters">
  <constructor-arg index="0">
    <map value-type="java.lang.String">
      <entry key="NormalizedAvailabilityInfo">
        <ref bean="availability"/>
      </entry>
      <entry key="NormalizedReachabilityInfo">
        <ref bean="reachability"/>
      </entry>
      <entry key="NormalizedPortInfo">
        <ref bean="port"/>
      </entry>
      <entry key="NormalizedCPUInfo">
        <ref bean="cpu"/>
      </entry>
      <entry key="NormalizedMemoryInfo">
        <ref bean="memory"/>
      </entry>
    </map>
  </constructor-arg>
</bean>

<bean id="availability" class="com.ca.im.oinet.connector.sources.kafka.dao.PMMetricFamilyFilter">
  <constructor-arg index="0"><value>availability</value></constructor-arg>
  <constructor-arg index="1"><value>{http://im.ca.com/normalizer}</value></constructor-arg>
</bean>

<bean id="reachability" class="com.ca.im.oinet.connector.sources.kafka.dao.PMMetricFamilyFilter">
  <constructor-arg index="0"><value>reachability</value></constructor-arg>
  <constructor-arg index="1"><value>{http://im.ca.com/normalizer}</value></constructor-arg>
</bean>

<bean id="port" class="com.ca.im.oinet.connector.sources.kafka.dao.PMMetricFamilyFilter">
  <constructor-arg index="0"><value>port</value></constructor-arg>
  <constructor-arg index="1"><value>{http://im.ca.com/normalizer}</value></constructor-arg>
  <property name="metricIncludeFilters">
    <set>
      <value>Availability</value>
      <value>BitsIn</value>
      <value>BitsOut</value>
      <value>BitsPerSecondIn</value>
      <value>BitsPerSecondOut</value>
      <value>UtilizationIn</value>
      <value>UtilizationOut</value>
    </set>
  </property>
</bean>

<bean id="cpu" class="com.ca.im.oinet.connector.sources.kafka.dao.PMMetricFamilyFilter">
  <constructor-arg index="0"><value>cpu</value></constructor-arg>
  <constructor-arg index="1"><value>{http://im.ca.com/normalizer}</value></constructor-arg>
</bean>

```

```
<bean id="memory" class="com.ca.im.oinet.connector.sources.kafka.dao.PMMetricFamilyFilter">
  <constructor-arg index="0"><value>memory</value></constructor-arg>
  <constructor-arg index="1"><value>{http://im.ca.com/normalizer}</value></constructor-arg>
</bean>
```

## Uninstall the OI Connector

If you no longer need the OI Connector, you can uninstall it without affecting the DX NetOps Performance Management installation.

### Follow these steps:

1. Navigate to the `<installation_directory>/uninstall` directory.
2. Issue the following command:  

```
./Uninstall_caperfcenter_oiconnector
```
3. (Optional) Manually remove the **OI Integration Group** group from DX NetOps Performance Management.

## Integrate with Application Delivery Analysis

The integration with CA Application Delivery Analysis (ADA) provides end-to-end performance monitoring through dashboards and views. These dashboards and views show historically-normal performance for users and metrics that cross acceptable performance thresholds. ADA gathers troubleshooting information. You can determine the origin of an application, network, or server performance problem using this information.

Integrate with ADA by registering an ADA data source in NetOps Portal.

For more information:

- About how to register an ADA data source, see [Configure a Data Source](#).
- About CA Application Delivery Analysis, see [the CA Application Delivery Analysis documentation](#).

## ADA Metrics

When registered as a data source, CA Application Delivery Analysis views report on the following metrics:

### Active Sessions

Measures the number of active TCP sessions reported by a monitor feed that match an application/server/network combination on the management console.

### Byte Loss Percentage

Measures the ratio of retransmitted data to total data, percentage of data lost on the monitored network, and loss rate in bytes per second.

### Completed Sessions

Measures the number of sessions completed during a 5-minute monitoring period.

### Connection Setup Time

Measures the amount of time that it takes to establish a TCP session between the client and server before data transfer can begin.

---

### **Data Transfer Time**

A Combined metric that measures the time that it takes to transmit a complete application response from the first response (the end of the Server Response Time) to the last packet sent in that request. Data Transfer Time excludes the initial server response time and includes NRTT if there is no more data to send than fits in the TCP window. The response time can be impacted by the design of the application or the performance of the server or network. The CA Application Delivery Analysis management console does not open an incident when the Data Transfer Time threshold is crossed.

### **Effective Network Round Trip Time**

Is a Network metric that consists of Network Round Trip Time plus Retransmission Delay. Note that Retransmission Delay is not the delay due to any retransmissions; it is the average amount of retransmission delay per round trip. It is important to note that the management console adds two averages, and combines two metrics. The CA Application Delivery Analysis management console opens a Network incident when the incident threshold for this metric is crossed.

### **Estimated Client Transaction Time**

Provides an approximation of Total Transaction Time in the absence of a Client segment. This metric is the summation of the Server segment Transaction Time and the WAN segment Network Round Trip Time. This is not an “engineering” level metric, but it can provide an indicator of the order of magnitude of response time for a given location or time of day. This metric is only available with a CA Application Delivery Analysis 9.3 data source.

### **Expired Sessions**

Measures the number of TCP sessions where the CA ADA Monitor service did not see the TCP session tear down (FIN or RST packet). Sessions which are inactive for a period of time are cleared out of memory and marked as Expired. The management console classifies a session as Expired if it does not observe any packets in a 15-minute period. Too many expired sessions left open can cause servers to become unresponsive.

### **From Server Bytes**

Measures the number of bytes that a server sent to a client.

### **From Server Packets**

Measures the number of packets that a server sent to a client.

### **Network Connection Time (NCT or NSCT)**

Is a Network metric that measures the amount of time between the Syn-Ack sent by the server and the Ack received back from the client. When a network is uncongested, it is a measurement of network latency that represents the minimum latency due to distance and serialization, and is the best possible round trip time for your network architecture. Sudden spikes in this value are commonly attributed to congestion, while a plateau (which goes up and stays up) typically indicates a path change. The CA Application Delivery Analysis management console opens a Network incident when the incident threshold for this metric is crossed.

### **Network Round Trip Time**

A Network metric that measures the time that a packet takes to travel across the network in both directions between the server and clients on a network, excluding loss. Application, server, and client processing time are excluded. The CA Application Delivery Analysis management console opens a Network incident when the incident threshold for this metric is crossed.

---

## **Observations**

The observation count measures the number of times during a 5-minute monitoring interval that a monitoring device calculated a performance metric for a particular application/server/network combination. Within a TCP transaction there can be different numbers of observations of different metrics. For example, there may be more observation counts for Network Round Trip Time than Server Response Time. Other metrics are links, and always have the same number of observations. For example, each TCP transaction has one Server Response Time observation and one Data Transfer Time observation. To rate a metric as Normal, Minor (yellow), or Major (orange), the metric must have a minimum number of observations.

## **Open Sessions**

Measures the number of sessions still open at the end of the data collection period. Open sessions might become Expired or Completed during subsequent reporting intervals.

## **Packet Loss Percentage**

A Network metric that measures the ratio of retransmitted data to total data within the network. This is measured from the vantage point of the monitoring device, which is next to the server. The monitoring device can identify packets retransmitted by the server because of data losses in the server-to-client direction along the network path. When data loss occurs in the client-to-server direction (in the network path before reaching the server, for example), the monitoring device cannot observe such packet loss, and that delay is not included in the Packet Loss Percentage. On the **Engineering** page of the CA Application Delivery Analysis management console, Packet Loss Percentage is part of the QoS report. The CA Application Delivery Analysis management console opens a Network incident when the incident threshold for this metric is crossed.

## **Rate from Server in Bytes**

Measures the rate of data from the server in bytes.

## **Rate from Server in Packets**

Measures the rate of data from the server in packets.

## **Rate to Server in Bytes**

Measures the rate of data to the server in bytes.

## **Rate to Server in Packets**

Measures the rate of data to the server in packets.

## **Refused Session Percentage**

A Server metric that measures the percentage of connection requests that the server explicitly rejected during the reporting interval. This metric is part of the Unfulfilled TCP/IP Session Requests report in the CA Application Delivery Analysis management console. The CA Application Delivery Analysis management console opens a Server incident when the incident threshold for this metric is crossed.

## **Refused Sessions**

Measures the number of connection requests that were explicitly rejected by the server during the three-way handshake.

---

**Retransmission Delay**

A Network metric that measures the elapsed time between sending the original packet and sending the last duplicate packet. The management console reports Retransmission Delay as an average across observations, and not just for the retransmitted packets. For example, if one packet in a set of 10 requires 300 ms of retransmission time, the Retransmission Delay is reported as 30 ms (300 ms/10 packets). The CA Application Delivery Analysis management console opens a Network incident when the incident threshold for this metric is crossed.

**Retransmitted Bytes**

Measures the amount of bandwidth used by retransmitted data.

**Retransmitted Packets**

Measures the increased load on the network infrastructure due to retransmitted packets.

**Server Connection Time (SCT)**

A Server metric that measures the amount of time that a server takes to acknowledge the initial client connection request by sending a Syn-Ack in response to the client's SYN packet. The CA Application Delivery Analysis management console opens a Server incident when the incident threshold for this metric is crossed.

**Server Response Time**

A Server metric that measures the time that it takes for a server to send an initial response to a client request or the initial server "think time." Increases in the Server Response Time generally indicate a lack of server resources such as CPU, memory, disk, or I/O, a poorly written application, or a poorly-performing tier in a multi-tier application. The CA Application Delivery Analysis management console opens a Server incident when the incident threshold for this metric is crossed.

**Session Duration**

Measures the duration of each TCP session.

**To Server Bytes**

Measures the number of bytes that a client sent to a server.

**To Server Packets**

Measures the number of packets that a client sent to a server.

**Total Bytes**

Measures the total bytes transmitted in and out.

**Total Packets**

Measures the total packets transmitted in and out.

**Total Sessions**

Indicates the total number of sessions that Completed or Expired in the sampling period. The sum of completed sessions and expired sessions is equal to the number of total sessions. It does not include Open, Unresponsive, or Refused sessions.

### **Transaction Time**

A Combined metric that measures the amount of time elapsed from when the client sends the request to when it receives the last packet in the response. Transaction Time is the sum of Server Response Time, Network Round Trip Time, Retransmission Delay, and Data Transfer Time. The CA Application Delivery Analysis management console does not open an incident when the Transaction Time threshold is crossed.

### **Unresponsive Session Percentage**

A Server metric that measures the percentage of sessions where a connection request was sent, but the server never responded. Part of the Unfulfilled TCP/IP Session Requests view. The CA Application Delivery Analysis management console opens a Server incident when the incident threshold for this metric is crossed.

### **Unresponsive Sessions**

Indicates the number of sessions where a connection request was sent, but the server never responded.

### **User Count**

Indicates the number of unique clients observed during the sample period.

### **User Throughput**

Indicates the number of bytes transmitted divided by the time required to transmit the information.

## **ADA Dashboards**

DX NetOps Performance Management offers CA Application Delivery Analysis (ADA) dashboards. ADA views are also included in NetOps Portal dashboards. The following CA ADA dashboards are available:

- Application Performance Dashboard
- Network Overview Dashboard (CA ADA)
- Network Performance Dashboard (CA ADA)
- Performance Events Dashboard
- Server Overview Dashboard (CA ADA)
- Server Performance Dashboard (CA ADA)

### **Application Performance Dashboard (ADA)**

The **Application Performance Dashboard** provides the following application performance and incident views:

- [Incident Count by Application.](#)
- [Performance by Application.](#)
- [Performance Map by Application.](#)

### **Performance Map by Application**

The Performance Map by Application view provides a Top N view into the worst performing applications based on a particular metric. Configure this view to measure application performance using any of the available CA Application Delivery Analysis metrics. If you receive an incident notification, or notice degraded performance in the Performance By Network view, use this view to drill into the Components reports on the Engineering report of the CA Application Delivery Analysis management console. Use the Components reports to begin troubleshooting the metric details for the selected application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing networks at the top of the graph.

You can also edit the filter criteria for the view to display performance information for a different metric. Filter criteria include:

- **Metric Type**  
Choose a metric to filter the list of networks.
- **Context Settings**  
Display the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Network Overview Dashboard (ADA)

Use the **Network Overview Dashboard** to determine the network device status. This dashboard gives you a broad view of network performance, and helps you to track discards and errors associated with networking infrastructure.

By default, this dashboard contains the following views.

### NOTE

Some views require a registered CA Application Delivery Analysis data source.

- Incident Count By Network
- Performance Scorecard
- Top CPU Utilization Routers/Switches Gauge/Table
- Top Enterprise Hosts by Volume
- Top Enterprise Protocols by Volume
- Top Interface Errors - Discards
- Top Interface Utilization - In - Trend/Table
- Top Interface Utilization - Out - Trend/Table
- Top Memory Utilization Routers/Switches - Gauge/Table

## Network Performance Dashboard (ADA)

The **Network Performance** dashboard provides network performance and incident views, and views of interface and router health and status. This dashboard emphasizes interface utilization, and includes lists of devices for easy drilldown into device performance and availability data.

By default, this dashboard contains the following views:

### NOTE

Some views require a registered CA Application Delivery Analysis data source.



- Groups and Sites
- [Incident Count By Network.](#)
- [Incident List by Network.](#)
- [Performance Maps.](#)
- Top Interface Errors - Discards
- Top Interface Utilization - In - Trend/Table
- Top Interface Utilization - Out - Trend/Table
- [Top Performance by Network.](#)
- [Top Performance Map by Network.](#)

## Performance Events Dashboard (ADA)

The **Performance Events** dashboard provides performance and incident views of network, server, and applications.

By default, this dashboard contains the following views:

- [Incident Count by Application.](#)
- [Incident Count By Network.](#)
- [Incident Count by Server.](#)
- [Incident List by Network.](#)
- [Incident List by Server.](#)
- [Performance by Application.](#)
- [Performance by Network.](#)
- [Performance by Server.](#)

## Server Overview Dashboard (ADA)

The **Server Overview Dashboard** shows an overview of server and application performance. Use the views on this dashboard to track memory and CPU utilization levels on critical servers.

By default, this dashboard contains the following views:

- Performance Scorecard
- Top CPU Utilization
- Top Disk Storage
- Top Disk Utilization
- Top Least Available Servers
- Top Least Reachable Servers
- Top Memory Utilization

## Server Performance Dashboard (ADA)

The **Server Performance** dashboard provides the following server performance and incident views, and health- and status-related data from servers and groups of servers:

- Groups and Sites
- Incident Count by Server
- Performance by Server
- Performance Map by Server
- Servers (Universal List)

## ADA Views

The CA Application Delivery Analysis (ADA) views that are available in NetOps Portal are different from the views that are available in CA Application Delivery Analysis. Each view includes the **Data Source** column when multiple data sources are available. Also, each view offers a context that filters the results. For example, you can configure CA Application Delivery Analysis views with a server context, which allows you to report on a particular server.

The following ADA views are available in NetOps Portal:

- [Engineering Trend](#)
- [Incident Counts](#)
- [Incident Lists](#)
- [Performance Maps](#)
- [Performance Scorecard](#)
- [Performance Views](#)

### Engineering Trend

If baseline data is applicable, the gray line indicates the number of observations. Not every metric supports baseline or observation data.

The **Engineering Trend** view provides an engineering chart for a particular server metric, network metric, or combined metric.

You can view data only from a single data source using this view. If you registered multiple instances of CA Application Delivery Analysis (ADA), edit the view to filter by a single data source. You can also change the dashboard group context by clicking the **Group** link. Using either method, you can select one data source in the Groups tree to serve as the view context.

You can measure application performance using an available ADA metric by configuring the **Engineering Trend** view. If you receive an incident notification, or notice degraded performance, you can drill to the Components reports on the Engineering report of the ADA management console using this view. Troubleshoot metric details for a selected application using the Components reports.

### Display Performance Information for a Different Metric

Edit the filter criteria for the view:

- **Metric Type**  
Choose an ADA metric to filter the list of applications.
- **Context Settings**  
Displays the view filter.  
**Context type:**
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

### Incident Count by Application

The Incident Count by Application view lists the applications that are impacted with the most incidents. Use this view to help you start troubleshooting reported issues by matching Network and Server incidents to trouble tickets.

By default, the incident count includes Open and Closed incidents where the performance threshold for a Server metric or Network metric was exceeded during more than one 5-minute reporting interval.

#### NOTE

This view does **not** include Combined Metrics (Data Transfer Time and Transaction Time) because CA Application Delivery Analysis (CA ADA) opens a Network or Server incident when the threshold for a Combined metric is exceeded.

Click an application to view its incidents in the CA ADA management console.

You may also edit the view to change the filter criteria for the view. Filter criteria include:

- **Scheduled Maintenance**

Choose whether to include incidents that are opened or closed during a scheduled maintenance period:

- **Excluded**

Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.

- **Included**

Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA ADA:

- Closes the existing incidents
- Opens new incidents, if performance degrades

- **Metric Type**

Filter the incident count by Server metric, Network metric, or Combined metric.

- **Minimum Elapsed Time (Minutes)**

Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.

- **Minimum Severity**

From least severe to most severe, include:

- **Minor:** Includes all incident severities in the Incident count. This is the default.
- **Major:** Includes incidents that have a Major or Unavailable severity.
- **Unavailable:** Includes Unavailable incidents.

- **Incident State**

For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

#### NOTE

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.

- **Context Settings**

Displays the view filter. A context type of:

- **Summary**  
Filters the view based on the selected group.
- **Server**  
Filters the view based on the selected server.

- **Apply Changes** Select the scope of your changes from the **Apply Changes** drop-down.

## Incident Count By Network

The Incident Count by Network view lists the networks that are impacted with the most incidents. Use this view to help you start troubleshooting reported issues by matching network incidents to trouble tickets. By default, the incident count includes open and closed incidents where the performance threshold for a Network metric was exceeded during more than one 5-minute reporting interval.

Click a network to view its incidents in the CA Application Delivery Analysis (CA ADA) management console.

You may also edit the view to change the filter criteria for the view. Filter criteria include:

- **Scheduled Maintenance**

Choose whether to include incidents that are opened or closed during a scheduled maintenance period:

- **Excluded**

Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.

- **Included**

Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA ADA:

- Closes the existing incidents
- Opens new incidents, if performance degrades

- **Metric Type**

Filter the incident count by the Network metric.

- **Minimum Elapsed Time (Minutes)**

Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.

- **Minimum Severity**

From least severe to most severe, include:

- **Minor:** Includes all incident severities in the Incident count. This is the default.
- **Major:** Includes incidents that have a Major or Unavailable severity.
- **Unavailable:** Includes Unavailable incidents.

- **Incident State**

For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

### NOTE

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.

- **Context Settings**

Displays the view filter. A context type of:

- **Summary**  
Filters the view based on the selected group.
- **Server**  
Filters the view based on the selected server.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.

## Incident Count by Server

The Incident Count by Server view lists the servers that are impacted with the most incidents. Use this view to help you start troubleshooting reported issues by matching server incidents to trouble tickets. By default, the incident count includes open and closed incidents where the performance threshold for a Server metric was exceeded during more than one 5-minute reporting interval.

Click a server to view its incidents in the CA Application Delivery Analysis (CA ADA) management console.

You may also edit the view to change the filter criteria for the view. Filter criteria include:

- **Scheduled Maintenance**  
Choose whether to include incidents that are opened or closed during a scheduled maintenance period:
  - **Excluded**  
Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.
  - **Included**  
Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA Application Delivery Analysis:
    - Closes the existing incidents
    - Opens new incidents, if performance degrades
- **Metric Type**  
Filter the incident count by the Server metric.
- **Minimum Elapsed Time (Minutes)**  
Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.
- **Minimum Severity**  
From least severe to most severe, include:
  - **Minor:** Includes all incident severities in the Incident count. This is the default.
  - **Major:** Includes incidents that have a Major or Unavailable severity.
  - **Unavailable:** Includes Unavailable incidents.
- **Incident State**  
For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

### NOTE

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Incident Counts

The Incident Count views show the networks, servers, and applications with the most incidents. The following Incident Count views are available:

### Incident Count by Network

The Incident Count by Network view lists the networks that are impacted with the most incidents. Use this view to help you start troubleshooting reported issues by matching network incidents to trouble tickets. By default, the incident count includes open and closed incidents where the performance threshold for a Network metric was exceeded during more than one 5-minute reporting interval.

Click a network to view its incidents in the CA Application Delivery Analysis (CA ADA) management console.

You may also edit the view to change the filter criteria for the view. Filter criteria include:

- **Scheduled Maintenance**  
Choose whether to include incidents that are opened or closed during a scheduled maintenance period:
  - **Excluded**  
Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.
  - **Included**  
Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA ADA:
    - Closes the existing incidents
    - Opens new incidents, if performance degrades
- **Metric Type**  
Filter the incident count by the Network metric.
- **Minimum Elapsed Time (Minutes)**  
Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.
- **Minimum Severity**  
From least severe to most severe, include:
  - **Minor:** Includes all incident severities in the Incident count. This is the default.
  - **Major:** Includes incidents that have a Major or Unavailable severity.
  - **Unavailable:** Includes Unavailable incidents.
- **Incident State**  
For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

### **NOTE**

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**

Filters the view based on the selected group.

- **Server**

Filters the view based on the selected server.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.

### **Incident Count by Server**

The Incident Count by Server view lists the servers that are impacted with the most incidents. Use this view to help you start troubleshooting reported issues by matching server incidents to trouble tickets. By default, the incident count includes open and closed incidents where the performance threshold for a Server metric was exceeded during more than one 5-minute reporting interval.

Click a server to view its incidents in the CA Application Delivery Analysis (CA ADA) management console.

You may also edit the view to change the filter criteria for the view. Filter criteria include:

- **Scheduled Maintenance**

Choose whether to include incidents that are opened or closed during a scheduled maintenance period:

- **Excluded**

Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.

- **Included**

Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA Application Delivery Analysis:

- • Closes the existing incidents
- • Opens new incidents, if performance degrades

- **Metric Type**

Filter the incident count by the Server metric.

- **Minimum Elapsed Time (Minutes)**

Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.

- **Minimum Severity**

From least severe to most severe, include:

- **Minor:** Includes all incident severities in the Incident count. This is the default.
- **Major:** Includes incidents that have a Major or Unavailable severity.
- **Unavailable:** Includes Unavailable incidents.

- **Incident State**

For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

### **NOTE**

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.

- **Context Settings**

Displays the view filter. A context type of:

- **Summary**

Filters the view based on the selected group.

- **Server**

Filters the view based on the selected server.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.

### **Incident Count by Application**

The Incident Count by Application view lists the applications that are impacted with the most incidents. Use this view to help you start troubleshooting reported issues by matching Network and Server incidents to trouble tickets.

By default, the incident count includes Open and Closed incidents where the performance threshold for a Server metric or Network metric was exceeded during more than one 5-minute reporting interval.

#### **NOTE**

This view does **not** include Combined Metrics (Data Transfer Time and Transaction Time) because CA Application Delivery Analysis (CA ADA) opens a Network or Server incident when the threshold for a Combined metric is exceeded.

Click an application to view its incidents in the CA ADA management console.

You may also edit the view to change the filter criteria for the view. Filter criteria include:

- **Scheduled Maintenance**

Choose whether to include incidents that are opened or closed during a scheduled maintenance period:

- **Excluded**

Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.

- **Included**

Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA ADA:

- Closes the existing incidents
- Opens new incidents, if performance degrades

- **Metric Type**

Filter the incident count by Server metric, Network metric, or Combined metric.

- **Minimum Elapsed Time (Minutes)**

Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.

- **Minimum Severity**

From least severe to most severe, include:

- **Minor:** Includes all incident severities in the Incident count. This is the default.
- **Major:** Includes incidents that have a Major or Unavailable severity.
- **Unavailable:** Includes Unavailable incidents.

- **Incident State**

For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

#### **NOTE**

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:



- Open
- Closed
- Open and Closed. This is the default.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Incident List by Network

The Incident List by Network view displays a summary of Network incidents. Use this view to drill into details on an incident in the CA Application Delivery Analysis (CA ADA) management console. By default, the incident list includes Open and Closed incidents where the performance threshold for a Network metric was exceeded during more than one 5-minute reporting interval.

Click an incident to view its details in the CA ADA management console.

You may also edit the view to change the filter criteria for the list of incidents. Filter criteria include:

- **Scheduled Maintenance**  
Choose whether to include incidents that are opened or closed during a scheduled maintenance period:
  - **Excluded**  
Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.
  - **Included**  
Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA ADA:
    - Closes the existing incidents
    - Opens new incidents, if performance degrades
- **Metric Type**  
Filter the incident count by the Network metric.
- **Minimum Elapsed Time (Minutes)**  
Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.
- **Minimum Severity**  
From least severe to most severe, include:
  - **Minor:** Includes all incident severities in the Incident count. This is the default.
  - **Major:** Includes incidents that have a Major or Unavailable severity.
  - **Unavailable:** Includes Unavailable incidents.
- **Incident State**  
For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

### NOTE

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Incident List by Server

The Incident List by Server view shows a summary of Server incidents, and lets you drill into details on an incident in the CA Application Delivery Analysis (CA ADA) management console. By default, the incident list includes Open and Closed incidents where the performance threshold for a Server metric was exceeded during more than one 5-minute reporting interval.

Click an incident to view its details in the CA ADA management console.

You may also edit the view to change the filter criteria for the list of incidents. Filter criteria include:

- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Metric Type**  
Filter the list of incidents by Server metric.
- **Minimum Elapsed Time (Minutes)**  
Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.
- **Minimum Severity**  
From least severe to most severe, include:
  - **Minor**: Includes all incident severities in the Incident count. This is the default.
  - **Major**: Includes incidents that have a Major or Unavailable severity.
  - **Unavailable**: Includes Unavailable incidents.
- **Incident State**  
For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

### NOTE

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Incident Lists

The Incident List views show summaries of network and server incidents. The following Incident List views are available:

### Incident List by Network

The Incident List by Network view displays a summary of Network incidents. Use this view to drill into details on an incident in the CA Application Delivery Analysis (CA ADA) management console. By default, the incident list includes Open and Closed incidents where the performance threshold for a Network metric was exceeded during more than one 5-minute reporting interval.

Click an incident to view its details in the CA ADA management console.

You may also edit the view to change the filter criteria for the list of incidents. Filter criteria include:

- **Scheduled Maintenance**  
Choose whether to include incidents that are opened or closed during a scheduled maintenance period:
  - **Excluded**  
Choose this option to exclude incidents that are opened or closed during a server's scheduled maintenance period. This is the default.
  - **Included**  
Choose this option to include incidents that are opened or closed during a server's maintenance period. This option increases the number of open and closed incidents because when a scheduled maintenance period begins, CA ADA:
    - Closes the existing incidents
    - Opens new incidents, if performance degrades
- **Metric Type**  
Filter the incident count by the Network metric.
- **Minimum Elapsed Time (Minutes)**  
Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.
- **Minimum Severity**  
From least severe to most severe, include:
  - **Minor:** Includes all incident severities in the Incident count. This is the default.
  - **Major:** Includes incidents that have a Major or Unavailable severity.
  - **Unavailable:** Includes Unavailable incidents.
- **Incident State**  
For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

### **NOTE**

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

### **Incident List by Server**

The Incident List by Network view displays a summary of Server incidents. Use this view to drill into details on an incident in the CA Application Delivery Analysis (CA ADA) management console. By default, the incident list includes Open and Closed incidents where the performance threshold for a Server metric was exceeded during more than one 5-minute reporting interval.

Click an incident to view its details in the CA ADA management console.

You may also edit the view to change the filter criteria for the list of incidents. Filter criteria include:

- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Metric Type**  
Filter the list of incidents by Server metric.
- **Minimum Elapsed Time (Minutes)**  
Choose a time period to specify a threshold for the total duration required to include an incident in the Incident Count. To include all incidents with the specified Minimum Severity, choose a Duration of Any. The default is 10 minutes.
- **Minimum Severity**  
From least severe to most severe, include:
  - **Minor:** Includes all incident severities in the Incident count. This is the default.
  - **Major:** Includes incidents that have a Major or Unavailable severity.
  - **Unavailable:** Includes Unavailable incidents.
- **Incident State**  
For the metrics that you have selected, specify whether to filter open or closed incidents in the Incident Count. By default, open and closed incidents are included based on the specified Minimum Severity and Duration.

#### **NOTE**

Closed incidents help you to identify incident trends that can lead to problem resolution. By default, CA ADA retains incidents for 3 months.

Choose one of the following options:

- Open
- Closed
- Open and Closed. This is the default.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Performance by Application

The Performance by Application view shows application performance data from the last 24 hours. If the time period for the report is longer than 24 hours, the view only displays performance data for the most recent 24 hours.

If you have enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, CA Application Delivery Analysis (CA ADA) classifies the metric using one of these ratings:

- **Acknowledged**  
Displays a severity state (diagonal stripes) to indicate a CA ADA or CA NetOps Portal (CAPC) user acknowledged an incident. When this happens, CAPC marks data that the incident covers as Acknowledged. CA ADA automatically marks future data covered by an acknowledged incident as Acknowledged.
- **No Data**  
Displays a severity state (blank) to indicate that no data is available.
- **Unrated**  
Displays a severity state (gray) to indicate that either there is insufficient past data (two full business days of data are needed) to establish a threshold, or there were not enough observations to exceed the minimum observations threshold.
- **Normal**  
Displays a severity state (green) to indicate that the metric value is between zero and the Minor threshold.
- **Minor**  
Displays a severity state (yellow) to indicate that the metric value exceeds the Minor threshold as defined in the CA ADA management console.
- **Major**  
Displays a severity state (orange) to indicate that the metric value exceeds the Major threshold as defined in the CA ADA management console.
- **Unavailable**  
Displays a severity state (red) to indicate that the application on a server is not running (unavailable). This rating appears when you click All Server Metrics in Settings. This rating is only applicable to a user-defined application where the CA ADA administrator has assigned servers to the application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing items at the top of the graph.

You may also the filter criteria for the view. By default, all data is included. Filter criteria include:

- **Metric Type**  
Select the Server metrics, Network metrics or Combined metrics you want to use to filter the list of networks.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes** Select the scope of your changes from the **Apply Changes** drop-down.

## Performance by Network

The Performance by Network view shows network performance data from the last 24 hours. If the time period for the report is longer than 24 hours, the view only displays performance data for the most recent 24 hours.

If you have enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, CA Application Delivery Analysis (CA ADA) classifies the metric using one of these ratings:

- **Acknowledged**

Displays a severity state (diagonal stripes) to indicate a CA ADA or CA NetOps Portal user acknowledged an incident. When this happens, the CA NetOps Portal (CAPC) marks data that the incident covers as Acknowledged. CA ADA automatically marks future data covered by an acknowledged incident as Acknowledged.

- **No Data**  
Displays a severity state (blank) to indicate that no data is available.
- **Unrated**  
Displays a severity state (gray) to indicate that either there is insufficient past data (two full business days of data are needed) to establish a threshold, or there were not enough observations to exceed the minimum observations threshold.
- **Normal**  
Displays a severity state (green) to indicate that the metric value is between zero and the Minor threshold.
- **Minor**  
Displays a severity state (yellow) to indicate that the metric value exceeds the Minor threshold as defined in the CA ADA management console.
- **Major**  
Displays a severity state (orange) to indicate that the metric value exceeds the Major threshold as defined in the CA ADA management console.
- **Unavailable**  
Displays a severity state (red) to indicate that the application on a server is not running (unavailable). This rating appears when you click All Server Metrics in Settings. This rating is only applicable to a user-defined application where the CA ADA administrator has assigned servers to the application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing items at the top of the graph.

You may also edit the filter criteria for the view. By default, all data is included. Filter criteria include:

- **Metric Type**  
Select the Server metrics, Network metrics or Combined metrics that you want to use to filter the list of networks.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Performance by Server

The Performance by Server view provides server performance data for the last 24 hours. If the time period for the report is longer than 24 hours, the view displays performance data only for the most recent 24 hours.

If you have enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, CA Application Delivery Analysis (ADA) classifies the metric using one of these ratings:

- **Acknowledged**  
Displays a severity state (diagonal stripes) to indicate an ADA or NetOps Portal user acknowledged an incident. When this happens, NetOps Portal marks data that the incident covers as Acknowledged. ADA automatically marks future data covered by an acknowledged incident as Acknowledged.
- **No Data**  
Displays a severity state (blank) to indicate that no data is available.
- **Unrated**

Displays a severity state (gray) to indicate that either there is insufficient past data (two full business days of data are needed) to establish a threshold, or there were not enough observations to exceed the minimum observations threshold.

- **Normal**  
Displays a severity state (green) to indicate that the metric value is between zero and the Minor threshold.
- **Minor**  
Displays a severity state (yellow) to indicate that the metric value exceeds the Minor threshold as defined in the ADA management console.
- **Major**  
Displays a severity state (orange) to indicate that the metric value exceeds the Major threshold as defined in the ADA management console.
- **Unavailable**  
Displays a severity state (red) to indicate that the application on a server is not running (unavailable). This rating appears when you click All Server Metrics in Settings. This rating is only applicable to a user-defined application where the ADA administrator has assigned servers to the application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing items at the top of the graph.

You can also the filter criteria for the view. By default, all data is included. Filter criteria include:

- **Metric Type**  
Select the Server metrics, Network metrics or Combined metrics you want to use to filter the list of networks.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Performance Map by Server

The Performance Map by Server view provides a Top N view into the worst performing servers based on a particular metric. Configure this view to measure server performance using any of the available CA Application Delivery Analysis (CA ADA) metrics. If you receive an incident notification, or notice degraded performance in the Performance By Server view, use this view to drill into the Components reports on the Engineering report of the CA ADA management console. Use the Components reports to begin troubleshooting the metric details for the selected server.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing servers at the top of the graph.

You may also edit the filter criteria for the view to display performance information for a different metric. Filter criteria include:

- **Metric Type**  
Choose a metric to filter the list of servers.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**

Filters the view based on the selected server.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.

## Performance Maps

If you receive an incident notification, or notice degraded performance in a Performance view, use this view to drill into the Components reports on the Engineering report of the CA Application Delivery Analysis (ADA) management console. Use the Components reports to begin troubleshooting the metric details for the selected application.

A performance map provides a Top N view into the worst performing networks, servers, or applications based on a particular metric.

If you receive an incident notification, or notice degraded performance in a Performance view, use this view to drill into the Components reports on the Engineering report of the CA Application Delivery Analysis (ADA) management console. Use the Components reports to begin troubleshooting the metric details for the selected application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing networks at the top of the graph.

### Top Performance Map by Network

The Top Performance Map by Network view provides a Top N view into the worst performing networks based on a particular metric. Configure this view to measure application performance using any of the available ADA metrics. If you receive an incident notification, or notice degraded performance in the Performance By Network view, use this view to drill into the Components reports on the Engineering report of the ADA management console. Use the Components reports to begin troubleshooting the metric details for the selected application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing networks at the top of the graph.

You may also edit the filter criteria for the view to display performance information for a different metric. Filter criteria include:

- **Metric Type**

Choose a metric to filter the list of networks.

- **Context Settings**

Display the view filter. A context type of:

- **Summary**

Filters the view based on the selected group.

- **Server**

Filters the view based on the selected server.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.

### Performance Map by Server

The Performance Map by Server view provides a Top N view into the worst performing servers based on a particular metric. Configure this view to measure server performance using any of the available ADA metrics. If you receive an incident notification, or notice degraded performance in the Performance By Server view, use this view to drill into the Components reports on the Engineering report of the ADA management console. Use the Components reports to begin troubleshooting the metric details for the selected server.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing servers at the top of the graph.



You may also edit the filter criteria for the view to display performance information for a different metric. Filter criteria include:

- **Metric Type**  
Choose a metric to filter the list of servers.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

### **Performance Map by Application**

The Performance Map by Application view provides a Top N view into the worst performing applications based on a particular metric. Configure this view to measure application performance using any of the available ADA metrics. If you receive an incident notification, or notice degraded performance in the Performance By Network view, use this view to drill into the Components reports on the Engineering report of the ADA management console. Use the Components reports to begin troubleshooting the metric details for the selected application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing networks at the top of the graph.

You may also edit the filter criteria for the view to display performance information for a different metric. Filter criteria include:

- **Metric Type**  
Choose a metric to filter the list of networks.
- **Context Settings**  
Display the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

### **Performance Scorecard**

The Performance Scorecard view provides a high-level view into the worst-performing applications by summarizing the percentage of time that an application is operating at different levels of performance, such as Normal, Minor, or Major.

The Performance Scorecard view provides a high-level view into the worst-performing applications by summarizing the percentage of time that an application is operating at different levels of performance, such as Normal, Minor, or Major.

If you have enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, CA Application Delivery Analysis (ADA) classifies the metric using one of these ratings:

- **Acknowledged**  
Displays a severity state (diagonal stripes) to indicate that an ADA or NetOps Portal user acknowledged an incident. When this happens, NetOps Portal marks data that the incident covers as Acknowledged. ADA automatically marks future data covered by an acknowledged incident as Acknowledged.
- **No Data**

Displays a severity state (blank) to indicate that no data is available.

- **Unrated**  
Displays a severity state (gray) to indicate that either there is insufficient past data (two full business days of data are needed) to establish a threshold, or there were not enough observations to exceed the minimum observations threshold.
- **Normal**  
Displays a severity state (green) to indicate that the metric value is between zero and the Minor threshold.
- **Minor**  
Displays a severity state (yellow) to indicate that the metric value exceeds the Minor threshold as defined in the ADA management console.
- **Major**  
Displays a severity state (orange) to indicate that the metric value exceeds the Major threshold as defined in the ADA management console.
- **Unavailable**  
Displays a severity state (red) to indicate that the application on a server is not running (unavailable). This rating appears when you click All Server Metrics in Settings. This rating is only applicable to a user-defined application where the ADA administrator has assigned servers to the application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing items at the top of the graph.

You may also edit the filter criteria for the view. By default, all data is included. Filter criteria include:

- **Metric Type**  
Choose a metric to filter the list.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Performance Views

Performance views show the performance data for servers, networks, and applications from the last 24 hours. If the time period for the report is longer than 24 hours, the view only displays performance data for the most recent 24 hours.

Performance views show the performance data for servers, networks, and applications from the last 24 hours. If the time period for the report is longer than 24 hours, the view only displays performance data for the most recent 24 hours.

If there are enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, CA Application Delivery Analysis (ADA) classifies the metric.

### Contents

#### Performance by Network

The Performance by Network view provides information about network performance. This view displays performance data from the last 24 hours. If the time period for the report is longer than 24 hours, the view only displays performance data for the most recent 24 hours.

If you have enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, ADA classifies the metric using one of these ratings:

- **Acknowledged**  
Displays a severity state (diagonal stripes) to indicate a ADA or NetOps Portal user acknowledged an incident. When this happens, the NetOps Portal marks data that the incident covers as Acknowledged. ADA automatically marks future data covered by an acknowledged incident as Acknowledged.
- **No Data**  
Displays a severity state (blank) to indicate that no data is available.
- **Unrated**  
Displays a severity state (gray) to indicate that either there is insufficient past data (two full business days of data are needed) to establish a threshold, or there were not enough observations to exceed the minimum observations threshold.
- **Normal**  
Displays a severity state (green) to indicate that the metric value is between zero and the Minor threshold.
- **Minor**  
Displays a severity state (yellow) to indicate that the metric value exceeds the Minor threshold as defined in the ADA management console.
- **Major**  
Displays a severity state (orange) to indicate that the metric value exceeds the Major threshold as defined in the ADA management console.
- **Unavailable**  
Displays a severity state (red) to indicate that the application on a server is not running (unavailable). This rating appears when you click All Server Metrics in Settings. This rating is only applicable to a user-defined application where the ADA administrator has assigned servers to the application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing items at the top of the graph.

You may also edit the filter criteria for the view. By default, all data is included. Filter criteria include:

- **Metric Type**  
Select the Server metrics, Network metrics or Combined metrics that you want to use to filter the list of networks.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

### **Performance by Server**

The Performance by Server view provides information about server performance. This view displays performance data from the last 24 hours. If the time period for the report is longer than 24 hours, the view only displays performance data for the most recent 24 hours.

If you have enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, ADA classifies the metric using one of these ratings:

- **Acknowledged**  
Displays a severity state (diagonal stripes) to indicate a ADA or NetOps Portal user acknowledged an incident. When this happens, NetOps Portal marks data that the incident covers as Acknowledged. ADA automatically marks future data covered by an acknowledged incident as Acknowledged.
- **No Data**

Displays a severity state (blank) to indicate that no data is available.

- **Unrated**  
Displays a severity state (gray) to indicate that either there is insufficient past data (two full business days of data are needed) to establish a threshold, or there were not enough observations to exceed the minimum observations threshold.
- **Normal**  
Displays a severity state (green) to indicate that the metric value is between zero and the Minor threshold.
- **Minor**  
Displays a severity state (yellow) to indicate that the metric value exceeds the Minor threshold as defined in the ADA management console.
- **Major**  
Displays a severity state (orange) to indicate that the metric value exceeds the Major threshold as defined in the ADA management console.
- **Unavailable**  
Displays a severity state (red) to indicate that the application on a server is not running (unavailable). This rating appears when you click All Server Metrics in Settings. This rating is only applicable to a user-defined application where the ADA administrator has assigned servers to the application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing items at the top of the graph.

You may also the filter criteria for the view. By default, all data is included. Filter criteria include:

- **Metric Type**  
Select the Server metrics, Network metrics or Combined metrics you want to use to filter the list of networks.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

### **Performance by Application**

The Performance by Server view provides information about application performance. This view displays performance data from the last 24 hours. If the time period for the report is longer than 24 hours, the view only displays performance data for the most recent 24 hours.

If you have enough observations, and if a threshold is available for each combination of application, server, network, and 5-minute period, ADA classifies the metric using one of these ratings:

- **Acknowledged**  
Displays a severity state (diagonal stripes) to indicate that an ADA or NetOps Portal user acknowledged an incident. When this happens, NetOps Portal marks data that the incident covers as Acknowledged. ADA automatically marks future data covered by an acknowledged incident as Acknowledged.
- **No Data**  
Displays a severity state (blank) to indicate that no data is available.
- **Unrated**  
Displays a severity state (gray) to indicate that either there is insufficient past data (two full business days of data are needed) to establish a threshold, or there were not enough observations to exceed the minimum observations threshold.
- **Normal**

Displays a severity state (green) to indicate that the metric value is between zero and the Minor threshold.

- **Minor**  
Displays a severity state (yellow) to indicate that the metric value exceeds the Minor threshold as defined in the ADA management console.
- **Major**  
Displays a severity state (orange) to indicate that the metric value exceeds the Major threshold as defined in the ADA management console.
- **Unavailable**  
Displays a severity state (red) to indicate that the application on a server is not running (unavailable). This rating appears when you click All Server Metrics in Settings. This rating is only applicable to a user-defined application where the ADA administrator has assigned servers to the application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing items at the top of the graph.

You may also the filter criteria for the view. By default, all data is included. Filter criteria include:

- **Metric Type**  
Select the Server metrics, Network metrics or Combined metrics you want to use to filter the list of networks.
- **Context Settings**  
Displays the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Top Performance by Network

The **Top Performance Map by Network** view provides a Top N view into the worst performing networks based on a particular metric. Configure this view to measure application performance using any of the available CA Application Delivery Analysis (ADA) metrics. If you receive an incident notification, or notice degraded performance in the Performance By Network view, use this view to drill into the Components reports on the Engineering report of the ADA management console. Use the Components reports to begin troubleshooting the metric details for the selected application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing networks at the top of the graph.

You can also edit the filter criteria for the view to display performance information for a different metric. Filter criteria include:

- **Metric Type**  
Choose a metric to filter the list of networks.
- **Context Settings**  
Display the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Top Performance Map by Network

The Top Performance Map by Network view provides a Top N view into the worst performing networks based on a particular metric. Configure this view to measure application performance using any of the available CA Application Delivery Analysis (ADA) metrics.

The Top Performance Map by Network view provides a Top N view into the worst performing networks based on a particular metric. Configure this view to measure application performance using any of the available CA Application Delivery Analysis (ADA) metrics. If you receive an incident notification, or notice degraded performance in the Performance By Network view, use this view to drill into the Components reports on the Engineering report of the ADA management console. Use the Components reports to begin troubleshooting the metric details for the selected application.

To isolate the component and metric contributing to a performance problem, click one of the worst-performing networks at the top of the graph.

You may also edit the filter criteria for the view to display performance information for a different metric. Filter criteria include:

- **Metric Type**  
Choose a metric to filter the list of networks.
- **Context Settings**  
Display the view filter. A context type of:
  - **Summary**  
Filters the view based on the selected group.
  - **Server**  
Filters the view based on the selected server.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

## Integrate with DX NetOps Network Flow Analysis

You integrate with DX NetOps Network Flow Analysis (NFA) by registering an NFA data source in NetOps Portal.

You can register NFA as a data source for an instance of NetOps Portal or NetQoS NetOps Portal, and perform initial configuration tasks. After you have registered NFA as a data source, you can perform some key administrative tasks, such as setting up SNMP profiles, user accounts, user permissions, and groups. You perform these tasks in NetOps Portal or NetQoS NetOps Portal. The settings that you specify are synchronized down to NFA.

**Prerequisites:** Verify that no one else is running a session of NFA. To prevent issues, ensure that only one user is writing to the database at any given time.

You must have the Administrator role in NFA and in NetOps Portal to complete the steps that are outlined in the following articles:

For more information:

- About the number of data sources that you can use, see the [Release Notes](#).
- About how to register an NFA data source in NetOps Portal, see [Configure a Data Source](#).

## Configure Network Flow Analysis in NetOps Portal

As an IT engineer or tools administrator, you can start monitoring network traffic composition and utilization statistics by deploying DX NetOps Network Flow Analysis (NFA).

This use case walks you through the steps to configure NFA as a data source for NetOps Portal or CA NetQoS NetOps Portal (CA NPC). This use case also describes how to configure NFA to receive flow data, and to configure the following elements, which are administered in NetOps Portal or CA NPC:

**Prerequisites:**

- You have installed NFA and NetOps Portal or CA NPC in your environment.
- You have administrative role rights for NFA and NetOps Portal.
- SNMP profiles
- IP domains
- User accounts, permissions, and roles
- Groups

## Configure Flow Collection

Configure the routers in DX NetOps Network Flow Analysis (NFA) to verify that they are sending data to the Harvesters.

NFA can begin to collect flows as soon as you complete the following tasks:

### Add a Harvester

Enable data to be processed and displayed by adding one or more Harvesters.

**Prerequisite:** If you have not already done so, register NFA, and set up the domains before you add any Harvesters.

**Follow these steps:**

1. From the page for the DX NetOps Network Flow Analysis data source, select **Harvester** from the **Configuration Settings** menu.  
The **Harvester** page appears and displays the current settings.
2. Click **Add**.  
The **Add Harvester** dialog opens.
3. Enter the following information:
  - **IP Address**  
Address of the Harvester server.
  - **Description**  
Identifying text about the Harvester, which appears in the Harvester page table.
  - **Domain**  
Parent tenant and domain combination for the Harvester and for any routers and interfaces that begin to supply data to the Harvester.  
Changing this setting affects the tenant-domain association for new routers that begin exporting flow data and any new interfaces that begin generating flow.  
In a multi-tenant environment, the Harvester tenant affects which SNMP profiles are available for routers to poll interfaces.  
The domain affects which operators and reports have access to the data from routers and interfaces.

**NOTE**

This option is visible only in an environment that contains multiple domains.

4. Click **Save**.

The new Harvester is added and appears in the Harvester list, provided that the IP address passes the connection test. If the test connection to the web service fails, an error message opens.

## Verify the Harvester Domain

Verify that each Harvester is associated with the appropriate domain before you set up routers to export flow data. If you have not already done so, set up any needed custom tenants and domains before you proceed.

### NOTE

The tenant feature is applicable only to deployments that include NetOps Portal. If your deployment uses CA NetQoS NetOps Portal, the tenant setting is always Default Tenant.

### Follow these steps:

1. From the page for the DX NetOps Network Flow Analysis data source, select **Harvester** from the **Configuration Settings** menu.  
The **Harvester** page appears and displays the current settings.
2. Select the row for the Harvester that you want to edit.
3. Click the **Edit** button.  
The **Edit Harvester** dialog opens.
4. (Optional) Change the Domain setting (tenant-domain combination) as needed.  
**Default:** Default Tenant \ Default Domain  
You can also change the IP Address and Description.

### NOTE

If no custom IP domains have been created, the Harvester table includes only the IP Address and Description columns.

5. Click **Save**.

Your changes are saved.

## Set Up the Routers

Enable NetFlow on each CA Network Flow Analysis router by completing the steps in this topic. You can configure routers to export any of the following flow protocols:

- NetFlow v5, v7, v9, and Sampled NetFlow
- sFlow version 5
- IPFIX, J-Flow, cFlow, and NetStream flow that complies with the standards for NetFlow v5, v7, or v9

Configure flow from each source to be exported to a single Harvester. If flow from one source is exported to multiple Harvesters, a number of problems result. If this occurs, contact Broadcom Support for help.

NetFlow provides a broad view of your network packet streams by creating flow records for all packets. The data from these flow records represents all packets. Sampled NetFlow/IPFIX and sFlow take samples from your packet streams, producing fewer flow records and lessening the impact to a collector. The lower your sampling rate, the less precise the data is likely to be.

In order for data from non-sampled flows to appear in reports of 15-minute (historical) data, the following minimum fields are required:

- One of the following: 1 - IN\_BYTES, 85 - IN\_PERMANENT\_BYTES, 231 - FW\_INITIATOR\_OCTETS, or 232 - FW\_RESPONDER\_OCTETS
- 4 - PROTOCOL
- 7 - L4\_SRC\_PORT
- 8 - IPV4\_SRC\_ADDR
- 10 - INPUT\_SNMP
- 11 - L4\_DST\_PORT
- 12 - IPV4\_DST\_ADDR
- 14 - OUTPUT\_SNMP



**Complete these tasks:**

1. Back up the current router configuration.
2. Configure NetFlow export for each interface individually:
  - a. Set the flow export version.
  - b. Set the flow source IP address. Cisco recommends that you configure a loopback source interface. The IP addresses of non-loopbacked interfaces can change.
  - c. Set the flow destination IP address and set the destination port to 9995. If you are using a custom value for the harvester listening port, use that value as the destination port. The port values must match or the Harvester does not receive flow data.
  - d. Set the flow expiration timeout to 1 minute.
3. Enable flow for each interface.
  - NetFlow v5 or v5-compatible flow:
    - Monitoring multiple interfaces on a router: Use either all ingress or all egress. Use the same option for all of the interfaces. Ingress and egress values may vary slightly due to routers dropping packets and changing ToS values as traffic travels between interfaces.
    - Monitoring a single known interface on a router: Use ingress and egress. This option results in fewer total flows from the router to the Harvester and puts less load on the network and the Harvester.
  - NetFlow v9 or v9-compatible flow:  
The Harvester identifies and deduplicates multiple flows on a single router, so you can use ingress and egress on multiple interfaces. You may find it most efficient to use this option for two or three interfaces. You have the option to enable ingress and egress across all interfaces, but this configuration may put an unnecessary burden on the Harvester.
4. Configure SNMP index persistence on each router that supports this feature.

The routers are set up.

**Add DSAs (Three-Tier Deployment)**

Add at least one DSA within 30 minutes of starting flow collection. Until you add a DSA to your three-tier deployment, 15-minute data is not available for reports. Reports that show a time range of more than 2 hours do not show any data.

**NOTE**

Do not add a DSA instance instead of editing the IP address of a retired DSA. In this case, routers continue to send data to the retired DSA--and that data is not available in reports. If you have to delete a DSA, contact Broadcom Support for assistance.

**Add a DSA****Follow these steps:**

1. Open the NFA console, logged in with Administrator rights.
2. Open the **DSA** page:
  - a. Select **Administration** from the NFA console menu.  
The **Administration** page opens.
  - b. Select **System: DSA** from the **Administration** menu.  
The **DSA** page opens and shows a list of the current DSAs.
3. Click **Add**.  
The **Add DSA** dialog opens.
4. Enter the IP address for the DSA server, and then click **Test Connection**.  
A connection test is performed to determine whether the NFA console server can contact the DSA server and can verify that MySQL is installed. If the test succeeds, a `Test success` message opens.

5. Respond to the test results:
  - a. If the `Test success` message opens, click **OK** to close it.
  - b. If an error message opens, close the error message and respond as described in [the "Troubleshoot Test Connection" section](#).
6. Click **Save** in the **Add DSA** dialog.  
The connection test is performed, followed by a test to locate DSA settings on the target server.
7. Note the test results:
  - If no error message opens, the tests have succeeded. The following events result:
    - The dialog closes and the DSA is added to the DSA list.
    - The Harvesters begin to include the new DSA in the destinations for new enabled interfaces that report 15-minute data.
    - The NFA console pushes settings down to the new DSA.
    - The DSA is configured to retrieve 15-minute data files from the NFA console.
    - Data from the DSA begins to be available for reports in approximately 30 minutes.
  - If an error message opens, close the error message and respond as described in [the "Troubleshoot Add Connection" section](#).

#### NOTE

- If a DSA does not begin to collect the 15-minute data within 30 minutes after flow collection begins, problems can result. The processed data accumulates on the NFA console server and processing slows. If the problem continues, the NFA console stops collecting the 15-minute data and unprocessed data accumulates on the Harvesters. If Watchdog traps are configured, the Watchdog sends out alerts that the Harvester or Reaper is falling behind. If the problem is left unchecked, the CA Network Flow Analysis services on the Harvesters may stop running.
- Each DSA is enabled to collect data for a maximum of 5,000 enabled interfaces that have reported data.
- For information about the data types, storage lifespan, minimum thresholds, and report types for the 15-minute data that is stored on DSA servers, see the topic 15-Minute Data in the *CA Network Flow Analysis Administrator Guide*.

### Troubleshoot Test Connection

Use the following tips for troubleshooting error messages that may open when you click **Test Connection** in the **Add DSA** dialog:

- "An invalid server IP address was entered":  
You entered an IP address in an invalid format. Make sure that you enter the IP address correctly.
- "System.Web.Services.Protocols.SoapException...":  
Verify that the NetQoS MySql service is running on the NFA console server. If the service is not running, start it.
- "Unable to connect to any of the specified MySQL hosts":  
Start the NetQoS MySql service on the DSA server. Verify that the NetQoS MySql service is running on the DSA server. If the service is not running, start it.
- "Unknown database 'nqrptr':"  
The DSA database nqrptr was not found on the target server. Verify that the DSA software installation was successful.

### Troubleshoot Add Connection

Use the following tips for troubleshooting error messages that might open when you click **Save** in the **Add DSA** dialog:

- An existing record is already in use:  
Enter the IP address of a DSA that is not already in the DSA list.
- Connection must be valid and open:

Verify that the NetQoS MySql service is running on the DSA server. If the service is not running, start it.

- `System.Web.Services.Protocols.SoapException...`:  
Verify that the target server is running and can be reached by the NFA console server. Verify that the NetQoS MySql service is running on the NFA console server. If the service is not running, start it.
- `Table 'nqrptr.settings' doesn't exist`:  
The DSA settings table was not found. The DSA software was not installed successfully on the target server.
  - Verify that you entered the correct IP address for the DSA—not the IP address for a Harvester or for the NFA console.
  - Verify that the DSA software installation was successful.

## Configure Traps

For more information, see [the DX NetOps Network Flow Analysis documentation](#).

## Results of Unregistering Network Flow Analysis Data Sources

You can unregister DX NetOps Network Flow Analysis. For example, you would unregister a DX NetOps Network Flow Analysis instance before registering it with a different NetOps Portal instance.

### IMPORTANT

Unregister only when necessary.

For more information about the description of the results of unregistering during an upgrade of DX NetOps Network Flow Analysis, see [the DX NetOps Network Flow Analysis documentation](#).

Understand the following rules before unregistering an DX NetOps Network Flow Analysis data source:

- **Users:** User accounts that do not have privileges to DX NetOps Network Flow Analysis in NetOps Portal are deleted from the DX NetOps Network Flow Analysis database. Existing User IDs remain unchanged. You cannot add new users or edit user account settings while unregistered.
- **Roles:** Roles are not deleted. Users continue to have their previous roles. Existing Role IDs remain unchanged. You cannot change the roles or permissions for existing users while unregistered.
- **Groups:**
  - Groups that do not exist in DX NetOps Network Flow Analysis are deleted. You cannot add or change groups while unregistered.
  - Nested groups that are associated with an interface are displayed as interface groups in the NFA console.
  - Groups that are not associated with an interface are displayed as permissions.
- **Single Sign-On and LDAP:** Single Sign-On and LDAP values remain unchanged.

## Set Up User Accounts

DX NetOps Network Flow Analysis (NFA) includes the admin predefined user account. This user account has full administrative privileges.

The Administrator must create a user account for each user using NFA, administrators and operators. Custom user accounts enhance security and take advantage of the narrowly-defined role rights that determine access to product features and data.

Custom user accounts are best deployed in a well-planned system that includes custom groups. Custom groups are assigned as permissions to let product operators view only the data, menus, and dashboards that they need to perform their daily tasks.

For more information about how to add user accounts, which includes how to assign a role to it, how to assign permission groups to it, how to assign group ownership to it, and how to grant product privileges to a registered NFA data source, see [Manage User Accounts](#).

## Set Up Groups

Add custom groups to help manage items in NetOps Portal. Custom groups are required to let operators see performance data from the routers they manage.

Properly configured, groups can prevent operators from viewing particular types of data for security reasons. The administrator can selectively grant users access to data in their area of responsibility, such as a physical location or subnet.

Before you start adding groups, plan a strategy and a structure. Consider the types of access permissions that operators require to perform their monitoring duties. If necessary, you can discuss your organizational and monitoring goals with a Broadcom technical representative.

For more information, see [Manage Groups](#).

## Test Data Source Connections (Register and Configure NFA Use Case)

In most cases, the status indicates that data source registration has completed successfully. If the status indicates an error, use the test feature on the Manage Data Sources page.

Click the Test button to run a test that confirms the proper registration and connection of a new data source. The test checks for version compatibility, and verifies that the data source is not registered with a different instance of the CA NetOps Portal software.

If the test fails, verify that the server name or IP address is accurate for the source type. For more information, see [Data Source Test Fails](#).

## Verify IP Domains

You can address potential IP address conflicts by verifying the IP domains. Domain identifiers indicate that two managed items that otherwise appear as duplicate IP addresses are actually two different managed items.

Global administrator in NetOps Portal can control which managed items are visible to and accessible by a particular administrator or user using IP domains.

Information about custom IP domains is sent down to the data sources during synchronization. Domains are available for use during configuration. You can use the DX NetOps Network Flow Analysis console Administration functions to add interfaces, custom virtual interfaces (CVIs), routers, and Harvesters to the custom domains that you create.

During the initial setup, verify that the existing IP domains are adequate to monitor your environment.

In this article:

### View the IP Domains List

IP domains are required for monitoring multiple environments with overlapping IP addresses. Set up all the domains that you need before you begin to export flow data.

### Follow these steps:

1. Log in to the NetOps Portal Console as a user with the required administrative role rights.
2. Display the current domains:
  - Hover over **Administration, Configuration Settings**, and then click **IP Domains**.  
The **Manage IP Domains** page opens and displays the current domains.
  - Select **Administration, Group Settings**, and then click **Groups**.  
The **Manage Groups** page opens and displays the domains in the All Groups tree.

If you have not created any custom IP domains, only the Default Domain appears in the list. This predefined domain has a 'null' setting for all parameters.

Any custom domains that you have created include values for the following parameters:

- **Name**  
Identifies the domain.
- **Description**  
(Optional) Describes this domain namespace, such as naming the enterprise that owns it.
- **Primary DNS Address**  
Is the IP address of the primary name server for this domain.
- **Primary DNS Port**  
Is the port number that the primary name server uses.
- **Secondary DNS Address**  
Is the IP address of the secondary name server for this domain. Can be the same as the primary address.
- **Secondary DNS Port**  
Is the port number that the secondary name server uses.

### Add Custom IP Domains

Administrators can set the domain assignment for Harvesters, routers, interfaces, and custom virtual interfaces (CVIs). Set up any custom tenants and domains that you need before you add the Harvesters.

Having the appropriate IP domains set up helps to achieve the following goals:

- Assign the correct tenant-domain when you add Harvesters so that their routers and interfaces inherit the correct associations. The routers have the appropriate SNMP profiles available to poll their interfaces.
- Make specific content accessible only to the operators who monitor the content.
- Enable Administrators to create domain-specific ToS labels, protocol groups, and Autonomous System (AS) names in DX NetOps Network Flow Analysis.
- Avoid IP address conflicts.  
For example, suppose that a router with a single IP address has interfaces that belong to different enterprises. The domain identifiers clarify that the interfaces are different managed items, even though they have the same IP address.

The Default Domain is created automatically. The Default Domain includes any items that are not assigned to a custom domain.

Repeat the steps as required to add more IP domains.

### **Follow these steps:**

1. Log in to the NetOps Portal Console as a user with the required administrative role rights.
2. Display the current domains.
3. Create a domain:
  - (CA PC) Click **New**.  
The IP Domains Administration dialog opens.
  - (NetOps Portal) Right-click **All Domains** and select **Add New Domain**.  
The Add Domain dialog opens.
4. Specify the appropriate parameters in the provided fields.  
The Device Name Alias (CA PC only) field indicates the alias to use for a managed device. A device alias is a user-configured name that is applied to the associated managed item in NetOps Portal. Click **Browse** to navigate to and import a CSV file of aliases. The CSV file contains a list of IP address-to-device alias mappings.  
Aliases that are associated with the primary IP address of a device take precedence over aliases that are associated with any secondary IP addresses. Look for the primary IP address in the Address column of the Inventory Devices list. Use the primary IP address of the device in the CSV file.

For example:

```
172.24.36.107,Austin Router
```

Browse to select the file and click Open.

If you include aliases for devices that you are managing, it can take up to 5 minutes to begin synchronizing these aliases with NetOps Portal.

#### NOTE

To remove an alias, import a CSV file that includes the IP address for the device and a blank alias column.

To change an alias, modify the alias entry in the CSV file and reimport the file.

### 5. Interface Description Override

(CA PC only) Indicates the alternate description to use for an interface. Interface descriptions appear in NetOps Portal already, but you can provide an alternate description. Click **Browse** to navigate to and import a CSV or TXT file of alternate descriptions. The file contains a comma-separated list of values that include the device IP address, interface name, interface description, and alternate interface description (alias) mappings.

For example:

```
172.24.36.107,ethernet_7,vmxnet3 Ethernet Adapter,Connection to Dallas
```

#### NOTE

Use the primary IP address of the associated device in the CSV or TXT file. Secondary IP addresses are not supported. Look for the primary IP address in the Address column of the Inventory Devices list.

Browse to select the file and click Open.

If you include alternate descriptions for interfaces you are managing already, it can take up to 5 minutes to begin synchronizing these descriptions with NetOps Portal.

#### NOTE

You can use the same alternate interface descriptions for more than one interface. To remove an alternate description, import a CSV or TXT file that includes the IP address for the device, the interface name, the interface description, and a blank alias column. When you remove an alternate description, the original interface description reappears in NetOps Portal views.

#### NOTE

If you use a spreadsheet program to remove all the alternate descriptions from a CSV file, include a column heading for the interface description override column in the imported file. If you do not include this column heading, the original interface descriptions do not reappear in NetOps Portal views.

To change a description, modify the alias entry in the CSV or TXT file and re-import the file.

#### – DNS Settings

(CA PC only) If selected, displays the Primary DNS/Port and Secondary DNS/Port options.

#### – Primary DNS Address

Is the IP address of the primary name server for this domain.

#### – Primary DNS Port

Is the port number that the primary name server uses.

#### – Secondary DNS Address

Is the IP address of the secondary name server for this domain. Can be the same as the primary address.

#### – Secondary DNS Port

Is the port number that the secondary name server uses.

#### – Enable DNS Proxy Address

(NetOps Portal only) Indicates whether the proxy address is enabled for this IP domain.

#### – DNS Proxy Address

(NetOps Portal only) Is the IP address of the DNS proxy server.

This setting is required only if your network is located behind a DNS proxy server.

6. Click **Save**. The new IP domain appears on the page.

## Verify SNMP Profiles

DX NetOps Network Flow Analysis sends secure Simple Network Management Protocol (SNMP) information to NetOps Portal at registration. This information is transformed into SNMP profiles that contain the information necessary to enable secure queries of device MIBs using SNMP. Profile information is updated at each synchronization.

During initial product setup, verify that the available SNMP profiles are adequate to monitor your environment. The available SNMP profiles are listed on the **Manage SNMP Profiles** page in NetOps Portal.

### View the SNMP Profiles

You can view a list of SNMP profiles that are defined. The list includes high-level information about the contents of each profile.

If tenant definitions have not been created, the SNMP profiles are shared among all registered data sources. The global administrator sees a list of SNMP profiles that are not explicitly associated with a tenant. Tenant administrators see only the items that are associated with their tenant.

#### **Follow these steps:**

1. Log in as a user with the Administrator role.
2. Hover over **Administration**, and then click **Configuration Settings, SNMP Profiles**.  
The **Manage SNMP Profiles** page opens, and the current list of SNMP profiles appears.

### Add an SNMP Profile

Administrators can create SNMP profiles in the NetOps Portal.

#### **Follow these steps:**

1. On the **Manage SNMP Profiles** page, click **New**.  
The **Add SNMP Profile** dialog opens.
2. Complete the fields and change the default settings as needed, and then click **Save**.  
You return to the list of SNMP profiles. The new profile appears in the list.  
NetOps Portal performs a global synchronization to send the profile information to CA Network Flow Analysis.

## Verify That Data Is Received

To verify the setup, complete the following tasks:

### Verify the IP Domains

Information about custom IP domains is sent down to DX NetOps Network Flow Analysis during synchronization. Domains are available for use during configuration. You can use the DX NetOps Network Flow Analysis console Administration functions to add interfaces, custom virtual interfaces (CVIs), routers, and Harvesters to the custom domains that you create.

Verify that the existing IP domains are adequate to monitor your environment.

For more information, see [Verify IP Domains](#).

### Verify the SNMP Profiles

DX NetOps Network Flow Analysis sends secure Simple Network Management Protocol (SNMP) information to NetOps Portal at registration. This information is transformed into SNMP profiles that contain the information necessary to enable secure queries of device MIBs using SNMP. Profile information is updated at each synchronization.



Verify that the available SNMP profiles are adequate to monitor your environment. The available SNMP profiles are listed on the **Manage SNMP Profiles** page in NetOps Portal.

For more information, see [Verify SNMP Profiles](#).

### **Verify That the Interfaces Are Enabled**

Verify that the expected interfaces are monitored.

#### **Follow these steps:**

1. Open the NFA console, logged in with Administrator rights.
2. Open the **Available Interfaces** page:
  - a. Select **Administration** from the NFA console menu.  
The **Administration** page opens.
  - b. Select **System: Enable Interfaces** in the **Administration** menu.  
The **Available Interfaces** page opens.
3. Expand the router details to display the interface list: Click the arrow to the left of the router name.
4. Review the list to review whether the interfaces are enabled.
5. Change the interface status: Select the check box next to one or more interfaces, then click **Enable** or **Disable**.  
The selected interfaces are immediately enabled or disabled.
6. Repeat these steps for each router.

### **Verify That the Interfaces Are Visible in the NFA Console**

Ensure that the configured interfaces are visible in the NFA console.

1. Verify that interfaces are visible on the **Enterprise Overview** page:
  - a. Log in to the NFA console.
  - b. Click **Enterprise Overview** in the main menu.
  - c. Ensure the report views show interfaces.
2. Verify that interfaces are visible in the **Interface Index**:
  - a. Click **Interfaces** in the NFA console menu.  
The **Interface Index** opens.
  - b. Ensure that the **Interface Index** includes the interfaces that you expect to see.  
You can locate interfaces by expanding router details to view interfaces. Alternatively, you can use the **Search** box to find routers or interfaces by entering all or part of a name or description.
3. If interfaces are not visible, perform the following preliminary troubleshooting tasks:
  - Verify that the DX NetOps Network Flow Analysis services are running on the NFA console server.
  - Review the logs. View the logs in the NFA console or open the logs from the <install\_path>\reporter\Logs directory.

## **Network Flow Analysis Views in NetOps Portal**

You can view DX NetOps Network Flow Analysis (NFA) data in NetOps Portal in the following ways:

- [Out-of-the-box Dashboards with Enterprise-Wide Data](#)
- [Custom NetOps Portal Dashboard Views with Interface-Specific Data](#)
- [Out-of-the-box NetOps Portal Page Views with DX NetOps Network Flow Analysis Data](#)
- [DX NetOps Network Flow Analysis Interface Context Views](#)

### **Out-of-the-box Dashboards with Enterprise-Wide Data**

- **Infrastructure Overview dashboard:**



- Interfaces Over Threshold
- Routers with the Most Flow Traffic
- Top Enterprise Hosts by Volume
- Top Enterprise Protocols by Volume
- Top Flows by Interface
- Top IP Interface Utilization (Flow)
- **Management: Management Overview dashboard:**
  - Top Flows by Interface
  - Top IP Interface Utilization (Flow)
- **Management: Network Overview dashboard:**
  - Top Enterprise Hosts by Volume
  - Top Enterprise Protocols by Volume
- **Capacity Planning: Router/Switch Capacity Watch Lists dashboard:**
  - Routers with the Most Flow Traffic

### **Custom NetOps Portal Dashboard Views with Interface-Specific Data**

Display interface-specific NFA data by adding the following views with an interface-context type:

- Calendar Heat Chart (Flow)
- Stacked Protocol Trend
- Stacked ToS Trend
- Top Conversations (Bar)
- Top Conversations (Pie)
- Top Conversations (Table)
- Top Hosts (Bar)
- Top Hosts (Pie)
- Top Hosts (Table)
- Top Protocols (Bar)
- Top Protocols (Pie)
- Top Protocols (Table)
- ToS Summary (Pie)
- ToS Summary (Table)

For more information about how to customize and add views to a dashboard, see [Manage Dashboards](#).

### **Out-of-the-box NetOps Portal Page Views with DX NetOps Network Flow Analysis Data**

- **IP Performance tab:**
  - Stacked Protocol Trend
  - Top Conversations (Pie)
  - Top Hosts (Pie)
  - ToS Summary (Pie)
- **CBQoS tab:**
  - Stacked Protocol Trend
  - Stacked ToS Trend

## DX NetOps Network Flow Analysis Interface Context Views

You can drill down NFA data for an interface within NetOps Portal without having to navigate to NFA. Interface context pages can include the following views:

- **IP Performance**

Provides a broad summary of information for the selected interface extracted from the flow data such as Top N Protocols, Top N Hosts, Top N Conversations and Top N ToS values. The interface also provides the interface utilization and discards details for the same interface provided through the data aggregator.

To navigate to the context page of the protocol or host, you can select the required protocol or host from any of the views.

**NOTE**

The links in the charts and tables on the **IP Performance** page open the corresponding network flow pages within NetOps Portal.

- The following interface pages are available:

- **Network Flow Protocol**

Provides trends and next level drill downs (hosts and conversations) for the selected Protocol in the table. Host and conversation trend tables provide utilization information in the context of the selected Protocol.

- **Network Flow Host**

Provides host trends and protocol stacked trends for the selected Host in the table.

- **Network Flow Conversation**

Provides host trends and protocol stacked trends for the selected Conversation in the table.

- **Network Flow ToS**

Provides trends and next level drill downs (hosts and conversations) for the selected Type of Service(ToS) in the table. Host and conversation trend tables provide utilization information in the context of the selected ToS.

## Enterprise-Level Views

You can view enterprise-wide data from DX NetOps Network Flow Analysis in several NetOps Portal dashboard views, which are described in the following topics.

- [Interfaces Over Threshold](#)
- [Routers with the Most Flow Traffic](#)
- [Top Enterprise Hosts by Volume](#)
- [Top Enterprise Protocols by Volume](#)
- [Top Flows by Volume](#)
- [Top IP Interface Utilization](#)

The top interfaces, hosts, protocols, or ToS have the highest traffic volume during the reporting period.

### Interfaces Over Threshold

The Interfaces Over Threshold view lists the most heavily used interfaces throughout the enterprise. A table summary shows the interfaces with utilization that exceeds the configured thresholds.

The Interfaces Over Threshold view shows the interfaces whose traffic exceeded the configured thresholds during the reporting period. The view includes the following information for up to ten top interfaces:

- **Status**

Identifies the interface status as Critical (Red - Meets or exceeds the user-defined Critical threshold) or Warning (Orange - Meets or exceeds the user-defined Warning threshold).

- **Interface Name**

Identifies the interface by its name. (Depending on the application setting for the name format, the name may be prefixed by the device name.)

- **Traffic Direction**  
Shows whether the data was inbound or outbound on the interface.
- **Speed**  
(CAPC) Records the data speed that is defined for the interface.
- **Average Utilization**  
Measures the average percentage of interface capacity that was used.
- **Percent Time Critical**  
Shows the percentage of the reporting period that the interface met or exceeded the Critical threshold.
- **Percent Time Warning**  
Shows the percentage of the reporting period that the interface met or exceeded the Warning threshold.

NetOps Portal views show the data from the time range that is defined for the page.

## Contents

### Opening the View

To see this view, go to one of the following locations:

- (CAPC) Infrastructure Overview dashboard; Summary context view in a custom dashboard
- (NPC) Enterprise, Traffic Analysis, Routers/Switches Overview, or custom dashboard

### Available Actions

You can perform several actions in this view:

- Change the thresholds, view name, and utilization settings as described in this topic.
- (CAPC) Change the columns that are shown in the table: click near a column border, click Columns, then choose the columns to display.
- Click an interface name to open the Interface context pages. You can review details or open additional views of interface data.

### Change the View Settings

#### Follow these steps:

1. Open the dialog for editing the view:
  - (CAPC) Click the Edit icon in the view title bar, and then click Edit.
  - (NPC) Click the arrow next to the title name, and select Edit from the menu.
 The dialog opens.
2. (Optional) Edit the text in the Title field to change the name in the view title bar.
3. (Optional) Edit the thresholds by changing any of the following values in the Interfaces Over Threshold Settings section:
  - **Critical - % Utilization:** Specify the utilization percentage for flagging interfaces with a status of Critical, the highest level of concern. If the utilization for an interface has met or exceeded this percentage, it is marked with a red (Critical) status symbol.
  - **Warning - % Utilization:** Specify the utilization percentage for flagging interfaces with a status of Warning. If the utilization for an interface has met or exceeded this percentage, but has not met the Critical threshold, the interface is marked with an orange (Warning) status symbol.
  - **Affected % of reporting period:** Specify the percentage of the reporting period that a utilization percentage must be violated in order for the threshold to be met.

For example, if the 'Affected % of reporting period' value is 25, the threshold is met for the interfaces that have a utilization level at or above the threshold level during 25% of the reporting period. With the default reporting period of 24 hours, the list includes interfaces at or above the threshold value for six hours or more during the previous 24 hours.

4. (Optional) (NPC) Define a new context to filter the interfaces that can appear in the view: select the Filter by value, and select a context type and setting in the Select Context dialog.  
Interfaces that are not in the selected group do not appear in the view, even if they violate a threshold. If you select a group, the defined context appears under the view title.
5. Select the scope of your changes from the **Apply Changes** drop-down.
6. Click **Save** to save your changes.

### **Find the Comparable View in the Network Flow Analysis (NFA) Console**

The Interfaces Over Threshold view is similar to the Interface Utilization view on the Enterprise Overview page in the NFA console.

### **Routers with the Most Flow Traffic**

The Routers with the Most Flow Traffic view displays the routers in your network that have the highest traffic. Traffic use is measured for both inbound and outbound traffic during the reporting period, as reported by CA Network Flow Analysis.

The view includes the following information for a maximum of 10 routers:

- **Name**  
Consists of the router's IP address and device name (Y-Axis). If an administrator defined an alias for the device item, the alias is displayed. Otherwise, the discovered device name is displayed.
- **Volume**  
Measures the total amount of traffic for the router expressed in megabytes, for example (X-Axis).

NetOps Portal views show the data from the time range that is defined for the page.

### **Opening the View**

To see this view, go to the following location:

- Infrastructure Overview and Router/Switch Capacity Watch Lists dashboards; Summary-type view in a custom dashboard

### **Available Actions**

You can perform several actions in this view:

- Change the view name by editing the view settings.
- Display details in a tooltip by hovering over a bar.
- Click a router bar to view details in the NetOps Portal Router pages.

### **Top Enterprise Hosts by Volume**

The Top Enterprise Hosts by Volume view shows the enterprise hosts that have the highest traffic volume, as reported by CA Network Flow Analysis.

The view shows a bar for each of a maximum of 10 hosts that have the highest traffic volume. The bar chart includes the following information:

- **Host**

Identifies the host server by its name and IP address (Y-Axis). If an administrator has defined an alias for the device, the alias is displayed. Otherwise, the discovered device name is displayed.

- **Volume**

Measures the total amount of data sent to or from the host, expressed in a scale that is appropriate for the highest-volume host (X-Axis).

Data appears from the time range that is defined for the page.

### **Opening the View**

To see this view, go to one of the following locations:

- (CAPC) Infrastructure Overview, Network Overview, or Summary context custom dashboard

### **Available Actions**

You can perform several actions in this view:

- Change the view name by editing the view settings.
- Display details in a tooltip by hovering over a bar.
- Click a name or bar to open the related views in the NFA console.

### **Find the Comparable View in the NFA Console**

The Top Enterprise Hosts by Volume view is similar to the Top Hosts view on the Enterprise Overview page in the NFA console.

## **Top Enterprise Protocols by Volume**

The **Top Enterprise Protocols by Volume** view shows the protocols with the highest volume of network traffic across the enterprise.

The view includes the following information for a maximum of ten protocols that are associated with the highest traffic during the reporting period:

- **Protocol**

Identifies the protocol by its keyword (Y-axis).

- **Volume**

Measures the total amount of data associated with the protocol expressed in a scale that is appropriate for the highest-volume protocol (X-axis).

Data appears from the time range that is defined for the page.

### **Opening the View**

To see this view, go to one of the following locations:

- (CA PC) Infrastructure Overview or Network Overview dashboard; Summary context view in a custom dashboard
- (NPC) Enterprise, Traffic Analysis, Network Overview, and custom dashboards

### **Available Actions**

You can perform several actions in this view:

- Change the view name in the view settings.
- Display details in a Tooltip by holding your cursor over a bar.
- Click a name or bar to open related views in the NFA console.

## **Find the Comparable View in the NFA Console**

The Top Enterprise Protocols by Volume view is similar to the Top Protocols view on the Enterprise Overview page in the NFA console.

## **Top Flows by Volume**

The Top Flows by Volume views show the interfaces across the enterprise that have the highest volume of inbound or outbound traffic.

The view shows the following information for a maximum of 10 top interfaces:

- **Name**  
Identifies the interface by its device name (such as its router name), followed by a colon (:) and the interface name (Y-Axis).
- **Volume**  
Measures the volume of flow data on the interface (X-Axis) expressed in a scale that is appropriate for the highest-volume interface.

NetOps Portal views show the data from the time range that is defined for the page.

## **Opening the View**

To see this view, go to one of the following locations:

- (CAPC) Infrastructure Overview and Management Overview dashboards; Summary context view in a custom dashboard
- (NPC) Custom dashboard

## **Available Actions**

You can perform several actions in this view, including the following ones:

- Change the data direction or the view name by editing the view settings.
- Display details in a tooltip by hovering over a bar. The tooltip identifies the interface position among the top 10, with Interface 0 as the one with the highest traffic volume.
- Click a name or bar to open related information on the interface context pages.

## **Find Flow Volume Data in the Network Flow Analysis (NFA) Console**

To see the flow volume of multiple top interfaces in the NFA console, create and run a Custom report. For example, you can view flow volume for the top interfaces in summary pie charts, summary tables, trend charts, and stacked trend charts. For instructions, see the topic "Set Up Custom Reports" in the *CA Network Flow Analysis Operator Guide*.

To see the flow volume of a single top interface in the NFA console, drill into details from the Enterprise Overview page:

1. Click an interface name or bar in one of the Top Interfaces views on the NFA console Enterprise Overview page.
2. Select Flows from the list labeled "For this interface, show me" on the Interface page that opens.
3. Click the gray bar on the left side of the page to change the presentation mode.
4. Click Volume in the Presentation menu that opens.  
The Flows views display a trend chart of inbound flow volume and outbound flow volume.  
To jump to the NetOps Portal Interface Pages data for the selected interface, click the arrow next to the Flows title, and select CAPC/NPC Interface Performance.

## Top IP Interface Utilization

The Top IP Interface Utilization (Flow) views show the high-utilization interfaces from across the enterprise.

The view includes the following information for a maximum of 10 top interfaces during the reporting period:

- **Name**  
Identifies the interface by its device name/interface name (Y-Axis).
- **Percent (Utilization)**  
Measures the percentage of interface capacity that was used (X-Axis). The view shows the utilization of either inbound or outbound capacity.

NetOps Portal views show the data from the time range that is defined for the page.

### Opening the View

To see this view, go to one of the following locations:

- (CAPC) Infrastructure Overview and Management Overview dashboards; Summary-type view in a custom dashboard
- (NPC) Traffic Analysis and custom dashboards

### Available Actions

You can perform several actions in this view:

- Change the data direction, view name, and context (the interfaces that are used) by editing the view settings.
- (NPC) Change the utilization thresholds.
- Display details in a tooltip by holding your cursor over a bar.
- Click a name or bar to open related information on the Interface context pages.

### Find the Comparable View in the Network Flow Analysis (NFA) Console

The Top IP Interface Utilization (Flows) view is similar to the Interface Utilization view on the Enterprise Overview page in the NFA console.

## Interface Stacked Trend View

Stacked Trend views show the top protocol or ToS values that are used for traffic on the interface that you have selected. The following Stacked Trend views are available:

- [Stacked Protocol Trend](#)
- [Stacked ToS Trend](#)


### Stacked ToS Trend

The **Stacked Type of Service (ToS) Trend** views show the interface traffic for the top ToS, including the time that the traffic occurred. A timeline of rates is included for each ToS value. You can configure the view to display rate, utilization, or volume information.

The stacked chart shows the value of rate, utilization, and volume as stacked lines. The value of the first metric is measured from the X axis. The value of each of the other metrics is measured from the top of the stacked line to the previous metric. For example, if the top of the first trend line indicates 10, and the second line indicates 25, the value of the second metric is 15.

The following chart shows an example of a stacked trend chart. At 18:55, the value of the metric that is represented by the blue color is approximately ten. At the same time, the value of the metric that is represented by the yellow color is

approximately seven. The value of the metric in yellow is calculated by subtracting the value at the bottom of the yellow trend line from the value at the top.



The view includes the following information:

- **Identifier**  
Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).  
(NPC) The identifier line also includes the interface speed.
- **ToS Bands**  
Show the data rate, data volume, or interface capacity utilization for each top ToS that is associated with traffic on the interface.
- **Time**  
Point in time during data transmission, expressed in hours and minutes (X-Axis).
- **Measurement Setting:**
  - **Rate:** Data transfer rate at each point in time, expressed in kilobits per second or a rate that is appropriate for the highest-volume ToS (Y-Axis). The rate is calculated by dividing the data volume by the elapsed transmission time.
  - **Bytes (Volume):** Data volume at each point in time, expressed in a scale that is appropriate for the highest-volume ToS (Y-Axis).
  - **Percent (Utilization):** Percentage of the total interface capacity that the ToS traffic uses (Y-Axis). The utilization percentage is calculated by dividing the data rate by the data speed.

Depending on the data direction, the view shows inbound, outbound, or total data on the interface:

- **Legend**  
Identifies the ToS for each color band by ToS number and label (bottom of the view).

NetOps Portal views show the data from the time range that is defined for the page.

In this article:

- [Open the View](#)
- [Available Actions](#)
- [Find ToS Trend Data in the NFA Console](#)

### **Open the View**

To see this view in console, go to one of the following locations:



- (CAPC) Custom dashboard; Interface Pages (with an interface selected): CBQoS tab
- (NPC) Interface Pages (with an interface selected): Interface QoS and custom tabs

#### NOTE

You can add **Multi-Interface Stacked ToS Trend** views to a custom dashboard or to a custom tab in the **Interface** pages of the console. This view consists of a group of interface-specific stacked ToS trend charts.

### Available Actions

You can perform several actions in this view, including:

- Change the traffic direction (In, Out, or Total), the type of measurement (Rate, Volume, or Utilization), and the view name by editing the view settings. If the view is on a custom interface context dashboard in the console, you can change the interface.
- (CAPC) Zoom in to narrow the time frame.
- (CAPC) **Display only the data for a single ToS:** Right-click a ToS in the legend at the bottom of the view, and then click **Focus**. This menu is available for a view that has multiple ToS values. (This option is active when the view contains multiple ToS values.)
- (CAPC) **Hide data for one of multiple ToS values:** Right-click a ToS in the legend at the bottom of the view, and then click **Hide**.
- (CAPC) Position your cursor over legend items to display explanatory tooltips.
- Jump to details on an NFA console Interface page by double-clicking a ToS value in the legend.

### Find ToS Trend Data in the NFA Console

You can display ToS volume in trend charts or stacked trend charts in the NFA console for a selected interface:

- **Overview**
  - **Report type:** Overview
  - **Presentation menu options:** Mixed Chart or Mixed Trend; Volume.
  - **Views:** Stacked ToS Trend (In and Out) for the Top N ToS, plus other overview views.
- **Top N ToS, Stacked Trends**
  - **Report type:** ToS
  - **Filter:** Top N ToS
  - **Presentation menu options:** Stacked Trend Chart; Volume.
  - **Views:** Stacked Trend for the Top N ToS (In, Out, and Total).
- **Top N ToS, Trends**
  - **Report type:** ToS
  - **Filter:** Top N ToS
  - **Presentation menu options:** Trend Chart; Volume
  - **Views:** Trend (In, Out, and Total) for each of the Top N ToS
- **Single ToS, Stacked Trends/Trends**
  - **Report type:** ToS
  - **Filter:** Single ToS value
  - **Presentation menu options:** Mixed Trend; Volume
  - **Views:** Trend (In and Out with baselines), Stacked ToS Trend (In and Out).
- **Single ToS, Trends**
  - **Report type:** ToS
  - **Filter:** Single ToS value. Presentation menu options: Mixed Chart; Volume.
  - **Views:** Stacked ToS Trend (In and Out).
- **Conversation**

- **Report type:** Conversations
- **Filter:** Single conversation source and destination.
- **Report subtype:** Protocols
- **Presentation menu options:** Volume.
- **Views:** Conversation Trend (maximum of seven views for different timespans).


To display Flow Forensics-level detail, click the **Flow Forensics** link and run a **Flow Forensics** report.

## Stacked Protocol Trend

The Stacked Protocol Trend views show the protocols that are used the most heavily for traffic on the selected interface. The views also show when the traffic occurred. A timeline of rates is included for each listed ToS value. You can configure the view to display rate, utilization, or volume information.

The stacked chart shows the value of rate, utilization, and volume as stacked lines. The value of the first metric is measured from the X axis. The value of each of the other metrics is measured from the top of the stacked line to the previous metric. For example, if the top of the first trend line indicates 10, and the second line indicates 25, the value of the second metric is 15.

The following chart shows an example of a stacked trend chart. At 18:55, the value of the metric that is represented by the blue color is approximately ten. At the same time, the value of the metric that is represented by the yellow color is approximately seven. The value of the metric in yellow is calculated by subtracting the value at the bottom of the yellow trend line from the value at the top.



The views include the following information:

- **Identifier**  
Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).  
(NPC) The identifier line also includes the interface speed.
- **Protocol Bands**  
Show the data rate, the data volume, or the interface capacity utilization for each top protocol that is associated with traffic on the interface.
- **Time (All Views)**  
Point in time during data transmission, expressed in hours and minutes (X-Axis).
- **Measurement Setting:**

- **Rate:** Data rate at each point in time. For example, the rate can be expressed in kilobits per second (Y-Axis). The rate is calculated by dividing the data volume by the elapsed transmission time.
- **Bytes (Volume):** Data volume at each point in time. For example, the volume can be expressed in kilobytes, (Y-Axis).
- **Percent (Utilization):** Percentage of the total interface capacity that the protocol uses (Y-Axis). The utilization percentage is calculated by dividing the data rate by the data speed.

Depending on the data direction, the view shows inbound, outbound, or total data on the interface.

- **Legend**

Identifies the protocol for each color band by protocol keyword and tcp/udp port (bottom of the view).

NetOps Portal views show the data from the time range that is defined for the page.

In this article:

- [Opening the Views](#)
- [Available Actions](#)
- [Find Protocol Trend Data in the NFA Console](#)

## **Opening the Views**

To see these views in the console, go to one of the following locations:

- (CAPC) **Interface Pages (with an interface selected):** Custom dashboard, IP Performance, and CBQoS tabs
- (NPC) **Interface Pages (with an interface selected):** Interface Capacity, Interface QoS, and custom tabs

### **NOTE**

You can add **Multi-Interface Stacked Protocol Trend** views to a custom dashboard or to a custom tab in the Interface pages in the console. This view consists of a group of interface-specific stacked protocol trend charts

## **Available Actions**

You can perform several actions in this view, including the following ones:

- Change the traffic direction, the type of measurement (Rate, Volume, or Utilization), and the view name by editing the view settings. If the view is on a custom interface context dashboard in the console, you can change the interface.
- (CAPC) Zoom in to narrow the time frame.
- (CAPC) **Display only the data for a single protocol:** Right-click a protocol in the legend at the bottom of the view, and click Focus. This menu is available for a view that has multiple protocols. (This option is active when the legend contains multiple protocols.)
- (CAPC) **Hide data for one of multiple protocols:** Right-click a protocol in the legend at the bottom of the view and, click Hide.
- (CAPC) Position your cursor over legend items to display explanatory Tooltips.
- Jump to details on an NFA console Interface page by double-clicking a protocol in the legend.
- (NPC) Jump to details on the corresponding Interface page by double-clicking a protocol band in the view. To select a destination tab on the Interface page, right-click the protocol band and select a tab from the menu.

## **Find Protocol Trend Data in the NFA Console**

You can display protocol volume in the NFA console in trend charts or stacked trend charts for a selected interface:

- **Overview**

- **Report type:** Overview
- **Presentation menu options:** Mixed Chart; Volume
- **Views:** Stacked Protocol Trend (In and Out) for the Top N Protocols, plus other overview views
- **Top N Protocols, Stacked Trends**
  - **Report type:** Protocols
  - **Filter:** Top N Protocols
  - **Presentation menu options:** Stacked Trend Chart; Volume
  - **Views:** Stacked Trend for the Top N Protocols (In, Out, and Total)
- **Top N Protocols, Trends**
  - **Report type:** Protocols
  - **Filter:** Top N ToS
  - **Presentation menu options:** Trend Chart; Volume
  - **Views:** Trend (In, Out, and Total) for each of the Top N Protocols
- **Single Protocol**
  - **Report type:** Protocols
  - **Filter:** Single protocol
  - **Views:** (Depending on the selected report subtype): trends, stacked trends, trend summaries, and multi-trend summaries for protocols, protocol hosts, and protocols in conversations.

To display Flow Forensics-level detail, click the **Flow Forensics** link, and run a **Flow Forensics** report.

## Interface ToS Summaries

The ToS Summary views show the Type of Service (ToS) values for traffic on the selected interface. The articles in this section describe these views.

### ToS Summary (Pie)

The ToS Summary (Pie) view shows an overview of the Type of Service (ToS) values for traffic on the selected interface.

The view includes a pie chart and table of information about the high-volume ToS values in use on the selected interface. The table includes the following information by default:

- **Identifier**  
Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).  
(NPC) The identifier line also includes the interface speed.
- **Type of Service**  
Name of the ToS values associated with high-volume traffic, identified by number and label.
- **Total**  
Shows the total data volume for the reporting period.
- **Percent**  
(CAPC) Lists the percentage of the total data volume for the Top N ToS.

NetOps Portal views show the data from the time range that is defined for the page.

### Opening the Views

To see these views in the NetOps Portal Console, go to one of the following locations:

- (CAPC) Interface Pages (with an interface selected): IP Performance tab; Custom dashboard
- (NPC) Interface Pages (with an interface selected): Custom tab

## Available Actions

You can perform several actions in this view, including the following ones:

- Change the traffic direction and view name by editing the view settings. If the view is on a custom interface context dashboard in the console, you can change the interface.
- (CAPC) Change the type of measurement.
- Jump to details on an NFA console Interface page by double-clicking a ToS name.

## Find ToS Summary Pie Charts in the NFA Console

You can display pie charts of ToS summary data in the NFA console for a selected interface:

- **Overview**
  - **Report type:** Overview
  - **Presentation menu option:** Pie Chart
  - **View:** ToS Summary (In and Out) for the Top N ToS
- **Top N ToS Summary**
  - **Report type:** ToS
  - **Filter:** Top N ToS
  - **Presentation menu option:** Pie Chart
  - **View:** ToS Summary (In, Out, and Total) for the Top N ToS
- **Single ToS Summaries**
  - **Report type:** ToS
  - **Filter:** Single ToS
  - **Report subtype:** Overview
  - **Presentation menu option:** Pie Chart
  - **Views:** ToS Protocol Summary (In and Out) for the single ToS; ToS Hosts Summary (From and To) for the single ToS; ToS Conversations Summary (Total) for the single ToS.

### NOTE

You can view additional versions of the summary pie charts by selecting Protocols, Hosts, or Conversations as the report subtype.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## ToS Summary (Table)

The ToS Summary (Table) views show rate, volume, or utilization for the top ToS values of the traffic on a particular interface. You can use this information to compare traffic for each of the top ToS values.

The example graphic shows the view in the console. The table shows the rate for each listed ToS value. You can configure the view to display rate, utilization, or volume information.

An interface identification string is shown under the view title. The table contains a row for each ToS with the Type of Service identifier (EF/AF, DSCP, and ToS values) and the following rate, volume, or utilization information:

- **Rate:**
  - (CAPC/NPC) Average rate of total, inbound, and outbound data for each ToS (Average Total, Average Out, and Average In)
  - (CAPC) Maximum rate of data that is outbound or inbound on the interface for each ToS (Maximum Out and Maximum In)

The rate is calculated by dividing the data volume by the elapsed transmission time.

- **Volume:** Volume of outbound, inbound, and all data for each ToS (Out, In, and Total), expressed in a scale that is appropriate for the highest-volume ToS.
- **Utilization:**
  - (CA PC/NPC) Average utilization of total, outbound, and inbound data that each ToS consumes (Average Total, Average In, and Average Out)
  - (CAPC) Maximum percentage of interface capacity that the outbound or inbound utilizes for each ToS (Maximum Out and Maximum In)

The utilization percentage is calculated by dividing the data rate by the data speed.

NetOps Portal views show the data from the time range that is defined for the page.

### **Opening the View**

To see this view in the console, go to one of the following locations:

- (CA PC) Custom dashboard
- (NPC) Interface Pages (with an interface selected): Interface QoS or custom tab

### **Available Actions**

You can perform several actions in this view:

- Change the type of measurement (Rate, Volume, or Utilization) and the view name by editing the view settings.
- (CAPC) Change the interface.
- Re-sort the table data by clicking a column heading. Click again to toggle between descending and ascending order.
- Change the Max Per Page value to show more or fewer items on each table page.
- (CAPC) Change the columns that are shown in the table: Click near a column border, click Columns, and then choose the columns to display.
- Click a Type of Service link to display more information about the ToS on Interface report pages in the NFA console.

### **Find ToS Summary Tables in the NFA Console**

You can display ToS summary tables in the NFA console for a selected interface:

- **Top N ToS Summary**
  - **Report type:** ToS.
  - **Filter:** Top N ToS.
  - **Presentation menu options:** Summary Table; Volume.
  - **View:** ToS Summary Table for the Top N ToS.
- **Protocol Summary for a Single ToS**
  - **Report type:** ToS.
  - **Filter:** Single ToS.
  - **Report subtype:** Protocols.
  - **Subtype filter:** Top N Protocols.
  - **Presentation menu options:** Summary Table; Volume.
  - **Views:** ToS Protocol Summary Table for the single ToS.
- **Host Summary for a Single ToS**

- **Report type:** ToS.
- **Filter:** Single ToS.
- **Report subtype:** Hosts.
- **Subtype filter:** Top N Hosts.
- **Presentation menu options:** Summary Table; Volume.
- **Views:** ToS Hosts Summary Table for the single ToS.
- **Conversation Summary for a Single ToS**
  - **Report type:** ToS.
  - **Filter:** Single ToS.
  - **Report subtype:** Conversations.
  - **Subtype filter:** Top N Conversations.
  - **Presentation menu options:** Summary Table; Volume.
  - **Views:** ToS Conversations Summary Table for the single ToS.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Interface Top Conversations

The Top Conversations views in this section show the conversations that generate the highest traffic on the currently selected interface.

### Top Conversations (Bar)

The Top Conversations (Bar) views show the conversations that have the highest traffic on the selected interface. A bar graph shows the volume for each conversation.

For example, use conversation information to determine the IP addresses of high-volume hosts. Contact the host owners or users to investigate the nature and purpose of the traffic.

You can view the conversations for incoming data, outgoing data, or all data.

The view includes a bar for each top conversation on the selected interface. A maximum of 10 conversations are shown. The view includes the following information:

- **Identifier**  
Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).  
(NPC) The identifier line also includes the interface speed.
- **Conversation Pair**  
Identifies the conversation source and destination servers by their names (the fully qualified DNS names, if they are available), followed by the IP addresses (Y-Axis).
- **Volume**  
Measures the total amount of data that was exchanged in the conversation expressed in a scale that is appropriate for the highest-volume conversation (X-Axis).

NetOps Portal views show the data from the time range that is defined for the page.

### Opening the View

To see this view in the NetOps Portal Console, go to one of the following locations:

- (CAPC) Custom dashboard
- (NPC) Interface Pages (with an interface selected): Interface Capacity or custom tab

## Available Actions

You can perform several actions in this view:

- Change the view name by editing the view settings.
- (CAPC) Change the traffic direction and the interface.
- Display details in a Tooltip by holding your cursor over a bar.
- Jump to conversation details on an NFA console Interface report page by clicking a bar or name.

## Find Conversation Data in the NFA Console

You can view conversation volume trend charts in the NFA console for any interface that you have selected:

- **Overview Multi-Trend**
  - **Report type:** Overview.
  - **Presentation menu options:** Mixed Trend; Volume.
  - **View:** Conversations Multi Trend Summary (Total) for the Top N Conversations, plus other views.
- **Top N Conversations Trend**
  - **Report type:** Conversations.
  - **Filter:** Top N Conversations.
  - **Presentation menu options:** Trend Chart; Volume.
  - **View:** Conversations Trend for the Top N Conversations.
- **Conversations for a Single Protocol**
  - **Report type:** Protocols.
  - **Filter:** Single protocol.
  - **Report subtype:** Conversations.
  - **Conversation Filter:** Top N Conversations.
  - **Presentation menu option:** Trend Chart.
  - **View:** Protocol Conversations Summary (Total) for a single protocol.
- **Conversations for a Single ToS**
  - **Report type:** ToS.
  - **Filter:** Single ToS.
  - **Report subtype:** Conversations.
  - **Conversation Filter:** Top N Conversations.
  - **Presentation menu options:** Trend Chart; Volume.
  - **View:** ToS Trend view for each conversation that uses the single ToS.

### NOTE

To see trend charts for a single conversation, click Top N Conversations and select a single conversation as the filter.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Top Conversations (Pie)

The Top Conversations (Pie) view includes a pie chart of the conversations that account for the most traffic on the selected interface.

The view includes a pie chart and a table of information about the high-volume conversations on the selected interface. A text string near the top of the view identifies the interface whose data is displayed. The table includes the following information by default:

- **Identifier**



Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).

(NPC) The identifier line also includes the interface speed.

- **Source - Destination Name**

Identifies the host servers that initiated and received the conversation data by their fully qualified DNS names (if available) and IP addresses.

- **Total**

Shows the total amount of data in the conversation expressed in a scale that is appropriate for the conversation with the highest volume.

- **Percent**

(CAPC) Records how much the conversation consumes out of the total traffic that is displayed.

NetOps Portal views show the data from the time range that is defined for the page.

### **Opening the Views**

To see these views in the NetOps Portal Console, go to one of the following locations:

- (CAPC) Custom dashboard; Interface Pages (with an interface selected): IP Performance tab
- (NPC) Interface Pages (with an interface selected): Interface QoS or custom tab

### **Available Actions**

You can perform several actions in this view, including the following ones:

- Change the view name by editing the view settings.
- (CAPC) Change the traffic direction and the type of measurement. If the view is on a custom interface context dashboard in the CA NetOps Portal Console, you can change the interface.
- (CAPC) Change the columns that are shown in the table: Click near a column border, click Columns, and then choose the columns to display.
- Jump to details on an NFA console Interface page by clicking a link.

### **Find Conversation Pie Charts in the NFA Console**

You can display conversation pie charts in the NFA console for any interface that you have selected:

- **Overview**
  - **Report type:** Overview
  - **Presentation menu option:** Pie Chart
  - **View:** Conversations Summary (Total) for the Top N Conversations, plus other overview views
- **Top N Conversations**
  - **Report type:** Conversations
  - **Filter:** Top N Conversations
  - **Presentation menu option:** Pie Chart.
  - **View:** Conversations Summary (Total) for the Top N Conversations
- **Conversations for a Single Protocol**

- **Report type:** Protocols
- **Filter:** Single protocol
- **Report subtype:** Conversations or Overview
- **Conversations Filter:** Top N Conversations. Presentation menu option: Pie Chart or Mixed Chart
- **View:** Protocol Conversations Summary (Total) for a single protocol.
- **Conversations for a Single ToS**
- **Report type:** ToS
- **Filter:** Single ToS
- **Report subtype:** Conversations
- **Conversation Filter:** Top N Conversations
- **Presentation menu option:** Pie Chart
- **View:** ToS Conversations Summary (Total) for a single ToS

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Top Conversations (Table)

The Top Conversations (Table) views show data for the top highest volume conversations on a particular interface. The maximum number of top conversations that are shown is 10. You can configure the view to show rate, utilization, or volume information.

The table contains a row for each conversation with the source and destination. Each conversation contains the fully qualified DNS host name (if available) and IP address of the servers that initiated and received the conversation data. The table also contains the following rate, volume, or utilization information:

- **Rate:**
  - (CAPC/NPC) For each conversation, the average rate of total data (Average Total), data that goes to the destination host (Average To), and data that comes from the source host (Average From).
  - (CAPC) For each conversation, the maximum rate of data that comes from the source host (Maximum From) and goes to the destination host (Maximum To).

The rate is calculated by dividing the data volume by the elapsed transmission time.

- **Volume:** For each conversation, the total amount of data (Total), data that comes from the source host (From), and data that goes to the destination host (To), expressed in a scale that is appropriate for the highest-volume conversation.

- **Utilization:**
  - (CAPC/NPC) For each conversation, the average utilization by data that comes from the source host (Average From), data that goes to the destination host (Average To), and total data (Average Total).
  - (CA PC) For each conversation, the maximum percentage of interface capacity that is used by the data that comes from the source host (Maximum From) or that goes to the destination host (Maximum To).

The utilization percentage is calculated by dividing the data rate by the data speed.

NetOps Portal views show the data from the time range that is defined for the page.

## Opening the View

To see this view in the NetOps Portal Console, go to one of the following locations:

- (CAPC) Custom dashboard
- (NPC) Interface Pages (with an interface selected): Custom tab

## Available Actions

You can perform several actions in this view:

- Change the type of measurement (Rate, Volume, or Utilization) and the view name by editing the view settings.
- (CAPC) Change the interface.
- (CAPC) Change the columns that are shown in the table: Click near a column border, click Columns, and then choose the columns to display.
- Click one of the links to jump to a pre-filtered Interface page report in the NFA console.

## Find Conversation Tables in the NFA Console

You can display tables with conversation volumes in the NFA console for a selected interface:

- **Top N Conversations**
  - **Report type:** Conversations.
  - **Filter:** Top N Conversations.
  - **Presentation menu option:** Summary Table; Volume.
  - **View:** Conversation Summary Table for the Top N Conversations.
- **Conversations for a Single Protocol**
  - **Report type:** Protocols.
  - **Filter:** Single protocol.
  - **Report subtype:** Conversations.
  - **Subtype filter:** Top N Conversations.
  - **Presentation menu option:** Summary Table; Volume.
  - **View:** Protocol Conversation Summary Table for a single protocol.
- **Conversations for a Single ToS**
  - **Report type:** ToS.
  - **Filter:** Single ToS.
  - **Report subtype:** Conversations.
  - **Subtype filter:** Top N Conversations.
  - **Presentation menu options:** Summary Table; Volume.
  - **View:** ToS Conversations Summary Table for a single ToS.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Interface Top Hosts

The Top Hosts views show the hosts that generate the highest traffic on the currently selected interface. The articles in this section describes these views.

### Top Hosts (Bar)

The Top Hosts (Bar) views show the top high-volume hosts for a particular interface. You can use this view to determine the IP addresses of hosts that are responsible for high volumes of network traffic. You can then contact the owner or user of each host to investigate the nature and purpose of the traffic. You can view the hosts for incoming flows, outgoing flows, or all flows.

The bar chart includes the following information for a maximum of 10 hosts:

- **Identifier**

Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).

- **Host Name**  
Identifies the host server by its fully qualified DNS name (if available) and IP address (Y-Axis).
- **Volume**  
Measures the total amount of data for the host on the interface, expressed in a scale that is appropriate for the highest-volume host (X-Axis).

NetOps Portal views show the data from the time range that is defined for the page.

### **Opening the View**

To see this view in the console, go to one of the following locations:

- (CA PC) Custom dashboard

### **Available Actions**

You can perform several actions in this view:

- Change the traffic direction and the view name by editing the view settings.
- (CAPC) Change the interface.
- Display details in a tooltip by holding your cursor over a bar.
- Jump to details for a specific host on an NFA console Interface page by clicking a bar.

### **Find Host Trend Views in the NFA Console**

The Enterprise Overview page in the NFA console shows traffic volume for the top hosts in a bar chart.

You also can display host volume in trend charts for a selected interface:

- **Overview**
  - **Report type:** Overview.
  - **Presentation menu options:** Mixed Trend; Volume.
  - **View:** Hosts Multi Trend Summary (From and To) for the Top N Hosts, plus other overview views.
- **Top N Host Trend**
  - **Report type:** Hosts.
  - **Filter:** Top N Hosts.
  - **Presentation menu options:** Trend Chart; Volume.
  - **View:** Host Trend for each of the Top N Hosts.

#### **NOTE**

To see trend charts for a single host, click Top N Hosts and select a host as the filter.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

### **Top Hosts (Pie)**

The Top Hosts (Pie) views show the hosts that account for the highest volumes of network traffic on the selected interface.

The table includes the following information by default:

- **Identifier**

Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).

- **Host Name**  
Identifies the host server by its fully qualified DNS name (if available) and IP address.
- **Total**  
Records the total amount of data for the host on the interface, expressed in a scale that is appropriate for the host with the highest volume.
- **Percent**  
(CAPC) Records the amount of traffic the host consumes out of the total traffic that is displayed.

NetOps Portal views show the data from the time range that is defined for the page.

### **Opening the Views**

To see these views in the NetOps Portal Console, go to one of the following locations:

- (CAPC) Custom dashboard; Interface Pages (with an interface selected): IP Performance tab

### **Available Actions**

You can perform several actions in this view, including the following ones:

- Change the traffic direction and the view name by editing the view settings. If the view is on a custom interface context dashboard in the console, you can change the interface.
- (CAPC) Change the columns that are shown in the table: Click near a column border, click Columns, and then choose the columns to display.
- Jump to details on an Interface report page in the NFA console by clicking a host link in the view.

### **Find Host Pie Charts in the NFA Console**

You can display pie charts with host volumes in the NFA console for a selected interface:

- **Overview**
  - **Report type:** Overview
  - **Presentation menu option:** Pie Chart
  - **View:** Host Summary (From and To) for the Top N Hosts, plus other overview views.
- **Top N Hosts Summary**
  - **Report type:** Hosts
  - **Filter:** Top N Hosts
  - **Presentation menu option:** Pie Chart
  - **View:** Host Summary (From, To, and Total) for the Top N Hosts
- **Hosts for a Single Protocol**
  - **Report type:** Protocols.
  - **Filter:** Single protocol.
  - **Report subtype:** Overview.
  - **Presentation menu option:** Pie Chart.
  - **View:** Protocol Hosts Summary (From, To, and Total) for the single protocol.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Top Hosts (Table)

The Top Hosts (Table) views show rate, volume, or utilization for the hosts that exchange the highest volume of data on a particular interface. The view is configured to show the rate for each listed host. You can configure the view to display rate, utilization, or volume information.

The view contains an interface identification string and a table. The table contains a row for each host with the fully qualified DNS host name (if available) and IP address, as well as the following rate, volume, or utilization information (by default):

- **Rate:**
  - (CAPC/NPC) For each host, the average rate of total data (Average Total), data that goes to the host (Average To), and data that comes from the host (Average From).
  - (CAPC) For each host, the maximum rate of data that comes from the host (Maximum From) and goes to the host (Maximum To).

The rate is calculated by dividing the data volume by the elapsed transmission time.

- **Volume:** For each host, the total amount of data (Total), data from the host (From), and data to the host (To), expressed in a scale that is appropriate for the highest-volume host.

- **Utilization:**
  - (CAPC/NPC) For each host, the average utilization by data that comes from the host (Average From), data that goes to the host (Average To), and total data (Average Total).
  - (CAPC) For each host, the maximum percentage of interface capacity that is used by the data from the host (Maximum From) or that goes to the host (Maximum To).

The utilization percentage is calculated by dividing the data rate by the data speed.

NetOps Portal views show the data from the time range that is defined for the page.

### Opening the View

To see this view in the NetOps Portal Console, go to one of the following locations:

- (CAPC) Custom dashboard
- (NPC) Interface Pages (with an interface selected): Custom tab

### Available Actions

You can perform several actions in this view:

- Change the type of measurement and the view name by editing the view settings.
- (CAPC) Change the interface.
- (CAPC) Change the columns that are shown in the table: Click near a column border, click Columns, then choose the columns to display.
- Click a name to jump to a pre-filtered Interface page report in the NFA console.

### Find Host Tables in the NFA Console

You can display tables with host volumes in the NFA console for a selected interface:

- **Top N Hosts Summary**
  - **Report type:** Hosts.
  - **Filter:** Top N Hosts.
  - **Presentation menu option:** Summary Table; Volume.
  - **View:** Host Summary Table for the Top N Hosts.
- **Hosts for a Single Protocol**

- **Report type:** Protocols.
- **Filter:** Single protocol.
- **Report subtype:** Overview.
- **Presentation menu option:** Summary Table; Volume.
- **View:** Protocol Host Summary Table for the single protocol.
- **Hosts for a Single ToS**
  - **Report type:** ToS.
  - **Filter:** Single ToS.
  - **Report subtype:** Hosts.
  - **Subtype filter:** Top N Hosts.
  - **Presentation menu options:** Summary Table; Volume.
  - **View:** ToS Hosts Summary Table for the single ToS.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Interface Top Protocols

The Top Protocols views show the protocols that are associated with the highest traffic on the interface that you have selected. The articles in this section describe these views.

### Top Protocols (Bar)

The Top Protocols (Bar) views show the top high-volume IP protocols for traffic on a particular interface. A bar chart shows which protocols account for the most traffic on the selected interface.

This view gives you an overall picture of the amount of data that is associated with particular protocols--and, therefore, with applications--on the interface. The view also lets you determine whether the application protocols are related to business-critical processes, or are related to low-priority or non-business related processes such as unauthorized web use. You can view protocol traffic for incoming flows, outgoing flows, or all flows.

The bar chart includes the following information for a maximum of 10 protocols:

- **Identifier**  
Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).  
(NPC) The identifier line also includes the interface speed.
- **Protocol**  
Identifies the protocol by its descriptor (Y-Axis).
- **Volume**  
Measures the total amount of protocol data expressed in a scale that is appropriate for the highest-volume protocol (X-Axis).

NetOps Portal views show the data from the time range that is defined for the page.

### Opening the View

To see this view in the console, go to one of the following locations:

- (CAPC) Custom dashboard
- (NPC) Interface Pages (with an interface selected): Custom tab

### Available Actions

You can perform several actions in this view, including the following ones:

- Change the traffic direction and the view name by editing the view settings.
- (CAPC) Change the interface.
- Display details in a Tooltip by holding your cursor over a bar.
- Jump to details for a specific protocol on an NFA console Interface page by clicking a bar or name.

### **Find the Comparable View in the NFA Console**

The Top Protocols bar charts in the console are similar to the Top Protocol view on the Enterprise Overview page in the NFA console.

### **Top Protocols (Pie)**

The Top Protocols (Pie) views show the protocols that are associated with the highest traffic volumes on the selected interface.

The table includes the following information by default:

- **Identifier**  
Identifies the interface that is used for the report. The identifier string consists of the router name, interface name, and interface description (under the view title).  
(NPC) The identifier line also includes the interface speed.
- **Protocol Name**  
Identifies the protocol by its keyword and TCP/UDP port assignment.
- **Total**  
Records the total volume of network traffic on the interface that is associated with the protocol
- **Percent**  
(CAPC) Records how much the protocol consumes out of the total traffic that is displayed.

NetOps Portal views show the data from the time range that is defined for the page.

### **Opening the Views**

To see these views in the console, go to one of the following locations:

- (CAPC) Custom dashboard
- (NPC) Interface Pages (with an interface selected): Interface QoS or custom tab

### **Available Actions**

You can perform several actions in this view:

- Change the traffic direction and the view name by editing the view settings.
- (CAPC) Change the interface.
- (CAPC) Change the columns that are shown in the table: Click near a column border, click Columns, and then choose the columns to display.
- Jump to details on an NFA console Interface page by clicking a protocol name.

### **Find Protocol Pie Charts in the NFA Console**

You can display pie charts with protocol traffic volumes in the NFA console for a selected interface:

- **Overview**



- **Report type:** Overview.
- **Presentation menu option:** Pie Chart.
- **View:** Protocol Summary (In and Out) for the Top N Protocols, plus other overview views.
- **Top N Protocol Summaries**
  - **Report type:** Protocols.
  - **Filter:** Top N Protocols.
  - **Presentation menu option:** Pie Chart.
  - **View:** Protocol Summary (In, Out, and Total) for the Top N Protocols.
- **Hosts or Conversations for Single Protocol**
  - **Report type:** Protocols.
  - **Filter:** Single protocol.
  - **Presentation menu option:** Pie Chart.
  - **Views:** Protocol Hosts Summary (From and To) for the single protocol; Protocol Conversations Summary (Total) for the single protocol.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Top Protocols (Table)

The Top Protocols (Table) views show the rate, volume, or utilization for the protocol traffic with the highest volume on a particular interface. For example, you can use this information to compare the data volume or utilization for particular protocols.

The view is set to show the traffic volume for each listed protocol. You can configure the view to display rate, utilization, or volume information.

The view contains an interface identification string and table. The table has a row for each protocol with the protocol name (keyword and TCP/UDP port assignment) and the following rate, volume, or utilization information (by default):

- **Rate:**
  - (CAPC/NPC) Average rate of total (Average Total), inbound (Average In), and outbound data (Average Out) for each protocol.
  - (CAPC) Maximum rate of data that is outbound, inbound, or both outbound and inbound (Maximum Out, Maximum In, and Maximum Total) on the interface for each protocol

The rate is calculated by dividing the data volume by the elapsed transmission time.
- **Volume:** Number of bytes/megabytes of outbound data (Out), inbound data (To), and all data (Total) for each protocol.
- **Utilization:**
  - (CAPC/NPC) Average utilization by inbound data (Average In), outbound data (Average Out), and total data (Average Total) for each protocol.
  - (CAPC) Maximum percentage of interface capacity that the outbound (Maximum Out) or inbound protocol data utilizes (Maximum In)

The utilization percentage is calculated by dividing the data rate by the data speed.

NetOps Portal views show the data from the time range that is defined for the page.

## Opening the View

To see this view in the NetOps Portal Console, go to one of the following locations:

- (CAPC) Custom dashboard
- (NPC) Interface Pages (with an interface selected): Custom tab

## Available Actions

You can perform several actions in this view:

- Change the type of measurement (Rate, Volume, or Utilization) and the view name by editing the view settings.
- (CAPC) Change the interface.
- Re-sort the table data by clicking a column heading. Click again to toggle between descending and ascending order.
- Change the Max Per Page value to show more or fewer items on each table page.
- (CAPC) Change the columns that are shown in the table: Click near a column border, click Columns, and then choose the columns to display.
- Click a name to jump to a pre-filtered Interfaces report in CA Network Flow Analysis.

## Find Protocol Tables in the NFA Console

You can use these ways to display tables of protocol volume data in the NFA console for a selected interface:

- **Top N Protocols**
  - **Report type:** Protocols.
  - **Filter:** Top N Protocols.
  - **Presentation menu options:** Summary Table; Volume.
  - **View:** Protocol Summary Table for the Top N Protocols, plus other overview views.
- **Protocols for a Single Host**
  - **Report type:** Hosts.
  - **Filter:** Single host.
  - **Report subtype:** Protocols.
  - **Presentation menu options:** Summary Table; Volume.
  - **View:** Host Protocol Summary Table for the single host.
- **Protocols for a Single Conversation**
  - **Report type:** Conversations.
  - **Filter:** Single conversation.
  - **Report subtype:** Protocols.
  - **Presentation menu options:** Summary Table; Volume.
  - **View:** Conversation Protocol Summary Table for the single conversation.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

## Calendar Chart (Flow)

The Calendar Heat Chart (Flow) view maps the utilization percentage of the selected interface over time to help you find recurring data patterns. Finding a pattern can help you to identify the source of high traffic rates and potential performance issues. For example, the view can show the hour of each day when utilization is the highest.

Each color represents a severity range that is calculated as a percentage of total capacity. High utilization is shown in orange and red. Low utilization is shown in green and blue.

The view includes the following information:

- **Identifier**  
Consists of the router name, interface name, and interface description (under the view title). The interface description consists of the ifDescr value by default, and can differ from the interface description that is shown in the NFA console.
- **Month, Date, and Day of the Week**  
Identify the day that the traffic occurred (X-Axis columns).
- **Hour**  
Identifies the hour of the day that the traffic occurred (Y-Axis).

To see the Calendar Heat Chart view in the CA NetOps Portal (CAPC) Console, add it to a custom dashboard.

You can perform several actions in this view:

- Change the data direction and view name, as described in this topic.
- (CAPC) Display details in a Tooltip by holding your cursor over a cell.
- (CAPC) Click Show All and choose a pattern-matching filter. For example, select Busy Hour to show only the data for the busiest hour of each day.

### **Find the Comparable View in the NFA Console**

To display Calendar Chart data for an interface in the NFA console, select an interface on the Interface page and select the following options:

- Report type: Utilization.
- Presentation menu option: Direction In or Direction Out.

To display Flow Forensics-level detail, click the Flow Forensics link and run a Flow Forensics report.

### **Change the View Settings**

You can adjust several settings in the Calendar Heat Chart.

#### **Follow these steps:**

1. (CAPC) Click the Edit icon in the view title bar and click Edit.
2. (Optional) Edit any of the following settings in the Calendar Heat Chart (Flow) Settings section:
  - **Title:** Change the name that appears in the view title bar.
  - **(CAPC) Time Display Format:** Select the time format for the chart, either 12 hours or 24 hours.
  - **(CAPC) Zone Start:** Set the starting value of each heat zone. The defaults are based on IT industry standards for performance. For example, the default Red Zone Start value is 70 percent utilization.  
**Defaults:** Green Zone Start = 0, Yellow Zone Start = 50, Orange Zone Start = 60, Red Zone Start = 70.
  - **(CAPC) Business Week Start:** Select the day that starts the business week.  
**Default:** Monday.
  - **(CAPC) Direction Settings:** Select the direction of traffic on the selected interface to include in the report:
    - **Out:** Outbound on the interface.
    - **In:** Inbound on the interface.
    - **Total:** Combination of inbound and outbound traffic.
3. (Optional) (CAPC) Change the context for the view data: Select a different interface from the Context Settings table.
4. Select the scope of your changes from the **Apply Changes** drop-down.
5. Click **Save** to save your changes.  
The settings dialog closes. The view refreshes to reflect your updates.

## **Integrate with CA Unified Communications Monitor**

The topics in this section include information about the views that show CA Unified Communications Monitor data in NetOps Portal.

### **Call Quality Breakdown**

The Call Quality Breakdown pie chart shows a rollup of systemwide call quality for the selected time frame. The percentage of all call minutes that fell into each severity range is displayed in the chart.

The call quality data in this chart comes from the Mean Opinion Scores (MOS) of every call leg that was measured within that time frame. MOS values are evaluated according to the Call Quality performance thresholds that are assigned to the

Locations where call activity occurred. The thresholds determine which MOS values are rated Normal, Minor, or Major. The threshold values in these severity categories can be customized per location or assigned automatically per codec.

MOS values do not apply to video streams. Network MOS is not included.

A unique color is assigned to each severity level shown in the pie chart. A legend explains the color assignments. Data that is unrated can indicate that the Minimum Observations threshold for that location was not met during the selected time frame.

## Call Quality Service Level Agreement

The Call Quality Service Level Agreement (SLA) view lets Managed Service Providers (MSPs) prove to a customer that they are meeting the SLA commitment for audio call quality by codec. MSPs create one Call Quality SLA view for each customer.

The Call Quality SLA view provides the following information:

- **Calls Meeting SLA**  
The number of calls that met the SLA commitment during the selected interval, based on your selections on the Settings dialog.  
If both legs of a call have MOS or Network MOS, then both legs must meet the SLA commitment for the call to be included in the Calls Meeting SLA value.  
Calls are not included in the Calls Meeting SLA value when one leg meets the commitment, but the other leg does not.
- **Calls Not Meeting SLA**  
The number of calls that did not meet the SLA commitment during the selected interval.
- **Total Calls**  
Total number of calls observed during the selected interval. Only Calls that had MOS or Network MOS for at least one of the two legs in the call are included in the total. The total includes the number of calls that met the SLA.

### NOTE

The total does not include calls that do not have MOS or Network MOS, such as short calls, abandoned calls, and calls with setup failures.

The number in this column is a link to the Calls Overview report in the CA UC Monitor management console. The Calls Overview report can contain up to a week's worth of data, and can therefore take a long time to open.

- **Percent**  
The percentage of calls that met the SLA commitment during the selected interval.
- **Importance**  
Indicates whether the SLA commitment was met, or indicates the importance of a failure to meet the SLA, based on your selections in the Settings dialog.

## Call Quality Trend

The Call Quality Trend view provides an enterprise-level view of Mean Opinion Score (MOS) data from all calls detected during the selected timeframe.

### About the Call Quality Trend View

The trend chart does not include an indication of the number of call minutes used to derive the metrics. To determine the number of observations behind the trend, navigate to the UC Monitor data source and view the Call Performance Overview report. Click the Metric Details link at the top of the report. MOS values do not apply to video streams. Network MOS is not included.

**NOTE**

This view does not support data from multiple data sources. If you registered multiple instances of UC Monitor, edit the view to filter by one data source. You can also change the dashboard group context by clicking the Group link. Using either method, you can select one data source in the Groups tree to serve as the view context.

**Call Quality Trend View Settings**

The Settings dialog provides filtering and display options for the view.

- **Title**  
The default title is Call Quality Trend. You can change the title as necessary.
- **Context Settings**  
Select another managed item to change the source of the data in the view.
- **Apply Changes**  
Select the scope of your changes from the **Apply Changes** drop-down.

**Performance Overview Dashboard**

The Performance Overview dashboard provides a network manager or higher-level executive with a daily or weekly summary of overall VoIP and video call performance.

You can click a link in a dashboard view and access the related report in the CA UC Monitor management console, with the appropriate context selected. The management console provides more details for the data that is summarized in the dashboard.

**Performance by Call Server**

The call servers shown in the Performance by Call Server view handled calls during the selected time frame. Their performance ratings are derived from the Locations that they served. The name of each server is a link that allows further investigation.

The Call Performance category includes both call quality and call setup metrics. Where available, video metrics are included. The bar graph represents the number of originated calls or call minutes that were rated Normal, Minor, or Major.

- **Calls Originated:** Calls from endpoints that are registered to a call server.
- **Call Minutes:** Minutes of call activity that were reported by a call server.

The Calls Originated column provides the total number of calls that were set up by the indicated call servers. Calls not included in the Calls Originated total were routed by the call servers in this view, but were set up by a different call server.

**Performance by Group**

The Performance by Group view is available only when custom groups are defined in CA NetOps Portal (CAPC). You can drill down into individual group members and their data by clicking a group name.

The view rates call performance in the incoming direction to gauge the listening quality for VoIP and video users in that group. Groups are sorted by worst call performance.

Click a Location link to see ratings for performance metrics and for the components participating in calls to the Locations in this view.

The Performance by Group view displays calls between Locations in the selected group. The view also identifies all associated call servers, including calls servers that are not explicitly part of the group definition.

## **Performance by Location**

The Performance by Location view lists all monitored Locations where endpoints had call activity. This view evaluates call performance in the incoming direction to gauge the listening quality for VoIP and video users.

Data that is unrated can indicate that the Minimum Observations threshold for that Location was not met during the selected time frame.

To see ratings for performance metrics and for the components participating in calls to the Locations listed in this view, click a Location link.

## **Performance by Media Device**

The Performance by Media Device view shows incoming calls from the PSTN and outgoing calls from the IP network that the device handled.

The media device category includes voice gateways and other devices that support call routing and processing. The metrics available for each device vary according to the device type and the environment. In a Microsoft environment, SNMP polling of media devices is not possible. As a result, fewer metrics are available for Microsoft media devices than for Cisco voice gateways.

The Calls Originated column provides the total number of calls that originated at points in the PSTN.

## **Performance Overview Dashboard View Settings**

The Settings dialog provides filtering and display options for each view in this dashboard:

- **Title**  
The default titles for the views in this dashboard are:
  - **Performance by Call Server**
  - **Performance by Group**
  - **Performance by Location**
  - **Performance By Media Device**
 You can change the titles as necessary.
- **Context Settings**  
Select another managed item to change the source of the data in the view.
- **Apply Changes** Select the scope of your changes from the **Apply Changes** drop-down.

## **Top Volume and Utilization Dashboard**

The Top Volume and Utilization Dashboard helps network engineers plan for network growth and track usage statistics for their unified communications systems. Data from CA UC Monitor can help engineers plan for capacity needs by determining the current operating levels of key unified communications components.

You can click a link in a view to access the related report in the CA UC Monitor management console with the appropriate context selected. The management console provides more details for the data that is summarized in the dashboard.

### **Contents**

#### **Top Groups**

The Top Groups Volume view compares the call volume of the groups of locations with the highest usage during the selected timeframe. The view provides a list of the top talkers at a particular point in time. All calls, including audio-only, video-only, and audio and video, are considered when compiling the list of top call volumes.

The Volume bar chart shows a comparison of call volumes among the groups of locations with the highest volumes. Each bar represents a relative activity level, which lets you easily compare call volumes among busy Locations.

The names in the Group column are links to the Top Volume report in the CA UC Monitor interface. You can see the volume statistics of individual group members in the context views included in this report.

By default, the Volume bar is calculated using the number of calls placed by locations in the indicated groups. You can select Call Minutes as the charted unit instead. Click the blue arrow to the left of the view name, and select Edit. In the dialog that opens, select Call Minutes from the “Calculate using” list.

### **Top Locations**

The Top Locations Volume view compares the call volume of the locations with the highest usage during the selected timeframe. The view provides a list of the top talkers at a particular point in time. All calls, including audio-only, video-only, or audio and video, are considered when compiling the list of top call volumes.

The Volume bar chart shows a comparison of call volumes among the locations with the highest volumes. Each bar represents a relative activity level, which lets you easily compare call volumes among busy locations.

The names in the Location/Media Device column are links to the Top Volume report. A related view of the busiest endpoints from the selected Location is also included in the report.

By default, the Volume bar is calculated using the number of calls placed by Locations in the indicated groups. You can select Call Minutes as the charted unit instead. Click the blue arrow to the left of the view name, and select Edit. In the dialog that opens, you can select Call Minutes from the “Calculate using” list.

### **Top Phones**

The Top Phones Volume view compares the call volume of the endpoints with the highest usage during the selected timeframe. The view provides a list of the top talkers at a particular point in time. All calls, including audio-only, video-only, and audio and video, are considered when compiling the list of top call volumes.

The Volume bar chart shows a comparison of call volumes among the endpoints with the highest volumes. Each bar represents a relative activity level so that you can compare call volumes among endpoints.

The numbers in the Phone Number column are links to the Phones report.

If no directory number is available for an endpoint, its IP address is provided in the IP Address column. The IP address can be a link to the endpoint web page. The endpoint name is also shown in the Name column.

### **Top Trunk Groups**

The Top Trunk Groups view helps you understand trunk group usage, and provides individual trunk usage statistics. Use this page to find underutilized trunk groups and overburdened trunks, where performance can deteriorate.

Statistics for all known trunk groups are included. The most heavily used trunks are shown first.

- **Name**

The name assigned to the trunk group. Click to drill down into an hourly breakdown of usage for this trunk group. The name of the trunk group is based on information from the monitored call traffic and the naming convention employed by the trunking equipment.

- **Utilization (%)**

The average usage during the time period, expressed as a percentage of capacity. Average usage is based on the capacity divided by the timeframe.

- *(Avaya only)* The trunk group capacity is discovered from SNMP polling of the Avaya Communication Manager. The usage for an Avaya trunk group is based on channel capacity, divided by the Centum Call Seconds (CCS) observed

for all active channels. Usage is a percentage of the trunk group channel capacity, divided by the observed CCS. The total is again divided by the selected timeframe.

- (*Cisco only*) Usage is the percentage of total voice interface capacity in use during the selected timeframe. Total interface capacity is derived from the known capacity of each voice interface in the group.

To identify underused trunk groups, sort this column to see the least used trunk group first.

- **Maximum Utilization**

The highest recorded usage during the selected time period, expressed as a percentage of capacity. Maximum usage is a good indicator of potential capacity issues. Even when the average usage is low, the maximum usage is significant because it often indicates the usage during the “busy hour” of the day. Maximum usage is based on the capacity, multiplied by the timeframe.

- (*Avaya only*) The capacity is based on information from the Avaya Communication Manager.
- (*Cisco only*) The capacity is discovered when voice gateway devices are discovered.

- **Grade of Service (GoS)**

(*Cisco only*) An estimation of the probability that a VoIP call will receive a busy signal. The GoS value (a decimal fraction) is always expressed with reference to the busy hour when the traffic intensity is the greatest. GoS is reported from the perspective of the origination Location or gateway device (the outgoing direction).

- **Call Minutes**

The number of call minutes used to calculate usage statistics. This value provides a sense of the scope of activity, and helps to determine the significance of the data, based on sample size.

- **Call Minutes Capacity**

The number of call minutes that were supported during the selected time period. The capacity is discovered during monitoring, or is manually supplied during configuration.

## **Top Voice Interfaces**

The Top Voice Interfaces view helps you to understand voice gateway usage, and provides individual gateway voice interface statistics. Use this view to look for overburdened interfaces, where performance can deteriorate. The minimum timeframe for this view is one day; no data appears if a narrower timeframe is selected. This view contains:

- Statistics for all known gateway voice interfaces.
- The most heavily used interfaces, which are shown first.

## **Top Volume and Utilization View Settings**

The Settings dialog provides filtering and display options for the views in this dashboard:

- **Title**

The default titles for the views in this dashboard are:

- **Top Groups**
- **Top Locations**
- **Top Phones**
- **Top Trunk Groups**
- **Top Voice Interfaces**

You can change the title as necessary.

- **Context Settings**

Select another managed item to change the source of the data in the view.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.



## Worst Performance Dashboard

As an IT staff, you can monitor the quality of unified communications system performance and summarize the performance data using the **Worst Performance Dashboard**. By default, this dashboard provides the **Worst Locations** and **Worst Phones** views, which focus on the locations and phones with the worst performance.

You can access the related report in the CA Unified Communications Monitor management console, with the appropriate context selected by clicking a link in a view. The management console provides more details of the data that is summarized in the dashboard.

In this article:

- [Worst Locations](#)
- [Edit the Settings for the Worst Locations View](#)
- [Worst Phones](#)
- [Edit the Settings for the Worst Phones View](#)

### Worst Locations

The **Worst Locations** view shows the pairs of locations that had the lowest-quality metrics for the data traveling between them. This view provides the following information:

- **Name**  
Displays the name of the location or media device that received the data stream with poor performance metrics. The **Name** can also refer to the name of the origination location for a data view filtered by a call setup metric.
- **Sending Name**  
Displays the name of the location or media device that sent the low-performing call data to the other location in the pair.
- **Call Server**  
Displays the call server that handled the calls for the shown severity breakdown.
- **Call Minutes**  
Displays the number of minutes that calls were active between this pair of locations.
- **Calls**  
Displays the number of distinct calls that ran between this pair of locations, and that also contributed to the worst performance metric displayed in the table.
- **Severity Breakdown**  
A stacked bar chart that shows applicable severity ratings as portions of the bars, which total 100 percent. Severity ratings are color-coded to match the severity indicators in other reports.
- **Unrated, Normal, Minor, Major**  
Shows the actual severity percentages for the selected quality metric type. Percentages always total 100 percent.

#### **NOTE**

An unrated metric indicates that a threshold is disabled, or that the threshold for minimum observations is too high for typical levels of call traffic.

### Edit the Settings for the Worst Locations View

The **Settings** dialog provides filtering and display options for the view. To view this dialog, click the **View Settings** (gear) icon on the view, and then select **Edit** from the menu:

- **Title**  
The default title is **Worst Locations**. You can change the title as necessary.
- **Metric Type**  
Select a different metric to change the type of data in the view.
- **Context Settings**

Select another managed item to change the source of the data in the view.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.

For more information about this field, see [Customize Views](#).

## **Worst Phones**

The **Worst Phones** view displays a list of endpoints with the lowest MOS for the selected timeframe. Endpoints are identified by the directory number or Session Initiation Protocol Uniform Resource Identifiers (SIP URI), and by their IP address.

- **Phone Number**

Displays the directory number or SIP URI of the endpoint.

- **Name**

Displays the name, MAC address, or host name of the endpoint.

- **IP Address**

Displays the IP address of the endpoint.

- **Call Minutes**

Displays the number of minutes that calls were active on this endpoint during the selected timeframe.

- **Calls**

Displays the number of distinct calls placed or received.

- **Average Metric**

Displays the average MOS or Network MOS for all calls to and from this endpoint during the selected timeframe. Unlike the severity reflected in the bar chart, the average is not weighted by the duration of the calls.

**NOTE**

The contents of this column are determined by your selection in the **Metric Type** field on the **Settings** dialog.

- **Severity Breakdown**

A stacked bar chart that shows each severity rating as a portion of the bar, which totals 100 percent. The severity ratings are color-coded to match the severity indicators in other reports.

- **Unrated, Normal, Minor, Major**

Depending on the selection in the **Show breakdown values as** field on the **Settings** dialog, these columns display the percentage of calls in a severity category, or the number of calls in a severity category.

**NOTE**

An unrated metric indicates that a threshold is disabled, or that the threshold for minimum observations is too high for typical levels of call traffic. Set a lower minimum for observations, or assign different thresholds to the endpoints in question.

## **Edit the Settings for the Worst Phones View**

The **Settings** dialog provides filtering and display options for the view. To view this dialog, click the **View Settings** (gear) icon on the view, and then select **Edit** from the menu:

- **Title**

The default title is **Worst Phones**. You can change the title as necessary.

- **Metric Type**

Select **MOS** or **Network MOS** to determine the contents of the **Average Metric** and **Severity Breakdown** columns.

- **Show breakdown values as**

Determines the contents of the severity categories (the **Unrated**, **Normal**, **Degraded**, and **Excessive** columns).

- **Calls:** Displays the number of calls in a severity category.

- **Percent:** Displays the percentage of calls in a severity category. Percentages always total 100 percent.

- **Context Settings**

Select another managed item to change the source of the data in the view.

- **Apply Changes**

Select the scope of your changes from the **Apply Changes** drop-down.

For more information about this field, see [Customize Views](#).

## Monitor Server Performance with DX Application Performance Management

You can configure monitoring NetOps Portal, the data aggregator, and the data collector for instrumentation in DX Application Performance Management (APM).

You can configure monitoring NetOps Portal, the data aggregator, and the data collector for instrumentation in DX Application Performance Management (APM). APM receives performance metrics about these DX NetOps Performance Management components and hosts. During the installation of these components, you can have the installer configure the APM agent. You can also configure the agent for these components post-installation.

### NOTE

After each configuration, ensure that your changes have taken effect by restarting the corresponding services.

### Follow these steps:

1. Log in to the host for the relevant component (NetOps Portal, the data aggregator, and the data collector DX NetOps Performance Management).
2. Change to the working directory for the component for which you want to configure monitoring by issuing the following command:

```
cd <installation_directory>
```

### Examples:

```
cd /opt/IMDataAggregator
```

```
cd /opt/CA/PerformanceCenter
```

### – **installation\_directory**

The installation directory of the component.

**Examples:** /opt/IMDataAggregator, /opt/CA

3. Run the `setEMHost.sh` script.
4. When prompted, enter the following parameters:
  - CA APM hostname
  - CA APM port

**Default:** 5001
5. On the NetOps Portal host, change to the working directory for APM by issuing the following command:

```
cd <APM_installation_directory>
```

### Example:

```
cd /opt/CA/PerformanceCenter/wily
```

### – **APM\_installation\_directory**

The installation directory for APM.

**Default:** /opt/CA/PerformanceCenter/wily

6. Configure monitoring for NetOps Portal by completing the following steps:
  - a. Configure monitoring for NetOps Portal services by issuing the following commands:
 

```
./addWily.sh <installation_directory>/PerformanceCenter
<installation_directory>/PerformanceCenter/PC/conf/wrapper-user.conf PC
./addWily.sh <installation_directory>/PerformanceCenter
<installation_directory>/PerformanceCenter/sso/conf/wrapper-user.conf sso
```

```
./addWily.sh <installation_directory>/PerformanceCenter
<installation_directory>/PerformanceCenter/DM/conf/wrapper-user.conf DM
./addWily.sh <installation_directory>/PerformanceCenter
<installation_directory>/PerformanceCenter/EM/conf/wrapper-user.conf EM
```

- **installation\_directory**

The installation directory for NetOps Portal.

**Default:** /opt/CA

- b. Restart the NetOps Portal services by issuing the following commands:
 

```
systemctl restart caperfcenter_console
systemctl restart caperfcenter_devicemanager
systemctl restart caperfcenter_eventmanager
systemctl restart caperfcenter_sso
```
7. (If you have configured monitoring for the data aggregator) Do the following:
  - a. Stop and restart ActiveMQ by issuing the following commands:
 

```
systemctl stop activemq
systemctl start activemq
```
  - b. Do one of the following steps based on your configuration:
    - Stop the Data Aggregator service by issuing the following command:
 

```
systemctl stop dadaemon
```
    - (Fault-tolerant environment) If the local data aggregator is running, issue the following command to shut it down and prevent it from restarting until maintenance is complete:
 

```
<da_installation_directory>/scripts/dadaemon maintenance
```
  - c. Do one of the following steps based on your configuration:
    - Start the Data Aggregator service by issuing the following command:
 

```
systemctl start dadaemon
```
    - (Fault-tolerant environment) Issue the following command to enable the fault-tolerant data aggregator so that it can start when necessary:
 

```
<installation_directory>/scripts/dadaemon activate
```
8. If you have configured monitoring for the data collector, stop and restart the Data Collector services by issuing the following commands:
 

```
systemctl stop activemq
systemctl start activemq
systemctl stop dcmd
systemctl start dcmd
```

## Generate Mediation Manager Device Packs

DX NetOps Mediation Manager (DX NetOps MM) uses device packs to monitor data from non-SNMP devices or obtain data from the Element Management System (EMS). DX NetOps MM includes device packs that support many vendors and device types. To support a new device type or a new vendor, create and deploy device packs in your environment. Use the DX NetOps MM device pack generator.

For more information, see the [DX NetOps Mediation Manager documentation](#).

## Integrate with Virtual Network Assurance

Display data for software-defined networking (SDN) and network functions virtualization (NFV) controllers and orchestrators by integrating DX NetOps Performance Management with DX NetOps Virtual Network Assurance (VNA).

Use the following process to integrate with VNA:

1. [Add a connection to VNA.](#)
2. [Configure a plug-in for each technology in your virtual network environment.](#)
3. Verify VNA data integration with NetOps Portal to confirm that VNA plug-in data is populating dashboards.

The following video shows how to verify the integration:

With DX NetOps Performance Management integrated with VNA and a plug-in configured, you can view and monitor VNA performance data related to the SDN and NFV controllers and orchestrators from the dashboards in NetOps Portal.

For more information about how to view and monitor VNA performance data from these dashboards, see [Modern Network Monitoring](#).

## Manage Connections to Virtual Network Assurance

You can manage the VNA Gateways, the plug-ins that are configured in the VNA Gateways, (23.3.3 and higher) the import rules, and (23.3.4 and higher) the user domains in the VNA Gateways.

Active *VNA Gateways* are the connections to DX NetOps Virtual Network Assurance (VNA).

You can manage (configure, edit, and delete) the *VNA Gateways* in the following ways:

- [Configure a VNA Gateway](#)
- [View a List of Configured VNA Gateways](#)
- [Edit a VNA Gateway](#)
- [Delete a VNA Gateway](#)

You can manage (configure, edit, and delete) the *plug-ins in a VNA Gateway* in the following ways:

- [Configure a Plug-in in a VNA Gateway](#)
- [View a List of the Plug-ins Configured in a VNA Gateway](#)
- [Edit a Plug-in](#)
- [Delete a Plug-in](#)
- (23.3.4 and higher) [Start a Plug-in](#)
- (23.3.4 and higher) [Stop a Plug-in](#)

(23.3.3 and higher) *Import rules* define to which IP domain the data aggregator maps groups of VNA inventory. They ensure the following:

- The data aggregator places VNA inventory into the correct IP domain and that it reconciles the inventory correctly with the SNMP device that matches it.
- Any alarms on devices in Spectrum are synched to the correct items/inventory in NetOps Portal.

Data collectors are assigned to an IP domain. The IP domain defines the IP space of the SNMP items/inventory that the data collector monitors. The data collector (the IP domain) to which the VNA Gateway is connected receives the VNA inventory from that VNA Gateway. VNAs monitor controllers, and those controllers can monitor SNMP items/inventory that span IP domains. With defined import rules, the VNA can monitor through the controller SNMP items/inventory outside the IP domain to which the data collector has access.

Based on the import rule, if the data aggregator moves a device from one IP domain to another, it also generates an event noting that. You can enable and disable the groups to which you can apply an import rule (add a group to an import rule).

For more information about how to manage the allowed groups for import rules, see [the DX NetOps Virtual Network Assurance documentation](#).

(23.3.3 and higher) You can manage (add, edit, and delete) the *import rules in a VNA Gateway* in the following ways:

- [Create an Import Rule in a VNA Gateway](#)
- [View a List of Import Rules in a VNA Gateway](#)
- [Edit an Import Rule in a VNA Gateway](#)
- [Delete an Import Rule in a VNA Gateway](#)

(23.3.4 and higher) You can manage (create, update, and view) *user domains in a VNA Gateway* in the following ways:

- [Create a User Domain in a VNA Gateway](#)
- [View a List of User Domains in a VNA Gateway](#)
- [Update the Name of a User Domain in a VNA Gateway](#)
- (23.3.5 and higher) [Get Inventory Statistics for a User Domain in a VNA Gateway](#)

(23.3.6 and higher) You can manage (add, edit, view details, delete) *the notification filters for a VNA Gateway* in the following ways:

- [View a List of Notification Filters in a VNA Gateway](#)
- [Add a Notification Filter to a VNA Gateway](#)
- [Edit a Notification Filter in a VNA Gateway](#)
- [View Notification Filter Details in a VNA Gateway](#)
- [Delete a Notification Filter from a VNA Gateway](#)

Global notification filters of alarms and events are a way for you to build filtering capabilities across plug-in.

## **Configure a VNA Gateway**

You can configure a VNA Gateway using the following methods:

- [Configure a VNA Gateway using NetOps Portal](#)
- [Configure a VNA Gateway through REST](#)

### **Configure a VNA Gateway using NetOps Portal**

**Follow these steps:**

1. Hover over **Administration, Monitored Items Management**, and then click **VNA Administration**.  
The **VNA Administration** page appears. The **<VNA Gateways>** view shows the VNA Gateways that are configured in NetOps Portal.
2. Click **New**.  
The **Configure VNA Gateway** dialog opens.
3. Complete the following fields, and then click **Save**:
  - **Host Name / IP Address**  
Specifies the host for the VNA Gateway.

**Required:** Yes

– **Port**

Specifies the port for the VNA Gateway.

**Default Port:** 8080

**Required:** Yes

– **Protocol**

Specifies the protocol for the VNA Gateway.

**Options:** http or https

**Default:** http

**Required:** No

– **Enable Authentication**

Defines whether the VNA Gateway is enabled for authentication. If you have configured VNA with a REST API password, accept the default.

**Default:** Selected

**Required:** No

In the **Authentication** section, complete the following fields:

- **Username**

Enter the username for the REST API user.

**Required:** Yes (if authentication is enabled)

- **Password**

Enter the REST API password.

**Required:** Yes (if authentication is enabled)

- **Confirm Password**

Enter the REST API password again.

**Required:** Yes (if authentication is enabled)

In the **Gateway Configuration** section, complete the following fields:

- **Data Collector**

Specifies the data collector that receives the VNA data (the data collector to which the VNA Gateway connects). You configure, or add, these VNA Gateways to a data collector. You can add multiple VNA Gateways to a data collector, however, only one associated VNA Gateway can be active at a time.

**Required:** Yes

**NOTE**

If you have configured the data collectors for fault tolerance, and this data collector is part of a standby group, and the active data collector fails over to a standby data collector, VNA will reconnect to the data collector that becomes active.

For more information about data collectors for fault tolerance, see [Configure the Data Collectors for Fault Tolerance](#).

- **Administrative State**

Defines the administrative state of the VNA Gateway. The administrative state determines whether the data collector receives inventory and performance data from the VNA Gateway.

**Options:**

- **Up:** The data collector registers the VNA Gateway and receives inventory and performance data from the VNA Gateway.

- **Down:** The data collector does not register the VNA Gateway and does not receive inventory and performance data from the VNA Gateway.

**Default:** Up

**Required:** Yes

- **Life Cycle State**

Specifies the life cycle state. The life cycle state determines whether the data collector processes the inventory and performance data that it receives from the VNA Gateway. Only one VNA Gateway can be "Active" at a time.

**Options:**

- **Active:** The data collector processes data from the VNA Gateway.
- **Retired:** The data collector does not process data from the VNA Gateway.

**Default:** Active

**Required:** Yes

The VNA Gateway is configured and added to the data collector.

### **Configure a VNA Gateway through REST**

Issue a POST to the following data aggregator REST web services `sdnGateways` endpoint, with the following Body:

`http://<DA_host>:8581/rest/tenant/<tenantID>/sdnGateways`

- **tenantID**

The ID of the tenant to which the data collector belongs. For the default tenant, the tenant ID is usually 1 .

**TIP**

You can get a list of the tenants, including the tenant ID, using the `http://<DA_host>:8581/rest/tenants` REST URL.

For more information, see [Tenants Web Service](#)

**Body:**

```
<SDNGateway version="1.0.0" >
  <Item version="1.0.0">
    <Name>GatewayName</Name>
    <MDRItemID>DC_ItemID</MDRItemID>
  </Item>
  <AdminStatus>status</AdminStatus>
  <LifeCycleState>state</LifeCycleState>
  <Hostname>host</Hostname>
  <Port>port</Port>
  <Protocol>protocol</Protocol>
  <Username>username</Username>
  <Password>password</Password>
</SDNGateway>
```

**Example:**

```
<SDNGateway version="1.0.0" >
  <Item version="1.0.0">
    <Name>SDN GW 2</Name>
    <MDRItemID>762</MDRItemID>
  </Item>
  <AdminStatus>UP</AdminStatus>
  <LifeCycleState>ACTIVE</LifeCycleState>
  <Hostname>sdn-gw2</Hostname>
  <Port>8080</Port>
  <Protocol>http</Protocol>
  <Username>admin</Username>
  <Password>admin password</Password>
</SDNGateway>
```



- **Name**  
Assigns a name to the VNA Gateway.
- **MDRItemID**  
Specifies the data collector that will receive the VNA data (the data collector to which the VNA Gateway connects).  
**NOTE**  
 If you have configured the data collectors for fault tolerance, and this data collector is part of a standby group, and the active data collector fails over to a standby data collector, VNA will reconnect to the data collector that becomes active.  
 For more information about data collectors for fault tolerance, see [Configure the Data Collectors for Fault Tolerance](#).
- **AdminStatus**  
Defines the administrative state of the VNA Gateway. The administrative state determines whether the data collector receives inventory and performance data from the VNA Gateway.  
**Values:**
  - **UP**: The data collector registers the VNA Gateway and receives inventory and performance data from the VNA Gateway.
  - **DOWN**: The data collector does not register the VNA Gateway and does not receive inventory and performance data from the VNA Gateway.
- **LifeCycleState**  
Specifies the lifecycle state. The lifecycle state determines whether the data collector processes the inventory and performance data that it receives from the VNA Gateway. Only one VNA Gateway can be "Active" at a time.  
**Values:**
  - **ACTIVE**: The data collector processes data from the VNA Gateway.
  - **RETIRED**: The data collector does not process data from the VNA Gateway.
- **Hostname**  
Specifies the host for the VNA Gateway.
- **Port**  
Specifies the port for the VNA Gateway.  
**Default:** 8080
- **Protocol**  
Specifies the protocol for the VNA Gateway.  
**Default:** http
- **Username**  
Specifies the username for the VNA Gateway.
- **Password**  
Specifies the password for the VNA Gateway.

### **View a List of Configured VNA Gateways**

View the list of configured VNA Gateways on the **VNA Administration** page, in the **<VNA Gateways>** view.

This view shows the following information:

- **Host**  
Displays the host for the VNA Gateway.
- **Port**  
Displays the port for the VNA Gateway.
- **Protocol**  
Displays the protocol for the VNA Gateway.
- **Data Collector**

- **Administrative Status**  
Displays the administrative state of the VNA Gateway. The administrative state determines whether the data collector is registered with the VNA Gateway and whether the data collector receives inventory and performance data from the VNA Gateway.
- **Operational Status**  
Displays the operational state of the VNA Gateway. The operational state specifies whether the VNA Gateway is up (running) or is down.
- **Life Cycle State**  
Displays the lifecycle state. The lifecycle state determines whether the data collector processes the inventory and performance data that it receives from the VNA Gateway. Only one VNA Gateway can be "Active" at a time.
- **Version**  
Displays the version of the VNA Gateway.
- **Last Inventory Update Time**  
Displays the date and time the data collector last processed the inventory and performance data that it received from the VNA Gateway.
- **Creation Time**  
Displays the date and time the VNA Gateway was configured.

(23.3.6 and higher)

## page

1038

VNA Administration page Notification

Home

Alarms

Performance

Inventory

Reports

System Health

Administration

VNA Administration

Monitored Inventory

Monitoring Configuration

Threshold Profiles

System Status

Connectors

VNA Gateways

Host	Port	Prot...	Data Collector	Admini...	Operati...	Life Cyc...	Ver
	8080	http	netops-long-dc4	Up	Up	Active	23
	8080	http	netops-long-dc3	Up	Up	Active	23
	8080	http	netops-long-dc2	Up	Up	Active	23

VNA Plugins

VNA Import Rules

VNA Domains Management

VNA Notification Filters

Notification Filters on netops-long-vna.netops.broadcom.net

No Data To Display

Filter  
(23.3.5 and higher)  
VNA Administration

Home

Alarms

Performance

Inventory

Reports

System Health

Administration

VNA Administration

Monitored Inventory

Monitoring Configuration

Threshold Profiles

System Status

Connectors

VNA Gateways

Host	Port	Prot...	Data Collector	Admini...	Operati...	Life Cyc...	Version
	8080	http	netops-long-dc4	Up	Up	Active	23.3.5-RELEASE Build:
	8080	http	netops-long-dc2	Up	Down	Active	23.3.5-RELEASE Build:
	8080	http	netops-long-dc3	Up	DC Con	Active	23.3.5-RELEASE Build:

VNA Plugins

VNA Import Rules

VNA Domains Management

User Domains on netops-long-vna.netops.broadcom.net

Domain Name	Domain ID
<input type="checkbox"/> Default Domain	
<input type="checkbox"/> Meraki Domain	
<input type="checkbox"/> VipsterDomain	
<input type="checkbox"/> vSphereDomain	
<input type="checkbox"/> ACIDomain	
<input type="checkbox"/> DNAC Domain	
<input type="checkbox"/> meraki_bandwidth	

page  
(23.3.4 and higher)

VNA Administration

DX NetOps

HomeAlarmsPerformanceInventoryReportsSystem HealthAdministration

JP

?

1

Search

VNA Administration

Monitored Inventory

Monitored Devices

Aggregated Components

Discovery Profiles

Monitoring Configuration

Monitoring Profiles

Metric Families

Vendor Certifications

Collections

VNA Administration

Flow Application Mapping

Threshold Profiles

Threshold Profiles

System Status

Data Aggregator

Data Collectors

Connectors

Syslog Configuration

VNA Gateways

Quick Filter

NewEditDelete

Host	Port	Prot...	Data Collector	Admini...	Operati...	Life Cyc...	Versi...
	8080	http	netops-long-dc	Up	Up	Active	23.3.2
	8080	http	netops-long-dc2	Up	Down	Active	23.3.2
	8080	http	netops-long-dc3	Up	Up	Active	23.3.2

VNA Plugins

VNA Import Rules

VNA Domains Management

VNA Plugins on

Quick Filter

NewStartStopEditDelete

Name	Status	Configuration Id	Description	Configuration
ACI Plugin	Running		Created by CAPM - PL...	"APIC_HOST_IP":"131...
Fortinet Plugin	Running		Fortinet	"DOMAIN_ID":"1","FO...
Meraki Plugin	Running		merakiDevnet	"APIGEE_PROXY_AUT...
Viptela Plugin	Running		viptela	"BATCH_COUNT":"10...
vSphere Plugin	Running		vSphere	"APPLY_FILTER":"fals...

10 per page

Page 1 of 1

Displaying 1 - 3 of 3

20 per page

Page 1 of 1

Displaying 1 - 5 of 5

BROADCOM

DX NetOps

Copyright © 2023 Broadcom. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. Privacy Terms of Use

page  
(23.3.3)

VNA Administration

DX NetOps

VNA Administration

Monitored Inventory

Monitoring Configuration

Threshold Profiles

System Status

Monitored Devices

Aggregated Components

Discovery Profiles

Monitoring Profiles

Metric Families

Vendor Certifications

Collections

VNA Administration

Flow Application Mapping

Threshold Profiles

Data Aggregator

Data Collectors

VNA Gateways

VNA Plugins on

Import Rules on

Quick Filter

New

Edit

Delete

Host	Port	Protocol	Data Collector	Administrative ...	Operational Sta...	Life Cycle State	Version	Last Inventory ...	Creation Time
	8080	http	netops-long-dc	Up	Up	Active	23.3.3-RELEASE ...	Oct 27, 2023 9:...	Oct 6, 2022 6:2...
	8080	http	netops-long-dc2	Up	Up	Active	23.3.3-RELEASE ...	Oct 27, 2023 2:...	Oct 6, 2022 6:2...
	8080	http	netops-long-dc3	Up	Up	Active	23.3.3-RELEASE ...	Oct 27, 2023 9:...	Oct 6, 2022 6:2...

10 per page

Page 1 of 1

Displaying 1 - 3 of 3

Quick Filter

New

Edit

Delete

Name	Configuration Id	Description	Configuration
Aruba Plugin		arubaBRCM	"ARUBA_ACCESS_TOKEN":"WdFSGUUr5uIKP...
AWS Plugin		AWS sandbox	"AWS_ACCESS_KEY":"KQ8DWQ56Og7GazTtaF...
Nuage Plugin		nuage_descr	"AVAILABILITY_DELTA_TIME":"300","AVAILABI...
One28T Plugin		one28T_descr	"CONDUCTOR_HOST":"servicesim.netops.bro...
SilverPeak Plugin		ds	"DOMAIN_ID":"48084","ENABLE_REST_API_EV...
VeloCloud Plugin		Created by CAPM - Plugin: VeloCloud Plugin	"DOMAIN_ID":"1","ENABLE_OUT_OF_BAND_M...

20 per page

Page 1 of 1

Displaying 1 - 6 of 6

Quick Filter

New

Edit

Delete

Name	Description	Status	Editor	Last Edit Date & Time
Aruba Rule		ACTIVE		Oct 24, 2023, 5:43:12 PM Greenwic...

page  
(23.3.2 and lower)

VNA Administration

Monitored Inventory

Monitored Devices

Aggregated Components

Discovery Profiles

Monitoring Configuration

Monitoring Profiles

Metric Families

Vendor Certifications

Collections

VNA Administration

Flow Application Mapping

Threshold Profiles

Threshold Profiles

System Status

Data Aggregator

Data Collectors

VNA Gateways

Quick Filter

New

Edit

Delete

Host	Port	Protocol	Data Collector	Administrative ...	Operational Sta...	Life Cycle State	Version	Last Inventory ...	Creation Time
	8080	http		Up	Up	Active	23.3.2-RELEASE ...	Oct 9, 2023 5:04...	Oct 6, 2022 6:23...
	8080	http		Up	Up	Active	23.3.2-RELEASE ...	Oct 4, 2023 5:11...	Oct 6, 2022 6:24...
	8080	http		Up	Up	Active	23.3.2-RELEASE ...	Oct 9, 2023 5:09...	Oct 6, 2022 6:24...

10 per page

Page 1 of 1

Displaying 1 - 3 of 3

VNA Plugins on netops-long-vna

Quick Filter

New

Edit

Delete

<input type="checkbox"/>	Name	Configuration Id	Description	Configuration
<input checked="" type="checkbox"/>	ACI Plugin	ACI	Created by CAPM - Plugin: ACI Plugin	"APIC_HOST_IP"
<input type="checkbox"/>	DNAC Plugin	DNAC	DNAC	"DNAC_HOST_I"
<input type="checkbox"/>	Fortinet Plugin	Fort	Fortinet	"DOMAIN_ID~"
<input type="checkbox"/>	Meraki Plugin	Mer	merakiDevnet	"APIGEE_PROX"
<input type="checkbox"/>	Viptela Plugin	Vipt	viptela	"BATCH_COUN"
<input type="checkbox"/>	vSphere Plugin	vSp	vSphere	"APPLY_FILTER"

page

Edit a VNA Gateway

Follow these steps:

1. From the **<VNA Gateways>** view, click the VNA Gateway (the connection) that you want to edit, and then click **Edit**.  
The **Configure VNA Gateway** dialog opens.

2. Edit the details, such as the life cycle state, and then save your changes.

Your changes to the VNA Gateway are saved.

Delete a VNA Gateway

Deleting an active VNA Gateway (the connection) deletes the VNA Gateway, the VNA historical performance data, and the following associated inventory (the VNA items):

- All VNA-only-managed devices and components from the data aggregator.
- VNA-only-managed components that are on SNMP-managed devices from the data aggregator.
- SD-WAN item types, such as tunnels, from the data aggregator.

TIP

As an alternative to deleting a VNA Gateway, you can preserve the inventory and performance data for the VNA Gateway but prevent the data collector from processing data from the VNA Gateway by setting the life cycle state of the VNA Gateway to "Retired".

1042

For more information, see [Edit a VNA Gateway](#).

#### NOTE

If you plan to uninstall VNA, delete related and active VNA Gateways (the connections) *after* uninstalling VNA. If you later reinstall VNA, add the VNA Gateway in NetOps Portal *after* reinstalling VNA.

For more information about how to uninstall VNA, see [the DX NetOps Virtual Network Assurance documentation](#).

#### Follow these steps:

1. From the **<VNA Gateways>** view, click the VNA Gateway (the connection) that you want to delete, and then click **Delete**.

The **Delete VNA Gateway** dialog opens.

2. Click **Yes**.

The VNA Gateway (the connection) is deleted.

#### Configure a Plug-in in a VNA Gateway

Configuring a plug-in configures it in a VNA Gateway, as a connection to VNA. When you configure a plug-in, VNA uses the configuration details to start the plug-in, to connect to the virtual network technology in your environment, and to begin collecting performance and inventory data.

You can configure a plug-in using the following methods:

- Use NetOps Portal
- [Use the DX NetOps Virtual Network Assurance REST API](#)

This procedure describes how to configure a plug-in using NetOps Portal.

#### Prerequisites:

- You have [set up the virtual network technology in your environment](#).
- You have [configured the associated VNA Gateway](#).

On the **VNA Administration** page, the **<VNA Gateways>** view shows the VNA Gateways that are configured in NetOps Portal. NetOps Portal queries the VNA Gateway for the plug-in configuration JSON for each configured plug-in.

#### Follow these steps:

1. From the **<VNA Gateways>** view, select the VNA Gateway in which you want to configure a plug-in. The list of plug-ins in the **VNA Plugins on <VNA gateway name>** view refreshes to show the plug-ins that are configured in the VNA Gateway.

2. Click **New**.

The **Configure VNA Plugin** dialog opens.

3. Complete the following fields, and then click **Save**:

##### – Plugin Name

Specifies the name of the virtual network technology in your environment to discover.

**Options:** (Cisco) ACI Plugin, (23.3.2 and higher) AppNeta Plugin, Aruba (Central) Plugin, (Amazon Web Services) AWS Plugin, BView (Broadcom BroadView) Plugin, Contrail (OpenContrail) Plugin, (Cisco) DNAC Plugin, Fortinet Plugin, (Cisco) Meraki Plugin, MoD (Broadcom Mirror on Drop) Plugin, NSXT (VMware NSX-T) Plugin, Nuage Plugin, One28T (128T SD-WAN) Plugin, OpenDaylight Plugin, Openstack (OpenStack) Plugin, SilverPeak (Silver Peak) Plugin, (VMware) VeloCloud Plugin, Versa (SD-WAN) Plugin, Viptela Plugin, (VMware) vSphere Plugin

##### – Plugin Config

Displays the configuration details that VNA uses to start the plug-in, to connect to the virtual network technology in your environment, and to begin collecting performance and inventory data. NetOps Portal populates this field with the plug-in configuration JSON that it retrieves from VNA.

**TIP**

You can get a description of the properties for the plug-in configuration JSON in the topic for virtual network technology in your environment.

For more information, see [the DX NetOps Virtual Network Assurance documentation](#).

The plug-in is configured in the VNA Gateway.

**View a List of the Plug-ins Configured in a VNA Gateway**

The configured plug-ins display on the **VNA Administration** page.

For more information about how to view this list, see [the "View a List of Configured VNA Gateways" section](#).

**Edit a Plug-in****Follow these steps:**

1. From the **VNA Gateways** view, click the VNA Gateway (the connection) associated to the plug-in that you want to edit. The associated plug-ins display in the **VNA Plugins on <VNA gateway name>** view.
2. Select the plug-in that you want to edit, and then click **Edit**. The **Configure VNA Plugin** dialog opens.
3. Edit the plug-in configuration details (the plug-in configuration JSON), and then save your changes.

The plug-in is edited.

**Delete a Plug-in**

You can delete a plug-in that is configured in a VNA Gateway (a connection).

**Follow these steps:**

1. From the **VNA Gateways** view, click the VNA Gateway (the connection) associated to the plug-in that you want to delete. The associated plug-ins display in the **VNA Plugins on <VNA gateway name>** view.
2. From the **VNA Plugins on <VNA gateway name>** view, select the plug-in that you want to delete, and then click **Delete**. The **Confirm Plugin Deletion** dialog opens.
3. Click **Yes**.

The plug-in is deleted.

**Start a Plug-in**

You can start stopped plug-ins ("Stopped" status) that are configured in a VNA Gateway.

**TIP**

You can check the status for a plug-in from the **Status** column in the **VNA Plugins on <VNA gateway name>** view.

**Follow these steps:**

1. From the **<VNA Gateways>** view, click the VNA Gateway (the connection) associated to the plug-in that you want to start. The associated plug-ins display in the **VNA Plugins on <VNA gateway name>** view.
2. Select the stopped plug-in that you want to start, and then click **Start**.

The stopped plug-in is started.



## Stop a Plug-in

You can stop running plug-ins ("Running" status) that are configured in a VNA Gateway.

### Follow these steps:

1. From the **<VNA Gateways>** view, click the VNA Gateway (the connection) associated to the plug-in that you want to stop.  
The associated plug-ins display in the **VNA Plugins on <VNA gateway name>** view.
2. Select the running plug-in that you want to stop, and then click **Stop**.
3. Click **Yes**.

The plug-in is stopped.

## Create an Import Rule in a VNA Gateway

(23.3.3 and higher)

By default, the IP domain for the data collector that is associated to the VNA receives inventory. If you have inventory that you want another IP domain to receive, create import rules for those groups and site groups that you want the data aggregator to place the group of VNA inventory into that IP domain.

### Follow these steps:

1. On the **VNA Administration** page, select the VNA Gateway to which you want to add an import rule.
2. In the **Import Rules on <VNA Gateway name>** view, click **New**.  
The **Create / Edit VNA Import Rule** page appears. A list of discovered groups and site groups for the selected VNA Gateway displays in the **Source Data** section of the page. These groups are for the plug-ins, and contain hierarchically-organized VNA devices.
3. In the **Source Data** section, select the groups or site groups to add to the import rule.  
By default, you can add only VNA domains (domain groups) and Viptela site groups to import rules. You can select only those groups or site groups that are not already added to an import rule.

#### TIP

If there is a group or site group that you cannot add to the import rule (it is not selectable), you can enable the setting for that group or site group.

For more information, see [the DX NetOps Network Flow Analysis documentation](#).

If an import rule is applied to the selected group based on a parent group, the **Override Import Rule?** dialog opens. Click **OK** to confirm that you want to override the applied import rule with this new import rule. The selected groups display under **Selected Groups** in the **Rule Definition** section.

4. In the **Rule Definition** section, complete the following fields:
  - **Import Rule Name**  
Defines the name of the import rule.
  - **Description**  
Defines the description for the import rule.
  - **IP Domain**  
Defines the IP domain into which you want the data aggregator to map the group of VNA inventory. The list of IP domains is from VNA.
5. Save your changes.

The group or site group is added to the import rule, and the import rule is added in the VNA Gateway.

## View a List of Import Rules in a VNA Gateway

(23.3.3 and higher)

The import rules in a VNA Gateway are listed on the **VNA Administration** page, in the **Import Rules on <VNA Gateway name>** view. To view the list, select the VNA Gateway for which you want to view the import rules.

The following information is available from the table in this view:

- **Name**  
Displays the name of the import rule.
- **Description**  
Displays the description for the import rule.
- **Status**  
Displays the status (Active or Inactive) for the import rule.
- **Editor**  
Displays the name of the user who last edited the import rule.
- **Last Edit Date & Time**  
Displays the data and time that the import rule was last edited.

**TIP**

You can also configure the view to show/display the **Gateway ID** and **Rule ID** columns.  
For more information, see [Customize Views](#).

### **Edit an Import Rule in a VNA Gateway**

(23.3.3 and higher)

**Follow these steps:**

1. From the **Import Rules on <VNA Gateway name>** view, select the import rule that you want to edit, and then click **Edit**.  
The **Create / Edit VNA Import Rule** page appears.
2. Do the following, and then save your changes:
  - To remove a group from the import rule, in the **Rule Definition** section, click the x next to the group that you want to delete.  
The selected group is removed from the import rule.
  - Edit the details of the import rule, such as the import rule name.

Your changes to the import rule are saved.

### **Delete an Import Rule from a VNA Gateway**

(23.3.3 and higher)

**Follow these steps:**

1. From the **Import Rules on <VNA Gateway name>** view, select the import rule that you want to delete, and then click **Delete**.  
The **Confirm Import Rule Deletion** dialog opens.
2. Confirm that you want to delete the import rule by clicking **Yes**.

The import rule is deleted from the VNA Gateway.

### **Create a User Domain in a VNA Gateway**

(23.3.4 and higher)

**Follow these steps:**

1. From the **VNA Gateways** view, select the VNA Gateway in which you want to create a user domain.  
The **VNA Domains Management on <VNA Gateway name>** view displays.
2. Click **New**.

The **Create Domain** dialog appears.

3. Enter a name for the user domain in the **Domain name** field, and then click **Create**.

The user domain is created in the VNA Gateway.

### **View a List of User Domains in a VNA Gateway**

(23.3.4 and higher)

The user domains in a VNA Gateway are listed in the **VNA Domains Management on <VNA Gateway name>** view.

The following information is available from the table in this view:

- **Domain Name**  
Displays the name of the user domain.
- **Domain ID**  
Displays the ID of the user domain (used in plug-in configuration).
- (23.3.6 and higher) The following columns display alarms/events grouped by severity for the user domain:
  - Critical Alarms
  - Major Alarms
  - Minor Alarms
  - Medium Alarms
  - Warning Alarms
  - Info Alarms
  - Unknown Alarms
  - Critical Events (hidden by default)
  - Major Events (hidden by default)
  - Minor Events (hidden by default)
  - Medium Events (hidden by default)
  - Warning Events (hidden by default)
  - Info Events (hidden by default)
  - Unknown Events (hidden by default)

### **Update the Name of a User Domain in a VNA Gateway**

(23.3.4 and higher)

You can update only the name of user domains in VNA Gateways.

#### **Follow these steps:**

1. From the **VNA Gateways** view, select the VNA Gateway in which you want to update the name of a user domain.  
The **VNA Domains Management on <VNA Gateway name>** view displays.
2. Select the user domain that you want to update, and then click **Update**.  
The **Update Domain** dialog appears.
3. Enter a name in the **Domain name** field, and then click **Update**.

The name of the user domain in the VNA Gateway is updated.

### **Get Inventory Statistics for a User Domain in a VNA Gateway**

(23.3.5 and higher)

**Follow these steps:**

1. From the **VNA Gateways** view, select the VNA Gateway in which you want to get inventory statistics for one of the following:

- For a specific user domain.

The **VNA Domains Management on <VNA Gateway name>** view displays. From this view, select the user domain for which you want to get inventory statistics, and then click **Stats**.

The **Inventory Stats** dialog with retrieved data for the user domain displays.

- For all user domains:

The **VNA Domains Management on <VNA Gateway name>** view displays. From this view, click **Stats**.

The **Inventory Stats** dialog with retrieved data for all user domains displays.

2. Click **Cancel** to close the dialog.

**Manage the Notification Filters for a VNA Gateway**

(23.3.6 and higher)

You can perform the following tasks from the **VNA Notification Filters** tab to manage notification filters for a VNA Gateway:

- [View a List of Notification Filters in a VNA Gateway](#)
- [Add a Notification Filter to a VNA Gateway](#)
- [Edit a Notification Filter in a VNA Gateway](#)
- [View Notification Filter Details in a VNA Gateway](#)
- [Delete a Notification Filter from a VNA Gateway](#)

For more information about notification filters, [the DX NetOps Virtual Network Assurance documentation](#).

**View a List of Notification Filters in a VNA Gateway**

(23.3.6 and higher)

View the list of notification filters by navigating to the **VNA Administration** page, **VNA Notification Filters**, **<VNA Gateway name>**.

The following information is available in this view:

- **Engine Id:** Displays the ID of the specific plugin for which the filter was created.
- **Notification Type:** Displays the type of notification (Alarm/Event) for which the filter was created.
- **Domain Id:** Displays the ID of the user domain (used in plug-in configuration).

**Add a Notification Filter to a VNA Gateway**

(23.3.6 and higher)

Add notification filters to apply them for notifications that are based on the configured parameters.

**Follow these steps:**

1. From the **VNA Administration** page, on the **<VNA Gateways>** view, click the VNA Gateway (the connection) associated with the notification filter for which you want to view the details.

The associated filters display in the **VNA Notification Filters** on the **<VNA Gateways>** view.

2. Click **New**.

The **Notification - Create Filter** dialog appears.

3. Complete the following fields:
  - **Engine ID**  
Defines the ID of the plugin engine that is associated with the current filter.
  - **Filter Type**  
Defines the type of the filter (Whitelist, or Blacklist). Based on the filter type, the notifications are filtered out or excluded.
  - **Notification Type**  
Defines the type of the notifications (Alarm, or Event) that the current filter processes.
  - **Clear Historical Alarms**  
Defines the clear historical alarms parameter for the notifications (should be true or false) to be filtered. This option is available only if you selected the **Notification Type** as Alarm.
  - **Severity**  
Defines the severity of the notifications to be filtered (Critical, Major, Medium, Minor, Warning, Info, Unknown).
  - **State**  
Defines the state of the notifications to be filtered (Created, Modified, Cleared, Unknown).
  - **Name**  
Defines the list of the names (separated by a comma) of the notifications to be filtered. The parameter is not case-sensitive.
  - **Cause**  
Defines the list of the causes (separated by a comma) of the notifications to be filtered. The parameter is not case-sensitive.
  - **Notification ID**  
Defines the list of the IDs (separated by a comma) of the notifications to be filtered. The parameter is not case-sensitive.
  - **Rule**  
Defines the list of the rules (separated by a comma) of the notifications to be filtered. The parameter is not case-sensitive. This option is available only if you selected the **Notification Type** as Alarm.
  - **Type from Source**  
Defines the list of the source types (separated by a comma) of the notifications to be filtered. The parameter is not case-sensitive.
  - **Acknowledged**  
Defines the acknowledged parameter for the notifications (true or false) to be filtered. This option is available only if you selected the **Notification Type** as Alarm.
4. Click **New**.

The notification filters are added to the VNA Gateway.

### **Edit a Notification Filter**

(23.3.6 and higher)

You can update the notification filter configuration on a VNA Gateway.

#### **Follow these steps:**

1. From the **VNA Administration** page, on the **<VNA Gateways>** view, click the VNA Gateway (the connection) associated with the notification filter for which you want to view the details.

The associated filters display in the **VNA Notification Filters** on the **<VNA Gateways>** view.

2. Select the filter that you want to edit, and then click **Edit**.

The **Notification - Create Filter** dialog appears.

3. Edit the required details of the notification filter, and then save your changes. You can skip some options, such as **Clear Historical Alarm**, if they are not required according to the filter **Notification Type**.

The changes to the notification filter on the VNA Gateway is saved.

### **View Notification Filter Details from a VNA Gateway**

(23.3.6 and higher)

You can view notification filter details from the VNA Gateway.

#### **Follow these steps:**

1. From the **VNA Administration** page, on the **&ltVNA Gateways>** view, click the VNA Gateway (the connection) associated with the notification filter for which you want to view the details.

The associated filters display in the **VNA Notification Filters** on the **&ltVNA Gateways>** view.

2. Select the filter the filter that you want to view details, and then click **Details**.

The **Notification - Create Filter** dialog appears.

The JSON configuration of the notification filter for the VNA Gateway displays.

### **Delete a Notification Filter from a VNA Gateway**

(23.3.6 and higher)

You can delete one ore more of the notification filters from a VNA Gateway.

#### **Follow these steps:**

1. From the **VNA Administration** page, in the **&ltVNA Gateways>** view, click the VNA Gateway (the connection) associated with the notification filter for which you want to view the details.

The associated filters display in the **VNA Notification Filters** in the **&ltVNA Gateways>** view.

2. Select the filter that you want to delete, and then click **Delete**.

The **Confirm Plugin Deletion** dialog opens.

3. Click **Yes**.

The notification filter is deleted from the VNA Gateway.

## **Secure Connections to Virtual Network Assurance**

You secure connections to DX NetOps Virtual Network Assurance (VNA) by enabling authentication and securing access to VNA over Hypertext Transfer Protocol Secure (HTTPS) in the VNA connection (the related and active VNA Gateway).

For more information about how to edit VNA Gateways (the connections to VNA), see [Manage Connections to Virtual Network Assurance](#).

## **Integrate with Splunk**

Review this article to learn how to integrate with a Splunk server.

**IMPORTANT**

Release level of this capability is Early Access. In Early Access, our focus is to verify functionality and gather feedback from a limited set of customers. If you are interested in trying out this capability, contact your account representative and create a ticket with Broadcom Support to request access to the capability. We plan to make Early Access capabilities generally available to all customers, and will determine the best time to do that on a per-capability basis.

**(23.3.4 and higher)**

Network operators access and analyze Syslogs from devices to troubleshoot network infrastructure issues. This can be a challenging process, as it involves correlating Syslogs across multiple nodes, switching between tools to find the right set of logs, and getting access to log analytics systems. These siloed workflows can have an adverse operational and business impact, and can result in a poorer mean-time-to-repair (MTTR) and a poorer experience. You could solve this issue with manual triage workflows and by manually searching for logs. This approach is inefficient and time-consuming.

To address the issue and to help you achieve your MTTR goals, integrate NetOps Portal with Splunk so that NetOps Portal can retrieve Syslogs in the context of alarms and devices using the out-of-the-box Syslog connector. The Syslog connector retrieves Syslogs from your Splunk instance, and then constructs the retrieved data in a format that NetOps Portal can consume.

The Syslog connector does the following:

- Reduces MTTR.
- Improves workflow efficiency.
- Eradicates the required lead time for the network operators to set up a Syslog-based triage in the context of an alarm or device.
- Offers a contextual view of Syslogs along with the support for text-based filtering, which makes the troubleshooting efficient for the network operators.

Use the Syslog connector to:

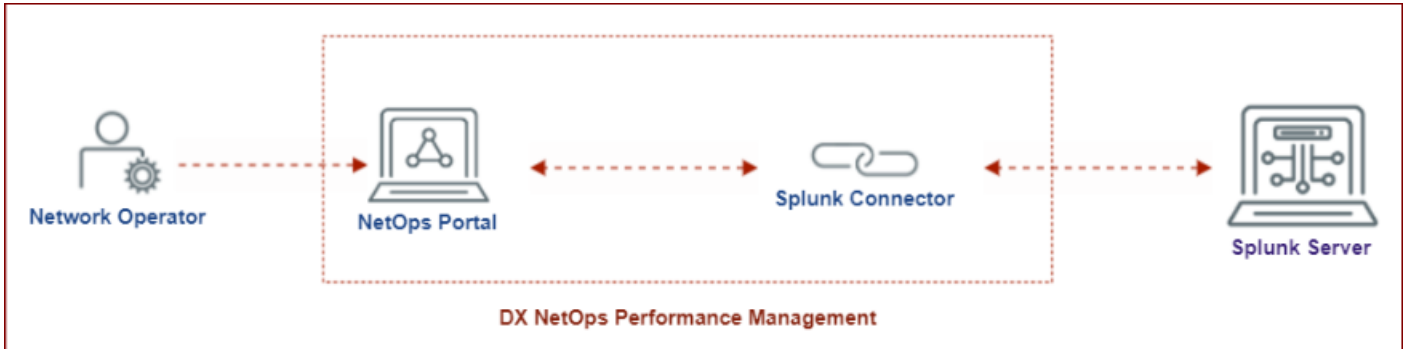
- Leverage your existing log monitoring investments.
- Improve accuracy of the log analysis.
- View Syslogs in the context of an alarm or device.

In this article:

- [High-Level Overview](#)
- [Review the Considerations](#)
- [Verify the Prerequisites](#)
- [Integrate with Splunk](#)
- [Create a Custom Query](#)
- [View Syslog Data](#)
- [View the Splunk Syslog Query](#)
- [Secure Connections to the Splunk Server](#)

**High-Level Overview**

The following illustration shows the high-level overview for the integration with Splunk:



### **Review the Considerations**

Before integrating with Splunk, consider that NetOps Portal retrieves Syslogs on-demand from the Splunk server by running an appropriate Splunk query. Furthermore, NetOps Portal does not store these Syslogs in any form.

### **Verify the Prerequisites**

Before you integrate with Splunk, ensure that you have met the following prerequisites:

- The data aggregator is configured as a data source in NetOps Portal.
- The data aggregator and Splunk servers are up and are running.

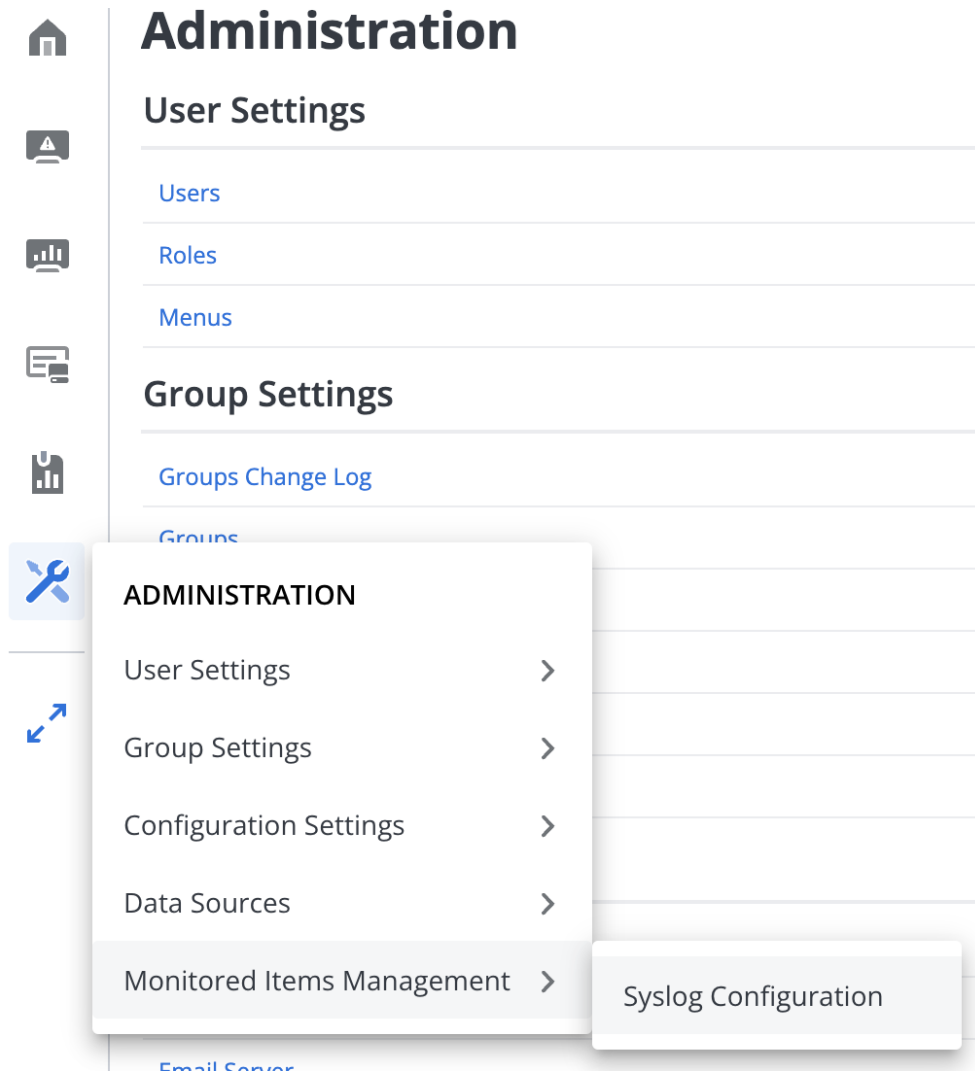
### **Integrate with Splunk**

As a network administrator, you integrate with Splunk by configuring the Splunk parameters in NetOps Portal.

#### **Follow these steps:**

1. Log in to the following NetOps Portal URL:  
<https://<hostname>:8181/pc/center/admin/debug/attrs>
2. Set the value of the `Syslog.Enable` parameter to `true`, click **Update**, and log out.
3. Log in to NetOps Portal.
4. Hover over **Administration, Monitored Items Management**, and then click **Syslog Configuration**.  
 The **Syslog Configuration** page opens.  
 The following image shows an example of how to navigate to this page:





5. In the **Splunk Connector** section, click **Enable Splunk Connector**.  
By default, the Syslog connector is enabled.
6. In the **Splunk Connector** section, complete the following fields:
  - **Protocol**  
Specifies the protocol that you want to use to secure the connection with the Splunk server.  
**Options:** http or https  
**Default:** https  
**Required:** yes
  - **Splunk Server Host Name**  
Specifies the host name or IP address of the Splunk server.  
**Example:** splunkhost.com  
**Required:** yes
  - **Management Port**  
Specifies the port number to connect to the Splunk server.  
**Default:** 8089  
**Required:** yes
  - **Splunk Access Token**  
Specifies the Splunk token required for authentication to access the Splunk server.

**Required:** yes

The following image shows an example of these fields:

### Splunk Connector



Enable Splunk Connector

Splunk Connector Status: ● Connected

Protocol \*

https ▼

Splunk Server Host Name \*

11.netops.broadcom.net

The host name or IP address of the Splunk server, e.g.,  
splunkhost.com

Management Port \*

8089

Splunk Access Token \*

[Redacted Access Token]

The access token (usually an org token) used to access your Splunk server.

For more information about these Splunk-related parameters, see [the Splunk documentation](#).

For example, for more information about authentication tokens, see [the "Create authentication tokens" section in the Splunk documentation](#).

7. In the **Splunk Connector** section, under **Field Mapping**, select *one* of the following options:

– **Identify by Name**

Specifies the event field mapping for a single index. In this section, you also define field mappings between the Syslog events and the corresponding fields:

**NOTE**

NetOps Portal uses these fields to generate a query (the Splunk Syslog query) to retrieve the required details. However, if you want to create a customized query to include additional information, select the **Identify by Query** option.

- **Splunk Index**

Specifies the Syslog index where the logs are stored in Splunk.

**Required:** yes

- **Splunk Source Types**

Specifies the source types of the Syslog in Splunk. For multiple entries, use a comma.

**Required:** no

- **Timestamp**

(Optional) Specifies the name of the field that maps to the event's timestamp.

**Default:** \_time

**NOTE**

If you specify another field, ensure that its value is in the correct format.

**Example:** 2023-10-31T14:00:18.527+05:30

**Required:** no

- **Severity**

Specifies the name of the field that maps to the event's severity.

**Required:** yes

- **Facility**

Specifies the name of the field that maps to the event's facility.

**Required:** yes

- **Message**

(Optional) Specifies the field that maps to the event's message.

**Default:** \_raw

**Required:** no

The following image shows an example of these fields:

#### Field Mapping

☒ Identify by Name ☐ Identify by Query

Define a mapping for the event fields in the syslogs.

#### Splunk Index \*

The index with the event data

#### Splunk Source Types

A comma separated list of one or more source types

#### Timestamp

The name of the field that maps to the event's timestamp.

#### Severity \*

The name of the field that maps to the event's severity.

#### Facility \*

The name of the field that maps to the event's facility.

#### Message

The name of the field that maps to the event's message.

#### – Identify by Query

Specifies a complex Splunk query. Select this option if you want NetOps Portal to retrieve information that it cannot retrieve using a value for **Identify by Name**. For example, retrieve data from multiple indexes or source types. Select this option for more flexibility to create and run complex queries.

For more information, see [the "Create a Custom Query" section](#).

The following example image shows an example of this option:

#### Field Mapping

☐ Identify by Name ☒ Identify by Query

#### Splunk Query

```
(index=netops OR index="sys_r5") | rename sc4s_syslog_facility as syslog_facility,
sc4s_syslog_severity as syslog_severity, _raw as syslog_message.
```

See the documentation for help with describing queries

**Default:** Identify by Query

8. Click **Save** to validate and save the configuration.

As part of the validation, NetOps Portal verifies the Splunk access token validity and host validity to ensure that the connection with the Splunk server is working. If the validation is successful, the entered configuration is saved and the **Splunk Connector Status** is "Connected". Otherwise, perform the necessary corrections, and then retry.

**NOTE**

The token might expire at a later stage, causing NetOps Portal to fail to retrieve data using the query. In such instances, NetOps Portal displays the error message that it receives from the Splunk server.

The Splunk parameters are configured. NetOps Portal is integrated with Splunk.

### **Create a Custom Query**

You can retrieve information from the Splunk server based on your unique requirements using a custom query. For example, this is helpful in scenarios where you want NetOps Portal to retrieve data from multiple indexes or use different source types.

**Prerequisite:**

- You have completed Step 1 through Step 4 in [the "Integrate with Splunk" section](#).
- You have validated the query by running it on the Splunk server. If it runs successfully, you can then use it in NetOps Portal.

**NOTE**

NetOps Portal does not validate the query.

**Follow these steps:**

1. In the **Splunk Connector** section, under **Field Mapping**, select the **Identify by Query** option. The **Splunk Query** field displays.
2. Add a custom query to this field. The mandatory fields are `index`, `syslog_facility`, and `syslog_severity`. Additionally, to enable NetOps Portal to parse the facility-, message-, and severity-related fields, ensure that when you map/rename them, that you use the syntax for the mapped names as `syslog_message`, `syslog_facility`, and `syslog_severity`.

**Example queries:**

```
index=netops OR index="sys_r5"| rename sc4s_syslog_facility as syslog_facility, sc4s_syslog_severity as
syslog_severity, _raw as syslog_message

index=netops OR index="\sys_r5\"| rename sc4s_syslog_facility as syslog_facility, sc4s_syslog_severity as
syslog_severity, _raw as syslog_message | search ((syslog_severity=*err*) OR (syslog_facility=*err*) OR
(syslog_message=*err*)) | fields syslog_severity,syslog_facility,syslog_message | sort -_time

(index="netops") OR ( index="cisco_logs" ) OR ( index="sys_r7" ) | eval syslog_severity =
coalesce( netops_severity, cisco_severity) | rename _raw as syslog_message
```

3. Click **Save**.

The custom query is created.

### **View Syslog Data**

You can view Syslog data retrieved from the Splunk server in the following contexts:

- [View data in the context of an alarm](#)
- [View data in the context of a device](#)

#### **View Syslog Data in the Context of an Alarm**

You can view Syslog data retrieved from the Splunk server for a specific alarm from the **Alarms** view. NetOps Portal retrieves the data from the Splunk server based on the creation time (plus or minus 5 minutes) of the alarm.

**Prerequisite:** The Syslog connector is enabled and the **Splunk Connector Status** is "Connected".

**Follow these steps:**

1. Hover over **Alarms**, and then click **Alarm Console**.  
The **Alarm Console** dashboard appears, and the **Alarms** view is displayed.
2. Select the alarm for which you want to view Syslog data from the view.  
The alarm context details display in the **Alarms** view.  
The following example image shows an example of a selected alarm in the **Alarms** view:

The screenshot shows the 'Alarm Console' dashboard. On the left is a sidebar with navigation icons. The main area is titled 'Alarms' and shows a table of alarm events. The first alarm is selected, and its details are shown on the right.

Severity	Date/Time	Item Name	Model Type
Major	Oct 19, 2023 9:50:40 A...	R5-...ca.com	Cisco7206VXR
Major	Oct 19, 2023 9:50:26 A...	R7-...3.ca.com	Cisco7206VXR

3. From the alarm context details, click the **Log Events** tab.  
The **Log Events** view displays. The data NetOps Portal retrieves from the Splunk server for the selected alarm is displayed.  
The following example image shows an example of this view:

The screenshot shows the 'Log Events' view for a selected alarm. The left sidebar shows the 'Log Events' tab selected. The main area displays a table of log messages.

Alarm Details: R5-...3.ca.com  
A BGP4 PEER 50-48-48-49.dr05.nrw.cny.frontiernet.net SESSION IS DOWN

Timestamp	Severity	Facility	Message
Oct 19, 2023 9:54:56 AM GMT	notice	local7	%BGP-5-ADJCHANGE: neig
Oct 19, 2023 9:54:56 AM GMT	notice	local7	%BGP-5-NBR_RESET: Neig
Oct 19, 2023 9:54:52 AM GMT	notice	local7	%BGP-5-ADJCHANGE: neig
Oct 19, 2023 9:54:52 AM GMT	notice	local7	%BGP_SESSION-5-ADJCHA

## **View Syslog Data in the Context of a Device**

### **Prerequisites:**

- The Syslog connector is enabled and the **Splunk Connector Status** is "Connected".
- Device Syslogs for the selected time range are present in the Splunk server.
- The Syslog host field of the Splunk server matches the DX NetOps device's IP address or hostname or FQDN.

### **NOTE**

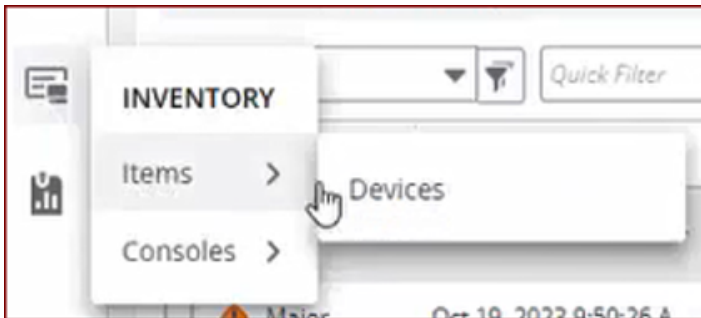
NetOps Portal retrieves and displays the data from the Splunk server based on the selected time range.

### **Follow these steps:**

1. Hover over **Inventory**, **Items**, and then click **Devices**.

The **Devices** page appears.

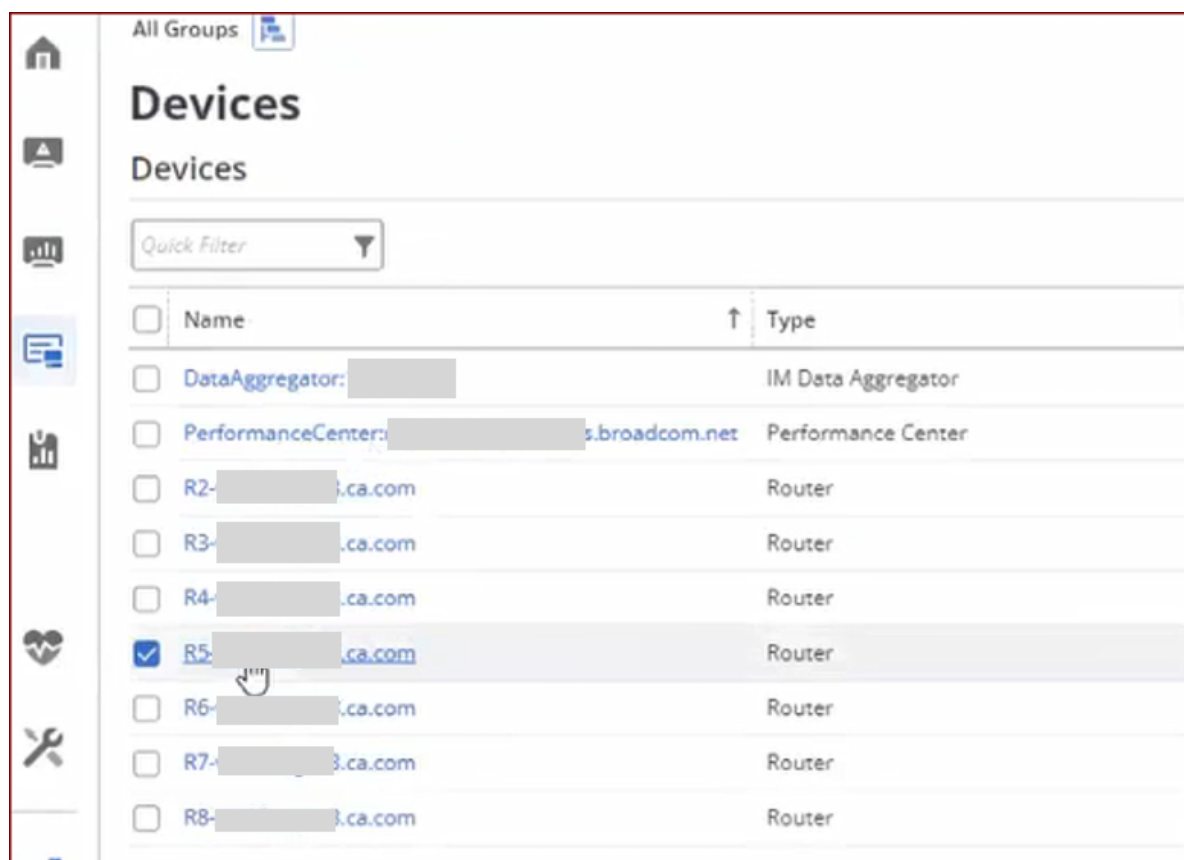
The following image shows an example of how to navigate to this page:



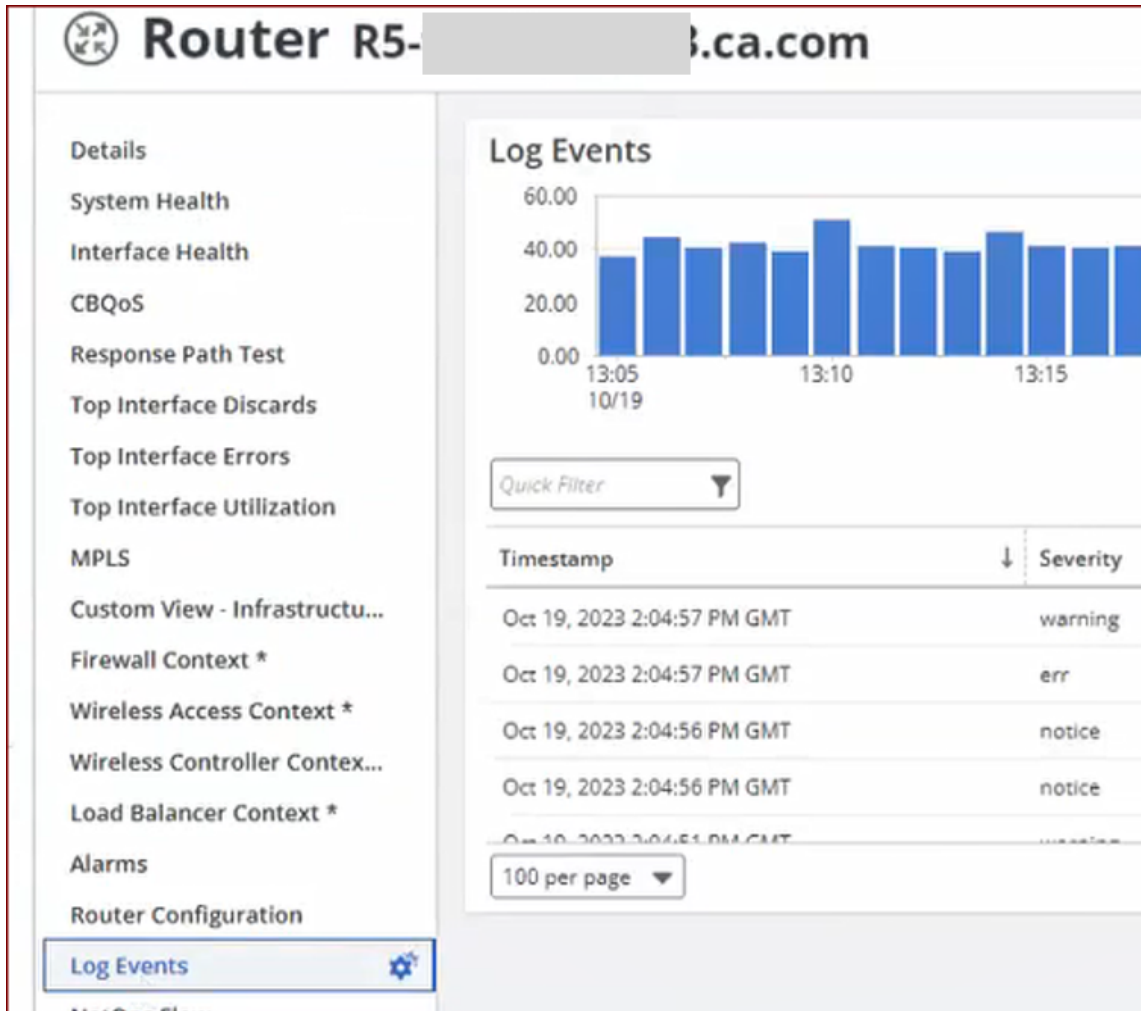
2. Click the device for which you want to view Syslog data from the view.

The device context page appears.

The following example image shows example links:



- Click the **Log Events** tab of the device context page.  
The **Log Events** view displays.  
The data NetOps Portal retrieves from the Splunk server for the selected device is displayed.  
The following example image shows an example of this view:



The bar chart displays the time range, and the table shows the data. When you click a bar on the bar chart, the related row is highlighted in the table.

You can also choose a time range from the top-right section of the page to view the historical data.

#### NOTE

The maximum limit to show the number of records in the table is 5000.

### View the Splunk Syslog Query

You can view the Splunk Syslog query that NetOps Portal dynamically generates to retrieve alarm/device-related data from the Splunk server. This query uses the fields that were configured while configuring the Splunk server. NetOps Portal generates the query, sends the query to the Splunk server, retrieves the data, and displays it in the appropriate context view.

You can use this dynamically-generated query as a starting point to delve deeper into analyzing the logs in Splunk. You can also copy and paste the query for troubleshooting or for sharing it with other stakeholders.

You can view the query in the following contexts:


- [View the query in the context of an alarm](#)
- [View the query in the context of a device](#)



## View the Splunk Syslog Query in the Context of an Alarm

**Prerequisite:** You have completed the steps in [the "View Syslog Data in the Context of an Alarm" section](#).

**Follow these steps:**

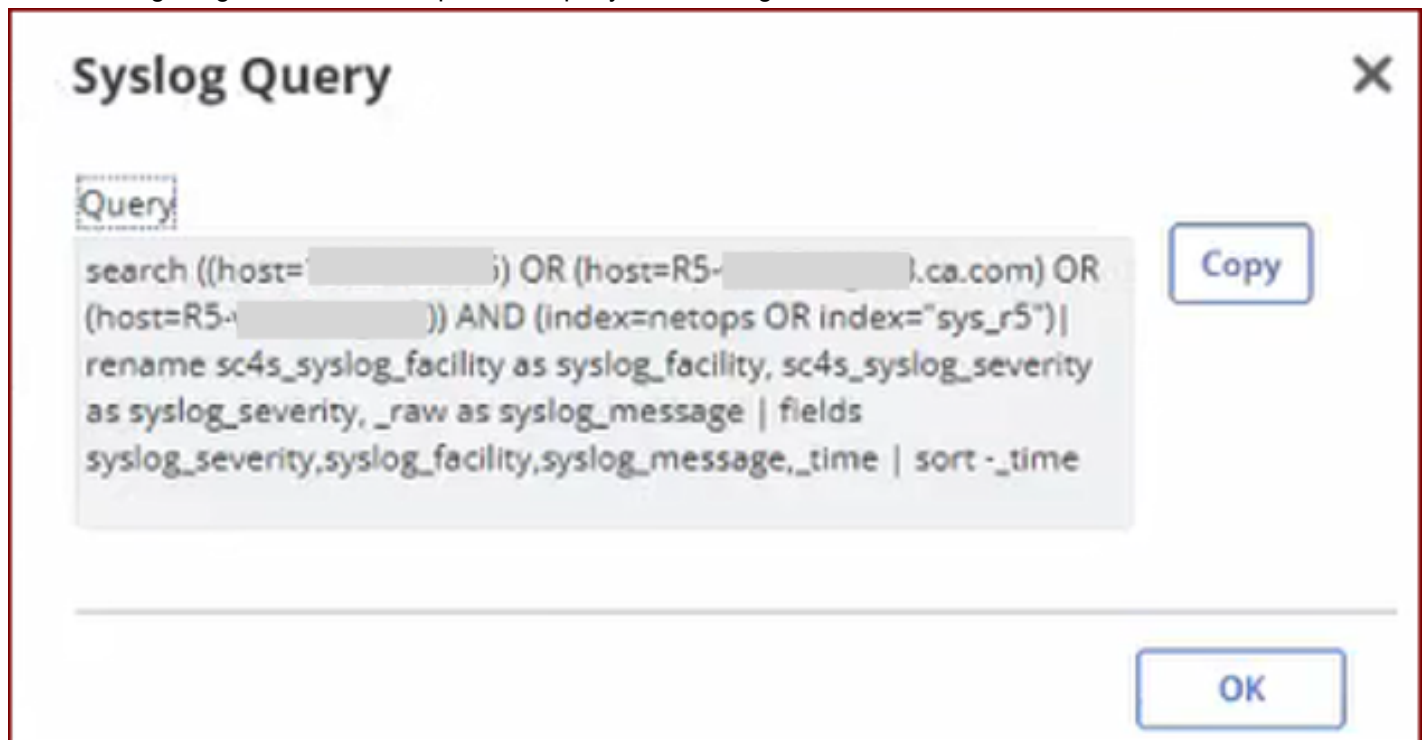
1. From the alarm context details, click the gear icon (  ), and then select the **Syslog Query** option.

**NOTE**

From this gear icon, you can also export the retrieved Syslog data in a CSV format using the **Export to CSV (scaled)** and **Export to CSV (unscaled)** options.

The **Syslog Query** dialog appears. The query that NetOps Portal uses to retrieve the alarm-related Syslog data from the Splunk server is displayed.

The following image shows an example of the query in the dialog:



2. To copy the query, click **Copy**.
3. Click **OK** to close the dialog.

## View the Splunk Syslog Query in the Context of the Devices

**Prerequisite:** You have completed the steps in [the "View Syslog Data in the Context of a Device" section](#).

**Follow these steps:**

1. From the alarm context details, click the gear icon, and then select the **Syslog Query** option.  
The **Syslog Query** dialog appears. The query that NetOps Portal uses to retrieve the device-related Syslog data from the Splunk server is displayed.
2. To copy the query, click **Copy**.
3. Click **OK** to close the dialog.

## **Secure Connections to the Splunk Server**

Secure the communication between NetOps Portal and the Splunk server (management port 8089 ) by adding certificates. To do so, you must import the Splunk server certificates into NetOps Portal.

### **Follow these steps:**

1. Access the following Splunk URL:

```
https://<hostname>:8089/
```

2. Click the lock icon on the browser page (before the address bar) and export the certificate.
3. Access the NetOps Portal server, and import the certificate into JAVA\_HOME on that server by issuing the following commands:

```
keytool -import -alias <splunkCert>  
-keystore $JAVA_HOME/lib/security/cacerts  
-file <exported.crt>
```

4. Log in to the following NetOps Portal URL:

```
https://<hostname>:8181/pc/center/admin/debug/attrs
```

5. Set the value of the `Syslog.SSL.validate` parameter to `true` , click **Update**, and log out.
6. Log in to NetOps Portal and verify that it can access the Splunk data in a secure mode.

You have secured the connection with the Splunk server (management port 8089 ).

---

# Troubleshooting

---

Use the troubleshooting section to diagnose and resolve issues with DX NetOps Performance Management.

Use the troubleshooting section to diagnose and resolve issues with DX NetOps Performance Management. Each page in this section contains one or more symptoms, and at least one solution to resolve the issue.

In this article:

- [Insufficient Permissions During Installation](#)
- [Enable NetOps Portal Logging](#)
- [Enable Data Collector Logging](#)
- [Download Data Collector Logs](#)
- [Disable Data Collector Logging](#)
- [SNMP Querying Command-Line Tools](#)
- [Enable the ActiveMQ Admin Console for the Data Aggregator or Data Collector](#)

Other troubleshooting topics are in this section.

## **Insufficient Permissions During Installation**

### **Symptom:**

After you specify an installation path, the following error appears:

```
Error: Insufficient permissions for installation path: Folder_PathPath must have executable permission for 'other'.Exiting the installer...
```

### **Solution:**

Ensure that the child and parent directories have the necessary executable permission by issuing the following command:

```
chmod 755 "Folder_Path"
```

## **Enable NetOps Portal Logging**

If you encounter an issue with a specific view in NetOps Portal, you can enable logging on that view. You can use this method to provide the necessary details to Broadcom Support.

### **Follow these steps:**

1. Log in as a user who has the Generate URLs from Views role right.
2. Open the dashboard that contains the problematic view.
3. Click the gear (**Edit**) icon on the view, and then select **Generate URL**.  
The Generate URL dialog opens.
4. Enable **Detailed View Logging**.
5. Click **Preview**.
6. Take a screenshot of the preview for Broadcom Support.
7. Run CARE on the NetOps Portal host and the data aggregator host.
8. Provide the archive files from CARE and the screenshot to Broadcom Support.

## **Enable Data Collector Logging**

If you encounter missing data, detailed poll logging is available from the data collector debug tool. When enabled for an IP address, the tool generates logs on the data collector. Before you download the logs, wait at least two poll cycles (ten minutes) after enabling logging to collect sufficient data.

The logs record the following details of every poll for the specified IP address:

- The requests that are sent to the device.
- The responses that are received from the device.
- The delta processing the data collector performs on the responses.
- The expression evaluation the data collector performs on the responses.

You cannot enable detailed poll logging for more than one IP address at a time. If a second IP address is enabled, detailed poll logging for the previous IP address is disabled. The logs for the previous IP address are discarded.

Detailed poll logging stores a maximum of 60 MB of logged messages by default. If the limit is reached, the oldest entries are discarded to make room for the newer messages. You can configure this limit with the cache setting.

Logging can be filtered by Poll Group ID, or by Component Item ID. Filtering restricts the logging to a specific Vendor Certification or a specific component.

### IMPORTANT

Detailed poll logging is stored in the runtime memory of the data collector. If the data collector is shut down, the log is lost. If the data collector restarts, the log resets.

#### Follow these steps:

1. Go to the following location:  
`<DA_host>:<port>/dcdebug/enablededebg.html`
2. Specify the IP address of the polled device.
3. Select an IP domain.
4. (Optional) Specify the **Filtering** and **Cache** settings.
5. Select **Enable Debug Logging**, and then click **Enable Debug Logging**.

### Download Data Collector Logs

If you encounter missing data, you can use the data collector debug tool to download the relevant logs. Before you download the logs, wait at least two poll cycles (ten minutes) after enabling logging to collect sufficient data. You can use this method to provide the necessary details to Broadcom Support.

#### Follow these steps:

1. Open the following URL:  
`<DA_host>:<port>/dcdebug/searchdebug.html`
2. Specify the IP address of the polled device.
3. Select an IP domain.
4. Select **Download all logs for IP**, and then click **Download As Zip**.

**Next step:** Disable the detailed poll logging.

### Disable Data Collector Logging

Disable the detailed poll logging after downloading the data collector logs.

#### Follow these steps:

1. Open the following URL:  
`<DA_host>:<port>/dcdebug/enablededebg.html`
2. Specify the IP address of the polled device.
3. Select an IP domain.
4. Select **Disable Debug Logging**, and then click **Disable Debug Logging**.

## SNMP Querying Command-Line Tools

DX NetOps Performance Management includes the `sapwalk2` and `sappoll` SNMP querying command-line tools. They are located in the `<DC_installation_directory>/scripts/` directory. The `sapwalk2` utility gathers an SNMP snapshot of network devices. Broadcom Support uses these snapshots to reproduce issues and verify SNMP values on devices. The `sappoll` utility retrieves SNMP data for a timeframe.

The following `sapwalk2` command can be helpful:

```
sapwalk2 -i ip_address -v snmp_version -s starting_oid -c community_name -
xv bridge_table_oid -o output_file.walk
```

### Example:

```
sapwalk2 -i 10.253.190.15 -v v2c -s 1.3.6.1 -c public -xv 1.3.6.1.2.1.17 -o
10.253.190.15.walk
```

The following `sappoll` command can be helpful:

```
sappoll -i ip_address -v snmp_version -c community_name -p snmp_port -r retries -d
timeout -f file_containing_oids -n number_of_polls -t poll_interval
```

### Example:

```
sappoll -i 138.42.110.45 -v v2c -c xxxxxx -p 161 -r 3 -d 3000 -f /tmp/OidList -n 5 -t 2
```

The file containing OIDs (for example, `/tmp/OidList`), should include the OIDs to poll in the following format:

```
I instance_oid
N node_oid
```

### Example:

```
I 1.3.6.1.2.1.1.1
```

## Enable the ActiveMQ Admin Console for the Data Aggregator or Data Collector

Generally, the ActiveMQ admin console should not be available on the network. However, if you must have access to the console to address a problem, you can enable it for the data aggregator or the data collector.

### Follow these steps:

1. Go to the following file based on the component:

#### – Data aggregator

```
<installation_directory>/IMDataAggregator/broker/<apache-activemq-*>/conf/activemq.xml
```

#### Example:

```
/opt/IMDataAggregator/broker/apache-activemq-5.18.3/conf/activemq.xml
```

#### • **installation\_directory**

The installation directory for the data aggregator.

**Default:** `/opt`

#### • **apache-activemq-\***

The installation directory of Apache ActiveMQ.

**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`

#### – Data collector

```
<DC_installation_directory>/IMDataCollector/broker/<apache-activemq-*>/conf/activemq.xml
```

#### Example:

```
/opt/IMDataCollector/broker/apache-activemq-5.18.3/conf/activemq.xml
```

- ***DC\_installation\_directory***  
The installation directory for the data collector.  
**Default:** /opt
  - ***apache-activemq-\****  
The installation directory of Apache ActiveMQ.  
**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2
2. Uncomment the `<import resource="jetty.xml"/>` line.
  3. (Optional) To update user access, edit the `jetty-realm.properties` property.
  4. (Optional) To encrypt the user passwords, issue the following command based on the component:
    - **Data aggregator**

```
<installation_directory>/broker/apache-activemq-*/lib/web run java -cp jetty-all-9.2.22.v20170606.jar org.eclipse.jetty.util.security.Password <password>
```

**Example:**

```
<installation_directory>/IMDataAggregator/broker/apache-activemq-5.18.3/lib/web run java -cp jetty-all-9.2.22.v20170606.jar org.eclipse.jetty.util.security.Password <password><password>
```
    - ***installation\_directory***  
The installation directory for the data aggregator.  
**Default:** /opt
    - ***apache-activemq-\****  
The installation directory of Apache ActiveMQ.  
**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2
    - **Data collector**

```
<DC_installation_directory>/IMDataCollector/broker/<apache-activemq-*>/lib/web run java -cp jetty-all-9.2.22.v20170606.jar org.eclipse.jetty.util.security.Password <password><password>
```

**Example:**

```
<DC_installation_directory>/IMDataCollector/broker/apache-activemq-5.18.3/lib/web run java -cp jetty-all-9.2.22.v20170606.jar org.eclipse.jetty.util.security.Password <password><password>
```
    - ***DC\_installation\_directory***  
The installation directory for the data collector.  
**Default:** /opt
    - ***apache-activemq-\****  
The installation directory of Apache ActiveMQ.  
**Example:** (23.3.4 and higher) apache-activemq-5.18.3 (23.3.1 - 23.3.3) apache-activemq-5.18.2
  5. Shut down the ActiveMQ broker on the data aggregator and on each data collector manually by issuing the following command based on your installation:
 

```
systemctl stop activemq
```
  6. Start the ActiveMQ broker on the data aggregator and on each data collector manually by issuing the following command based on your installation:
 

```
systemctl start activemq
```

## Access Denied to MySQL Utilities

### Symptom:

Access is denied when you try to access a MySQL utility from the command line.

### Example:

If you try to run `mysqldump` without a password, the following error returns:

```
mysqldump: Got error: 1045: Access denied for user 'root'@'localhost' (using password: NO) when trying to connect
```

**Solution:**

As a step toward enhanced security, the MySQL utilities require a MySQL password. In previous releases, if you were running the utilities as the root user, the password had a default. We strongly recommend you upgrade or at least change the default password with the `mysqladmin` utility.

**Example:**

To change the default 'root' password from the default ('*default*' in the following example) to your custom password ('*newpassword*' in the following example), use the following command:

```
mysqladmin -uroot -pdefault password newpassword
```

## Automatic Rediscovery Does Not Run After Updating Vendor Group Priority

**Symptom:**

DX NetOps Performance Management does not run rediscovery after I updated the vendor certification priorities and updated the order of vendor certifications.

**Solution:**

Updating the order of vendor certifications and the `<PriorityGroup>` tag in the same REST operation can disrupt automatic rediscovery. Manually trigger rediscovery for the affected metric family.

**Follow these steps:**

1. Navigate to the data aggregator data source in NetOps Portal.  
The **Monitored Devices** page appears.
2. Under **Monitoring Configuration**, click **Metric Families**.  
The **Metric Families** page appears.
3. Select the affected metric family, and then click **Update Metric Family**.  
The **Rediscover Associated Devices** window opens.
4. Click **Yes** to proceed.

Discovery runs for all devices in the metric family.

## Browser Shows Error when Logging In

**Symptom:**

When I enter my password on the Login page, I am redirected to an error page in the web browser.

**Solution:**

This symptom does not indicate that you entered incorrect SAML credentials. Instead, the browser error (such as 401 or 500) indicates that Single Sign-On redirected the browser to the login URL, but the Identity Provider (IdP) server is down.

Follow these steps to correct the issue:

- Verify that the IdP server is running.
- Test the network connection between the CA NetOps Portal server and the IdP server.

## Cannot Create a Vendor Certification

### Symptom:

I cannot create a vendor certification and I receive an error message.

### Solution:

Open the karaf log files in the Data Aggregator installation directory, and follow these steps:

1. Look for the MIB name string or the name of the metric family that you selected.
2. Review the stack trace of the exception to find the CertManagerException and the reason for the error. The reason for the error follows the exception.

**Example:** The expression parser did not expect the token after ++, as shown:

```
Caused by: com.ca.im.dm.certmgr.interfaces.CertManagerException: Tech Cert: {http://
im.ca.com/
normalizer}NormalizedCPUInfo, Unable to compile expression: [Error: expected end of statement

[Near : {... stemID ++ extremeSystemBoardID ....}]
```

3. Fix the error based on the reason that is provided. Verify that the following requirements are met:
  - The expression group does not contain a mix of scalar and table entries.
  - Expressions contain valid syntax.
  - At least one expression is defined for a metric family variable.
  - At least two metric family variables are defined. Names and Indexes are required (except for scalar-only metrics).
  - The vendor certification variable used in the expression is from the chosen MIB table (valid in the user interface).

## Cannot Remove a Custom Vendor Certification

### Symptom:

I want to stop using a custom vendor certification. I cannot find a method to remove or deactivate the certification.

### Solution:

Removing a custom vendor certification is not supported. Move the vendor certification to the bottom of the vendor certification priority list.

## Cannot Find the Data Aggregator RIB Document

### Symptom:

I cannot locate the Data Aggregator RIB document.

### Solution:

Follow these steps:

1. Verify that Data Aggregator has been successfully added to CA NetOps Portal. For more information, see [Manage Data Sources](#).
2. Verify that the RIB web service for Data Aggregator is running.
3. Verify that the RIB web service for Data Aggregator is publishing the RIB document:
  - a. Use the Data Source's RIB web service to request the list of available RIB documents.

**Example:**



```
http://da_host:8581/rib/doclist
```

4. If you know the document ID, check the document:

```
http://da_host:8581/rib/doc/docId
```

**Example:**

```
http://da_host:8581/rib/doc/CA.IM.DA.NormalizedPortInfo
```

## Cannot Remove a Metric Family

**Symptom:**

I want to stop using a custom metric family. I cannot find a method to remove the metric family from my instance of DX NetOps Performance Management.

**Solution:**

Removing a custom metric family is not supported. To stop using the custom metric family, remove it from all monitoring profiles. For more information, see [Manage Monitoring Profiles](#).

## Cannot View More than 5000 Device Components in Inventory List

**Symptom:**

I have more than 5000 devices registered in NetOps Portal. However, the list of device components in the inventory or search results shows only 5000 items.

**Resolution:**

By default, NetOps Portal limits the number of devices in the inventory to 5000, but you can change this default value.

**Follow these steps:**

1. Go to the following URL:

```
http://<PC_host>:8181/pc/center/admin/debug
```

– **PC\_host**

The hostname of the NetOps Portal installation.

2. Log in as an administrative user.
3. Select **Global Attributes**.
4. Specify a value for the `UniversalList.Limit` attribute, and then click **Update**.

**WARNING**

Increasing the default device limit of 5000 can cause performance issues in NetOps Portal.

The maximum number of registered devices that are visible in the inventory is equal to the value of `UniversalList.Limit`. A restart of the services is not required for the changes to take effect.

## Clean Up After a Failed NetOps Portal Installation

**Symptom:**

The NetOps Portal installation can fail to complete. For example, if the `tmp/` directory lacks sufficient space for the installation files, the installation fails.

**Solution:**

If you experience an installation failure, clean up the directories before you reinstall NetOps Portal.

**Follow these steps:**

1. Clean up the `tmp/` directory by removing unnecessary files.
2. Remove the `/opt/CA` directory by issuing the following command:  

```
rm -rf /opt/CA
```
3. Remove the installer registration file by issuing the following command:  

```
rm /var/.com.zerog.registry.xml
```
4. Remove the NetOps Portal service files by issuing the following command:  

```
rm service caperfcenter_*
```
5. [Restart the server.](#)
6. [Install NetOps Portal again.](#)

## Data Aggregator Disk Space is Decreasing

**Symptom:** The available disk space on the data aggregator host is decreasing unexpectedly.

**Solution:** Errors can cause orphaned rollup messages to build up in Apache ActiveMQ. These message files are large. Verify the presence of the messages and purge them if necessary.

**Follow these steps:**

1. Navigate to the following directory on the data aggregator host:  

```
<DC_installation_directory>/broker/<apache-activemq-*>/data/kahhadb
```

  - **`DC_installation_directory`**  
The installation directory of the data collector.  
**Default:** `/opt/IMDataCollector`
  - **`apache-activemq-*`**  
The installation directory for Apache ActiveMQ.  
**Example:** (23.3.4 and higher) `apache-activemq-5.18.3` (23.3.1 - 23.3.3) `apache-activemq-5.18.2`
2. Search for "not removing data file" by issuing the following command:  

```
grep "not removing data file" db-*.log | more
```
3. Compare the return to the following example:  

```
2015-10-13 09:30:31,411 [eckpoint Worker] TRACE MessageDatabase
- not removing data file:
```

This return indicates that the ActiveMQ broker has orphaned messages that are blocking the queue. If the return in your system is similar, continue to the next step.
4. Navigate to the following URL to access the ActiveMQ broker:  

```
http://DA_host:8161/admin/queues.jsp
```

ActiveMQ might require login with the Admin account.  
**Default Username:** `admin`  
**Default Password:** `admin`
5. Locate ActiveMQ.DLQ. If the pending message count is greater than 0, click **purge**.  
ActiveMQ clears the orphaned messages.

## Data Aggregator Fails to Synchronize

**Symptom:**

When I try to synchronize Data Aggregator with NetOps Portal, I see a 'Synchronization failure' message in the Status column.

### Solution:

Data Aggregator could not handle the data that is sent to it during synchronization. Review the Device Manager application log file, called DMSERVICE.log. This file appears in the CA/PerformanceCenter/DM/logs directory. The log entry shows a general SOAP exception if Data Aggregator is unable to handle data that was received from CA NetOps Portal during synchronization.

Look for an exception and stack trace within the following phases of synchronization:

- Pull
- Global Sync
- Bind (only executes when initially synchronizing with a data source)
- Push

Contact CA Technical Support with this information.

## Data Aggregator or Data Collector Does Not Initialize

### Symptom:

On certain systems, the data aggregator or the data collector processes start successfully, but the component is never initialized. The status shows that the process is running. This problem typically occurs on systems with underpowered CPUs.

This problem occurs only in the following situations:

- When you upgrade the data aggregator or the data collectors.
- When you restart them after clearing out the Apache `karaf` data directory.

### Solution

To verify that this issue occurred, complete the following verification steps:

1. Look at the `karaf.log` file on the relevant host:
  - **(Data aggregator)** The `<installation_directory>/apache-karaf/data/log/karaf.log` file.
    - **`installation_directory`**  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
  - **(Data collector)** The `<DC_installation_directory>/apache-karaf/data/log/karaf.log` file.
    - **`DC_installation_directory`**  
The installation directory for the data collector.  
**Default:** `/opt/IMDataCollector`

This issue produces an error message in the log. The following message is an example of this error:

```
2016-01-12 02:32:46,303 | WARN | Event Dispatcher | AetherBasedResolver
| mvn.internal.AetherBasedResolver 583 | 3 - org.ops4j.pax.logging.pax-
logging-api - 1.8.3 | | Error resolving artifact com.ca.im:data-
mgmt.provision.xml:features:2.7.0-RELEASE-137:Could not find artifact com.ca.im:data-
mgmt.provision.xml:features:2.7.0-RELEASE-137 in central (http://repol.maven.org/
maven2/)
shaded.org.eclipse.aether.resolution.ArtifactResolutionException: Could not find
artifact com.ca.im:data-mgmt.provision.xml:features:2.7.0-RELEASE-137 in central
(http://repol.maven.org/maven2/)
```

```

at
  shaded.org.eclipse.aether.internal.impl.DefaultArtifactResolver.resolve(DefaultArtifactReso
at
  shaded.org.eclipse.aether.internal.impl.DefaultArtifactResolver.resolveArtifacts(DefaultArt
...
...
2016-01-12 02:32:46,327 | WARN    | Event Dispatcher | FeaturesServiceImpl
  | res.internal.FeaturesServiceImpl 1367 | 7 - org.apache.karaf.features.core - 2.4.3
  | | Unable to add features repository mvn:com.ca.im/data-mgmt.provision/2.7.0-
RELEASE-137/xml/features at startup
java.io.IOException: Error resolving artifact com.ca.im:data-
mgmt.provision:xml:features:2.7.0-RELEASE-137: Could not find artifact
  com.ca.im:data-mgmt.provision:xml:features:2.7.0-RELEASE-137 in central (http://
repo1.maven.org/maven2/)
at
  org.ops4j.pax.url.mvn.internal.AetherBasedResolver.resolve(AetherBasedResolver.java:584)
at
  org.ops4j.pax.url.mvn.internal.AetherBasedResolver.resolve(AetherBasedResolver.java:528)
at
  org.ops4j.pax.url.mvn.internal.AetherBasedResolver.resolve(AetherBasedResolver.java:506)
at
  org.ops4j.pax.url.mvn.internal.AetherBasedResolver.resolve(AetherBasedResolver.java:481)
at org.ops4j.pax.url.mvn.internal.Connection.getInputStream(Connection.java:123)

```

2. Confirm that the `data-aggregator` or `data-collection-manager` features are not resolved in Apache Karaf:

a. Log in to the Karaf console on the relevant host, and then issue the following commands based on the host:

- **(Data aggregator)**

```
ssh -p8501 karaf@localhost
```

- **(Data collector)**

```
ssh -p8601 karaf@localhost
```

b. List the features by issuing the following command:

```
features:list | grep data
```

When this issue occurs, this list does not include the following features:

- `data-aggregator`
- `data-collection-manager`

Apply the fix for the relevant host:

- [Data aggregator](#)
- [Data collector](#)

## **Data Aggregator**

### **Follow these steps:**

1. Identify the features repository URL to be added from the `<installation_directory>/apache-karaf/etc/org.apache.karaf.features.cfg` file by issuing the following command:

```
[root@sapv1 ~]# grep --color provision <installation_directory>/apache-karaf/etc/
org.apache.karaf.features.cfg
featuresRepositories = mvn:org.apache.karaf.assemblies.features/standard/<version>/
xml/features,mvn:org.apache.karaf.assemblies.features/spring/<version>/xml/
features,mvn:org.apache.karaf.assemblies.features/enterprise/<version>/xml/
features,mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/features
#featuresRepositories=mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/
features
```

**Example:**

```
[root@sapv1 ~]# grep --color provision /opt/IMDataAggregator/apache-karaf/etc/
org.apache.karaf.features.cfg
featuresRepositories = mvn:org.apache.karaf.assemblies.features/standard/4.3.3/
xml/features,mvn:org.apache.karaf.assemblies.features/spring/4.3.3/xml/
features,mvn:org.apache.karaf.assemblies.features/enterprise/4.3.3/xml/
features,mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/features
#featuresRepositories=mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/
features
```

**– *installation\_directory***

The installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

2. Select the content from this file according to the error message in the `karaf.log` file. The number after `RELEASE-` can be different.
3. Connect through Karaf by issuing the following command:
 

```
ssh -p8501 karaf@localhost
```

The password is **karaf**.
4. Manually add the features by issuing the following commands:
 

```
features:addurl mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/features
features:install data-aggregator
```
5. Verify that the `karaf.log` file does not include a repeat of the error message.  
During normal operation, the log shows the various bundles in the application initializing.
6. Verify that the data aggregator has initialized. For example, verify that the data aggregator REST page loads from the following URL:
 

```
da_host:8581/rest/
```
7. [Restart each data collector.](#)
8. If the error message also appears in the `karaf.log` file on the data collector, [apply the data collector fix.](#)

**Data Collector****Follow these steps:**

1. Identify the features repository URL to be added from the `<DC_installation_directory>/apache-karaf/etc/org.apache.karaf.features.cfg` file by issuing the following command:
 

```
[root@sapv1 ~]# grep --color provision <DC_installation_directory>/apache-karaf/etc/
org.apache.karaf.features.cfg
```

```
featuresRepositories = mvn:org.apache.karaf.assemblies.features/standard/<version>/
xml/features,mvn:org.apache.karaf.assemblies.features/spring/<version>/xml/
features,mvn:org.apache.karaf.assemblies.features/enterprise/<version>/xml/
features,mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/features
#featuresRepositories=mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/
features
```

**Example:**

```
[root@sapv1 ~]# grep --color provision /opt/IMDataCollector/apache-karaf/etc/
org.apache.karaf.features.cfg
featuresRepositories = mvn:org.apache.karaf.assemblies.features/standard/4.3.3/
xml/features,mvn:org.apache.karaf.assemblies.features/spring/4.3.3/xml/
features,mvn:org.apache.karaf.assemblies.features/enterprise/4.3.3/xml/
features,mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/features
#featuresRepositories=mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/
features
```

**– DC\_installation\_directory**

The installation directory for the data collector.

**Default:** /opt/IMDataCollector

2. Select the content from this file according to the error message in the `karaf.log` file. The number after `RELEASE-` can be different.
3. Connect through Karaf by issuing the following command:
 

```
ssh -p8601 karaf@localhost
```

The password is **karaf**.
4. Manually add the features:
 

```
features:addurl mvn:com.ca.im/data-mgmt.provision/2.7.0-RELEASE-133/xml/features
features:install data-collection-manager
```
5. Verify that the `karaf.log` file does not include a repeat of the error message. During normal operation, the log shows the various bundles in the application initializing.
6. Verify that the data collector has initialized. For example, in NetOps Portal, hover over **Administration, Monitored Items Management**, and then click **Data Collectors**. Verify that the **Status** is "Collecting Data".

## Data Collector Dropped Polling Event Message

**Symptom:**

A "Data Collector Dropped Polling" event appeared in my events list.

**Solution:**

Clock drift can sometimes cause polling to stop on some components and devices. When this occurs, restart the Data Collector.

For more information, see [Restart the Data Collector](#).

## Data Collector Installs But Does Not Appear in the Data Collector List Menu

**Symptom:**

The data collector is installed successfully, but it does not appear in the **Data Collector List** menu.

**Solution:**

Do the following steps:

1. Review the `<DC_installation_directory>/apache-karaf/shutdown.log` file to ensure that the data collector was not shut down automatically.
  - **`DC_installation_directory`**  
The default installation directory for the data collector.  
**Default:** `/opt/IMDataCollector`

The data collector is shut down automatically when you incorrectly specify the data aggregator host, tenant, or IP domain when you installed the data collector. The `shutdown.log` file provides error information as to why the data collector was shut down. The data collector shuts down for one of the following reasons:

  - The data aggregator host information, tenant, or IP domain that was specified during the data collector installation were incorrect:
    - If you specified the data aggregator host information incorrectly, uninstall and reinstall the data collector.
    - If you specified the tenant incorrectly, uninstall and reinstall the data collector.
    - If you specified the IP domain incorrectly, uninstall and reinstall the data collector.
  - Contact with the data aggregator could not be established.
2. Ensure that an established connection to the data aggregator exists by issuing the following command:
 

```
netstat -a | grep 61616
```
3. If a connection to the data aggregator does not exist, do the following steps:
  - a. View the `<DC_installation_directory>/broker/apache-karaf/conf/activemq.xml` file on the Data Collector host. This file contains the hostname or IP address of the data aggregator host that you specified when you installed Data Collector.
    - **`DC_installation_directory`**  
The default installation directory for the data collector.  
**Default:** `/opt/IMDataCollector`
  - b. Search for the “networkConnector” section of the `activemq.xml` file. This section should contain a line as follows:
 

```
<networkConnector name="manager"
  uri="static:(tcp://test:61616) "
  duplex="true"
  suppressDuplicateTopicSubscriptions="false"/>
```

Ensure that the data aggregator hostname that is specified in the `networkConnector` section is correct and resolves through DNS or `/etc/hosts` entries. The data collector cannot communicate with the data aggregator if you entered the data aggregator hostname incorrectly during the Data Collector installation.
  - c. Type the following command help ensure that the connection opens successfully when you open a telnet connection to the data aggregator host on port 61616:
 

```
telnet <dahostname> 61616
```

This command confirms that the data aggregator is listening in on that port.
  - d. If the telnet connection does not open successfully, the reasons could be as follows:
    - The data aggregator is not running. Ensure that it is running. Open a console, and issue the following command:
 

```
systemctl status dadaemon
```
    - If Data Aggregator is not running, start the data aggregator. Log on to the data aggregator host computer as the root user or a sudo user with access to a limited set of commands. Do one of the following steps:
      - Start the Data Aggregator service:

```
systemctl start dadaemon
```

- (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:

```
<installation_directory>/scripts/dadaemon activate
```

- **installation\_directory**

The default installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

The request to initiate the connection is not making it from the data collector to the data aggregator successfully. Ensure that the port that is specified in the `networkConnector` section of the `broker.xml` file is open for incoming connections on the data aggregator. Be sure that there are no firewall rules preventing this connection.

## Data Collector Shows Polling Status Not Connected

### Symptom:

The data collector loses its connection to the data aggregator for some time (its **Polling Status** is "Not Connected"), starts to cache polling data, and does not send this data to the data aggregator. The data collector and data aggregator are not in sync with each other for time, as seen with the `date +%s` command. This command shows Unix time (num of secs since 1970), which ignores time zone differences.

### Solution:

1. Have the systems administrator adjust the time on the data aggregator or data collector, whichever is not correct.
2. On the server with the adjusted time, restart the ActiveMQ service and the corresponding process (either the data aggregator or the data collector).  
For more information, see [Restart Performance Management Component Services](#).
3. Ensure that `chrony` is configured to run often on all DX NetOps Performance Management servers, and that it is running.  
For more information, see [the Red Hat documentation](#).

## Data Is Missing from Views

### Symptom:

Some of the table columns in the interface views are empty. For example, interface and device names, interface speeds, and utilization data are missing from views.

### Solution:

Some data sources do not support authentication passwords or privacy passwords that are below the minimum length.

SNMP profiles that use the SNMPv3 format let you enable authentication and privacy options. When you create a valid SNMPv3 profile, specify an authentication password that is eight characters or more in length. These profiles may not be successful in communicating with devices. In this case, SNMP data are missing for the affected interfaces.

Similarly, blank passwords are not supported for SNMP v3 profiles with MD5 or SHA as the Authentication Protocol.

## Data Source Registration Fails

### Symptom:

I attempt to add a new data source, but the registration fails. The following message appears: 'Create Data Source Failed: Data source communication failure.'

### Solution 1:



This message indicates that the data source is unreachable. Do the following:

- Verify that the data source is running.
- Verify that the DNS hostname or IP address of the server where the data source database is installed is correct. You can edit the data source to view this information.
- Check intervening firewalls. Make sure they are configured to let CA NetOps Portal communications reach the data sources. For more information about the ports to open, see the *Installation Guide*.

### Solution 2:

If the failure occurred with a CA Infrastructure Management Data Aggregator data source, verify that it is running. Access the following URL:

```
http://<host>:<port_number>/rest
```

where 'host' is the IP address of the server where the Data Aggregator is installed, and 'port\_number' is the port used to access the RESTful web service, usually 8181.

The web service status indicates whether the Data Aggregator is running.

### Solution 3:

Check the Device Manager application.log file. The file is written to the following directory:

```
CA\PerformanceCenter\PC\logs
```

The log entry references the URI used by CA NetOps Portal to communicate with the data source, along with a stack trace.

## Data Source Synchronization Fails

### Symptom:

When attempting to synchronize a data source, the `Synchronization failure` message appears.

### Solution:

This message can indicate one of the following:

- The data source is unreachable. Do the following:
  - Verify that the data source is running.
  - Verify that the Domain Name System (DNS) hostname or IP address of the server where the database for the data source is installed is correct.  
For more information, see [Configure a Data Source](#).
- The data source could not handle the data NetOps Portal sent it during synchronization. Check the log files by completing the following steps:
  - a. Check the data source log for the data source. As a user with the Administrator role, on the **Manage Data Sources** page, select the data source for which you want to view the log, and then click **Log**.  
The **Data Source Log** page opens displaying the data source log for the data source.  
For more information, see [Synchronize Data Sources](#).
  - b. If you still cannot determine the source of the problem, check the `CA\PerformanceCenter\PC\logs\application.log` Device Manager log file. The data source could not handle the data NetOps Portal sent to it during synchronization if the log entry shows a general SOAP exception.
- NetOps Portal encountered an issue during the attempted synchronization.  
[Check the log files](#). Look for an exception and stack trace within the following phases of synchronization:

- Pull
- Global Sync
- Bind (only executes when initially synchronizing with a data source)
- Push

The log contains detailed information about the steps that NetOps Portal performed during each phase. This information can help pinpoint the cause for the synchronization failure.

- The system times are not synchronized. Check the Network Time Protocol (NTP) server or the system time on each DX NetOps Performance Management server (including data source servers and the NetOps Portal server).

## Data Source Test Fails

### Symptom:

I test a data source during the registration process, but the test fails.

### Solution 1:

Do the following:

- Verify the DNS hostname or IP address of the server where the database for the data source is installed.
- Try running the data source registration. The data source registration might succeed even if the test failed.
- Check the logs for registration failure information. For more information, see [Data Source Registration Fails](#).

### Solution 2:

If the failure occurred with a Data Aggregator data source, verify that it is running. Access the following URL:

```
http://<host>:<portnumber>/rest
```

#### NOTE

This URL does not open the correct page in Mozilla Firefox. Use another supported browser.

The web service status indicates whether the Data Aggregator is running.

### Solution 3:

If the failure occurred with a data source other than a Data Aggregator, check the application log file (PC/logs/application.log) for a corresponding event. The log entry includes the URL that CA NetOps Portal uses to communicate with the data source, as well as a stack trace.

## Discovery Does Not Start

### Symptom:

Select discovery profiles, and then click **Run** to run a discovery, but discovery fails to start, or you cannot click **Run**.

### Solution:

Possible reasons for a discovery failure or for a disabled Run button include the following:

- The IP domain previously specified in the discovery profile has been deleted. Assign the discovery profile to an IP domain.
- A data collector is not installed for the IP domain that is specified in the selected discovery profile.  
For more information about how to install data collector hosts, see [Installing](#).
- One or more Data Collector hosts are installed for the IP domain that is specified in the selected discovery profile. However, all of the Data Collector hosts that are installed for the IP domain are stopped. Start the Data Collector hosts.
- The tenant is deactivated. Activate the tenant.

## Gaps Appear in Reports or Views

### Symptom:

A view or report shows a gap in the value for a metric.

### Solution:

By default, DX NetOps Performance Management shows gaps in the data for counter values for counter wraps, bad counter values, and missing data. This behavior protects the integrity of report data.

For more information, see [Trend Views](#) and [Configure Counter Behavior](#).

## Gaps in Data Appear during Throttling

### Symptom:

Certain devices are missing data in DX NetOps Performance Management. You notice a gap in polling or missed polls even though the device is running.

### Solution:

The throttling of outgoing SNMP requests can cause gaps in data. Throttling occurs when some polls in a poll group have responded with REQUEST\_TIMED\_OUT in a single poll cycle. The device is assumed to be too busy to respond to polls for that poll group. Polling is attempted again in the next poll cycle.

To address this issue, verify that polling is being throttled. To prevent gaps in data, reduce the poll rate or configure a monitoring profile poll filter. For more information, see [Poll Sensitive and Critical Devices Without a Performance Impact](#) and [Manage Monitoring Profiles](#).

To verify that polling is being throttled, use the following method that is most convenient:

- Select the **Details** tab for a device. Look for any events related to throttling. "Polling Stopped" appears in the event list when a throttling event occurs on the device that is missing data.  
For more information about these events, see [Polling Stopped Event Message](#).
- Hover over **System Health**, and then click **Data Aggregator Polling**:
  - a. Review the Device Polling Statistic chart, which shows a systemwide total of stopped polls. Investigate any spikes above zero.
  - b. Review the Stopped Polls by Device chart, which shows the devices with the most stopped polls. Investigate the top devices struggling to keep up with the polling load.
- Query the `NormalizedDevicePollingStatistics` metric family. Determine the number of polls that were stopped during the poll cycle. Also, determine the number of poll groups that were stopped.
- Review the Poll Summary log for `badPollRequestTimedOutCount` and `notSentPollRequestCount`.
- If you have Detailed Poll Logging from dcdebug, review the dcdebug Detailed Poll Log for the following string. If the value is above zero, then polls have been throttled.

```
DCMResponseVariable [name={http://im.ca.com/
normalizer}NormalizedDevicePollingStatistics.NumPollsStoppedDueToPriorTimeouts,
value=
```

## Group Membership Is Not Updated During Synchronization

### Symptom:

After synchronization, your group memberships are not updated completely.

### Solution:

Your system might have many rule groups to evaluate each synchronization cycle. Each rule group is all the rules for a single group. DX NetOps Performance Management processes as many rule groups as possible within the configured time. By default, the processing time is 60 seconds. If processing takes longer than the configured time, DX NetOps Performance Management finishes processing the current rule group. During the next synchronization cycle, it processes the oldest unprocessed rule groups first.

To evaluate more rules groups during each cycle, increase the processing time.

#### Follow these steps:

1. Log in to the NetOps Portal host.
2. Set a new processing time by issuing the following command:
 

```
mysql netqosportal -unetqos -p password -e "replace into general values('Rules.MaxTimeSeconds', 'time');"
```

#### – Time

Defines the rule-processing time in seconds.

#### TIP

Set the processing time to 120 seconds and evaluate the system performance.

**Default:** 60

DX NetOps Performance Management allocates the specified time to rules processing during synchronization.

## Insecure Connection Message in Firefox

#### Symptom:

When logging in to NetOps Portal, the following warning message appears when you click inside the **Username** or **Password** field:

This connection is not secure. Logins entered here could be compromised.

#### Solution:

Some versions of Firefox show this warning for any login screen that is not using HTTPS. NetOps Portal uses HTTP by default because it is hosted as an internal application behind corporate firewalls. To configure more security, [configure NetOps Portal to Use HTTPS](#).

For more information about this message, see the Firefox help.

## Inventory is Empty After a Data Source is Registered

#### Symptom:

I installed a data source and registered it, but I do not see any managed items in the Inventory.

#### Solution 1:

Check to make sure that the data source is registered and has an active status. Do the following:

1. Log in as a user with administrative privileges.
2. Select **Administration**, **Data Sources**, and click **Data Sources**.  
The Manage Data Sources page opens.

#### Solution 2:

One of the following may have occurred:

- Data source registration failed. For more information, see [Data Source Registration Fails](#).
- Data source synchronization failed. For more information, see [Data Source Synchronization Fails](#).

**Solution 3:**

Check the permissions for the user account that you used to log in. If the user account has no assigned permission groups, you see no managed items.

Make sure that you have not logged in as a user associated with the Default Tenant, who sees no managed items.

## Low Data Aggregator Disk Space

**Symptom:**

The data aggregator has run out of disk space.

**Solution:**

Move the data aggregator to another location on the same host with more disk space.

**Follow these steps:**

- Do one of the following steps:
  - Stop the ActiveMQ and Data Aggregator services by issuing the following commands:
 

```
systemctl stop activemq
systemctl stop dadaemon
```
  - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:
 

```
<installation_directory>/scripts/dadaemon activate
```

    - installation\_directory**  
The installation directory for the data aggregator.  
**Default:** /opt/IMDataAggregator
- Uninstall the ActiveMQ service and the Data Aggregator service from the following data aggregator locations by issuing the following commands:
 

```
<installation_directory>/scripts/activemq uninstall
<installation_directory>/scripts/dadaemon uninstall
```

  - installation\_directory**  
The installation directory for the data aggregator.  
**Default:** /opt/IMDataAggregator
- Move the contents from the current installation directory to the new installation directory by issuing the following command:
 

```
cp -rfp <installation_directory>
<new_installation_directory>
```

  - installation\_directory**  
The installation directory for the data aggregator.  
**Default:** /opt/IMDataAggregator
  - new\_installation\_directory**  
The new installation directory for the data aggregator.
- Delete the current installation directory by issuing the following command:
 

```
rm -rf <installation_directory>
```

  - installation\_directory**  
The installation directory of the data aggregator.  
**Default:** /opt/IMDataAggregator

5. Clean up the data directory by issuing the following command:

```
rm -rf <new_installation_directory>/apache-karaf/data
```

- **new\_installation\_directory**

The new installation directory of the data aggregator.

6. Update the following files to point to the new installation directory by issuing the following command:

```
/etc/DA.cfg file
/var/.com.zerog.registry.xml<new_installation_directory>/scripts/
activemq <new_installation_directory>/scripts/dadaemon <new_installation_directory>/
apache_karaf/bin/setenv<new_installation_directory>/apache_karaf/bin/restart
```

- **new\_installation\_directory**

The new installation directory of the data aggregator.

7. Install the ActiveMQ service and Data Aggregator service from the following locations:

```
<new_installation_directory>/scripts/activemq
install
<new_installation_directory>/scripts/dadaemon install
```

- **new\_installation\_directory**

The new installation directory of the data aggregator.

8. Do one of the following steps based on your installation:

- Start the ActiveMQ and Data Aggregator services:

```
systemctl start activemq
systemctl start dadaemon
```

- (Fault tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:

```
<new_installation_directory>/scripts/dadaemon activate
```

- **new\_installation\_directory**

The new installation directory of the data aggregator.

## Metric Family is Incomplete

### Symptom:

I successfully imported a custom metric family, but later found a defective metric definition. For example, the Name property has a maximum length of 32 characters. If this limit is exceeded, it can cause synchronization problems.

### Solution:

Delete the custom metric family.

#### WARNING

Use this procedure only in emergency situations. Otherwise, do not use this procedure.

### Follow these steps:

1. Locate the following directory:

```
<installation_directory>/apache-karaf/certifications/custom/deploy
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

2. Delete the XML files that were created and deployed for the metric family. They are named as follows:

- im.ca.com-normalizer-<technology>.xml
- im.ca.com-inventory-<technology>.xml

If applicable, also delete the im.ca.com-certifications-snmp-<vendor>.xml file that was created for the vendor certification.

3. Do one of the following steps:

- Start the Data Aggregator service by issuing the following command:

```
systemctl start dadaemon
```

- (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:

```
<installation_directory>/scripts/dadaemon activate
```

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** /opt/IMDataAggregator

The data aggregator restarts.

4. Verify that the previously imported metric family or vendor certification does not appear in NetOps Portal. All previously discovered components for this custom certification are also deleted.
5. Hover over **Administration**, **Data Sources**, and then click **Data Sources**.  
The **Managed Data Sources** page appears.
6. Select the data aggregator data source, and then click **Resync**.  
The **Resync Data Source** dialog opens.
7. Click **Resync**.  
The components for remaining metric families synchronize between data aggregator and NetOps Portal.
8. Edit and correct the custom metric family XML file.
9. Import the corrected metric family XML file.

## Metric Family is Not Supported

### Symptom:

I created a monitoring profile to poll metric families on a collection of devices. However, in the Polled Metric Families table, one of those metric families has a status of 'Unsupported.'

### Solution:

To correct the problem, follow these steps:

1. Verify that the polled device responds to SNMP queries.
2. Navigate to the unsupported metric family.
3. Verify that a vendor certification supports the metric family. If no vendor certification is defined, create a custom vendor certification.
4. Verify that all key vendor certification attributes are supported on the device. If all key vendor certificate attributes are supported:
  - a. Navigate back to the device
  - b. Select the metric family for which you added a custom vendor certification
  - c. Click Update Metric Family.  
Your device configuration is updated.

## Metric Values Do Not Appear in Table in OpenAPI

### Symptom:

I ran a query to generate a table with metric values. However, when I open the table, no metric values appear in the results.

### Solution:

The OpenAPI does not currently support the aggregation of multiple data samples into a single value for a particular time range.

To preview the results of your query, select the appropriate output format, such as:

- HTML table with extra metrics
- JSON
- XML

## MIB Fails to Compile

### Symptom:

When I review the list of MIBs in the Select MIB page of the Create Vendor Certification wizard, I receive an error message that a MIB did not compile.

### Solution:

If a MIB failed to compile, follow these steps:

1. Check the error message in the Select MIB page.
2. Perform one of these actions, depending on the error type:
  - Syntax error -- Using details in the error message, correct the syntax error in your MIB and import the corrected MIB.
  - Dependency error -- Upload the required MIB to resolve the dependency issue.

When a new MIB is imported, any existing MIB that failed to compile is recompiled in addition to the new or modified MIB.

## Multiple SNMP Devices Trigger Intrusions Alarms

### Symptom:

I have many SNMP devices behind a more restricted firewall configuration (such as DMZ networks). For security reasons, the SNMP devices have different community strings. I defined an SNMP profile for each different community string, but now I am getting intrusion alarms and have been logged out of CA NetOps Portal.

### Solution:

To find the correct SNMP profile for a device, CA NetOps Portal tries all of the SNMP profiles. This behavior can trigger intrusion alarms and can log you out of CA NetOps Portal.

To resolve this issue, follow this process:

1. Create a separate discovery profile for a critical SNMP device.
2. Assign the SNMP profile with the correct community string to the discovery profile.
3. Repeat steps one and two for each critical SNMP device.

When discovery is run, only the assigned SNMP profile is used.



## No Charts or Images are Visible in IE with HTTPS

### Symptom:

Some charts or images are not appearing in NetOps Portal when using Internet Explorer (IE) with HTTPS. A red X appears to indicate that the chart or image is broken.

### Solution:

By default, in IE, the Transport layer security (TLS) encryption setting that is required for HTTPS is set to TLS 1.0. NetOps Portal charts and images require TLS 1.1.

For more information about how to change the TLS setting in IE, see [the Microsoft Internet Explorer documentation](#).

## 'No Data to Display' Message in Views

### Symptom:

Some of the views on a dashboard are empty. A message states, "No Data to Display".

### Solution:

A graph or table view on a dashboard can show "No Data to Display" for the following reasons:

- You have not installed or registered the data source that is related to data the view displays.  
Views receive data from a single data source. Some view containers appear on dashboards even if the corresponding data source is not registered. They are always empty until the data source is registered.  
You can prevent such views from displaying in dashboards by hiding, or suppressing, these views.  
For more information, see [Customize Your User Settings](#).
- The data source is registered, but it has been temporarily disabled.  
NetOps Portal polls only enabled data sources for data. If the data source is disabled, Administrators can enable it by editing it.  
For more information, see [Configure a Data Source](#).
- The view type has not yet been customized.  
Some type of views do not have default settings, and display data only after you have customized it. For example, the MultiView and MultiTrend views require customization before they display data.  
For more information, see [Views](#).
- No data is available for the selected time range. To test this theory, select a different time range.
- Not enough time has transpired since polling started on the devices that are selected for reporting.  
If the polling interval is fairly long, the first data point can take a little longer to appear. Polling rates are set in the data sources.  
A service is not running.  
If the Device Manager service is not running, the "no data" message can appear.
- The current group does not contain items of the required type for this view.  
The group of items whose data is reported on the dashboard is shown above the Time Period selector. Check the view: Is this Server report trying to show data from a group of routers?
- The group is new or has recently been changed.  
Check group membership. A group rule might be misconfigured. If your user account has the required role right, edit the view to select another group context. Or click the **Group Filter** link above the Time Period selector, and then select another group context for the dashboard.
- The user account of the logged-in user does not have permission to view monitored items that have reported data.  
For more information, see [Manage User Accounts](#).
- The data source has not properly synchronized with NetOps Portal.  
For more information, see [Data Source Synchronization Fails](#).
- Components or managed items were not discovered.

This problem is data source-specific. For the data aggregator data source, you can check inventory discovery history. On the **Discovery Profiles** page, select the discovery profile that you created for the initial discovery, and then click **History**.

For more information about how to view discovery profile history results, see [Run Device Discovery](#).

- You have nbot configured or enabled metric families.  
The data aggregator data source automatically applies out-of-the-box monitoring profiles to the predefined collections, such as the All Routers collection. However, custom groups and custom collections are subject to misconfigured custom monitoring profiles.
- The database query timed out.  
Network connectivity issues between the NetOps Portal server and the data source can cause this problem.

## No Output is Generated After Running the Device Pack Generator

### Symptom:

I ran the Device Pack Generator but no output files were generated.

### Solution:

To verify that validation of the devicePackConfig file is successful, execute the following command:

```
createIMDevicePack <path_of_devicePackConfig.xml> -v
```

If the validation is unsuccessful, review the error message in the console. To fix the problem, make the recommended changes in the devicePackConfig file and rerun the script using the -v option.

You can also verify that the directory information provided in <Path> is correct and that it contains the CSV files. Also verify that the header information provided matches the CSV files.

If the problem persists, contact CA Technologies Support.

## No Performance Data for a Device Pack

### Symptom:

I generated and deployed my device pack in NetOps Portal. I do not see any performance data for my device pack.

### Solution:

Data becomes available after two performance polls.

## OpenAPI Query Results in Empty Table

### Symptom:

After you run a query, the resulting Query URL shows an empty table.

### Solution:

When selecting to show metric columns in a table format, select at least one configuration column, such as item name. Another solution is to select a format other than 'HTML table with export'.

## NetOps Portal Cannot Contact the Data Aggregator

### Symptom:

The data aggregator is installed successfully, but its status on the **Manage Data Sources** page is 'Unable to Contact'.

#### Solution:

Do the following steps:

1. Log on to the data aggregator host computer.
2. From a console, verify that the data aggregator is running by issuing the following command:
3. If the data aggregator is running, a network issue is most likely preventing NetOps Portal from contacting the data aggregator. Resolve all network problems.
4. If the data aggregator is not running, start the data aggregator. Log on to the data aggregator host computer as the root or sudo user with access to a limited set of commands. Do one of the following steps:

```
systemctl status dadaemon
```

- Start the data aggregator by issuing the following command:

```
systemctl start dadaemon
```

- (Fault-tolerant environments) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:

```
<installation_directory>/scripts/dadaemon activate
```

- **installation\_directory**

The installation directory for the data aggregator.

**Default:** /opt/IMDataAggregator

## Polling Does Not Complete for My Sensitive Device

#### Symptom:

Polling on my critical device cannot complete in a single polling cycle. The large amount of network traffic sometimes completely stops my device. How can I reliably poll this device to help ensure quality performance?

#### Solution:

Polling is vital for monitoring a device. However, excessive polling can cause a large amount of network traffic, which can degrade your ability to monitor a device successfully. If too much network traffic is overwhelming your sensitive device, you can try the following adjustments to reduce overall traffic to the device:

- Adjust your monitoring profile to remove unnecessary metric families from polling.
- Apply a filter in your monitoring profile to reduce the number of polled interfaces.
- Adjust your monitoring profile to poll less often (for example, change the SNMP Poll Rate to 15 minutes, instead of the default 5 minutes).
- Adjust the SNMP traffic threshold to lower the number of SNMP requests that are sent to the device at a time.
- Adjust the SNMP timeouts threshold to control how many polling timeouts cause polling to suspend for the current polling cycle.

## Polling Has Stopped on Discovered Metric Family

#### Symptom:

I select a device from the Monitored Devices page and see that polling has stopped for a metric family that the device supports. I did not intend for polling to stop for that metric family.

#### Solution:

Follow this process to determine why polling has stopped and perform the appropriate steps to address the cause:

1. Verify that a monitoring profile is defined and is set to poll the desired metric family.  
If this requirement is not already met, create or edit a monitoring profile with the desired metric family defined in it.

2. Verify that the device is associated with the device collection.  
If the device is not associated with the device collection, add the device to the device collection.

**NOTE**

For information about adding a device to a device collection, see the *CA NetOps Portal Administrator Guide*.

3. Verify that the monitoring profile is associated with the device collection and the device.  
If the monitoring profile is not associated, create the relationship between the monitoring profile and the device collection.

After you complete one of these actions to restart polling, select the device on the Monitored Devices page to verify:

- The status of the metric family on the Polled Metric Families tab has changed.
- The status in the Interface Components table has changed to Active.

Polling resumes automatically on existing devices.

New devices can be discovered using one of the following methods:

- Select the polled metric family on the Monitored Devices page, and click Update Metric Family.
- Set the Change Detection rate in the monitoring profile for that metric family, with Automatic Discovery set to True.

## Polling Safety Valve Event Message

**Symptom:**

A "Polling Safety Valve" event appeared in my events list.

**Solution:**

Clock drift can sometimes cause polling to stop on some components and devices. When this occurs, restart the Data Collector.

For more information, see [Restart the Data Collector](#).

## Polling Stopped Event Message

**Symptom:**

A "polling stopped" event appeared in my events list.

**Solution:**

By default, the data aggregator controls SNMP polling, helping to ensure that too many poll requests do not overwhelm a device. One method for controlling poll traffic is the SNMP timeouts threshold. The default threshold value is 15. When 15 or more SNMP requests timeout for a polling group, polling is suspended for the remainder of the current polling cycle for that polling group. Other polling groups on the device continue to poll. An event is generated, informing you of the situation.

**NOTE**

Polling resumes at the beginning of each poll cycle. When no timeouts occur in a complete 5-minute poll cycle, a "clear" event is generated.

For more information, see [Poll Sensitive and Critical Devices Without a Performance Impact](#) and [Gaps in Data Appear during Throttling](#).

## PrimaryIPAddress ATTRIBUTE\_VALUE\_NOT\_ALLOWED Error in Karaf Log

**Symptom**

The following error message appears in the karaf.log file:

```
ERROR | 0c98d87-thread-1 | date time | AbstractReconciler |
 iliation.impl.AbstractReconciler 495 | .im.aggregator.discovery | |
  Got unexpected error for detect change only write for attr {http://im.ca.com/
  inventory}Device.PrimaryIPAddress on item Item_ID#: ATTRIBUTE_VALUE_NOT_ALLOWED
```

### Solution

This error is benign and occurs when the system detects a duplicate IP address. One possible cause is that the vCenter has a decommissioned VM and the IP address has been reused. Both VMs are created in the DX NetOps Performance Management environment.

## QueryBuilder Certificate Warning

### Symptom:

In an environment with HTTPS enabled for the data aggregator, a certificate warning appears when you try to run QueryBuilder.

### Solution:

QueryBuilder requires Single Sign-On authentication. The QueryBuilder URL redirects to NetOps Portal for the Single Sign-On website. The Single Sign-On URL uses the NetOps Portal host IP address instead of its host name. This scenario leads to a certificate warning because when you generate Secure Sockets Layer (SSL) certificate using the openssl command line utility, the certificate is generated for the host name, not the host IP address. Add the NetOps Portal host IP address as a Subject Alternative Name (SAN) in the v3 certificate extension.

For more information, see [Configure the Data Aggregator to Use HTTPS and ActiveMQ JMX to Use SSL](#).

## Report on All Pages Times Out

### Symptom:

When you run a report on All Pages, it times out. When you click the link for a timed out report, the PDF contains a "Query failed" error message.

### Solution:

If you experience memory issues, ensure all your servers, especially the data repository, meet the minimum requirements and sizing guidelines. For information about the sizing requirements, see the [DX NetOps Sizing Tool](#). If a report on All Pages times out and the minimum requirements and sizing guidelines are met, contact Broadcom Support.

## Services Do Not Start After Installing the Data Collector and VNA

### Symptom:

After installing the data collector and DX NetOps Virtual Network Assurance, the Data Collector service (dcmd ) service or the WildFly service in DX NetOps Virtual Network Assurance fails to start or run properly.

### Solution:

Other processes are consuming more memory. Update the memory that they use:

- To update the memory that the Data Collector service (dcmd ) service uses, modify the maximum memory for the data collector host.

For more information, see [Modify Maximum Memory Usage for Data Aggregator and Data Collector Components](#).

- To update the memory that WildFly in DX NetOps Virtual Network Assurance uses, modify the maximum memory for DX NetOps Virtual Network Assurance.

For more information about how to configure memory usage, see [DX NetOps documentation](#).

## Spectrum Alarms Are Missing from the Alarms View

Use the following troubleshooting topic to address DX NetOps Spectrum (Spectrum) alarm issues.

Address the following Spectrum alarm issues:

- [The Event Manager service cannot create an alarm subscription](#)
- [Alarms are missing when reporting on a group context in NetOps Portal](#)
- [The Alarms in the Alarms View are not Showing the Most Recent Items Displayed in OneClick](#)
- [The Alarms in the Alarms View do not Match the Alarms Displayed in OneClick](#)
- [Spectrum Alarms can take up to a minute to show up in Alarms Views](#)

### The Event Manager service cannot create an alarm subscription

#### Symptom:

The **Alarms** views in NetOps Portal do not display any alarms.

The alarm service process within the Event Manager cannot create an alarm subscription, and is unable to retrieve alarms from the Spectrum data source.

Review the following possible causes and solutions:

- **Possible cause:**  
The Spectrum or Event Manager data source have data source synchronization issues.

#### Solution:

Review the data source synchronization failures.

For more information, see [Synchronize Data Sources](#).

- **Possible cause:**  
The alarm service process in the Event Manager cannot create an alarm subscription.

#### Solution:

- a. Check the Event Manager log, `EMService.log`, for the noted errors.

For more information about the Event Manager log, see [NetOps Portal Logs](#).

- b. If the log includes WARNING messages, such as `Failed to Retrieve Alarms`, or ERROR messages, such as `Handled unexpected Exception while processing Task`, then increase the value for the `RequestTimeoutSec` property in the alarm service configuration settings.

For more information about this property, see [Update the Configuration of the Alarm Service](#).

### Alarms are missing when reporting on a group context in NetOps Portal

#### Symptom:

Alarms are missing when reporting on a group context in NetOps Portal. To view Spectrum alarms in the context of a group in NetOps Portal, synchronize the items from Spectrum.

#### Solution:

Ensure that the applicable device exists within the correct IP Domain collection in Spectrum.

For more information about how to synchronize Spectrum devices with NetOps Portal, see [the DX NetOps Spectrum documentation](#).

### **The Alarms in the Alarms View are not Showing the Most Recent Items Displayed in OneClick**

#### **Symptom:**

The alarms that are displayed in the **Alarms** views in NetOps Portal do not show the most recent alarms that are displayed in OneClick.

#### **Solution:**

1. Determine if there are Spectrum data source synchronization issues.  
For more information, see [View System Status](#).
2. If there are synchronization issues, determine the cause of the issue.  
For more information, see [Synchronize Data Sources](#).

### **The Alarms in the Alarms View do not Match the Alarms Displayed in OneClick**

#### **Symptom:**

The alarms in the **Alarms** views in NetOps Portal do not match the alarms that are displayed in OneClick.

#### **Solution:**

NetOps Portal can report on a maximum of 20,000 alarms. If the group selection in NetOps Portal results in more than 20,000 alarms, some alarms will not display in the **Alarms** views in NetOps Portal.

Unlike Spectrum, NetOps Portal displays symptomatic and normal-severity alarms by default.

Do one of the following:

- Make NetOps Portal to be more consistent with the types of alarms that OneClick displays by hiding symptomatic and normal-severity alarms using an alarm filter.  
For more information about how to define alarm filters, see [Alarms View](#).
- Change the group selection in NetOps Portal so that the view displays a smaller number of alarms.

#### **NOTE**

NetOps Portal reports on the newest alarms.

### **Spectrum Alarms can take up to a minute to show up in Alarms Views**

#### **Symptom:**

Alarms that are generated in Spectrum are taking up to a minute or more to show up in **Alarms** views in NetOps Portal.

#### **Solution:**

By default, the alarm service process in the Event Manager polls for raised, updated, and cleared alarms every 30 seconds. You can reduce the frequency to every 10 seconds by updating the `SubscriptionHeartbeatIntervalSec` property in the default alarm service configuration settings.

For more information about how to update these settings, see [Update the Configuration of the Alarm Service](#).

## **Unable to Back Up the Data Repository**

Use this topic to help troubleshoot data repository back up issues.

#### **Symptom:**

When I back up the data repository by running the `vbr.py` script, the following message appears:

```
"Another vbr instance is already running"
```

#### **Solution:**

A previous backup attempt failed (for example, you did not set up password-less SSH correctly).

To back up the data repository, remove the `/tmp/.initiator.mutex` file from the system where the data repository is installed. The next scheduled backup occurs normally.

### Symptom:

You cannot back up the data repository and issue the `copycluster` command within Vertica using the `vbr` utility. The backup fails with this message:

```
[dradmin@hostname backup]$ /opt/vertica/bin/vbr.py --task backup --config-file /home/dradmin/backup/
backup_snapshot.ini
Error: Errors connecting to remote hosts: <IP>. Check SSH settings, and that the same Vertica version is
installed on all nodes.
Backup FAILED.
```

Checking the ssh configuration in the `/etc/ssh/sshd_config` file, `AllowTcpForwarding = No`. Enabling TCP forwarding by setting `AllowTcpForwarding = Yes` in the `/etc/ssh/sshd_config` file on the Vertica hosts, on both source and destination systems, is a prerequisite to backing up the data repository. The utility cannot forward connections from the database hosts to the backup or copycluster hosts.

### Solution:

Enable TCP forwarding. SSH connections to backup hosts do not require SSH forwarding.

1. Edit the `/etc/ssh/sshd_config` file settings on the Vertica hosts, on both source and destination systems by completing the following steps:
  - a. Ensure that the `AllowTcpForwarding` option is uncommented.
  - b. Set `AllowTcpForwarding = Yes`.
  - c. Save the file changes.
2. Restart the ssh service for it to read in the new configuration by issuing the following command:

```
systemctl restart ssh
```

TCP forwarding is enabled.

For more information, see [the Vertica documentation](#).

## Unable to Resolve Issue

Broadcom Support can triage issues using the troubleshooting information that the CARE tool collects. The tool includes configuration files in its scripts directory (`/RemoteEngineer`) for the NetOps Portal, the data aggregator, the data collectors, and the data repository components.



DX NetOps Performance Management includes the `re.sh` script in the following directories for the following components:

Component	
NetOps Portal	<code>&lt;installation_directory&gt;/PerformanceCenter/RemoteEngineer</code> <b>Example:</b> <code>/opt/CA/PerformanceCenter/RemoteEngineer</code> <ul style="list-style-type: none"> <li>• <b><i>installation_directory</i></b> The default installation directory for NetOps Portal. <b>Default:</b> <code>/opt/CA</code></li> </ul>
The data aggregator	<code>&lt;installation_directory&gt;/RemoteEngineer</code> <b>Example:</b> <code>/opt/IMDataAggregator/RemoteEngineer</code> <ul style="list-style-type: none"> <li>• <b><i>installation_directory</i></b> The default installation directory for the data aggregator. <b>Default:</b> <code>/opt/IMDataAggregator</code></li> </ul>
The data collectors	<code>&lt;installation_directory&gt;/RemoteEngineer</code> <b>Example:</b> <code>/opt/IMDataCollector/RemoteEngineer</code> <ul style="list-style-type: none"> <li>• <b><i>installation_directory</i></b> The default installation directory for the data collector. <b>Default:</b> <code>/opt/IMDataCollector</code></li> </ul>
The data repository	<code>&lt;installation_directory&gt;/IMDataRepository_verticaVers</code> <b>Example:</b> <code>/opt/CA/IMDataRepository_vertica10.1.1/RemoteEngineer</code> <ul style="list-style-type: none"> <li>• <b><i>installation_directory</i></b> The default installation directory for the data repository. <b>Default:</b> <code>/opt/CA</code></li> </ul>

### Verify the Prerequisites

Before running the CARE tool for a component, ensure that you have the required packages for the installer of that component.

For more information, see [Installation Requirements and Considerations](#).

### Run the CARE Tool

When prompted by Broadcom Support, run the CARE tool.

#### **Follow these steps:**

1. At a command prompt, navigate to the directory by issuing the following command:

```
cd <component_installation_directory>
```

- ***component\_installation\_directory***

Specifies the directory where CARE is installed.

**Example:** `/opt/CA/PerformanceCenter/RemoteEngineer`

2. Run the CARE tool by issuing the following command:

```
./re.sh
```

**NOTE**

The first time that you issue the command, a license agreement appears.

The CARE tool collects troubleshooting information for Broadcom Support triage purposes.

## Unexpected Data Aggregator Shutdown

**Symptom:**

The data aggregator shuts down unexpectedly.

**Solution:**

The data aggregator shuts down if it loses contact with the data repository. If contact with the data repository is lost, an audit message is logged in the `<installation_directory>/apache-karaf/shutdown.log` file.

- **installation\_directory**

The installation directory of the data aggregator.

**Default:** `/opt/IMDataAggregator`

**NOTE**

The `shutdown_details.log` logs heartbeat messages between the data aggregator and the data repository, as well as any data aggregator shutdowns for debugging purposes.

For more information about these logs, see [Data Aggregator Logs](#).

**To resolve connectivity or other data repository issues, perform the following steps:**

1. Verify that the data repository process is running. Do the following actions:
  - a. Log in to the database server you use for the data repository as the database administrator user, not as the root user.
  - b. Open the Vertica admintools utility from the `/opt/vertica/bin/adminTools` directory.  
The Administration Tools dialog opens.
  - c. Select **(1) View Database Cluster State**.  
The returning window should state: "ALL" for Host and "UP" for State.
2. If the data repository is not running, attempt to start it by performing the following steps:
  - a. Log in to the database server you use for the data repository.
  - b. Open the Vertica admintools utility from the `/opt/vertica/bin/adminTools` directory.  
The Administration Tools dialog opens.
  - c. Select **(3) Start Database**.
  - d. Press the **Space** bar next to the database name, select **OK**, and then press the **Enter** key on your keyboard.  
You are prompted for the database password.
  - e. Enter the database password, and then press the **Enter** key on your keyboard.  
The data repository database starts.

**NOTE**

If you see an error message stating that you cannot connect because of a username or password error, it is possible that a database password change is why the data aggregator has disconnected from the data repository.

- f. Select **(E) Exit**, and then press the **Enter** key on your keyboard.  
If the data repository does not start, contact Broadcom Support.
3. If the data repository is running, you have a network connection problem, such as a network latency issue. Address your network connectivity problem.
  4. After the data aggregator is running again, set up an automatic recovery of the data aggregator process.

## Vendor Certification Expression is Erroneous

### Symptom:

The MVEL compiler might not give an evaluation exception (error) for bad expressions. This situation can happen for some syntax errors, including missing or open parentheses, and multiple asterisks.

The incorrect expression is compiled without an error condition until an expression evaluation is performed with the appropriate variables. Database columns that are the target for the intended expression are not populated.

### Solution:

Turn on debug logging for the ExpressionEvaluator using the following steps:

1. Locate the `<installation_directory>/apache-karaf/etc/org.ops4j.pax.logging.cfg` file, create the following entry, and then save your changes:
 

```
log4j.logger.com.ca.im.core.expressionevaluator=DEBUG
```

  - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
2. Do one of the following steps:
  - Start the Data Aggregator service by issuing the following command:
 

```
systemctl start dadaemon
```
  - (Fault-tolerant environment) Enable the fault-tolerant data aggregator so that it can start when necessary by issuing the following command:
 

```
<installation_directory>/scripts/dadaemon activate
```

    - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
3. Search for evaluation exceptions in the `<installation_directory>/apache-karaf/data/log/karaf.log` file.
  - **installation\_directory**  
The installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`

## Vertica Fails to Install in a Cluster Environment

### Symptom:

Vertica fails to install in my cluster environment.

### Solution:

The hosts in a data repository cluster require that the root or sudo passwordless SSH user is set up prior to installing or upgrading the data repository. Verify that the passwordless SSH user is set up for the Vertica Linux database administrator user between the nodes in the cluster.

For more information about this prerequisite to installing the data repository, see [Prepare to Install the Data Repository](#).

## Vertica Fails to Start

### Symptom:

The data repository does not restart after you stopped the database through Vertica admintools.

### Solution:

A leap second occurred on 30 June 2015. The extra second causes Vertica to fail to start on certain versions of the Linux operating system.

To resolve this issue, restart each node before the first time you start the database.

The supported versions of Red Hat Enterprise Linux (RHEL) contain a fix for this issue.

## Vertica Fails to Install due to 'iptables' Error

### Symptom:

When you run the `dr_install.sh` script, Vertica fails to install due to the following 'iptables' error:

```
Prerequisites not fully met during local (OS) configuration for  
verify-ipAddress.xml:  
  
WARN (N0010): https://my.vertica.com/docs/version/HTML/index.htm#cs hid=N0010  
  
Linux iptables (firewall) has some non-trivial rules in tables: filter
```

### Solution:

If you encounter this Vertica known issue, reach out to CA Support for a workaround.

## View Shows Invalid RIB Query Syntax Error

### Symptom

A view does not load, and the view shows the following error message:

Invalid RIB query syntax. See logs for details.

### Solution

Percentile or Projection attributes are disabled. Verify that the attributes for Percentiles and Metric Projections are configured correctly in the metric family XML. Enable the calculation, or remove the reporting field from the view.

For more information, see [Metric Family XML Structure](#).

# APIs

---

You can automate provisioning and configuration tasks or extract data from the DX NetOps Performance Management database using the APIs that DX NetOps Performance Management offers. These REST web services exposes the most frequently repeated or time-consuming tasks.

Some APIs consist of RESTful web services. You can access a set of resources by a fixed set of operations using the REST model. This model takes advantage of widely deployed HTTP features that are supported by common hardware, such as gateway devices.

The global administrator can perform many administrative tasks using the RESTful web services.

The following categories of APIs are available:

- **NetOps Portal REST Web Services**  
Use these web services to automate tasks that you manually perform in NetOps Portal.  
For more information, see [NetOps Portal REST Web Services](#).
- **Data aggregator REST Web Services**  
Use these web services to manage administrative operations, such as retrieving data or managing relationships between profiles and tenants or groups.  
For more information, see [Data Aggregator REST Web Services](#).
- **OpenAPI**  
Use the flexible Open API tool to extract data from the DX NetOps Performance Management database. OpenAPI enables integration between DX NetOps Performance Management data and external applications.  
For more information, see [OpenAPI](#).

## Interact with the REST Web Services

You can interact with the available web services using a REST client or using cURL.

### **NOTE**

Due to known limitations, set up the REST client using a browser other than Firefox.

## NetOps Portal REST Web Services

Programmatically manage dashboards, data sources, group, roles, and more using the NetOps Portal RESTful web services.

You can programmatically perform the following tasks using the NetOps Portal RESTful web services:

- Access the NetOps Portal RESTful web services.  
For more information, see [Use NetOps Portal Web Services](#).
- Manage REST API tokens.  
For more information, see [Generate a REST API Token](#).
- Manage dashboards.  
For more information, see [Dashboards Web Service](#).
- Manage data sources.  
For more information, see [Data Sources Web Service](#).
- Manage devices and device life cycles.  
For more information, see [Devices Web Service](#).
- Manage IP domain definitions.  
For more information, see [Domains Web Service](#).
- Manage groups.

- For more information, see [Groups Web Service](#).
- Manage roles.  
For more information, see [Roles Web Service](#).
- Manage user accounts.  
For more information, see [Users Web Service](#).
- Manage tenants.  
For more information, see [Tenants Web Service](#).
- Retrieve a list of the configuration items, such as custom user accounts, roles, or groups, that are already in the system.  
For more information, see [Console Info Web Service](#).
- Manage SNMP profiles.  
For more information, see [Automate Tenant Configuration with REST Web Services](#).
- Manage business hours definitions.  
For more information, see [Business Hours Web Service](#).
- Manage maintenance indicators.  
For more information, see [Maintenance Indicators Web Service](#).
- Manage alarm attributes.  
For more information, see [Alarm Attributes Web Service](#).

For more information about how to use these web services, see [Use NetOps Portal Web Services](#).

## Use NetOps Portal Web Services

Access the API components that DX NetOps Performance Management includes as RESTful web services.

In this article:

- [Access the NetOps Portal Web Services](#)
- [Access the NetOps Portal REST Web Services Documentation](#)
- [The Endpoints and the XML Schema](#)
- [Connect a REST Client to the NetOps Portal REST Web Services](#)

### Access the NetOps Portal Web Services

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can access the API components that DX NetOps Performance Management includes.

Access these components as RESTful web services using the following methods:

- [Use a web browser](#)
- [Use cURL](#)

### Access the NetOps Portal REST Web Services from a Web Browser

The NetOps Portal REST Web Services endpoints require authentication.

#### **NOTE**

You can access the full list of NetOps Portal RESTful web services at the following URL:

`http://<PC_host>:8181/pc/center`

The resulting API launch page includes a list of the available web services, endpoint addresses, and Web Application Description Language (WADL) and Web Service Definition Language (WSDL) URLs. From the launch page, you can access the WADL of each web service for testing. If you use a testing utility to run web service calls, you receive feedback that is useful for debugging purposes. Using a testing utility can be a timesaver. You can supply username and password parameters as service endpoints for automatic authentication of all service calls. Testing utilities require a WSDL file that

describes the service being tested. WSDL files are XML files that conform to the WSDL. In the REST format, use the simpler WADL. To access the SOAP web services, use the link to the WSDL.

### **Access the NetOps Portal REST Web Services using cURL**

You can access the NetOps Portal REST web services using the standard cURL utility. cURL is a popular command-line utility to interact with various network protocols, including HTTP and HTTPS. It is pre-installed on the Mac. You must install it on Windows. The following are a few examples of using cURL.

#### **Examples:**

Request the password from the user using the `-u` option with only the username:

```
curl -s -k -u admin https://<PC_host>:8181/rest
```

Have a shell script request the username and password, without the password appearing in 'ps' output:

```
echo "DA REST requires authentication."
read -p "Username: " user
read -s -p "Password: " pass
echo
curl --config - -s -k https://<PC_host>:8181/rest <<< "user = \"${user}:${pass}\""
```

### **Access the NetOps Portal REST Web Services Documentation**

Most of the web services provide their own documentation, including lists and descriptions of the available parameters and operations. You can access this documentation in HTML format from the launch page:

```
http://<PC_host>:8181/pc/center/rest
```

### **The Endpoints and the XML Schema**

The endpoints act on a common set of data. You can represent the data in different data formats (for example, MIME types). The format depends on the endpoint that consumes or produces the data. An XML schema describes the data and other supported data formats, such as JSON.

The articles in this section of the documentation describe the basic terms and parameters of the XML schema to create scripts for the NetOps Portal RESTful web services. You can group the data by namespace. A schema describes the types and elements of each namespace. *Types* define the structure of the data, while *elements* are instances of a type. For example, REST endpoints produce or consume elements, and the structure of each element is described by its type.

### **Connect a REST Client to the NetOps Portal REST Web Services**

Connect the REST client that you will use to invoke the NetOps Portal web services to the NetOps Portal RESTful web services using one of the following authorization types:

- [Token authentication](#)  
Use this method when you want to protect your password and *not* send it in clear text from the REST client.
- [Basic authentication](#)

### **Connect a REST Client to the NetOps Portal REST Web Services Using Token Authentication**

**Prerequisite:** You have a REST API token. If you do not have a REST API token, you can [generate one](#).

#### **Follow these steps:**

1. Launch a REST client.

2. Enter the following URL for the NetOps Portal RESTful web services API in the **URL** field, selecting **GET** for the **HTTP Method**:  
`http://<PC_host>:8181/pc/center/webservice/<Web_Service_Name>`
  - **Web\_Service\_Name**  
The web service name.  
**Examples:** `tenants`, `users`, `devices`, `groups`, `roles`  
For a list of the web services, see [NetOps Portal REST Web Services](#).
3. Select **Bearer Token** as the authorization type.
4. Paste the REST API token in the **Token** field.
5. Click **Send**.

The REST client is connected to the NetOps Portal REST web services by way of the REST API token.

### Connect a REST Client to the NetOps Portal REST Web Services Using Basic Authentication

Follow these steps:

1. Launch a REST client.
2. Enter the following URL for the NetOps Portal RESTful web services API in the **URL** field, selecting **GET** for the **HTTP Method**:  
`http://<PC_host>:8181/pc/center/webservice/<Web_Service_Name>`
  - **Web\_Service\_Name**  
The web service name.  
**Examples:** `tenants`, `users`, `devices`, `groups`, `roles`  
For a list of the web services, see [NetOps Portal REST Web Services](#).
3. Select **Basic Auth** as the authorization type.
4. Enter a valid username and password for a user account that has global administrator access, and then click **OK**.
5. Select **Custom Header** from the **Headers** menu.  
The **Request Header** dialog opens.
6. Enter `Content-Type` as the value for the **Name** parameter.
7. Enter `application/xml` in the **Value** field, and then click **OK**.  
The **Headers** section now shows the following updated values:  
`Authorization: Basic YWRtaW46YWR...`  
`Content-Type: application/xml`

The REST client is connected to the NetOps Portal REST web services by way of basic authentication.

## Dashboards Web Service

Perform dashboard management tasks using the `dashboards` endpoint for the NetOps Portal web service.

Users with the Administer Users and the Allow Access to REST Services role rights, or LDAP-authenticated user accounts with the Admin role right, can use the `dashboards` endpoint to do the following tasks:

- Manage (create and update) dashboards.
- Import and export dashboards.
- Display data views.
- Build the dashboards once and deploy them in an additional tenant. During the import process, you can choose to deploy extra dashboards without manually recreating them.

In this article:



- [Verify the Prerequisite](#)
- [Access the dashboards Web Service Documentation](#)
- [Export a Dashboard](#)
- [Import a Dashboard](#)
- [Update a Dashboard](#)

### **Verify the Prerequisite**

- You have connected the REST client that you will use to invoke the web service to the NetOps Portal RESTful web services.

### **Access the dashboards Web Service Documentation**

Issue the following call to see the parameters for the web service:

```
http://<PC_host>:8181/pc/center/rest/dashboards/documentation
```

### **Export a Dashboard**

Export a dashboard to an XML file that you will use to import into another instance of NetOps Portal. This process requires the internally-assigned page ID. Use NetOps Portal to find the ID.

#### **Follow these steps:**

1. Navigate to the dashboard that you want to export.
2. In the browser window, find the page ID in the URL.

#### **Example:**

```
http://<PC_host>:8181/pc/desktop/page?pg=2000040
```

The page ID is 2000040 .

3. In the REST client, set the URL to the following:

```
http://<PC_host>:8181/pc/center/webservice/dashboards/<pageId>
```

#### **Example:**

```
http://<PC_host>:8181/pc/center/webservice/dashboards/2000040
```

4. Select **GET** for **HTTP Method**.  
An error or success message appears in the response.

An XML file is created that represents the exported dashboard.

### **Import a Dashboard**

#### **WARNING**

Exporting a dashboard in XML format using the web service, and then importing it from the XML file can result in a broken dashboard.

**Best Practice:** Create a dashboard by copying an existing one using NetOps Portal.

For more information, see [Manage Dashboards](#).

#### **Follow these steps:**

1. (If you are importing a dashboard for a specific tenant) Log in as the administrator for that tenant.
2. Browse to the XML file that represents a dashboard that you have exported, or paste the exported text of the XML file into the **Body** text field.
3. Update the `dashboardMenu` , `menuItem` , and `dashboardTitle` attributes in the XML body for the dashboard that you want to import:

```
<dashboardMenu>MyCustomDashboard_DashboardMenuName</dashboardMenu>
<menuItem>MyCustomDashboard_MenuItemName</menuItem>
<dashboardTitle>MyCustomDashboard_DashboardTitleName</dashboardTitle>
```

– **<dashboardMenu>**

The title that appears at the top of the dashboard menu. Specify an existing top-level dashboard menu title for this attribute.

**Examples:** Infrastructure Health, Capacity Planning

– **<menuItem>**

The title of the dashboard that appears in the dashboards menu. This property is unique to each tenant.

– **<dashboardTitle>**

The title of the dashboard that appears at the top of the dashboard page. This property is unique to each tenant.

4. (Optional) Test the import of the dashboard that you want to import before you import it. Perform the following **POST** operation:

```
http://<PC_host>:8181/pc/center/webservice/dashboards/test/
```

The following message appears to indicate that the test was successful:

```
<?xml version="1.0" encoding="UTF-8"?>
<msg>Validation successful.</msg>
```

Checking is performed for the unique `menuItem` and `dashboardTitle` attributes.

5. Set the URL to the following:

```
http://<PC_host>:8181/pc/center/webservice/dashboards/import
```

6. Perform a **POST** operation.

An error or success message appears in the response.

The web service assigns a page ID to the imported dashboard.

## **Update a Dashboard**

You update a dashboard based on a page ID. Edit the XML to modify the dashboard.

### **Follow these steps:**

1. Set the URL to the following:

```
http://<PC_host>:8181/pc/center/webservice/dashboards/<pageId>
```

2. Paste the XML file that represents the dashboard into the **Body** text field, and then edit the XML as needed.

3. Perform a **PUT** operation.

An error or success message appears in the response.

## **Data Sources Web Service**

You can manage data sources using the `datasources` web service.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services.

You can perform the following tasks to manage data sources using the `datasources` web service:

- Get item IDs from local IDs.
- Get local IDs from the item IDs.
- Manage the data source log file.
- Modify data source settings, such as the authentication method to use.
- Register new data sources.
- Remove data sources.
- Synchronize data sources.
- View the current registered data sources.
- Get the current PC version.

### **Access the datasources Web Service Documentation**

Issue the following call to see the parameters for the `datasources` web service:

```
http://PC_host:8181/pc/center/webservice/datasources
```

Issue the following call to see a list of supported operations:

```
http://PC_host:8181/pc/center/rest/datasources/documentation
```

### **Retrieve the Values for the Data Source Parameters**

You can retrieve the current values for the following data source parameters using the GET command:

- **id**  
An internally assigned identifier for the data source.
- **enabled**  
Indicates whether the data source is enabled.  
**Default:** true
- **name**  
The hostname of the data source.
- **authtype**  
The type of authentication to use for this data source.  
**Values:** NONE, BASIC
- **type**  
**Values:**
  - **REPORTER:** DX NetOps Network Flow Analysis
  - **SUPER\_AGENT:** CA Application Delivery Analysis
  - **VOIP\_MONITOR:** CA Unified Communications Monitor
  - **EVENT\_MANAGER:** The Event Manager service
  - **SPECTRUM\_IM:** DX NetOps Spectrum
  - **APM:** DX Application Performance Management
  - **DATA\_AGGREGATOR:** The default DX NetOps Performance Management data source
  - **CATALYST\_CONNECTOR:** The CA Catalyst Connector
  - **SERVICE\_OPERATIONS\_INSIGHT:** CA Service Operations Insight
  - **RESERVED\_CUSTOMER\_N:** An enum that has been reserved for the unspecified custom uses.
- **consoleSameAsDataSource**  
Indicates whether the data source web console address is the same as the hostname. Use this parameter in cases where network address translation is deployed.  
**Default:** true
- **consoleAddress**  
The IP address where DX NetOps Performance Management contacts the data source console.

## Data Sources Web Service Examples

The following examples demonstrate some of the operations for the `datasources` web service:

- **get item ids**

Retrieves the NetOps Portal item ids given the data source local id. If an id is not found, it is omitted from the returned data.

**URL:** `http://hostname:8181/pc/center/webservice/datasources/{idName}/{idValue}/itemids`

**Where:**

- **{idName}**

One of the property name values that the `get id names` method of this web service returns.

`{idName}` can be any of the following values:

- **dataSourceId**

Example: `http://PC_address:8181/pc/center/webservice/datasources/dataSourceId/3/itemids`

- **dataSourceGUID**

Example: `http://PC_address:8181/pc/center/webservice/datasources/dataSourceGUID/7b0ef6b1e9094d599821b2b07d78d83d/itemids`

- **dataSourceConsoleName**

Example: `http://PC_address:8181/pc/center/webservice/datasources/dataSourceConsoleName/Data%20Aggregator%40PC_DA.ca.com/itemids`

- **{idValue}**

A value for the property denoted by `idName`.

**HTTP method:** POST

**XSD for the provided XML:** `http://hostname:8181/pc/center/rest/datasources/xsd`

To translate an array of local ids into item ids, supply XML in the following format:

```
<LocalIDs>
  <LocalID ID="4514"/>
  <LocalID ID="4705"/>
  <LocalID ID="4501"/>
  <LocalID ID="4540"/>
  <LocalID ID="4511"/>
  <LocalID ID="4499"/>
</LocalIDs>
```

The function returns an array of `ItemIDResult` objects.

**Example:**

```
<ItemIDResults>
  <ItemIDResult ItemID="406" LocalID="4514"/>
  <ItemIDResult ItemID="407" LocalID="4705"/>
  <ItemIDResult ItemID="408" LocalID="4501"/>
  <ItemIDResult ItemID="409" LocalID="4540"/>
  <ItemIDResult ItemID="410" LocalID="4511"/>
  <ItemIDResult ItemID="411" LocalID="4499"/>
</ItemIDResults>
```

- **get local ids**

Retrieves the data source local id given the NetOps Portal item ids. If an id is not found, it is omitted from the returned data.

**URL:** `http://hostname:8181/pc/center/webservice/datasources/{idName}/{idValue}/localids`

**Where:**

- **{idName}**

One of the property name values that the `getIdNames` method of this web service returns.

`{idName}` can be any of the following values:

- **dataSourceId**

Example: `http://PC_address:8181/pc/center/webservice/datasources/dataSourceId/3/localids`

- **dataSourceGUID**

Example: `http://PC_address:8181/pc/center/webservice/datasources/dataSourceGUID/7b0ef6b1e9094d599821b2b07d78d83d/localids`

- **dataSourceConsoleName**

Example: `http://PC_address:8181/pc/center/webservice/datasources/dataSourceConsoleName/Data%20Aggregator%40PC_DA.ca.com/localids`

- **{idValue}**

A value for the property denoted by `idName`.

**HTTP method:** POST

**XSD for the provided XML:** `http://hostname:8181/pc/center/rest/datasources/xsd`

To translate an array of the specified NetOps Portal item ids to local ids, supply XML in the following format:

```
<ItemIDs>
  <ItemID ID="412"/>
  <ItemID ID="413"/>
</ItemIDs>
```

The function returns an array of `ItemIDResult` objects.

**Example:**

```
<ItemIDResults>
  <ItemIDResult ItemID="412" LocalID="4590"/>
  <ItemIDResult ItemID="413" LocalID="4760"/>
</ItemIDResults>
```

## Devices Web Service

Perform device management tasks using the NetOps Portal RESTful web services.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services.

Use the `devices` web service to:

- View the devices that belong to a tenant.
- Retrieve a list of interfaces from a device (a router or switch) in the inventory of managed items.

Devices are associated with a tenant. If you are not deploying multi-tenancy, devices are associated with the Default Tenant. Retrieve a list of devices that are associated with your tenant by running the `getList` method to the root `/devices/`.

### Verify the Prerequisites

- You have connected the REST client that you will use to invoke the PC web services to the PC RESTful web services.
- You are a local admin user.

### Access the devices Web Service Documentation

Issue the following call to see the parameters for the `devices` web service:

```
http://PC_host:8181/pc/center/webservice/devices
```

Issue the following call to see a list of supported operations:

```
http://PC_host:8181/pc/center/rest/devices/documentation
```

### Available GET Methods

Issue these calls as a local admin user.

- **get list**  
Gets a list of devices that belong to the tenant of the logged-in user.  
`http://PC_host:8181/pc/center/webservice/devices`
- **get interfaces**  
Gets a list of all interfaces that belong to a specified device.  
`http://PC_host:8181/pc/center/webservice/devices/{idName}/{idValue}/interfaces`

You can retrieve a list of internally assigned identifiers that you can use in other methods to identify devices using the following URL:

```
http://PC_host:8181/pc/center/webservice/devices/idNames
```

Or you can retrieve a list that is filtered by tenant and by device subtype as follows:

```
http://PC_host:8181/pc/center/webservice/devices/subtype/tenant/tenantIdName/  
tenantIdValue
```

The `subtype` parameter further describes the device. For example, the `server`, `switch`, and `router` subtypes identify devices, while an interface can have subtype "physical" or "virtual".

#### TIP

You can retrieve a list of group members that includes their subtype by issuing a GET operation to the following URL for the `groups` web service:

```
http://PC_host:8181/pc/center/webservice/groups/{idName}/{idValue}/items
```

For more information about the `groups` web service, see [Groups Web Service](#).

## Set Device Alias Names Using the Devices Web Service

As an Administrator, you can set an alias name for a monitored device using the `devices` endpoint for the NetOps Portal REST web service. The alias name appears in the inventory list of devices and in the inventory list of interfaces.

#### TIP

You can also set the alias name for a monitored device using the `update_alias_name.sh` script. You can use this script to set the alias names for multiple monitored devices. Device alias names set from the REST web service take precedence over the alias names that are set using the `update_alias_name.sh` script.

For more information, see [Set Alias Names Using a Script](#).

The following URL shows the syntax for the `devices` endpoint for the NetOps Portal REST service:

```
http://<PC_host>:<port>/pc/center/webservice/devices
```

- **PC\_host**  
Specifies the NetOps Portal host name.
- **port**  
Specifies the NetOps Portal required port number.

**Default: 8181**

For more information about the NetOps Portal server ports that should be open to allow DX NetOps Performance Management communications to function properly, see [Installation Requirements and Considerations](#).

**Follow these steps:**

1. Determine the devices for which you want to set an alias name by issuing a GET request to the following `devices` endpoints for the NetOps Portal REST service:

```
http://<PC_host>:<port>/pc/center/webservice/devices
```

A list of monitored devices is shown in the response.

2. Issue a PUT request to the following `devices` endpoint for the NetOps Portal REST web service:

```
http://<PC_host>:<port>/pc/center/webservice/devices/deviceItemId/<device_id>/
aliasName/<alias_name>
```

**Example:**

```
http://<PC_host>:<port>/pc/center/webservice/devices/deviceItemId/119/aliasName/Alias
name for 96.24
```

- **device\_id**  
The identification number for the monitored device for which you want to set an alias name.
- **alias\_name**  
The alias name for the monitored device.

**Set Interface Alias Names Using the Devices Web Service**

As an Administrator, you can set an alias name for an interface using the `devices` endpoint for the NetOps Portal REST web service. The alias name appears in dashboards and in views depending on a user's role rights.

**NOTE**

You can also set an alias name for an interface using the `update_alias_name.sh` script. You can use this script to set the alias names for multiple interfaces. Interface alias names set using REST web services take precedence over the alias names that are set using the `update_alias_name.sh` script.

For more information, see [Set Alias Names Using a Script](#).

The following URL shows the syntax for the `devices` endpoint for the NetOps Portal REST service:

```
http://PC_host:8181/pc/center/webservice/devices
```

**Follow these steps:**

1. Determine the interfaces for which you want to set an alias name by issuing a GET request to *one* of the following `devices` endpoint for the NetOps Portal REST service:

```
http://PC_host:8181/pc/center/webservice/devices/deviceItemId/<device_id>/interfaces
```

- **device\_id**  
The device item identification number for the monitored device with which the interface is associated.

```
http://PC_host:8181/pc/center/webservice/devices/domainItemId/<domain_id>/<device_IP>/interfaces
```

- **domain\_id**  
The identification number for the IP domain that the device is in.
- **device\_IP**  
The IP address of the device.

A list of interfaces for the device is shown in the response.

2. Issue a PUT request to *one* of the following `devices` endpoints for the NetOps Portal REST service, with the following body:

```
http://PC_host:8181/pc/center/webservice/devices/deviceItemId/<device_id>/setInterfaceNameAlias
```

– **device\_id**

The identification number for the device item with which the interface is associated.

```
http://PC_host:8181/pc/center/webservice/devices/domainItemId/<domain_id>/<device_IP>/
setInterfaceNameAlias
```

– **domain\_id**

The identification number for the IP domain that the device is in.

– **device\_IP**

The IP address of the device.

**Body:**

```
<interfaces>
  <interface>
    <itemId>interface_item_id</itemId>
    <nameAlias>alias_for_interface</nameAlias>
  </interface>
  ...
</interfaces>
```

**Example:**

```
<interfaces>
  <interface>
    <itemId>164</itemId>
    <nameAlias>Et0 - alias</nameAlias>
  </interface>
  <interface>
    <itemId>165</itemId>
    <nameAlias>Se0 - alias</nameAlias>
  </interface>
</interfaces>
```

– **interface\_item\_id**

The identifier of the interface for which you want to set an alias name.

– **alias\_for\_interface**

The alias name that you want to set for the interface.

## Set Component Alias Names Using the Devices Web Service

As an Administrator, you can set the alias name for a component using the `devices` endpoint for the NetOps Portal REST web service. The alias name appears in dashboards and in views depending on a user's role rights.

**NOTE**

You can also set an alias name for a component using the `update_alias_name.sh` script. You can use this script to set the alias names for multiple components. Component alias names set using REST web services take precedence over the alias names that are set using the `update_alias_name.sh` script.

For more information, see [Set Alias Names Using a Script](#).

The following URL shows the syntax for the `devices` endpoint for the NetOps Portal REST service:

```
http://<PC_host>:8181/pc/center/webservice/devices
```



**Follow these steps:**

1. Determine the components for which you want to set an alias name by issuing a GET request to *one* of the following `devices` endpoints for the NetOps Portal REST service:

```
http://<PC_host>:8181/pc/center/webservice/devices/deviceItemId/<device_id>/components
```

- **`device_id`**

The identification number for the monitored device with which the component is associated.

```
http://<PC_host>:8181/pc/center/webservice/devices/domainItemId/<domain_id>/<device_IP>/components
```

- **`domain_id`**

The identification number for the IP domain that the device is in.

- **`device_IP`**

The IP address of the device.

A list of components for the device is shown in the response.

2. Issue a PUT request to *one* of the following `devices` endpoints for the NetOps Portal REST service, with the following body:

```
http://<PC_host>:8181/pc/center/webservice/devices/deviceItemId/<device_id>/setComponentNameAlias
```

- **`device_id`**

The identification number for the device item with which the component is associated.

```
http://<PC_host>:8181/pc/center/webservice/devices/domainItemId/<domain_id>/<device_IP>/
setComponentNameAlias
```

- **`domain_id`**

The identification number for the IP domain that the device is in.

- **`device_IP`**

The IP address of the device.

**Body:**

```
<deviceComponents>
  <deviceComponent>
    <itemId>component_item_id</itemId>
    <nameAlias>alias_for_component</nameAlias>
  </deviceComponent>
  ...
</deviceComponents>
```

**Example:**

```
<deviceComponents>
  <deviceComponent>
    <itemId>164</itemId>
    <nameAlias>Et0 - alias</nameAlias>
  </deviceComponent>
  <deviceComponent>
    <itemId>165</itemId>
    <nameAlias>Se0 - alias</nameAlias>
  </deviceComponent>
</deviceComponents>
```

- **`component_item_id`**

The identifier of the component for which you want to set an alias name.

- **`alias_for_component`**

The alias name that you want to set for the component.

## Manage Device Life Cycles Using the Devices Web Service

You can manage the device life cycle state.

A monitored device's life cycle state defines the monitoring behavior for the device, such as whether the device is active, retired, or in maintenance. Components inherit the life cycle state of the associated device. You can define the usage state of SNMP and ICMP devices by managing the life cycle.

You can manage the device life cycle state using the NetOps Portal API (using the `devices` endpoint) or using NetOps Portal. This article describes how to manage the state using the NetOps Portal API.

For more information about how to manage the state using NetOps Portal, see [Manage Device Life Cycles](#).

### TIP

In environments where DX NetOps Spectrum manages the life cycle state, manage the state using the DX NetOps Spectrum REST API instead of this REST API.

For more information, see the [DX NetOps Spectrum documentation](#).

### NOTE

You cannot manage the life cycle state of DX NetOps Mediation Manager devices.

Issue a PUT request to the following `devices` endpoint:

```
http://<PC_host>:8181/pc/center/webservice/devices/deviceItemId/<itemID>/lifeCycleState/<newState>
```

### Example:

The following example sets the life cycle state for NetOps Portal device item ID 1037 to "ACTIVE".

```
http://<PC_host>:8181/pc/center/webservice/devices/deviceItemId/1037/lifeCycleState/ACTIVE
```

- **itemID**  
Specifies the NetOps Portal device item ID.
- **newState**  
Specifies the life cycle state, in ALL CAPS.

#### Options:

- **ACTIVE**  
Specifies the normal operating status of a device.
- **RETIRED**  
Specifies that the device is no longer in use and that no monitoring occurs. This state disables polling, threshold monitoring, notifications, and change detection. DX NetOps Performance Management does not update the SNMP Profile or change the hostname.

#### NOTE

In this state, virtual devices (Virtual Machine or ESX) discovered with SystemEDGE are not polled for its vCenter statistics.

- **MAINTENANCE**  
Specifies that the device is temporarily under maintenance. You can configure the behavior of this state.

#### NOTE

In this state, DX NetOps Performance Management does not update the device during discovery or rediscovery.

## Domains Web Service

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services.

Managed items are associated with IP domains by the data sources during data collection. Consult the documentation for each registered data source to determine the steps to take to associate items with domains.

In this article:

### **Verify the Prerequisite**

- You have connected the REST client that you will use to invoke the NetOps Portal web services to the NetOps Portal RESTful web services.

### **Access the domains Web Service Documentation**

Issue the following call to see the parameters for the `domains` web service:

```
http://PC_host:8181/pc/center/webservice/domains
```

Issue the following call to see a list of supported operations:

```
http://PC_host:8181/pc/center/rest/domains/documentation
```

### **Available GET Methods**

- **get list**  
Gets a list of all of the IP domains that belong to the tenant for the logged-in user.  
Run the `getList` method to the root `/domains/` to retrieve a list of all IP domain IDs for the tenant of the currently logged-in administrator:
  - If you are logged in as a global administrator, you see a list of domain IDs for the Default Tenant only.
  - If you are logged in as a tenant administrator, you see a list of domain IDs for your tenant.
- **get**  
Retrieves information about a specified IP domain:  

```
http://PC_host:8181/pc/center/webservice/domains/idName/idValue
```

#### **NOTE**

The administrator sees only the IP domains that fall within the Default Tenant. A tenant administrator sees only the IP domains within that tenant.

- **get domain for group**  
Gets the IP domain associated with a specified group:  

```
http://PC_host:8181/pc/center/webservice/domains/group/groupIdName/groupIdValue
```
- **get id names**  
Retrieves a list of identifiers that can be used to identify IP domains in other web service operations:  

```
http://PC_host:8181/pc/center/webservice/domains/idNames
```
- **get list**  
Gets all of the IP domains that belong to the tenant for the logged-in user:  

```
http://PC_host:8181/pc/center/webservice/domains
```
- **get list by tenant**  
Retrieves the list of all of the IP domains that are associated with the specified tenant ID:  

```
http://PC_host:8181/pc/center/webservice/domains/tenantItemId/tenantId
```

- **get list with translation**

Gets all of the IP domains that belong to the tenant for the logged-in user. Any localized text is translated to the specified language:

```
http://PC_host:8181/pc/center/webservice/domains/cultureId
```

#### **Available POST Method**

- **create**

Creates an IP domain:

```
http://PC_host:8181/pc/center/webservice/domains
```

#### **Available PUT Method**

- **update**

Updates a specified IP domain.

```
http://PC_host:8181/pc/center/webservice/domains/idName/idValue
```

#### **Available DELETE Method**

- **delete**

Deletes an IP domain definition.

```
http://PC_host:8181/pc/center/webservice/domains/{idName}/{idValue}
```

#### **Basic IP Domain Parameters**

The current values for the following domain parameters are available from the GET command:

- **cultureID**

Specifies a language (locale). Supply a language identifier from the following list:

- en-US (English, United States)
- ja-JP (Japanese)
- zh-CN (Simplified Chinese)
- fr-FR (French, France)

- **dnsProxyAddress**

Is the IP address of the DNS proxy server.

- **domainItemID**

Is an internal (database) identifier for a tenant definition.

- **groupItemID**

Is an internal (database) identifier for the group definition associated with a domain.

- **itemDesc**

Describes this domain namespace, such as naming the enterprise that owns it.

- **itemName**

Identifies the domain.

tenantID is an internal (database) identifier for a tenant definition. Identifies the tenant associated with this domain.

- **primaryDNSAddress**

Is the IP address of the primary name server for this domain.

- **primaryDNSPort**  
Is the port number that the primary name server uses.
- **secondaryDNSAddress**  
(Optional) Is the IP address of the secondary name server for this domain.
- **secondaryDNSPort**  
(Optional) Is the port number that the secondary name server uses.
- **isDNSProxyEnabled**  
Indicates whether the proxy address is enabled for this IP domain.
- **deviceAlias**  
Indicates the alias to use for a managed item. A device alias is a user-configured name that is applied to the associated managed item in NetOps Portal.
- **deviceAliasList**  
Identifies a CSV or TXT file of alternate interface descriptions. A comma-separated list of values that include the device IP address, interface name, interface description, and alternate interface description (alias) mappings.
- **interfaceDescriptionOverride**  
Indicates the alternate description to use for an interface. Overrides the interface descriptions that appear in NetOps Portal by default.

## Groups Web Service

Perform common group management tasks using the `groups` web service, such as creating and managing groups of monitored items.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services. You can create new groups and can add items to them manually. You can also write rules to create and populate groups that are based on item attributes.

### Verify the Prerequisite

- You have connected the REST client that you will use to invoke the NetOps Portal web services to the NetOps Portal RESTful web services.

### Access the Documentation

Issue the following call to see the parameters for the `groups` web service:

```
http://PC_host:8181/pc/center/webservice/groups
```

### Access the Available Methods

Issue the following call to see a list of supported operations for the `groups` web service:

```
http://PC_host:8181/pc/center/rest/groups/documentation
```

## Manage Groups Using the Groups Web Service

You can add and delete multiple groups or sub-groups, add items to groups and create group rules to populate groups, and you can export and import group definitions.

You can manage (create, edit, delete groups) groups or collections using NetOps Portal or using the `groups` endpoint for the NetOps Portal API. This article describes how to manage them using the `groups` web service.

For more information about how to manage groups using NetOps Portal, see [Manage Groups](#).

In this article:

- [Create a Group](#)
- [Export a Group](#)
- [Identify a Group](#)
- [Look Up the ID of a Group](#)
- [Retrieve a List of Groups or Sub-Groups](#)
- [Syntax for Site Group Management](#)
- [groups Web Service Example Syntax](#)
- [Retrieve a List of Group Members](#)
- [Add an Item to a Group](#)
- [Add a Group Rule to a Group](#)
- [Group and Group Rule Attributes](#)
- [About Modifying or Deleting Groups](#)
- [Remove an Item from a Group](#)
- [Delete a Group](#)
- [Migrate a Group to Another NetOps Portal Database](#)
- [About Group Deletion](#)

#### NOTE

The procedures in this article contain URIs that are constructed using the default server port, 8181.

### Create a Group

The global Administrator can create and populate groups. Groups are always associated with a tenant definition. If you are not deploying multi-tenancy, groups are associated with the default tenant. Otherwise, group management tasks apply only to the current tenant.

You can add monitored devices to groups automatically by creating group rules and applying them to groups. You can also add sub-groups to your groups.

This includes:

- [Create a Group in the Default Tenant](#)
- [Create a Group in a Custom Tenant](#)

### Create a Group in the Default Tenant

Create and configure a group associated with the default tenant by supplying the group ID or the group path as parameters, in the `/All Groups` path. You can also use this procedure to create a site group.

#### Follow these steps:

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:

```
http://PC_host:8181/pc/center/webservice/groups/<useIds>/<allowDeletes>
```

#### Example:

```
http://PC_host:8181/pc/center/webservice/groups/false/false
```

#### – useIds

Indicates whether the `groups` web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported.

#### Values:

- **true**: The groups web service uses the IDs of the groups.
- **false**: The groups web service does not use the IDs of the groups, and instead, creates new groups.

#### – **allowDeletes**

Indicates whether the groups web service can delete the group that you are creating. Allows the groups web service to update the group rules that are defined in any existing groups that are overwritten by this XML.

##### **Values:**

- **true**: Honor the allowDeletes values at the group level.
- **false**: Ignore the allowDeletes values at the group level.

For more information about this parameter, see the ["Group and Group Rule Attributes"](#) section.

2. Select **POST** for **"HTTP" Method**.
3. Provide a valid Username and Password in the request header for a user account that has global administrator access to NetOps Portal.
4. Select **'application/xml'** as the **'Body Content-type'**.
5. Add the following XML within the "Body" text section, replacing the values with the values that you want to use for the new group:

```
<GroupTree path="/All Groups">
  <Group name="Group Name" desc="Description of the group"
    inherit="true" type="user group"/>
</GroupTree>
```

#### **XML for a site group:**

```
<Group desc="Site group to represent the entire United States"
  inherit="true" location="North America" name="USA-Site" type="site"/>
```

#### **XML for a sub-group:**

In this example, the Raleigh group is added as a sub-group of the All Groups\USA group:

```
<GroupTree path="/All Groups/USA">
  <Group name="Raleigh" desc="This is the group for
    managed items in Raleigh, NC" inherit="true"
    type="user group"/>
</GroupTree>
```

#### – **inherit**

Indicates whether the group includes child items of group members. For example, if the `inherit` attribute is set to `true`, device interfaces are group members if the device has been added to the group.

**Values:** `true`, `false`

#### – **type**

Indicates the type of group. The following values are supported.

##### **Values:**

- **user group** (default): A group that a user has created.
- **site**: A user-created group that represents a physical site.

For example, supply the following XML:

```
<GroupTree path="/All Groups">
  <Group name="USA" desc="Group to represent the entire
    United States" inherit="true" type="user group"/>
</GroupTree>
```

6. Run the method.

In this example, the USA group is added under the default **All Groups** group in the **Groups** tree.

### **Create a Group in a Custom Tenant**

Create and configure a group associated with a custom tenant by supplying the group ID or the group path that you retrieved in the ["Create a Group in the Default Tenant"](#) section as parameters, in the `/All Groups/Tenants/` path.

The following shows an example of how to add groups to a non-default tenant.

### Follow these steps:

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:

```
http://PC_host:8181/pc/center/webservice/groups/<useIds>/<allowDeletes>
```

#### Example:

```
http://PC_host:8181/pc/center/webservice/groups/false/false
```

#### – useIds

Indicates whether the `groups` web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported.

##### Values:

- `true`: The `groups` web service uses the IDs of the groups.
- `false`: The `groups` web service does not use the IDs of the groups, and instead, creates new groups.

#### – allowDeletes

Indicates whether the `groups` web service can delete the group that you are creating. Allows the `groups` web service to update the group rules that are defined in any existing groups that are overwritten by this XML.

##### Values:

- `true`: Honor the `allowDeletes` values at the group level.
- `false`: Ignore the `allowDeletes` values at the group level.

For more information about this parameter, see the ["Group and Group Rule Attributes"](#) section.

2. Select **POST** for **"HTTP" Method**.
3. Provide a valid Username and Password in the request header for a user account that has global administrator access to NetOps Portal.
4. Select **'application/xml'** as the **'Body Content-type'**.
5. Add the following XML within the "Body" text section, replacing the values with the values that you want to use for the new group:

```
<GroupTree path="/All Groups/Tenants/MyCustomTenant/Groups">
  <Group name="North Carolina" desc="NC in the Southeast Region" inherit="true" type="user group"/>
</GroupTree>
```

6. Run the method.

The North Carolina group is created under the MyCustomTenant custom tenant. Only the global Administrator and the Administrator for the custom tenant can manage this group, and only MyCustomTenant users can see this group.

A group is created in a custom tenant.

## Export a Group

You can view the details and rule definitions of a given group by issuing a GET request to the parent of that group.

### Follow these steps:

1. Set up a REST client with a connection to the NetOps Portal server.
2. Use the following format for the URL in the REST client:

```
http://PC_host:8181/pc/center/webservice/groups/<idName>/<idValue>
```

**Example:** Retrieve the details and rule definitions for the groups under the `ParentGroup` group that is under `All Groups` :

```
http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2fParentGroup
```

#### – idName

The value that is used to specify the group.



**TIP**

You can retrieve the possible values by performing a GET operation on the following URL:

`http://PC_host:8181/pc/center/webservice/groups/idNames`

**Parameters:**

- **groupId:** The internally-assigned ID of the group.
- **groupPath:** The path of the group, with each group delimited by encoded front slashes (%2F).

– **idValue**

The value that is used to specify a group, based on the nature of the ID name. Dependent on the value of the `idName` field, as follows:

- **groupId:** The ID of the group is expected.
- **groupPath:** The path of the group, with each group delimited by encoded front slashes (%2F) is expected.

3. Select **GET** for the **HTTP method**.
4. Provide a valid username and password in the request header for a user account that has global administrator access to NetOps Portal.
5. Run the method.

The XML that is returned contains the group item identifier.

In the following example, the XML that is returned includes information about the **Inventory** system group:

```
<?xml version="1.0" encoding="UTF-8"?>
<GroupTree id="5" inheritDefault="true" path="Inventory">
  <Group desc="This group contains groups for the various item types
that are associated with all data sources." id="12" inherit="false"
location="" name="All Items" type="automatic group">
    <Group desc="Includes networks from all Application Delivery
Analysis data sources." id="40" inherit="true" location=""
name="Application Delivery Analysis Networks" type="automatic
group" />
    <Group desc="This group contains all the applications reported by
each data source." id="41" inherit="true" location="" name=
"Applications" type="automatic group" />
    <Group desc="This group contains all the device components reported
by each data source." id="113" inherit="true" name="Device
Components" type="automatic group" />
    <Group desc="This group contains all the VMware ESX hosts reported
by each data source." id="35" inherit="true" location="" name="ESX
Hosts" type="automatic group" />
    <Group desc="Includes interfaces from all data sources." id="50"
inherit="true" location="" name="Interfaces" type="automatic group" />
    <Group desc="This group contains all the pingable devices reported
by each data source." id="114" inherit="true" name="Pingable
Devices" type="automatic group" />
    <Group desc="Includes routers from all data sources." id="31"
inherit="true" location="" name="Routers" type="automatic group" />
    <Group desc="Includes servers from all data sources." id="32"
inherit="true" location="" name="Servers" type="automatic group" />
    <Group desc="Includes switches from all data sources." id="33"
inherit="true" location="" name="Switches" type="automatic group" />
    <Group desc="This group contains all the virtual machines reported
by each data source." id="34" inherit="true" location=""
name="Virtual Machines" type="automatic group" />
    <Group desc="This group contains all the voice interfaces reported
```

```

    by each data source." id="51" inherit="true" location="" name=
    "Voice Interfaces" type="automatic group" />
<Group desc="This group contains all the voip locations reported
    by each data source." id="52" inherit="true" location="" name=
    "VoIP Locations" type="automatic group" />
</Group>
<Group desc="Includes every data source that has reported
    configuration information to the performance center." id="4"
    inherit="false" location="" name="Data Sources" type="automatic
    group">
    <Group desc="Contains configuration information reported to the
        performance center by Data Aggregator" id="115" inherit="true"
        name="da" type="system group" />
    <Group desc="Contains configuration information reported to the
        performance center by Event Manager" id="100" inherit="true"
        name="EventManager@servername.domain.com" type="system group">
        <Group desc="Devices created by event manager" id="101" inherit=
            "true" name="Devices" type="system group" />
    </Group>
</Group>

```

The above results show that the group ID for the Inventory group is 5:

```
<GroupTree id="5" inheritDefault="true" path="Inventory">
```

#### TIP

You can get the complete group path using this value. Enter the following URL:

```
http://PC_host:8181/pc/center/webservice/groups/groupItemId/5
```

From this result, you can see that the complete path to the **Inventory** group is `/All Groups/Inventory`. All sub-groups of the **Inventory** system group are also returned.

You can then use the values that are returned for the `group name` parameter to construct a group path for a sub-group, as in the following example:

```
/All%20Groups/Inventory/Device%20Components
```

The URL to get the group ID from the group path would resemble the following example:

```
http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2FInventory
%2FDevice%20Components
```

#### NOTE

When specifying a group path in the URL, use the percent encoded value (`%2F`) to delimit each group in the path. In the URL syntax, forward-slash (`/`) characters are reserved as the delimiting character for path segments.

### Identify a Group

Identify groups when you want to update them. Use one of the following methods:

- Use the full `group path`, which identifies the group as a sub-group of the default first-level or root group (All Groups). The group path shows the position of a group (as a node) within the **Groups** tree, which extends hierarchically from the root group.
- Use the `group ID`, an internally-assigned numeric value that identifies a node. If the groups are nested within multiple containers, use this method to avoid having to enter the full `group path`.

### **Look Up the ID of a Group**

Many NetOps Portal API calls require the numeric `groupId` for a group. You can look up the IDs for groups using any of the following three methods:

- [Look up the ID of a single group path.](#)
- [Look up the ID of several groups.](#)
- [Retrieve a list of the child groups of a specified group.](#)

### **Look up the ID of a Single Group Path**

Use the `groupId` lookup service to look up the ID of a single group path using a simple HTTP GET.

With this method, you look up the ID of All Groups/Inventory. The syntax is as follows:

```
http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2FInventory/
groupId
```

Use URL text encoding. When specifying a group path in the URL, replace spaces with `%20` and use the percent encoded value (`%2F`) to delimit each group in the path.

The entire group path is not required. For example, you can find the All Groups/Tenants/Acme/Groups/Sites/Boston group by searching for `groupPath/Acme%2FGroups%2FSites%2FBoston/groupId`.

The XML that is returned is as follows:

```
<groupId>
  <groupId>5</groupId>
</groupId>
```

#### **NOTE**

The search might match more than one group, but returns only one group.

### **Look up the ID of Several Groups**

Use the bulk group `find` service to look up the ID of several groups with a single HTTP POST. The syntax is as follows:

#### **TIP**

You can look up the IDs for groups using this method for any number of groups. For each group path that is found, the corresponding group ID and the group Name are returned. If a path does not match a group, an ID is not returned.

```
http://PC_host:8181/pc/center/webservice/groups/find
```

The XML to POST is as follows:

```
<groups>
  <group Path='/All Groups/Tenants/Acme/Groups/Sites/San Francisco'/>
  <group Path='/All Groups/Tenants/Acme/Groups/Sites/Austin'/>
  <group Path='/All Groups/Tenants/Acme/Groups/Sites/Boston'/>
</groups>
```

The XML that is returned is as follows:

```
<?xml version="1.0"
encoding="UTF-8"?>
<groups>
  <group ID="10503"Name="San Francisco" Path="/All Groups/Tenants/Acme/Groups/Sites/San Francisco"/>
  <group Path="/All Groups/Tenants/Acme/Groups/Sites/Austin"/>
  <group ID="10505"Name="Boston" Path="/All Groups/Tenants/Acme/Groups/Sites/Boston"/>
</groups>
```

#### NOTE

Since the Austin group does not exist, a group ID is not returned for this group.

### Retrieve a List of the Child Groups of a Specified Group

Use the `groups` web service to retrieve a list of all child groups of a specified group.

#### NOTE

This call takes several minutes to return a large group hierarchy. This method is the slowest as it returns more data.

For example, to look up the ID of the `All Groups/Inventory` group, the syntax is as follows:

```
http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2FInventory
```

#### NOTE

When specifying the group path, use URL text encoding. Replace spaces with `%20` and separate the groups by `%2F`.

### Retrieve a List of Groups or Sub-Groups

You can:

- [Retrieve a list of the groups under all groups.](#)
- [Retrieve a list of the sub-groups under a group.](#)

### Retrieve a List of the Groups Under All Groups

You can retrieve a list of the groups under `All Groups` using one of the following options:

- By identifying the default group. Issue the call with the `groupPath` parameter, for example:

```
http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups
```

#### NOTE

When using some REST clients, the `All Groups` syntax is required rather than `All%20Groups`. But in general, blank spaces are not valid in URLs.

- By getting the identifier (`siteId`) for a site group. Issue the call with the `groupId` parameter, for example:

```
http://PC_host:8181/pc/center/webservice/groups/groupItem/siteId
```

The following XML shows an example of the return for the site ID:

```
<GroupTree siteId="118" inheritDefault="true" path="Austin, TX">
```

### Retrieve a List of the Sub-Groups Under a Group

You can retrieve a list of the sub-groups under a group that you specify using one of the following options:

- Issue the call with the `groupPath` parameter, for example:

`http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2FInventory`

- Issue the call with the `groupId` parameter, for example:

**NOTE**

The `groupId` of the default **Inventory** group is 5.

`http://PC_host:8181/pc/center/webservice/groups/groupId/5`

The XML that is returned includes the syntax of any group rules that are applied to that group.

## Syntax for Site Group Management

You can create group rules with multiple comparisons, in addition to regular expressions. For example, use the following syntax in the XML to post a group rule that adds devices whose name begins with the word 'Cisco':

```
<Match>
  <Compare readOnly="true" using="MEMBER_OF">
    <Property name="ItemID" type="device"/>
    <Value reference="/All Groups">1</Value>
  </Compare>
  <Compare readOnly="false" using="STARTS_WITH">
    <Property name="AlternateName" type="device"/>
    <Value>Cisco*</Value>
  </Compare>
</Match>
```

For group path syntax, forward-slash (/) characters are appropriate for the XML that you post. This example assumes that you already have a group structure of `All Groups\Texas\Austin`:

```
<GroupTree inheritDefault="true" path="/All Groups/Texas/Austin">
  <Group desc="" inherit="true" location="" name="CA Officetype="custom group">
    <Group desc="" inherit="true" location="" name="Austin Lab" type="custom group"/>
  </Group>
  <Group desc="" inherit="true" location="" name="Austin Data Center" type="custom group"/>
</GroupTree>
```

In the URL for the `groups` web service request, use a backslash character for a group path. Do not use forward-slash (/) characters in the URL.

## groups Web Service Example Syntax

You can retrieve a list of all groups under the highest-level group in the **Groups** tree (the default, **All Groups**) using one of the following options:

- Identify the default group by issuing the call with the `groupPath` parameter, for example:

**NOTE**

When using some REST clients, the `All Groups` syntax is required rather than `All%20Groups`. But in general, blank spaces are not valid in URLs.

`http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups`

- Identify the default group (whose `groupId` value is 1) by issuing the call with the `groupId` parameter, for example:

`http://PC_host:8181/pc/center/webservice/groups/groupId/1`

This includes:

- [Sub-group Syntax](#)
- [Create a Site Group](#)
- [Group Rules](#)

### **Sub-group Syntax**

You can get a list of all sub-groups under a group that you specify using one of the following options:

- Issue the call with the `groupPath` parameter, for example:  
`http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2FInventory`
- Issue the call with the `groupId` parameter, for example:

#### **NOTE**

The `groupId` of the default **Inventory** group is 5.

`http://PC_host:8181/pc/center/webservice/groups/groupId/5`

The XML that is returned includes the syntax of the group rules that are applied to that group.

### **Create a Site Group**

#### **TIP**

You can use a procedure similar to the following to create sub-groups. Add the following XML within the **Body text** section, and then run the method:

```
<GroupTree path="/All Groups/USA">
  <Group name="Raleigh" desc="This is the group for
    managed items in Raleigh, NC" inherit="true"
    type="custom group"/>
</GroupTree>
```

In this example, the Raleigh group is added as a sub-group of the All Groups\USA group.

### **Follow these steps:**

1. As a local admin user, create a site group using the following URL:

`http://PC_host:8181/pc/center/webservice/groups/<useIds>/<allowDeletes>`

#### **Example:**

`http://PC_host:8181/pc/center/webservice/groups/true/false`

#### – **useIds**

Indicates whether the `groups` web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported. In this example, the XML does not contain a group ID, so the value is `false`.

#### **Values:**

- **true:** The `groups` web service uses the IDs of the groups.
- **false:** The `groups` web service does not use the IDs of the groups, and instead, creates new groups.

#### – **allowDeletes**

Indicates whether the `groups` web service can delete the group that you are creating. Allows the `groups` web service to update the group rules that are defined in any existing groups that are overwritten by this XML.

#### **Values:**

- **true:** Honor the `allowDeletes` values at the group level.
- **false:** Ignore the `allowDeletes` values at the group level.

For more information about this parameter, see the "[Group and Group Rule Attributes](#)" section.

2. Select **POST** for **HTTP Method**.
3. Provide a valid **Username** and **Password** in the request header for a user account that has administrator access to NetOps Portal.

4. Select **application/xml** as the '**Body Content-type**'.
5. Add the following XML within the **Body text** section, replacing the values with the values that you want to use for the new site group:

```
<GroupTree path="/All Groups">
  <Group name="East Coast USA" desc="This is a site group"
    inherit="true" type="site group" location="North America"
    bHourID="99990" timeZone="EST"/>
</GroupTree>
```

- **type**  
Indicates the type of group. The following values are supported.  
**Values:**
  - **user group** (default): A group that a user has created.
  - **site**: A user-created group that represents a physical site.
- **bHourID**  
(Optional) The internally assigned identifier of the business-hour definition that you want to associate with this site group.
- **timeZone**  
(Optional) The time zone to associate with this site group. Time zones can only be associated with site groups, not with user groups.
- **inherit**  
Indicates whether the group includes child items of group members. For example, if you set this attribute to `true`, device interfaces are group members if the device has been added to the group.

For example, supply the following XML:

```
<Group name="East Coast" desc="Site group to represent the
  entire United States" inherit="true" type="site group" location="North
  America" bHourID="99990" timeZone="EST"/>
```

6. Run the method.

The USA site group is created under the default **All Groups** group in the **Groups** tree.

## Group Rules

Group rules support multiple comparisons, in addition to regular expressions. For example, you want a group rule that adds devices in the Routers group whose name begins with 'Cisco'. The group ID of the Router group is 31. Use the following syntax in the XML:

```
<Match>
  <Compare readOnly="true" using="MEMBER_OF">
    <Property name="ItemID" type="device"/>
    <Value reference="/All Groups/Inventory/All Items/Routers">31</Value>
  </Compare>
  <Compare readOnly="false" using="STARTS_WITH">
    <Property name="AlternateName" type="device"/>
    <Value>Cisco*</Value>
  </Compare>
</Match>
```

### NOTE

Scope group rules as narrowly as possible. Do not scope group rules to All Groups.

## Retrieve a List of Group Members

Retrieve a list of sub-groups and items that are direct members of a specified group using the `groups` web service. The XML that is returned displays information about the specified group as the parent element, and separates groups and items into their own separate elements.

The list that is returned does not include items that were added to the group as children of a managed item that was directly added to the group. For example, if a router is a direct member of a group, the list does not include the interfaces that belong to that router.

### Follow these steps:

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:

```
http://PC_host:8181/pc/center/webservice/groups/<idName>/<idValue>/items
```

– **idName**

The ID for the group.

**TIP**

You can retrieve the possible values by performing a GET operation on the following URL:

```
http://PC_host:8181/pc/center/webservice/groups/idNames
```

**Parameters:**

- **groupId**: The internally-assigned ID of the group.
- **groupPath**: The path of the group, with each group delimited by encoded front slashes (%2F ).

– **idValue**

The value that is used to specify a group, based on the nature of the ID name. Dependent on the value of the **idName** field, as follows:

- **groupId**: The ID of the group is expected.
- **groupPath**: The path of the group, with each group delimited by encoded front slashes (%2F ) is expected.

2. Select **GET** for the **HTTP** method.
3. Provide a valid Username and Password in the request header for a user account that has global administrator access to NetOps Portal.
4. Run the method.

The XML that is returned lists sub-groups and managed items that are members of the group. It resembles the following:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<group id="787" name="Test Group" description="" type="group" subType="user">
  <groups>
    <group id="788" name="Test Child Group 1" description="" type="group" subType="user"/>
    <group id="789" name="Test Child Group 2" description="" type="group" subType="user"/>
  </groups>
  <itemTypes>
    <itemType type="Devices">
      <items>
        <item id="121" name="Austin Switch" description="Cisco IOS Software, Switch 192.168.1.1" type="device"
subType="switches" addedBy="BY_USER"/>
        <item id="124" name="Austin Router" description="Cisco IOS Software, Router 192.168.0.1" type="device"
subType="router" addedBy="BY_USER"/>
        <item id="127" name="Austin Server" description="Linux" type="device" subType="server"
addedBy="BY_USER"/>
      </items>
    </itemType>
    <itemType type="Interfaces">
      <items>
```



```

    <item id="417" name="eth0/1/0:7" description="Ethernet0/1/0:7" type="interface" subType="physical"
addedBy="BY_USER"/>
    <item id="418" name="eth0/1/0:8" description="Ethernet0/1/0:8" type="interface" subType="physical"
addedBy="BY_USER"/>
    <item id="420" name="eth0/1/0:9" description="Ethernet0/1/0:9" type="interface" subType="physical"
addedBy="BY_USER"/>
  </items>
</itemType>
</itemTypes>
</group>

```

## Add an Item to a Group

Add individual managed items to groups using the `groups` web service. The item ID for each item is required.

Start by getting a list of managed items in the database and their IDs using the `devices` web service. The `get idNames` method returns a list of identifiers that you can use in other methods to identify devices. A submethod, `get interfaces`, returns a list of device interfaces. You can filter the results by item subtype.

For more information, see [Devices Web Service](#).

The web service checks to avoid item duplication and to make sure the group is valid. If any of these steps fail, then the service exits out with an error:

- The group exists.
- The user who authenticated with the web service has access permissions to the group.
- The group can have items added to it. You can add items of the following subtypes to groups: user, site, service provider-defined items.

Each item that is specified in the list is also validated according to the following criteria:

- The list of items to add does not contain duplicate IDs. Nor does the list contain the IDs of any items that are already members of the target group.
- An item that corresponds to the specified ID exists.
- The item is NOT a group. If a user wants to add a group to an existing group, there are other services for that.
- The user has access to (permission to view) that item.

The XML that is returned shows the results of the validation. They include a report of the items that were added and the items that were not added.

## Follow these steps:

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:

```
http://PC_host:8181/pc/center/webservice/devices/<idNames>
```

2. In the XML that is returned, locate the device IDs for the devices to add to the target group.

3. Enter the following URL:

```
http://PC_host:8181/pc/center/webservice/groups/<idName>/<idValue>/items
```

### Example:

```
http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups%2FmyGroup/
items
```

#### — idName

The ID for the group.

#### TIP

You can retrieve the possible values by performing a GET operation on the following URL:

```
http://PC_host:8181/pc/center/webservice/groups/idNames
```

Use one of the following parameters for the **idName**:

- **groupId**: The internally assigned ID of the group.
- **groupPath**: The path of the group, with each group delimited by encoded front slashes (%2F).

– **idValue**

The value that is used to specify a group, based on the nature of the ID name. Dependent on the value of the **idName** field, as follows:

- **groupId**: The ID of the group is expected.
- **groupPath**: The path of the group, with each group delimited by encoded front slashes (%2F) is expected.

4. Select **POST** for the **HTTP method**.
5. Provide a valid username and password in the request header for a user account that has global administrator access to NetOps Portal.
6. Select **'application/xml'** as the **'Body Content-type'**.
7. Add the following XML within the "Body" text section:

```
<items>
  <item id="value"/>
</items>
```

For example, the following XML adds a single router with managed item ID 684 to the Austin group:

```
<items>
  <item id="684"/>
</items>
```

**TIP**

You can also submit a list of managed items to add to a single group.

8. For example, add XML like the following to add five items:

```
<items>
  <item id="123"/>
  <item id="234"/>
  <item id="345"/>
  <item id="456"/>
  <item id="567"/>
</items>
```

The items that you specified in your list are added to the group.

9. In NetOps Portal, [verify group membership](#).

The item is added to the group.

## Add a Group Rule to a Group

Group rules determine the managed items/devices DX NetOps Performance Management adds to those groups/collections whose rules apply to those items/devices as it discovers them.

**TIP**

To get the syntax for type of rule, create a rule using NetOps Portal, and then [exporting the group](#). For more information about how to create rules using NetOps Portal, see [Manage Group Rules](#).

### Follow these steps:

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:  
`http://PC_host:8181/pc/center/webservice/groups/<useIds>/<allowDeletes>`

**Example:**

```
http://PC_host:8181/pc/center/webservice/groups/false/true
```

– **useIds**

Indicates whether the `groups` web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported. In this example, the XML does not contain a group ID, so the value is `false`.

**Values:**

- `true`: The `groups` web service uses the IDs of the groups.
- `false`: The `groups` web service does not use the IDs of the groups, and instead, creates new groups.

– **allowDeletes**

Indicates whether the `groups` web service can delete the group that you are creating. Allows the `groups` web service to update the group rules that are defined in any existing groups that are overwritten by this XML.

**Values:**

- `true`: Honor the `allowDeletes` values at the group level.
- `false`: Ignore the `allowDeletes` values at the group level.

For more information about this parameter, see the ["Group and Group Rule Attributes"](#) section.

2. Select **POST** for the **HTTP method**.
3. Provide a valid Username and Password in the request header for a user account that has global administrator access to NetOps Portal.
4. Select **'application/xml'** as the **'Body Content-type'**.
5. Modify the XML within the "Body" text section with the values that you want to use for the group rule, and then run the method:

For more information about the attributes that you can add to the XML, such as `allowDeletes`, see the ["Group and Group Rule Attributes"](#) section.

**Example:**

The following group rule adds the Cisco-3345 device to the Austin group:

```
<GroupTree path="/All Groups/USA">
  <Group name="Austin" desc="The group for items in Austin, TX">
    <Rules allowDeletes="true" saveRules="true">
      <Rule add="device" name="Add Devices">
        <Match>
          <Compare readOnly="true" using="MEMBER_OF">
            <Property name="ItemID" type="device"/>
            <Value reference="/All Groups">1</Value>
          </Compare>
          <Compare readOnly="false" using="EQUALS">
            <Property name="AlternateName" type="device"/>
            <Value>Cisco-3345</Value>
          </Compare>
        </Match>
      </Rule>
    </Rules>
  </Group>
</GroupTree>
```

The `groups` web service creates the group rule by adding XML to the group. If the group does not already exist, the web service also adds the group.

## **Group and Group Rule Attributes**

When creating a group by way of the `groups` web service, you can add the following attributes to the XML:

- `allowDeletes`.
- `saveRules`.

## The allowDeletes Attribute

The `allowDeletes` attribute defines whether the `groups` web service can delete groups or rules, and can have the following values:

- `true`
  - At the group level, the groups that exist in the **Groups** tree but are not specified in the XML are deleted.
  - At the rule level, when the rules run, items that exist in the group but no longer satisfy the group rules are deleted.
- `false`
  - At the group level, the groups that exist in the **Groups** tree but are not specified in the XML are not deleted.
  - At the rule level, when the rules run, items that exist in the group but no longer satisfy the group rules are not deleted.
- `inherit`
  - At the group level, the value for the `allowDeletes` attribute of an ancestor group in the **Groups** tree determines if the groups that exist in the **Groups** tree but are not specified in the XML are deleted. This is the default if you do not specify a value.
  - At the rule level, when the rules run, the value for the `allowDeletes` attribute of an ancestor group determines if the items that exist in the group but no longer satisfy the group rules are deleted. This is the default if you do not specify a value.

**Best practice:** Define a `GroupTree` level `allowDeletes` value to use as the default for any group that does not define an `allowDeletes` value.

## The saveRules Group Rule Attribute

The `saveRules` group rule attribute can have the following values:

- `true`  
If the rule is changed, the web service saves the changes to the database.
- `false`  
If the rule is changed, the web service does not save the changes to the database.

## About Modifying or Deleting Groups

You can modify or delete only the groups that you own. If you created a group, you are the owner of that group. The global administrator is the owner of the groups in the **Groups** tree. In addition, the global administrator or tenant administrator can edit your user account to give you ownership of a branch of the **Groups** tree.

### TIP

You can retrieve a list of the groups of which you are the owner using the `users` web service.

For more information about the `users` web service, see [Users Web Service](#).

## Remove an Item from a Group

You can remove individual items from a group using the `groups` web service. First, view the current group membership (from the **Manage Groups** page in NetOps Portal). Then apply a new rule to the group that specifies the `allowDeletes` attribute. When you post the rule, it deletes items that are already in the group but that do not meet the criteria that the rule specifies.

### Follow these steps:

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:  
`http://PC_host:8181//pc/center/webservice/groups/<useIds>/<allowDeletes>`

### Example:

`http://PC_host:8181//pc/center/webservice/groups/false/true`

– **useIds**

Indicates whether the `groups` web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported. In this example, the XML does not contain a group ID, so the value is `false`.

**Values:**

- `true`: The `groups` web service uses the IDs of the groups.
- `false`: The `groups` web service does not use the IDs of the groups, and instead, creates new groups.

– **allowDeletes**

Indicates whether the `groups` web service can delete the group that you are creating. Allows the `groups` web service to update the group rules that are defined in any existing groups that are overwritten by this XML.

**Values:**

- `true`: Honor the `allowDeletes` values at the group level.
- `false`: Ignore the `allowDeletes` values at the group level.

For more information about this parameter, see the ["Group and Group Rule Attributes"](#) section.

2. Select **POST** for the **HTTP method**.
3. Provide a valid Username and Password in the request header for a user account that has global administrator access to NetOps Portal.
4. Select **'application/xml'** as the **'Body Content-type'**.
5. Add the following XML within the "Body" text section, replacing the values with the values that you want to use for the group rule:

```
<GroupTree path="/All Groups/USA">
  <Group name="Austin" desc="The group for items in Austin, TX">
    <Rules allowDeletes="true" saveRules="true">
      <Rule add="device" name="Add Devices">
        <Match>
          <Compare readOnly="true" using="MEMBER_OF">
            <Property name="ItemID" type="device"/>
            <Value reference="/All Groups">1</Value>
          </Compare>
          <Compare readOnly="false" using="EQUALS">
            <Property name="subType" type="device"/>
            <Value>server</Value>
          </Compare>
          <Compare readOnly="false" using="NOT_EQUALS">
            <Property name="subType" type="device"/>
            <Value>VM</Value>
          </Compare>
        </Match>
      </Rule>
    </Rules>
  </Group>
</GroupTree>
```

This example creates a rule that adds devices with subtype 'server' but removes devices with subtype 'VM' (virtual machine) from the Austin group.

6. Run the method.
7. In NetOps Portal, [verify group membership](#).

The item is removed from the group.

## Delete a Group

The groups web service does not offer a DELETE method to delete a user group or a custom collection based on the group ID. Deletion of groups is possible using the update method. You perform another POST of the entire XML for a group with the group elements that need to be deleted removed from the XML. Their absence, and the `allowDeletes` parameter, removes them.

Set `allowDeletes` on a group that is one level higher in the tree to `true` to ensure that a sub-group is deleted if no entry for that group is included in the XML. This attribute is inherited by all sub-groups. It is not applied to the parent group itself.

You can only delete user groups. You cannot delete system groups and out-of-the-box collections. To delete groups that you defined in **My User Groups**, use NetOps Portal.

### WARNING

The user groups that you delete can only be restored by recreating them.

You can supply the group ID or the group path as parameters.

In the "Create Groups Using Web Services" procedure, the example XML created the Raleigh group that was a sub-group of the USA group. The full XML for this task would be as follows:

```
<GroupTree path="/All Groups">
  <Group name="USA" desc="Group to represent the entire United
States" inherit="true" type="user group">
    <Group name="Raleigh" desc="This is the group for managed items
in Raleigh, NC" inherit="true" type="user group"/>
  </Group>
</GroupTree>
```

Enable future deletion of these groups by POSTing the XML to the following URL:

```
http://PC_host:8181/pc/center/webservice/groups/<useIds>/<allowDeletes>
```

### Example:

```
http://PC_host:8181/pc/center/webservice/groups/false/true/
```

- **useIds**  
Indicates whether the groups web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported. In this example, the XML does not contain a group ID, so the value is `false`.  
**Values:**
  - `true`: The groups web service uses the IDs of the groups.
  - `false`: The groups web service does not use the IDs of the groups, and instead, creates new groups.
- **allowDeletes**  
Indicates whether the groups web service can delete the group that you are creating. Allows the groups web service to update the group rules that are defined in any existing groups that are overwritten by this XML.  
**Values:**
  - `true`: Honor the `allowDeletes` values at the group level.
  - `false`: Ignore the `allowDeletes` values at the group level.

For more information about this parameter, see the "Group and Group Rule Attributes" section.

And then delete the Raleigh sub-group.

### Follow these steps:

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:

```
http://PC_host:8181/pc/center/webservice/groups/<useIds>/<allowDeletes>
```

**Example:**

```
http://PC_host:8181/pc/center/webservice/groups/false/true/
```

– **useIds**

Indicates whether the `groups` web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported. In this example, the XML does not contain a group ID, so the value is `false`.

**Values:**

- `true`: The `groups` web service uses the IDs of the groups.
- `false`: The `groups` web service does not use the IDs of the groups, and instead, creates new groups.

– **allowDeletes**

Indicates whether the `groups` web service can delete the group that you are creating. Allows the `groups` web service to update the group rules that are defined in any existing groups that are overwritten by this XML.

**Values:**

- `true`: Honor the `allowDeletes` values at the group level.
- `false`: Ignore the `allowDeletes` values at the group level.

For more information about this parameter, see the ["Group and Group Rule Attributes"](#) section.

2. Select **POST** for **"HTTP" Method**.
3. Provide a valid Username and Password in the request header for a user account that has global administrator access to NetOps Portal.
4. Select **'application/xml'** as the **'Body Content-type'**.
5. Add the following XML within the "Body" text section. Replace the values with values to identify the groups that contain the groups to be deleted:

```
<GroupTree path="/All Groups">
  <Group name="USA" desc="Group to represent the entire United
    States" allowDeletes="true" type="user group"/>
</GroupTree>
```

6. Run the method.

The web service updates the XML for the specified `All Groups\USA` group to remove the Raleigh sub-group. The Raleigh group is thus deleted from the **Groups** tree.

**Migrate a Group to Another NetOps Portal Database**

You can import a list of group definitions using the `groups` web service. To import group definitions into a new database, you must first retrieve, and then export, all of your existing definitions.

You can supply the group ID or the group path that you retrieved in a previous procedure as parameters.

The basic URL to retrieve group definitions is as follows:

```
http://PC_host:8181/pc/center/webservice/groups/groupPath/All%20Groups
```

That method uses the group path name to identify the root group. You can also retrieve the "All Groups" structure by group item ID, as follows:

```
http://PC_host:8181/pc/center/webservice/groups/groupItemId/1
```

Prepare to export all of the groups that you have defined in NetOps Portal by performing a retrieval of the All Groups group. Retrieving this root group calls up the entire tree in the XML that is returned.

**Follow these steps:**

1. Using a REST client with a connection to the NetOps Portal server, use the following format for the URL:

```
http://PC_host:8181/pc/center/webservice/groups/<idName>/<idValue>
```

**Example:**

```
http://PC_host:8181/pc/center/webservice/groups/groupItemId/1
```

– **idName**

The ID for the group.

**TIP**

You can retrieve the possible values by performing a GET operation on the following URL:

```
http://PC_host:8181/pc/center/webservice/groups/idNames
```

Use one of the following parameters for the **idName**:

- **groupItemId**: The internally assigned ID of the group.
- **groupPath**: The path of the group, with each group delimited by encoded front slashes (%2F ).

– **idValue**

The value that is used to specify a group, based on the nature of the ID name. Dependent on the value of the **idName** field, as follows:

- **groupItemId**: The ID of the group is expected.
- **groupPath**: The path of the group, with each group delimited by encoded front slashes (%2F ) is expected.

2. Select **GET** for "HTTP" Method.
3. Provide a valid username and password in the request header for a user account that has global administrator access to NetOps Portal.
4. Select 'application/xml' as the 'Body Content-type'.
5. Run the method.  
The XML that is returned defines your entire group structure.
6. Import the group definitions into another NetOps Portal database by completing the followig steps:

- a. Using a connection to the NetOps Portal server from a REST client, use the following format for the URL:

```
http://PC_host:8181/pc/center/webservice/groups/<useIds>/<allowDeletes>
```

**Example:**

```
http://PC_host:8181/pc/center/webservice/groups/false/true
```

• **useIds**

Indicates whether the `groups` web service uses the `id` attribute of a group to identify it. Because they are internally assigned, the IDs are a less reliable way to identify groups that have been exported and reimported. In this example, the XML does not contain a group ID, so the value is `false`.

**Values:**

- `true`: The `groups` web service uses the IDs of the groups.
- `false`: The `groups` web service does not use the IDs of the groups, and instead, creates new groups.

• **allowDeletes**

Indicates whether the `groups` web service can delete the group that you are creating. Allows the `groups` web service to update the group rules that are defined in any existing groups that are overwritten by this XML.

**Values:**

- `true`: Honor the `allowDeletes` values at the group level.
- `false`: Ignore the `allowDeletes` values at the group level.

For more information about this parameter, see the "[Group and Group Rule Attributes](#)" section.

- b. Select **POST** for "HTTP" Method.
- c. Run the method.



The groups that you exported are added as sub-groups of the `All Groups` root group.

7. In NetOps Portal, [verify the structure of the imported groups](#).

The group definitions are imported.

### About Group Deletion

At the group level, to delete a group and retain the groups under the `GroupTree` path (unless excluded from the XML body), set the `allowDeletes` parameter to `true`. The groups web service deletes the groups under the `GroupTree` path that are excluded from the XML body. When the `allowDeletes` parameter is set to `false` for a group, the groups web service retains the groups under the `GroupTree` path regardless of whether they are included in the XML body.

Apply and set this attribute to `true` for a container group when you want to delete a sub-group.

**Example 1:** The following example XML deletes any unlisted groups under Austin:

```
<GroupTree inheritDefault="true" path="/All Groups/Texas/Austin" allowDeletes="true">
  <Group desc="" inherit="true" location="" name="CA Office" type="user group">
  <Group desc="" inherit="true" location="" name="Austin Lab" type="user group"/>
</Group>
  <Group desc="" inherit="true" location="" name="Austin Data Center" type="user
group"/>
</GroupTree>
```

**Example 2:** The following example XML deletes unlisted groups under Texas, including the Austin sub-group from the previous example:

```
<GroupTree path="/All Groups/Texas">
  <Group name="USA" desc="Group to represent the entire United States"
allowDeletes="true" type="user group"/>
</GroupTree>
```

For group path syntax, forward-slash (/) characters are appropriate for the XML documents that you post.

In the URL for the web service request, use a backslash (\) character for a group path. Do not use forward-slash (/) characters in the URL.

## Roles Web Service

Manage (view, create, and modify) roles using the `roles` web service.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services. `role` is a parameter assigned to user accounts that controls user access to DX NetOps Performance Management features and dashboards. Based on user job functions, the role grants administrative access to DX NetOps Performance Management configuration using `role rights`. With assigned roles, users can access data and NetOps Portal features that they require to perform their duties and restrict access to features that they do not require. To assign custom or out-of-the-box roles to user accounts, use [the users RESTful web service](#).

For more information about how to manage roles using NetOps Portal, see [Manage Roles and User Accounts](#).

### Verify the Prerequisite

- You have connected the REST client that you will use to invoke the NetOps Portal web services to the NetOps Portal RESTful web services.

## **Access the roles Web Service Documentation**

Issue these requests as a user with the Administrator role.

Retrieve a list of parameters for the `roles` web service by issuing a GET request to the following URL:

```
http://<PC_host>:8181/pc/center/webservice/roles
```

Retrieve a list of operations for the `roles` web service by issuing a GET request to the following URL:

```
http://<PC_host>:8181/pc/center/rest/roles/documentation
```

## **Basic roles Parameters**

The current values for the following user account role settings are available from GET requests:

- **accessRights**  
The role rights that are allocated to this role. To retrieve a list of the available rights for this role, the **categoryId** parameter is required. The `categoryId` parameter with value 1 corresponds to both NetOps Portal and the data aggregator.
- **culture**  
Specifies a language (locale). Supply a language identifier from the following list:
  - en-US (English, United States)
  - ja-JP (Japanese)
  - fr-FR (French, France)
- **description**  
(Optional) Describes the role to help you identify it.
- **enabled**  
Determines whether the role is enabled for use (activated). Values are true or false.
- **name**  
The name for the role. The name is limited to 50 characters.
- **selections**  
Provides sets of access rights that you can selectively grant to this role, organized into categories.
- **userCount**  
The number of users who have this role assigned to their user account.
- **userID**  
The internally-assigned value for the role.

## **Roles Web Service Example Syntax**

Issue the following call to see the available operations and parameters for the roles web service:

```
http://<PC_host>:8181/pc/center/rest/roles/documentation
```

## **Available GET Methods**

- **get access rights**  
Retrieves a list of the rights that are assigned to a specified role. Use the following syntax:  

```
http://<PC_host>:8181/pc/center/webservice/roles/idName/idValue/rights/cultureId
```
- **get access rights by category**  
Retrieves a list of all available role rights for a specified category. The category is either NetOps Portal or a data source. A `categoryId` of 1 applies to both NetOps Portal and the data aggregator.

`http://<PC_host>:8181/pc/center/webservice/roles/rights/categoryid/cultureId`

- **get category**

Retrieves an XML document that includes the `categoryId`:

`http://<PC_host>:8181/pc/center/webservice/roles/rights/categories/en-US`

- **get by tenant**

Retrieves a list of all roles that are associated with the tenant for the logged-in user. Use the following syntax:

`http://<PC_host>:8181/pc/center/webservice/roles/idName/idValue/rights/tenant/tenantIdName/tenantIdValue/cultureId`

- **get categoryId**

Retrieves XML that shows the available category IDs:

`http://<PC_host>:8181/pc/center/webservice/roles/rights/categories/cultureId`

A value of 1 corresponds to NetOps Portal and Data Aggregator. Role rights in that category only apply to a Data Aggregator data source.

- **get id names**

Retrieves a list of identifiers that can be used to identify roles in other web service methods. Use the following syntax:

`http://<PC_host>:8181/pc/center/webservice/roles/idNames`

### **Available PUT Methods**

- **copy**

Copies a role. This method creates a copy of a specified role and associates it with the tenant for the logged-in user. Use the following syntax:

`http://<PC_host>:8181/pc/center/webservice/roles/idName/idValue/copy/roleName/description/enabled/cultureId`

- **update**

Modifies a specified role. The role name and tenant ID are required elements of the role parameters:

`http://<PC_host>:8181/pc/center/webservice/roles/idName/idValue/cultureId`

### **Available POST Methods**

- **create**

Creates a role. The new role is associated with the tenant for the logged-in user:

`http://<PC_host>:8181/pc/center/webservice/roles/`

- **create for tenant**

Creates a role and assigns it to the specified tenant:

`http://<PC_host>:8181/pc/center/webservice/roles/tenant/tenantIdName/tenantIdValue`

### **Create a User Account Role**

#### **NOTE**

**Best Practice:** Create user account roles before creating user accounts.

Use the roles web service to create user account roles. The `role` is a parameter of the user account that grants the user access to certain administrative features and to perform selected tasks. The following procedure invokes the `roles` web service using a REST client.

Issue the following call to see the parameters for the roles web service:

```
http://<PC_host>:8181/pc/center/webservice/roles
```

### Follow these steps:

1. In the REST client interface, enter the following URL in the **URL** field, and then select **POST** for the method:

```
http://<PC_host>:8181/pc/center/webservice/roles/
```

2. Provide a valid Username and Password for a user account that has Administrator access to NetOps Portal.
3. Select '**application/xml**' as the '**Body Content-type**' in the Body settings.
4. Paste XML in the Body field that resembles the following example:

```
<role>
  <name>TestRoleName</name>
  <description>Test Role Description</description>
  <enabled>true</enabled>
  <accessRights>
    <accessRight>
      <accessRightName>ViewToS</accessRightName>
      <categoryId>1</categoryId>
      <enabled>true</enabled>
    </accessRight>
  </accessRights>
</role>
```

- **accessRights**

Correspond to role rights.

- **accessRightName**

The name of the role right. For example, the `administerGroups` role right allows the user with this role to manage a limited section of the Groups tree.

- **categoryId**

Identifies the category of role rights, such as NetOps Portal role rights or data-source-specific role rights.

5. Run the method.

The user account role is created.

## Users Web Service

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the `users` endpoint for the NetOps Portal REST web service. The Administrator and User Administrator roles have these rights.

You can manage user accounts in the following ways:

- Retrieve a current list of user accounts.
- Create user accounts.
- Modify user accounts.
- Change user account passwords.
- Delete user accounts.

You can manage user accounts using this endpoint and using NetOps Portal. This article is about managing user accounts using the endpoint.

For more information about how to manage user accounts using NetOps Portal, see [Manage User Accounts](#).

The following topics provide information about related parameters.

### **Basic User Account Parameters**

The following is a list of the parameters for the `users` endpoint:

#### **NOTE**

You can get the parameters and available operations for this endpoint by issuing a GET request to the following URL:

```
http://<PC_host>:8181/pc/center/rest/users/documentation
```

- **userId**  
The internally-assigned value for the user account.
- **name**  
The login name for the user account. The name is limited to 50 characters.
- **description**  
(Optional) The description of the user account to help you identify it.
- **enabled**  
Defines whether the user account is enabled for use (activated).
- **removable**  
Defines whether you can remove, or delete, the user account from the database.  
**Values:** true or false
- **timezone (tz)**  
Corresponds to the time zone in which the user views report data.  
**Default:** UTC (Coordinated Universal Time)
- **userLevel**  
Identifies the product privilege assigned to the user account.
- **role**  
Specifies the role assigned to the user account.
- **permissionId**  
Specifies the ID for the **My Assigned Groups** for the user account.
- **tenantId**  
The internal (database) identifier for the tenant with which the user account is associated.
- **culture**  
Specifies the language (locale) code for the user.  
**Options:**
  - en-US (English, United States)
  - fr-FR (French, France)
  - ja-JP (Japanese)

For more information about the languages that DX NetOps Performance Management supports, see [Language Support](#).

#### **TIP**

You can view the current values for the `culture` parameter for users that are associated with the same tenant as the account that is used to run this command by issuing the following GET call:

```
http://<PC_host>:8181/pc/center/webservice/users/<cultureId>
```

- **cultureId**  
The language code for the language in which you would like to view the output, such as `en-US`.

## Users Web Service Example Syntax

The following HTTP methods are available with the `users` endpoint:

### Available GET Methods

- **get groups owned by user**

Retrieves a list of groups for the specified user. The groups that are returned are groups that are owned by the specified user, meaning that this user can modify or delete these groups:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/groupsOwnedByUser
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

- **get groups**

Retrieves a list of groups to which the specified user has view access:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/groups
```

The groups that are returned are groups that are within the permission set of the specified user. The user cannot modify or delete these groups.

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

- **get administered groups**

Retrieves a list of all administered groups associated with a specified user account:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/administeredGroups
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

- **get id names**

Retrieves a list of identifiers that can be used to identify users in other web service methods:

```
http://<PC_host>:8181/pc/center/webservice/users/userId/<idValue>
```

- **idValue**

A value for the property denoted by `userId`.

- **idValue**

The value for the identifying category. For example, if `idName` is `userId`, provide the user ID. If `idName` is `userName`, provide the login name for the user account.

- **get authentication types**

Returns a list of identifiers that can be used to assign authentication types to users:

```
http://<PC_host>:8181/pc/center/webservice/users/authenticationTypes
```

- **get user name**

Retrieves the user account by name:

```
http://<PC_host>:8181/pc/center/webservice/users/userName/<userName>
```

- **userName**

The name of the user account to retrieve.

### Available PUT Methods

- **update password**

Changes the password of a specified user account:

**NOTE**

Passwords are sent in cleartext to avoid publicizing the encryption key that the web service uses. As a result, to protect the password privacy, change the password using the `users` endpoint for the NetOps Portal REST web service *only* on the NetOps Portal host.

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/password/<newPassword>
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

- **newPassword**

The new URL-encoded password for the user. For example, if the password contains a question mark (?), URL-encode that character as "%3f".

- **update role**

Updates the role assignment of a specified user account:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/role/<roleIdName>/<roleIdValue>
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

- **roleIdName**

The name for the role ID.

- **roleIdValue**

The value for the property denoted by `roleIdName`.

- **update time zone**

Updates the time zone of a specified user account:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/timeZone/<newTimeZone>
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

- **newTimeZone**

The new time zone for the user.

- **set groups**

Updates the permission groups that have been granted to a specified user account:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/groups
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

**XML Format:**

```
<group>
  <group ID="5245"/>
  <group ID="5246"/>
  <group ID="5247"/>
  ...
</group>
```

- **add groups**

Adds the specified groups:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/addGroups
```

- **idName**  
A property name value that the `get id names` method of this web service returns.
- **idValue**  
A value for the property denoted by `idName`.

**XML Format:**

```
<group>
  <group ID="5245"/>
  <group ID="5246"/>
  <group ID="5247"/>
  ...
</group>
```

- **remove groups**

Removes the specified groups:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/removeGroups
```

- **idName**  
A property name value that the `get id names` method of this web service returns.
- **idValue**  
A value for the property denoted by `idName`.

**XML Format:**

```
<group>
  <group ID="5245"/>
  <group ID="5246"/>
  <group ID="5247"/>
  ...
</group>
```

- **set administered groups**

Updates the branches of the groups tree to which the specified user has administrative access:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/administeredGroups
```

- **idName**  
A property name value that the `get id names` method of this web service returns.
- **idValue**  
A value for the property denoted by `idName`.

**XML Format:**

```
<group>
  <group ID="5245"/>
  <group ID="5246"/>
  <group ID="5247"/>
  ...
</group>
```

- **set default group**

Updates the default group for the specified user:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/defaultGroup/<groupId>
```

**Example:**

```
http://PC.mycompany.net:8181/pc/center/webservice/users/userId/7/defaultGroup/11
```

- **idName**  
A property name value that the `get id names` method of this web service returns.
- **idValue**  
A value for the property denoted by `idName`.
- **groupId**



The ID for the default group that you want to set for the user.

- **add administered groups**

Adds the specified groups to a user account:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/addAdministeredGroups
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

**XML Format:**

```
<group>
  <group ID="5245"/>
  <group ID="5246"/>
  <group ID="5247"/>
  ...
</group>
```

- **remove administered groups**

Removes the specified groups from a user account:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/removeAdministeredGroups
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

**XML Format:**

```
<group>
  <group ID="5245"/>
  <group ID="5246"/>
  <group ID="5247"/>
  ...
</group>
```

- **update product privilege per datasource**

Updates the product privilege of a specified user account to enable access to the user interface of a specific data source:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/ds/dsId/
productPrivilege/newProductPrivilege
```

- **idName**

A property name value that the `get id names` method of this web service returns.

- **idValue**

A value for the property denoted by `idName`.

- **dsId**

The ID for the data source that you want to grant access to for the user.

## **Available POST Methods**

- **create**

Creates and configures a user account. This user is associated with the tenant for the logged-in user. If you are deploying multi-tenancy, this user is assigned to the tenant of the authenticated user account making the REST service call. If you are not deploying multi-tenancy, this association is transparent to you; new user accounts are associated with the default tenant. The parameters include a role assignment:

```
http://<PC_host>:8181/pc/center/webservice/users/role/<roleIdName>/<roleIdValue>
```

- **roleIdName**

**TIP**

If you do not know the `roleIdName`, issue a GET call to the `roles` endpoint, for example:

```
http://PC_host:8181/pc/center/webservice/roles/idNames
```

**Examples:**

- `roleName`
- `roleId`

**– roleIdValue**

This value depends on the `roleIdName` that you selected. For example, if `roleName` is used, substitute a valid role name for `roleIdValue`.

**NOTE**

This role must be available within the tenant.

Provide a valid **Username** and **Password** for a user account that has host or tenant administrator access to NetOps Portal.

**NOTE**

The password is automatically set to be the same as the user name.

**XML Format:**

```
<user>
  <name>{UserName}</name>
  <description>{UserDescription}</description>
  <enabled>{UserEnabled}</enabled>
  <removable>{UserRemovable}</removable>
  <timezone>{UserTimeZone}</timezone>
  <culture>{UserCulture}</culture>
  <administeredGroups>
    <group ID="{groupID}"/>
    <group ID="{groupID}"/>
  </administeredGroups>
</user>
```

Replace values with the values that you want to use for this user account.

For example, supply the following parameters:

```
<user>
  <name>Jane Doe</name>
  <description>User associated with the John Doe Corporation tenant.</description>
  <enabled>true</enabled>
  <removable>true</removable>
  <timezone>CST6CDT</timezone>
  <culture>en-US</culture>
  <administeredGroups>
    <group ID="105"/>
    <group ID="367"/>
  </administeredGroups>
</user>
```

**NOTE**

The `administeredGroups` tag is optional. To create a user without administered groups, exclude the tag.

- **create in tenant**

Creates a user account in the specified tenant, with the specified role assignment:

```
http://<PC_host>:8181/pc/center/webservice/users/tenant/<tenantIdName>/<tenantIdValue>/
role/<roleIdName>/<roleIdValue>
```

**– tenantIdName**

The tenant ID name.

- **tenantIdValue**

The tenant ID value for the tenant in which to create the new user account.

- **roleIdName**

The role ID name for the role to assign to the new user account.

- **clear administered groups**

Clears administered groups with a special ID that indicates no groups are selected:

```
http://<PC_host>:8181/pc/center/webservice/users/<idName>/<idValue>/administeredGroups
```

**XML Format:**

```
<group>
  <group ID="2"/>
</group>
```

## User Account Product Privilege Settings

Determine the product privilege of a specified user account using NetOps Portal. The *product privilege* is a type of permission set associated with a user account. The product privilege grants user access to features in selected data sources and does not apply to NetOps Portal functionality.

For more information:

- About how to grant product privileges to registered data sources to user accounts using NetOps Portal, see [Manage User Accounts](#).
- About the product privileges that are available for registered data sources, see [Product Privilege](#).

## Available DELETE Methods

- **delete by user name**

Deletes a user account by name:

```
http://<PC_host>:8181/pc/center/webservice/users/userName/<userName>
```

- **userName**

The name of the user account to delete.

**NOTE**

- You can delete all user accounts except the **admin** and **user** predefined user accounts.
- Deleting a user account by way of the `users` endpoint deletes their user-owned items (notifications, on-demand reports, and scheduled reports), with no choice to move, or assign, these items to another user. To assign the user-owned items to another owner, delete the user account using NetOps Portal.

## Tenants Web Service

Manage (view, create, and modify) tenant definitions using the `tenants` web service.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the `tenants` endpoint for the NetOps Portal web services. The basic tenant definition contains a few parameters to identify the tenant. The infrastructure—devices, networks, servers—and the monitoring parameters for a customer's monitored systems must be associated with the tenant definition. Each tenant must contain at least one IP domain, plus as many of the following definitions as required to manage the associated enterprise infrastructure and applications:

- User accounts
- Roles
- Custom and system groups
- Custom reports
- Custom menus

To associate definitions and monitoring parameters with an existing tenant definition, log in as the tenant administrator and create the required definitions using the required web services. The definitions are then associated with the tenant definition and available to users logged in with this tenant's user accounts.

For more information about how to create a tenant using the `tenants` web service, see [Automate Provisioning and Configuration using the tenants Web Service](#).

### **Tenants Web Service Parameters**

You can list the available parameters and operations for the `tenants` web service using the following REST URL:

```
http://<PC_host>:8181/pc/center/rest/tenants/documentation
```

#### **Parameters**

- **tenantDescription**  
(Optional) Describes the tenant.
- **idName**  
Is a name for the tenant.
- **status**  
Is the status of this tenant.  
**Values:**
  - **Activated:** Enables tenant user accounts for use.
  - **Disabled:** Prevents any actions by user accounts that are associated with this tenant.
- **removable**  
States whether the item can be deleted (removed from the database).  
**Values:** `true` or `false`
- **theme**  
Specifies the format—the theme that controls the appearance of the page in the browser window—to use for this tenant. All operators whose user account is associated with this tenant see this same theme. The CA-Blue and CA-White themes are available.  
**Default:** `CA-Blue`
- **defaultCulture**  
Specifies a language (locale). Supply a language identifier from the following list:
  - `en-US` (English, United States)
  - `ja-JP` (Japanese)
  - `fr-FR` (French, France)
- **accountId**  
Identifies this tenant; usually corresponds to the MSP account number. If a value is supplied as input, it must be unique across all defined tenants.
- **tenantID**  
The internal (database) identifier for a tenant definition.

### **Tenants Web Service Example Syntax**

You can retrieve a current list of tenant definitions, create tenant definitions, and modify these definitions by changing their parameters using the `tenants` web service.

## HTTP Methods

You can use the following HTTP methods with the `tenants` web service:

- **GET**

Returns a list of tenant definitions sorted by name. Available on the `tenantID` endpoint. Use the following syntax:

```
http://<PC_host>:8181/pc/center/webservice/tenants/
```

- **POST**

Creates a custom tenant. Use the following syntax:

```
http://<PC_host>:8181/pc/center/webservice/tenants/
```

- **PUT**

Updates an existing tenant definition. Use the following syntax:

```
http://<PC_host>:8181/pc/center/webservice/tenants/
```

## Tenants Web Service Examples

The following are `tenants` web service examples:

### Example: Update the Description of a Tenant

The following is an example of a PUT that updates the `tenantDescription` parameter of a tenant:

```
http://<PC_host>:8181/pc/center/webservice/tenants/  
<idName>/<idName>/<tenantDescription>/<NewDescription>
```

For more information about how to connect the REST client that you will use to invoke the NetOps Portal web services to the NetOps Portal RESTful web services, see [Use NetOps Portal Web Services](#).

Substitute the desired values for the italicized terms. The following parameters are required:

- **idName**  
The name of the tenant that you want to edit.
- **NewDescription**  
The new description to identify this tenant.

### Example: Return a List of Tenant IDs and Names

The following is an example of a GET that returns a list of tenant IDs and names using the `tenants` web service:

```
http://<PC_host>:8181/pc/center/webservice/tenants/idNames
```

The following XML is returned:

```
<?xml version="1.0" encoding="UTF-8"?>  
<idNames>  
  <idName value="tenantAccountId" />  
  <idName value="tenantItemId" />  
  <idName value="tenantName" />  
</idNames>
```

The response text of the HTTP request is either the expected result or an error message to indicate a problem. The following HTTP response code ranges are used:

- 200 - Command status is 'OK'. Indicates a successful response.
- 400 - A user error has occurred. Errors in this range indicate a problem with the input text (400) or the user credentials (403) and can usually be easily corrected.
- 500 - A system error occurred. Errors in this range typically indicate a system fault. These errors can require assistance from Broadcom Technical Support to resolve them.

For more information about HTTP status codes, see the [IETF website](#).

## Console Info Web Service

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services.

In this article:

### Consoleinfo Web Service

Retrieve information about NetOps Portal console configuration using the consoleinfo web service that the API provides. You can pass this information into other web service methods, such as time zone information that is required to create business hour definitions.

### Consoleinfo Web Service Example Syntax

Issue the following call to see the parameters for the `consoleinfo` web service:

```
http://PC_host:8181/pc/center/rest/consoleinfo/documentation
```

### Available GET Methods

- **get all time zones**  
Gets a list of all of the time zones that are available for use in business hour definitions:  
`http://PC_host:8181/pc/center/web/service/consoleinfo/allTimezones/cultureID`
- **cultureId**  
Specifies a language (locale). Supply a language identifier from the following list:
  - en-US (English, United States)
  - ja-JP (Japanese)
  - zh-CN (Simplified Chinese)
  - fr-FR (French, France)
- **get installed language packs**  
Gets a list of language packs that are installed on the server:  
`http://PC_host:8181/pc/center/web/service/consoleinfo/languagepacks/cultureID`
- **get time zones assigned to sites**  
Retrieves a list of all of the time zones that are assigned to site groups to which the logged-in user has access:  
`http://PC_host:8181/pc/center/web/service/consoleinfo/timezonesAssignedToSites`

### Event Web Service

You can retrieve a list of events that have been raised in your environment using the event web service.

Perform a GET to the following URL to see a list of events for the specified managed item:

```
http://PC_host:8181/pc/eventId/item/itemId
```

- **eventId**  
Event ID of the event to retrieve properties for.
- **itemId**  
Item ID of the item that the event is on.

For more information, see [Devices Web Service](#).

## Automate Provisioning and Configuration using the tenants Web Service

Automate provisioning and configuration tasks using the `tenants` web service.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services. The most frequently repeated or time-consuming tasks are exposed to you with web services. Some of these APIs consist of web services that conform to the Representational State Transfer (REST) model.

This use case illustrates a procedure that an administrator can deploy to create multiple tenant definitions using the `tenants` web service. Because each tenant has its own user accounts to provide access to NetOps Portal, we also describe user account creation within tenants. In this use case, we describe the steps to take when using a REST client, a generic web services user interface application. The examples in this use case contain URIs that are constructed using the default server port, 8181.

### Create Tenants Programmatically

Create tenants using the `tenants` web service. The basic tenant definition contains a few parameters to identify the MSP customer and let other operators access managed items and configuration for the customer. An administrator account is a required component of the tenant definition so that the customer can perform some tenant setup.

You can associate monitored devices and product settings with the tenant definition in separate steps. Each tenant must contain at least one IP domain. You and the tenant administrator can then configure other product settings that are required to manage the enterprise infrastructure and applications for that customer.

### Create Tenants Using Web Services

Use any REST client to create and configure a tenant using the `tenants` web service.

#### Follow these steps:

1. Set up a REST client with a connection to the NetOps Portal server.
2. Enter a URL for the web services API in the REST client. Use the following format:  

```
http://<PC_host>:8181/pc/center/webservice/tenants/
```
3. Select **POST** for **"HTTP" Method**.
4. Provide a valid Username and Password for the user account who will run REST calls and who has global administrator access to NetOps Portal.
5. Select **'application/xml'** as the **'Body Content-type'** in the Body settings.
6. Add the following XML within the "Body" text section:

```
<tenant>
  <tenantName>Name of tenant</tenantName>
  <tenantDesc>Description of the tenant</tenantDesc>
  <accountIdentifier>unique string for this tenant</accountIdentifier>
  <status>{activated or disabled}</status>
```

```

    <removable>{true or false}</removable>
    <theme>{CA-Blue or CA-White}</theme>
    <defaultCulture>culture</defaultCulture>
  </tenant>

```

7. Replace any values with the values that you want to use for the new tenant.  
For example, supply the following parameters:

```

<tenant>
  <tenantName>John Doe</tenantName>
  <tenantDesc>John Doe Corporation tenant</tenantDesc>
  <accountIdentifier>JD1234</accountIdentifier>
  <status>Enabled</status>
  <removable>false</removable>
  <theme>CA-Blue</theme>
  <defaultCulture>en-US</defaultCulture>
</tenant>

```

For more information about tenant parameters, see [Tenants Service Example Syntax](#).

8. Run the method.
9. Repeat the preceding steps until you have created as many tenants as you require.

The new tenant definition is created, and the admin and user account are created.

### IMPORTANT

The passwords for the admin and user account are set to the username, and are one-time passwords. These password expire immediately. On first login, when prompted to change this password, these users must change them by clicking the change password link.

## Create Users Using Web Services

Use any REST client to create and configure a user account using the users web service.

Every user account is automatically associated with a tenant. If you are deploying multi-tenancy, the new user is assigned to the tenant of the authenticated user account that was used to make the REST service call. If you are not deploying multi-tenancy, this association is transparent to you; new user accounts are associated with the Default Tenant.

### Follow these steps:

1. Set up a REST client with a connection to the NetOps Portal server.
2. Enter a URL for the NetOps Portal RESTful web services API in the REST client. Use the following format:  
`http://<PC_host>:8181/pc/center/webservice/users/role/{roleIdName}/roleIdValue/`

#### – {roleIdName}

Use values that are specified in `http://PC_host:8181/pc/center/webservice/roles/idName s`.

**Examples:** 'roleName' and 'roleId'.

#### – roleIdValue

This value depends on the `roleIdName` that you selected. For example, if 'roleName' is used, substitute a valid role name for `roleIdValue`.

This role must be available within the tenant.

3. Select **POST** for **"HTTP" Method**.
4. Provide a valid Username and Password for a user account that has host or tenant administrator access to NetOps Portal.
5. Select **'application/xml'** as the **'Body Content-type'** in the Body settings.
6. Add the following parameters within the "Body" text section:

```

<user>

```



```

<name>{UserName}</name>
<description>{UserDescription}</description>
<enabled>{UserEnabled}</enabled>
<removable>{UserRemovable}</removable>
<timezone>{UserTimeZone}</timezone>
<culture>{UserCulture}</culture>
  <administeredGroups>
    <group ID="{GroupID}" />
    <group ID="{GroupID}" />
  </administeredGroups>
</user>

```

7. Replace any values with the values that you want to use for the new user account. For example, supply the following parameters:

```

<user>
  <name>Jane Doe</name>
  <description>User associated with the John Doe Corporation tenant.</description>
  <enabled>true</enabled>
  <removable>true</removable>
  <timezone>CST6CDT</timezone>
  <culture>en-US</culture>
  <administeredGroups>
    <group ID="105" />
    <group ID="367" />
  </administeredGroups>
</user>

```

For more information about user parameters, see [Users Web Service](#).

#### NOTE

The `administeredGroups` parameter is optional. To create a user without administered groups, exclude the tag.

8. Run the method.
9. Repeat the preceding steps until you have created as many users as you require.

The new user definition is created.

### Basic User Account Parameters

The current values for the following user account settings are available from a GET operation:

```
http://<PC_host>:8181/pc /center/webservice/users/cultureId
```

#### NOTE

- This URL returns information about users associated with the same tenant as the account that is used to execute this command.
- The GET method does not return password information. When you create a new user account, the password is automatically set to be the same as the user name.  
You can update the password of a specified user account using the PUT method. The password is sent in cleartext to avoid publicizing the encryption key for the web service to use. As a result, to protect the password privacy, use the method for changing the password *only* on the server where DX NetOps Performance Management is installed.

For *cultureId*, supply the language code for the language in which you would like to view the output, such as 'en-US'.

- **userID**

Is an internally assigned value for the user account.

- **name**  
Is a login name for the user account. The name is limited to 50 characters.
- **description**  
(Optional) Describes the user account to help you identify it.
- **enabled**  
Determines whether the user account is enabled for use (activated).
- **removable**  
States whether the item can be deleted (removed from the database).  
**Values:** true or false

#### NOTE

You cannot delete the two predefined user accounts, **admin** and **user**.

- **timezone (tz)**  
Corresponds to the time zone in which the user will view report data.  
**Default:** UTC (Coordinated Universal Time).
- **userLevel**  
Identifies the product privilege assigned to this user account.
- **role**  
Specifies the role assigned to the user account.
- **tenantId**  
Is an internal (database) identifier for the tenant with which the user account is associated.
- **culture**  
Specifies a language (locale).  
**Values:**
  - en-US (English, United States)
  - ja-JP (Japanese)
  - fr-FR (French, France)

### User Account Product Privilege Settings

You can determine the DX NetOps Performance Management privilege of a specified user account using NetOps Portal. *product privilege* is a type of permission set associated with a user account. This permission grants user access to features in selected data sources and does not apply to NetOps Portal functionality.

For more information, see [Product Privilege](#).

## Business Hours Web Service

Create and manage business hours definitions and deploy them to filter views using the `businesshours` web service.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the `businesshours` endpoint for the NetOps Portal RESTful web services. Enhance reporting by creating sets of business hours definitions. Business hours help product operators prioritize their troubleshooting workload by highlighting events that occur during business-critical time periods.

Use the groups web service to create site groups, and use the business hours web service to manage business hours definitions. If you add business hours first, you can then associate them with site groups during site-group creation.

You can add and delete multiple business hours definitions and assign them to site groups.

In this use case, we describe the steps to take when using a REST client, a generic web services user interface application. The examples in this use case contain URIs that are constructed using the default server port, 8181.

In this article:

### Verify the Prerequisite

- You have connected the REST client that you will use to invoke the NetOps Portal web services to the NetOps Portal RESTful web services.

### Access the businesshours Web Service Documentation

Issue the following call to see the parameters for the `businesshours` web service:

```
http://PC_host:8181/pc/center/webservice/businesshours
```

Issue the following call to see a list of supported operations:

```
http://PC_host:8181/pc/center/rest/businesshours/documentation
```

### Create Site Groups Using Web Services

Create and configure a site group that is associated with the Default Tenant using the `groups` web service. You can create rules and apply them to site groups so that items are added automatically. The steps to create groups within a custom tenant are slightly different. You can supply the group ID or the group path as parameters.

For more information about how to create site groups, see [Groups Web Service](#).

### Manage Time Zone Associations

You can manage associations of time zones with site groups using web services.

Assign a time zone to a site group using the following URL in a PUT operation:

```
http://PC_host:8181/pc/center/webservice/businesshours/assign/timezone/site/siteGroupId
```

Remove the assignment using the following syntax in a PUT operation:

```
unassign/timezone/site/siteGroupId
```

Get a list of time zones that are assigned to site groups using the following syntax:

```
http://PC_host:8181/pc/center/webservice/businesshours/timezonesAssignedToSites
```

Remove a time zone association by running the same method with the `unassign/timezone` syntax.

### **Follow these steps:**

1. Using the REST client with a connection to the NetOps Portal server, enter a URL for the NetOps Portal RESTful web services API in the REST client. Use the following format:

```
http://PC_host:8181/pc/center/webservice/businesshours/assign/timezone/site/siteGroupId
```

2. Select **PUT** for **"HTTP" Method**.
3. Provide a valid Username and Password for a user account that has administrator access to NetOps Portal.
4. Select **application/xml** as the **'Body Content-type'** in the Body settings.
5. Add the following XML within the **Body text** section:

```
<GroupTree path="/All Groups">
  <Group name="East Coast USA" desc="This is a site group"
    inherit="true" type="site group" location="North America"
    bHourID="99990" timeZone="EST"/>
</GroupTree>
```

- **bHourID**  
The internally assigned identifier of the business hours definition that you created previously.
- **timeZone**

The time zone to associate with this site group.

6. Run the method.

If it does not already exist, the USA site group is created under the default All Groups group in the Groups tree. The Eastern Standard timezone (EST) is assigned to this site group.

7. Repeat the preceding steps until you have associated time zones with all site groups to which you plan to apply business hours.

### Create a Business Hours Definition

You can create business hours definitions using the business hours web service and any REST client. For this procedure, you can log in as a global administrator or as a tenant administrator. The global administrator creates the business hours definitions within the Default Tenant, while the tenant administrator creates the definition within that tenant.

The steps to create business hours within a custom tenant are slightly different. You can supply the tenant ID as a parameter.

Business hours definitions comprise a starting hour and an ending hour. The `startHour` and `endHour` parameters in the XML must be the same for all days to avoid an error in any POST or PUT operation.

For systems with multiple tenants, specify the tenant in the URL. To get a list of tenant IDs, perform a GET to the following URL:

```
http://PC_host:8181/pc/center/webservice/tenants/idNames
```

#### NOTE

You can also modify business hours within a custom tenant by performing a PUT operation to the following URL:

```
/businesshours/tenantId/tenant_Id/id/id
```

where `id` is the ID of the business hours definition.

#### Follow these steps:

1. Using the REST client with a connection to the NetOps Portal server, enter a URL for the NetOps Portal RESTful web services API in the REST client. Use the following format:

– **Default Tenant:**

```
http://PC_host:8181/pc/center/webservice/businesshours/
```

– **Specific Tenant**

```
http://PC_host:8181/pc/center/webservice/businesshours/tenantId/tenant_ID
```

2. Select **POST** for **"HTTP" Method**.

3. Provide a valid Username and Password for a user account that has global administrator or tenant administrator access to NetOps Portal.

4. Select **application/xml** as the **'Body Content-type'** in the Body settings.

5. Add the following XML within the **Body text** section:

```
<BusinessHour>
  <Name>Bakery</Name>
  <Description>HEB Bakery</Description>
  <Monday>
    <HourRange startHour="5" endHour="12"/>
  </Monday>
  <Tuesday>
    <HourRange startHour="5" endHour="12"/>
  </Tuesday>
  <Wednesday>
    <HourRange startHour="5" endHour="12"/>
  </Wednesday>
  <Thursday>
```

```

<HourRange startHour="5" endHour="12"/>
</Thursday>
    <Friday>
<HourRange startHour="5" endHour="12"/>
</Friday>
    <Saturday/>
    </Sunday>
</BusinessHour>

```

In this example, a business hours definition named Bakery is created. The business hours start at 5 a.m. and end at noon. Saturday and Sunday are excluded.

6. Run the method.
7. Repeat the preceding steps until you have created as many business hours definitions as you require.

### **Manage Business Hours Associations**

Apply business hours filter by associating a business hours definition with a site group. The following procedure details how to assign business hours to a site group.

Remove a business hours assignment to a site group using the following syntax in a PUT operation:

```
unassign/businesshour/site/siteGroupId
```

Remove a business hours definition association with a site group by running the method with the `unassignbusinesshour` syntax.

#### **Follow these steps:**

1. Using the REST client with a connection to the NetOps Portal server, enter a URL for the NetOps Portal RESTful web services API in the REST client. Use the following format:

```
http://PC_host:8181/pc/center/webservice/businesshours/assign/businesshour/businessHourId/site/siteGroupId
```

#### **– businessHourId**

The internally assigned identifier of the business hours definition that you created previously.

2. Select **PUT** for **"HTTP" Method**.
3. Provide a valid Username and Password for a user account that has administrator access to NetOps Portal.
4. Select application/xml as the **'Body Content-type'** in the Body settings.
5. Add the following XML within the **Body text** section:

```

<GroupTree path="/All Groups">
    <Group name="East Coast USA" desc="This is a site group"
    inherit="true" type="site group" location="North America"
    bHourID="99990" timeZone="EST"/>
</GroupTree>

```

6. Run the method.

If it does not already exist, the USA site group is created under the default All Groups group in the Groups tree. The business hours definition with ID 99990 is assigned to this site group.

#### **NOTE**

NetOps Portal validates the site group for an associated time zone by performing a verification. If the site group does not have a time zone assignment, an error message appears.

7. Repeat the preceding steps until you have associated time zones with all site groups to which you plan to apply business hours.

## Query the RIB to Return a View with Business Hours Filtering

You can return data for a specific metric by entering queries into a web browser by querying the Report Information Base (RIB). This example presents a NetOps Portal RIB query that returns Top Discards data from a data aggregator data source.

Precede the NetOps Portal RIB queries with the following URL:

```
http://PC_host:8481/dm/rib/query/
```

You can append URL parameters to specify property values:

```
http://PC_host:8481/dm/rib/query/ribquery/?property1=value1&property2=value2
```

The following RIB query returns Top Discards data from a data aggregator data source:

```
http://<server IP address>:port/dm/rib/query/
SELECT .PollItem.ID, .PollItem.DevDisplayName, .Item.DisplayName, .Discards.Sum, .DiscardsIn.Sum, .DiscardsOut.Sum
FROM CA.IM.DA.MF.NormalizedPortInfo.IFSTATS WHERE .Group.GroupID = 1039 AND .EndTime(300) > 1366208760
AND .EndTime(300) <= 1366212360 GROUPBY .PollItem.ID, .Item.DisplayName, .PollItem.DevDisplayName
ORDERBY .Discards.Sum DESC LIMIT 10
```

### TIP

If necessary, you can escape the RIB query and parameters in your web browser, for example:

```
http://PC_host:port/dm/rib/query/SELECT%20.PollItem.ID,
%20.PollItem.DevDisplayName,%20.Item.DisplayName,
%20.Discards.Sum,%20.DiscardsIn.Sum,%20.DiscardsOut.Sum%20FROM
%20CA.IM.DA.MF.NormalizedPortInfo.IFSTATS%20WHERE%20.Group.GroupID%20=
%201039%20AND%20.EndTime(300)%20%3E%201366208760%20AND%20.EndTime(300)%20%3C=
%201366212360%20GROUPBY%20.PollItem.ID,%20.Item.DisplayName,
%20.PollItem.DevDisplayName%20ORDERBY%20.Discards.Sum%20DESC%20LIMIT%2010
```

Add the following URL parameters to return Top Discards data for a set of business hours in a specific time zone. NetOps Portal Administrators configure sets of business hours.

### NOTE

Some queries support only data filtering by time zone and business hours.

- **RIB.TimeZone**  
Is the string identifier of the time zone used to filter data results.
- **RIB.BusinessHours**  
Is the NetOps Portal ID of the business hour definition used to filter data results. Include this parameter in the `propertiesToTranslate` value to ensure that the ID is translated. IDs that are not translated are submitted unchanged to each applicable data source.
- **propertiesToTranslate**  
Is a list of parameter names whose values contain a NetOps Portal ID to translate to a local data source ID.

## Example 1

To return data filtered by time zone, add the time zone parameter (shown in bold text) to the URL. In the following example, the data is filtered to include only data for items in sites configured for the America/New\_York time zone:

```
http://pghost:8481/dm/rib/query/
SELECT .PollItem.ID, .PollItem.DevDisplayName, .Item.DisplayName, .Discards.Sum, .DiscardsIn.Sum, .DiscardsOut.Sum
FROM CA.IM.DA.MF.NormalizedPortInfo.IFSTATS WHERE .Group.GroupID = 1039 AND .EndTime(300) > 1366208760
AND .EndTime(300) <= 1366212360 GROUPBY .PollItem.ID, .Item.DisplayName, .PollItem.DevDisplayName
ORDERBY .Discards.Sum DESC LIMIT 10?RIB.TimeZone=America/New_York
```

## Example 2

To return data filtered by time zone and business hours, add the time zone and business hours parameters (shown in bold text) to the URL. In the following example, the data is filtered to include only data for items in sites configured for the America/New\_York time zone and business hours matching the NetOps Portal definition for ID 6434:

```
http://pchost:8481/dm/rib/query/
SELECT .PollItem.ID, .PollItem.DevDisplayName, .Item.DisplayName, .Discards.Sum, .DiscardsIn.Sum, .DiscardsOut.Sum
FROM CA.IM.DA.MF.NormalizedPortInfo.IFSTATS WHERE .Group.GroupID = 1039 AND .EndTime(300) > 1366208760
AND .EndTime(300) <= 1366212360 GROUPBY .PollItem.ID, .Item.DisplayName, .PollItem.DevDisplayName
ORDERBY .Discards.Sum DESC LIMIT 10?RIB.TimeZone=America/
New_York&RIB.BusinessHours=6434&propertiesToTranslate=RIB.BusinessHours
```

## Troubleshooting

Errors are returned if valid syntax invokes a definition that is not itself valid. For example, you attempt to create a site group. The syntax includes the `businessHourId` for an invalid business hours definition. In such a case, the HTTP Response XML includes an error message similar to the following text:

```
<Group bHourID="99990" desc="This is a site group" inherit="true" location="North America" name="East Coast
USA" result="Error with validating business hour ID: Business hour definition with an ID of '99990' not
found!" timeZone="EST" type="site group"/>
```

Business hours definitions comprise a starting hour and an ending hour. The `startHour` and `endHour` parameters in the XML must be the same for all of the days that are included in the definition. Otherwise, you see an error in POST or PUT operations.

## Maintenance Indicators Web Service

Automate the management (create, delete, and update) of the maintenance indicators using the `maintenanceindicators` web services.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services. Maintenance indicators represent times when maintenance is occurring. After you schedule maintenance indicators, views indicate maintenance with shading.

Maintenance indicators apply to all the devices and components in a site group. When the associated site group is selected in the context, maintenance indicators appear in applicable views as you navigate between dashboards. The subtitle of each view indicates whether maintenance indicators apply to the view.

Subgroups do not directly inherit maintenance indicators from the site groups. Associate the maintenance indicators with each relevant subgroup. However, when rendering views, these filters apply to all items based on the selected site group. The filters of the selected site group apply to all items in that group and in any subgroups. When you change the selected site group to a subgroup, the filters of the parent group are no longer applicable.

Reference groups inherit associated maintenance indicators from the original site group.

## Maintenance Indicators Web Service Operations

Issue the following call to see the available operations for the maintenance indicators web service:

```
http://PC_host:8181/pc/center/rest/maintenanceindicators/documentation
```

- **PC\_host:8181**  
The NetOps Portal IP address. 8181 is the required port.

## Operations

- **get list**

Get a list of all maintenance indicators definitions for the authenticated user.

**URL:**

`http://PC_host:8181/pc/center/webservice/maintenanceindicators/`

**HTTP method:** GET

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/maintenanceindicators/xsd`

- **add**

Create a maintenance indicators definition from the provided XML.

**URL:**

`http://PC_host:8181/pc/center/webservice/maintenanceindicators/tenantId/{tenantId}`

– **{tenantId}**

Specify the ID for the desired tenant.

**HTTP method:** POST

**XSD for the provided XML:** `http://PC_host:8181/pc/center/rest/maintenanceindicators/xsd`

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/maintenanceindicators/xsd`

- **delete**

Delete the maintenance indicators definition using the specified ID.

**URL:**

`http://PC_host:8181/pc/center/webservice/maintenanceindicators/id/{maintenanceIndicatorId}`

– **{maintenanceIndicatorId}**

Specify the ID for the desired maintenance indicator.

**HTTP method:** DELETE

Return type is a string.

- **get**

Get a specific maintenance indicators definition.

**URL:**

`http://PC_host:8181/pc/center/webservice/maintenanceindicators/id/{maintenanceIndicatorId}`

**HTTP method:** GET

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/maintenanceindicators/xsd`

- **get id names**

Return a list of identifiers that can be used in other methods to identify certain objects. For this web service, an empty list is returned.

**URL:**

`http://PC_host:8181/pc/center/webservice/maintenanceindicators/idNames`

**HTTP method:** GET

- **get list for tenant**

Get a list of all maintenance indicators definitions for the specified tenant.

**URL:**

`http://PC_host:8181/pc/center/webservice/maintenanceindicators/tenantId/{tenantId}`

– **{tenantId}**

Specify the ID for the desired tenant.

**HTTP method:** GET

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/maintenanceindicators/xsd`

- **update**

Update the specified maintenance indicators definition from the provided XML.

**URL:**

`http://PC_host:8181/pc/center/webservice/maintenanceindicators/tenantId/{tenantId}/  
id/{maintenanceIndicatorId}`



**HTTP method:** PUT

**XSD for the provided XML:** `http://PC_host:8181/pc/center/rest/maintenanceindicators/xsd`

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/maintenanceindicators/xsd`

### **Create Maintenance Indicators Using Web Services**

Use any REST client to create and configure maintenance indicators using the maintenance indicators web service.

#### **Follow these steps:**

1. Set up a REST client with a connection to the NetOps Portal server.
2. Use the following format for the URL in the REST client:  
`http://PC_host :8181/pc/center/webservice/maintenanceindicators/tenantId/8`

#### **NOTE**

8 is the ID for the Default Tenant.

3. Select `POST` for the '**HTTP' Method**.
4. Provide a valid Username and Password for a user account that has global administrator access to NetOps Portal.
5. Select '**application/xml**' as the '**Body Content-type**'.
6. Add the following XML within the "Body" text section, replacing any values with the values that you want to use for the new maintenance indicators:

```
<MaintenanceIndicator>
  <Name>Maintenance Indicator Name</Name>
  <Description>Maintenance Indicator Description</Description>
  <MaintenanceYear>Year</MaintenanceYear>
  <MaintenanceMonth>Month</MaintenanceMonth>
  <MaintenanceDay>Day</MaintenanceDay>
  <StartHour>StartHour</StartHour>
  <EndHour>EndHour</EndHour>
  <SelectedSites>
    <SelectedSite>SiteID</SelectedSite>
    <SelectedSite>SiteID</SelectedSite>
  </SelectedSites>
</MaintenanceIndicator>
```

**Example:** In this example, 3569 is the site ID for the **Framingham** site:

```
<MaintenanceIndicator>
  <Name>Framingham router maintenance</Name>
  <Description>Upgrade of the network for framingham</Description>
  <MaintenanceYear>2016</MaintenanceYear>
  <MaintenanceMonth>8</MaintenanceMonth>
  <MaintenanceDay>5</MaintenanceDay>
  <StartHour>19</StartHour>
  <EndHour>22</EndHour>
  <SelectedSites>
    <SelectedSite>510</SelectedSite>
    <SelectedSite>511</SelectedSite>
  </SelectedSites>
</MaintenanceIndicator>
```

- **Name**  
Specify a name for this maintenance indicators definition.
  - **Description**  
Specify a description for this maintenance indicators definition.
  - **MaintenanceYear**  
Specify a four-digit year for when the related maintenance takes place.
  - **MaintenanceMonth**  
Specify a month (1 - 12) for when the related maintenance takes place.
  - **MaintenanceDay**  
Specify a day (1 - 31) for when the related maintenance takes place.
  - **StartHour**  
Specify an hour (0 - 23) for when the related maintenance takes place.
  - **EndHour**  
Specify an hour (1 - 24) for when the related maintenance takes place.
  - **SelectedSites**  
(Optional) Specify one or more tags containing the site IDs to assign this maintenance indicators definition to. To create a definition without sites that are assigned, omit this parameter.
7. Run the method.
  8. Repeat the preceding steps until you have created as many maintenance indicators as you require.

## Alarm Attributes Web Service

Automate the management (create, edit, and delete) of the alarm attributes using the `alarmattributes` web services.

Users with the Administer Users and the Allow Access to REST Services rights, or LDAP-authenticated user accounts with the Admin role can use the NetOps Portal RESTful web services. You can create alarm attributes for a specific tenant or for all tenants (global).

### TIP

You can also manage alarm filter attributes using NetOps Portal.  
For more information, see [Alarms View](#).

### NOTE

- You can create the same attribute at the tenant level and the global level. The tenant-level attribute applies to the tenant user and the global-level attribute applies only to all other tenant users.
- You can create a maximum of five attributes in each type (String, Boolean, and Integer) to the **Alarm Console** dashboard. For example, you can create up to five String, five Boolean, and five Integer custom attributes. The Alarms view or alarm filters do not display the attributes that you create over this limit.
- Do not manage external attributes using these web services. External attributes directly correspond to specific MIB variables in DX NetOps Spectrum.

## Alarm Attributes Web Service Operations

Issue the following call to see the available operations for the `alarmattributes` web service:

```
http://PC_host:8181/pc/center/rest/alarmattributes/documentation
```

- **PC\_host:8181**  
Specifies the NetOps Portal host name. 8181 is the required port.

## Operations

- **delete**

Delete the attribute using the specified ID.

**URL:**

`http://PC_host:8181/pc/center/webservice/alarmattributes/tenantId/{tenantId}/attributeId/{attributeId}`

- **{tenantId}**

The id for the desired tenant

- **{attributeId}**

The hexadecimal id of the alarm attribute

**HTTP method:** DELETE

Return type is a string.

- **delete global**

Delete the global attribute using the specified ID.

**URL:**

`http://PC_host:8181/pc/center/webservice/alarmattributes/global/attributeId/{attributeId}`

- **{attributeId}**

The hexadecimal id of the alarm attribute

**HTTP method:** DELETE

Return type is a string.

- **get id names**

Return an empty list. In other web services, it can potentially return a list of identifiers that can be used in other methods to identify certain objects.

**URL:**

`http://PC_host:8181/pc/center/webservice/alarmattributes/idNames`

**HTTP method:** GET

- **get list for tenant**

Get a list of all attribute definitions for the specified tenant.

**URL:**

`http://PC_host:8181/pc/center/webservice/alarmattributes/tenantId/{tenantId}`

- **{tenantId}**

The id for the desired tenant

**HTTP method:** GET

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/alarmattributes/xsd`

- **add or update**

Add or update an attribute for the specified tenant.

**URL:**

`http://PC_host:8181/pc/center/webservice/alarmattributes/tenantId/{tenantId}`

- **{tenantId}**

The id for the desired tenant

**HTTP method:** POST

**XSD for the provided XML:** `http://PC_host:8181/pc/center/rest/alarmattributes/xsd`

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/alarmattributes/xsd`

- **get global list**

Get a list of all global attribute definitions.

**URL:**

`http://PC_host:8181/pc/center/webservice/alarmattributes/global`

**HTTP method:** GET

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/alarmattributes/xsd`

- **add or update global**

Add or update a global attribute.

**URL:**

`http://PC_host:8181/pc/center/webservice/alarmattributes/global`

**HTTP method:** POST

**XSD for the provided XML:** `http://PC_host:8181/pc/center/rest/alarmattributes/xsd`

**XSD for the returned XML:** `http://PC_host:8181/pc/center/rest/alarmattributes/xsd`

## Create Alarm Attributes Using Web Services

Use any REST client to create and configure alarm attributes using the alarm attributes web service.

### Follow these steps:

1. Set up a REST client with a connection to the NetOps Portal server.
2. Use the following format for the URL in the REST client:

`http://PC_host :8181/pc/center/webservice/alarmattributes/tenantId/8`

**NOTE**

8 is the ID for the Default Tenant.

3. Select POST for the **HTTP Method**.
4. Provide a valid Username and Password for a user account that has global administrator access to NetOps Portal.
5. Select '**application/xml**' as the '**Body Content-type**'.
6. Add the following XML within the **Body** text section, replacing any values with the values that you want to use for the new alarm attribute:

```
<AlarmAttribute>
  <AttributeID>HEX_ID</AttributeID>
  <Name>Attribute_Name</Name>
  <Description>Attribute_Description</Description>
  <Type>Value_Type</Type>
  <AddAsFilter>true</AddAsFilter>
  <AddAsColumn>true</AddAsColumn>
</AlarmAttribute>
```

**Example:**

**Example:**

```
<AlarmAttribute>
  <AttributeID>0x129e7</AttributeID>
  <Name>Topology Model Name</Name>
  <Description>The Topology Model Name String</Description>
  <Type>STRING</Type>
  <AddAsFilter>true</AddAsFilter>
  <AddAsColumn>true</AddAsColumn>
</AlarmAttribute>
```

- **AttributeID**  
Specify the hexadecimal ID of the alarm attribute.
- **Name**  
Specify a name for this alarm attribute.
- **Description**  
Specify a description for this alarm attribute.
- **Type**  
Specify one of the following allowable types:

- **STRING**  
A String field, for example, the Device Name.
  - **BOOLEAN**  
A Boolean value, for example, Acknowledged.
  - **ADDRESS\_RANGE**  
An IP address field, for example, the IP Address.
  - **INTEGER**  
A numeric value, for example, Number of Occurrences.
  - **HEX**  
A hexadecimal value, for example, Cause Code.
  - **OCTET\_STRING**  
A string value that is treated as a numeric value.
  - **AddAsFilter**  
Specify whether to add the attribute to the filter list.  
**Default:** `true`
  - **AddAsColumn**  
Specify whether to add the attribute as a column in an Alarm view. The column is hidden by default.  
**Default:** `true`
7. Run the method.
  8. Repeat the preceding steps until you have created as many alarm attributes as you require.

## Data Aggregator REST Web Services

You can manage administrative operations, such as retrieving data or managing relationships between profiles and tenants or groups, using the data aggregator REST web services.

You can interact with the REST APIs using a REST client tool that can send requests and can receive responses or cURL.

These REST web services include two types:

- **Generic**  
Use this type to manage metric families and SNMP vendor certification.
- **Data-driven**  
Use this type for most Data Aggregator REST web services, such as to read and modify data aggregator configuration information, such as monitoring profiles and groups.

### Access the Data Aggregator REST Web Services

Access the data aggregator REST web services using one of the following methods:

- [From a web browser](#)
- [Using cURL](#)

### Access the Web Services from a Web Browser

The data aggregator REST web services endpoints require authentication and require that Single Sign-On (SSO) be up and running.

The endpoints map to types of items, such as groups, monitoring profiles, tenants, and device certifications. Use specific endpoints to return a list of results, or create, update, or delete an item.

#### **NOTE**

Specify individual metric family items and SNMP vendor certification items using name instead of ID.

**IMPORTANT**

- As a security mechanism, provide HTTP basic authentication credentials with a valid username and password over HTTP/HTTPS when making calls to the data aggregator REST services. This is the same username and password that you use to log in to NetOps Portal. The REST logins are logged on the data aggregator in the `$KARAF_HOME/data/logs/AuthenticationAudit.log` file.
- Set the Context-type field to application/xml when you perform operations using data aggregator REST web services.

**Basic Data Aggregator REST web services operations:**

- `GET http://.../endpoint`  
Returns a list of all items of the specified type. The `getlist.xsd` schema defines the format for the return data.
- `GET http://.../endpoint/[id | name]`  
Returns the details for a single item with the specified ID or certification name. The XSD schema defines the format for the return data.
- `POST http://.../endpoint`  
Creates an object of the specified type with specified facets. The XSD schema defines the format for the return data.  
**NOTE**  
With POST operations, you can create objects with the same name but different IDs. This convention is allowed because it is often valid when the objects are scoped to different tenants.
- `PUT http://.../endpoint/[id | name]`  
Updates the attributes of the specified item. The `update.xsd` schema defines the format and expected fields.  
**NOTE**  
With PUT operations, you can create objects with the same name but different IDs. This convention is allowed because it is often valid when the objects are scoped to different tenants.
- `DELETE http://.../endpoint/[id | name]`  
Deletes the item that is specified using the ID or certification name.

**Access the Web Services using cURL**

You can access the data aggregator REST web services using the standard cURL utility. cURL is a popular command-line utility to interact with various network protocols, including HTTP and HTTPS. It is pre-installed on the Mac. You must install it on Windows. The following are a few examples of using cURL.

**Examples:**

Request the password from the user using the `-u` option with only the username:

```
curl -s -k -u admin https://localhost:8582/rest
```

Have a shell script request the username and password, without the password appearing in 'ps' output:

```
echo "DA REST requires authentication."
read -p "Username: " user
read -s -p "Password: " pass
echo
curl --config - -s -k https://localhost:8582/rest <<< "user = \"$user:$pass\""
```

**Generic REST Web Services**

Use *generic* REST web services to manage metric families and for limited SNMP vendor certification support. Generic REST web services are self-filtering and do not use an argument within the URL to manage relationships.

You can display the details about the user-facing, generic REST web services using the following URL:

```
http://DA_host:8581/genericWS
```

This URL includes detailed information about generic web services, including XSDs, URIs, supported HTTP methods, attributes, and relationships.

To view detailed documentation about a specific endpoint, access the following URL:

```
http://DA_host:8581/genericWS/endpoint/documentation
```

### **Manage Relationships using Generic REST Web Services**

When managing relationships, the generic REST web services do not use an argument within the URL. Instead, generic REST web services rely on the basic operations alone to manage relationships. The endpoints filter on themselves to expose the information. These methods are used for managing relationships between the metric families and SNMP vendor certifications.

You can view, create, and delete relationships for generic REST web services using the GET, PUT, and DELETE methods.

These methods use the following URL:

```
http://DA_host:8581/genericWS/endpoint/name/endpoint
```

### **Example: List Metric Families Related to an SNMP Vendor Certification**

To return a list of metric families for a specified SNMP vendor certification, use the following method:

- **Method:** GET
- **URL:**

```
http://DA_host:8581/genericWS/certifications/snmp/name/metricfamilies
```

### **Data-Driven REST Web Services**

Use *data-driven* REST web services to read and modify data aggregator configuration information, such as monitoring profiles and groups. You can display the details about the data-driven REST web services using the following URL:

```
http://DA_host:8581/rest
```

This URL includes detailed information about endpoints, and data-driven REST web services, including XSDs, URIs, supported HTTP methods, attributes, and relationships.

To view detailed documentation about a specific endpoint, access the following URL:

```
http://DA_host:8581/rest/endpoint/documentation
```

### **View XSD Schemas using Data-Driven REST Web Services**

Do the following steps before performing an HTTP request:

1. Verify the XML schema definition (XSD) for the endpoint.
2. Review the format of the return or upload XML that the service provides.

Each item of content that is placed in an XML document must adhere to the description of the endpoint.

The XSD files for each operation contain tags that describe the attributes and the purpose of the metric families. To obtain the XSD for an endpoint, use the following paths with the data-driven web services URL:

```
http://DA_host:8581/rest/endpoint/XSD/operation.xsd
```

- ***operation***

Specifies the type of operation to execute.

**Values:**

- **get**  
The XSD for a single item get.
- **getlist**  
The XSD for a list of the endpoint items.
- **filterselect**  
The XSD for advanced filter criteria and return XML format to be specified using GET Tunneling.
- **create**  
The XSD that any input XML must match when trying to create.
- **update**  
The XSD that any input XML must match when trying to update.

**NOTE**

Not all endpoints support all operations. If it does not support an operation, the web service fails and returns the following message:

```
405 Method Not Allowed
```

### **Filter on Attributes in the XSD Schema for Data-Driven REST Web Services**

You can filter on attributes such as the item name, description, and other such attributes. For example, filter monitoring profiles by the metric families they contain. You can use this information to determine whether to add or remove metric families to or from a monitoring profile.

**Follow these steps:**

1. Enter the following URL in a web browser:  
`http://DA_host:8581/rest/`  
A list shows the available data-driven web services.
2. Click a web service.  
The documentation page for that web service opens.
3. Click the URL under the **filtered get list** method.  
The XSD schema opens.
4. Look for the elements that include the following property:  
`substitutionGroup="AttributeFilterTypeSubstitution"`  
Use this information to determine on which attribute you want to filter.
5. Open a REST client editor or HTTP tool that sends requests and gets responses, and set the Content-type to **application/xml**.
6. Enter the following filter criteria:
  - **HTTP method:** POST
  - **URL:**  
`http://DA_host:port/rest/endpoint/filtered/`
  - **Body:**

```
<FilterSelect xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="filter.xsd">
  <Filter>
    <elementName type="CONTAINS">filter-criteria</elementName>
  </Filter>
</FilterSelect>
```

- ***filter-criteria***



Specifies the actual value of the attribute.

- **elementName**

Specifies the element name (attribute) on which to filter.

**NOTE**

You can specify selection criteria also, such as poll rates only. This method is also known as GET Tunneling.

For more information, see the following example.

Results are returned in the **Body** tab of the HTTP Response pane.

**Example: Return a List of Monitoring Profiles that Contain a Metric Family Using Filter and Selection Criteria (GET Tunneling)**

You can return a list of the monitoring profiles that contain a metric family using poll rate as the selection criteria (GET Tunneling) using the following method and URL:

- **Method:** POST

- **URL:**

`http://DA_host:8581/rest/monitoringprofiles/filtered/`

- **Body:**

```
<FilterSelect xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="filter.xsd">
  <Filter>
    <MonitoringProfile.FacetTypes type="CONTAINS">{http://im.ca.com/
normalizer}NormalizedPortInfo</MonitoringProfile.FacetTypes>
  </Filter>
  <Select use="exclude" isa="exclude">
    <MonitoringProfile use="exclude">
      <PollRate use="include"/>
    </MonitoringProfile>
  </Select>
</FilterSelect>
```

**Example: Return a List of Monitoring Profiles that Contain an Item Using Filter and Selection Criteria**

You can return the monitoring profiles that contain an item using name as the selection criteria using the following method and URL:

- **Method:** POST

- **URL:**

`http://DA_host:8581/rest/monitoringprofiles/filtered/`

- **Body:**

```
<FilterSelect xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="filter.xsd">
  <Filter>
    <MonitoringProfile.FacetTypes type="CONTAINS">{http://im.ca.com/
normalizer}NormalizedPortInfo</MonitoringProfile.FacetTypes>
  </Filter>
  <Select use="exclude" isa="exclude">
    <Item use="exclude">
      <Name use="include"/>
    </Item>
  </Select>
</FilterSelect>
```

```

        </Item>
    </Select>
</FilterSelect>

```

### **Manage Basic Relationships with Data-Driven REST Web Services**

You can create or delete relationships using the following methods and URL:

- **Method:** PUT or DELETE
- **URL:**

```
http://DA_host:8581/rest/endpoint/id/relatesto/endpoint/id
```

You can view nested relationships using the following URL:

```
http://DA_host:8581/rest/endpoint/id/relatesto/endpoint/endpoint
```

### **Return a List of Associated Groups and Devices for a Specific Monitoring Profile**

You can return a list of the related groups and devices for a specified monitoring profile ID with the attributes in the current `getlist.xsd` file using the following method and URL:

- **Method:** GET
- **URL:**

```
http://DA_host:8581/rest/monitoringprofile/781/relatesto/groups/devices
```

### **Limit Data-Driven Scope to Tenant Domains**

You can limit the data-driven scope of your operations to a specific tenant domain for some of the data aggregator features, rather than accessing information in the entire global repository.

You can access information for an endpoint with a specific tenant domain using the following basic method and URL:

- **Method:** GET
- **URL:**

```
http://DA_host:port/rest/tenant/id/endpoint
```

#### **NOTE**

You can use some, but not all, web services with tenant domains.

## **Change When Same Day, Same Hour Baseline Averages Are Calculated**

After a limited amount of data has been collected, the baseline average is calculated for the same hour for every preceding day of the week.

When more data is available, a switchover in the calculation method occurs automatically and Data Aggregator establishes "normal" by averaging hourly samples across available preceding same days of the week.

By default, this automatic switchover occurs when at least 3 same day of the week, same hour data samples are available for the past 12 weeks. You can change when this automatic switchover occurs.

Consider the following information about changing when same day of the week, same hour baseline averages are calculated:

- You have the option of changing both attributes, or just one of the attributes.
- Both attribute values must be a numeric value that is greater than or equal to 1.
- No upper limit is enforced. However, the retention policy of the hourly roll ups defines the upper limit. By default, the hourly retention rate is 90 days (which is approximately 12 weeks of data). If you increase the maximum number of preceding weeks, increase the hourly roll-up retention rate also.
- The MinimumNumberOfRequiredDataPoints attribute value must be less than or equal to the MaximumNumberOfWeeks value.

### Follow these steps:

1. Enter the following information in a web browser:  
`http://DA_host:port/rest/sdshbaselineconfig`
  - **DA\_host:port**  
 Specifies the Data Aggregator host name and the port number.  
**Default port:** 8581  
 The sdshbaselineconfig webservice endpoint URL opens.
2. Review the current values for the minimum required number of data points and the maximum number of preceding weeks.
3. Open a REST client editor or HTTP tool that sends requests and gets responses and set the Content-type to application/xml.
4. Enter the following criteria:
  - HTTP method = PUT
  - Enter the minimum required number of data points within the maximum number of preceding weeks (to trigger the switchover in the baseline calculation method) that you want to change on the Body tab in the HTTP Request pane. For example:

```
<SdshBaselineConfiguration version="1.0.0">

  <SDSHSettings>

    <MinimumNumberOfRequiredDataPoints>5</MinimumNumberOfRequiredDataPoints>

    <MaximumNumberOfWeeks>10</MaximumNumberOfWeeks>

  </SDSHSettings>

</SdshBaselineConfiguration>
```

In this example, the minimum number of data points to be available for baseline average calculation has been changed to 5. The number of preceding weeks to look for these data points has been changed to 10.

## Manage Polling Behavior for Components

You can use this Data Aggregator REST web service to disable or enable polling on other components. This feature allows for more granular polling control than monitoring profile filters alone.

You can also disable or enable polling on specific interfaces using [monitoring profiles](#).

By default, polling is enabled for all new components. You can use a Data Aggregator REST web service to disable polling for all new components associated with specific metric families.

For more information, see [Manage Default Polling Behavior](#).

**Follow these steps:**

1. Set up a REST client with a connection to the Data Aggregator server.
2. Query the following REST URL to find the `pollable` item ID:  
`http://DA_host:8581/rest/pollable`
3. Use the following format for the URL in the REST client:  
`http://DA_host:8581/rest/pollable/<itemID>`
  - ***itemID***  
Specify the `pollable` item ID.
4. Select `PUT` for the **'HTTP' Method**.
5. Provide a valid Username and Password for a user account that has global administrator access to the Data Aggregator.
6. Select **'application/xml'** as the **'Body Content-type'**.
7. Add the following XML within the "Body" text section:

```
<Pollable version="1.1.1">
  <IsPollEnabled>false</IsPollEnabled>
</Pollable>
```

- ***IsPollEnabled***  
To disable polling, specify `false`. To enable polling, specify `true`.
8. Run the method.

## Manage Default Polling Behavior

By default, polling is enabled for all new components. You can use the `discoverydefaultconfig` REST web service to disable polling for all new components of specific metric families.

**NOTE**

You can disable or enable polling on specific interfaces using NetOps Portal.

For more information, see [Manage Interface Polling Behavior](#).

**NOTE**

You can disable or enable polling on other components using the `pollable` REST web service.

For more information, see [Manage Polling Behavior for Components](#).

Using the Data Aggregator REST web service to disable polling allows for more granular polling control than monitoring profile filters alone.

For more information, see [Manage Monitoring Profiles](#).

In this article:

- [Disable Polling for New Components](#)
- [Re-enable Polling for New Components](#)
- [Manage Filtered Components](#)

### Disable Polling for New Components

You can disable polling for all new components of specific metric families.

**Follow these steps:**

1. Set up a REST client with a connection to the Data Aggregator server.

2. Query the following REST URL to find the `DiscoveryDefaultConfig` item ID by issuing a GET to the `discoverydefaultconfig` web service:  
`http://DA_host:8581/rest/discoverydefaultconfig`
3. Use the following format for the URL in the REST client:  
`http://DA_host:8581/rest/discoverydefaultconfig/<itemID>`
  - **itemID**  
Specify the `DiscoveryDefaultConfig` item ID.
4. Select **PUT** as the **HTTP Method**.
5. Provide a valid Username and Password for a user account that has global administrator access to the data aggregator.
6. Select **'application/xml'** as the **'Body Content-type'**.
7. Add the following XML within the **Body** text section:  

```
<DiscoveryDefaultConfig version="1.0.0">
  <DisablePollingOfNewComponentsList>
    <DisablePollingOfNewComponents>{http://im.ca.com/
normalizer}Metric_Family_Name</DisablePollingOfNewComponents>
  </DisablePollingOfNewComponentsList>
</DiscoveryDefaultConfig>
```

  - **Metric\_Family\_Name**  
Specify the metric family that is associated with the components for which you want to disable polling.
8. Edit the XML to include the metric families that are associated with the components for which you want to disable polling.  
**Example:**  

```
<DiscoveryDefaultConfig version="1.0.0">
  <DisablePollingOfNewComponentsList>
    <DisablePollingOfNewComponents>{http://im.ca.com/
normalizer}NormalizedPortInfo</DisablePollingOfNewComponents>
    <DisablePollingOfNewComponents>{http://im.ca.com/
normalizer}NormalizedCPUInfo</DisablePollingOfNewComponents>
  </DisablePollingOfNewComponentsList>
</DiscoveryDefaultConfig>
```
9. Run the method.

### **Re-enable Polling for New Components**

If you want to re-enable polling for new components of a specific metric family, simply repute the XML excluding the metric family.

#### **Follow these steps:**

1. Query the following REST URL to find the `DiscoveryDefaultConfig` item ID by issuing a GET to the `discoverydefaultconfig` web service:  
`http://DA_host:8581/rest/discoverydefaultconfig`
2. Use the following format for the URL in the REST client:  
`http://DA_host:8581/rest/discoverydefaultconfig/<itemID>`
  - **itemID**  
Specify the `DiscoveryDefaultConfig` item ID.
3. Select **PUT** as the **HTTP Method**.
4. Provide a valid Username and Password for a user account that has global administrator access to the data aggregator.

5. Select **'application/xml'** as the **'Body Content-type'**.
6. Add the following XML within the **Body** text section:
 

```
<DiscoveryDefaultConfig version="1.0.0">
  <DisablePollingOfNewComponentsList>
    <DisablePollingOfNewComponents>{http://im.ca.com/
normalizer}Metric_Family_Name</DisablePollingOfNewComponents>
  </DisablePollingOfNewComponentsList>
</DiscoveryDefaultConfig>
```

  - **Metric\_Family\_Name**  
Specify the metric family that is associated with the components for which you want to disable polling.
7. Edit the XML to exclude the metric families that are associated with the components for which you to re-enable polling.  
**Example:** You want to re-enable polling for new components of the NormalizedCPUInfo metric family.
 

```
<DiscoveryDefaultConfig version="1.0.0">
  <DisablePollingOfNewComponentsList>
    <DisablePollingOfNewComponents>{http://im.ca.com/
normalizer}NormalizedPortInfo</DisablePollingOfNewComponents>
  </DisablePollingOfNewComponentsList>
</DiscoveryDefaultConfig>
```
8. Run the method.

### Manage Filtered Components

You can control whether filtered components are created with the `DoNotCreateFilteredItems` attribute.

#### **Follow these steps:**

1. Query the following REST URL to find the `DiscoveryDefaultConfig` item ID by issuing a GET to the `discoverydefaultconfig` web service:  
`http://DA_host:8581/rest/discoverydefaultconfig`
2. Use the following format for the URL in the REST client:  
`http://DA_host:8581/rest/discoverydefaultconfig/<itemID>`
  - **itemID**  
Specify the `DiscoveryDefaultConfig` item ID.
3. Select **PUT** as the **HTTP Method**.
4. Provide a valid Username and Password for a user account that has global administrator access to the Data Aggregator.
5. Select **'application/xml'** as the **'Body Content-type'**.
6. Add the following XML within the **Body** text section:
 

```
<DiscoveryDefaultConfig version="1.0.0">
  <DoNotCreateFilteredItems>true</DoNotCreateFilteredItems>
</DiscoveryDefaultConfig>
```

  - **DoNotCreateFilteredItems**  
Specifies whether to create filtered items.  
**Options:**
    - **true:** Filtered items are not created.
    - **false:** Filtered items are created.
7. Run the method.

## Poll Sensitive and Critical Devices Without a Performance Impact

Critical devices are sensitive to too many polls, which can lead to performance problems. To throttle the SNMP poll requests and avoid overwhelming your sensitive devices, configure the SNMP polling controls.

By default, SNMP polling is controlled in three ways:

- **SNMP traffic threshold**  
No more than 15 SNMP requests can be sent to a device at a time. Poll and discovery SNMP requests over 15 are queued and sent to a device when possible during the polling cycle. Up to 600 requests can be queued.
- **SNMP timeouts threshold**  
When 15 or more SNMP requests timeout, polling is suspended for the remainder of the current polling cycle. An event is generated, informing you of the situation.  
Polling resumes at the beginning of each poll cycle. When the timeouts do not exceed the 15 threshold, a clear event is generated.
- **SNMP PDU Segmentation**  
SNMP packets can be limited to contain a maximum number of varbinds by splitting a larger SNMP request into multiple smaller requests and reassembling the responses.

These thresholds are designed to prevent overwhelming a device with too many poll requests or PDUs that are too large. You can override these SNMP polling thresholds defaults by modifying the following parameters:

- **TimeoutFailSafeThrottleDefault**  
Specifies the default maximum number of timeouts that trigger the fail-safe throttle.
- **MaxOutstandingRequestsDefault**  
Specifies the default maximum number of outstanding requests.
- **MaxRequestSizeDefault**  
Specifies the default maximum number of varbind that a single SNMP request may contain.

These thresholds can also be modified per device so that a specific device can be assigned individual threshold values.

For example, your older router is exceptionally sensitive to polling. But, this router is critical and must be polled as frequently as possible. You already adjusted your monitoring profile to remove unnecessary metric families from polling. You also applied a filter in your monitoring profile to reduce the number of polled interfaces. However, polling still causes this router to crash. Therefore, your only option is to adjust the default SNMP polling parameters for your sensitive router.

You can add any of the following parameters to the policy for individual IPs or IP ranges in an `IPRange` section within the `IPRangeList`:

- **TimeoutFailSafeThrottle**  
The maximum number of timeouts that are applied on the devices within this IP range.
- **MaxOutstandingRequests**  
The maximum number of outstanding requests that are sent to the devices within the indicated IP range.
- **MaxRequestSize**  
Limits the number of varbinds in an outgoing SNMP request.  
If the number of varbinds in the SNMP request exceeds the value of `MaxRequestSize`, the outgoing request is split into two or more smaller requests.  
Some IP ranges are not covered in the `IPRange` sections. For global settings, use the `MaxRequestSizeDefault` parameter to set the varbind limit.

### NOTE

If `MaxRequestSize` is 0, the original request is sent regardless of its size.

For related troubleshooting information, see [Gaps in Data Appear during Throttling](#) and [Polling Stopped Event Message](#).

### Follow these steps:

1. Find the ID for your IP Domain (that contains your sensitive router) by opening:

`http://<DA_host:port>/rest/ipdomains`

– **DA\_host:port**

Specifies the data aggregator hostname and the port number where you are accessing the REST web services from.

2. Locate your IP Domain ID in the following SNMP throttle policy list, and note the corresponding policy ID:

`http://<DA_host:port>/rest/snmpthrottlepolicies`

3. Determine the number of varbinds that you want to include in a single outgoing SNMP request. Some devices ignore requests that are too large without sending an error. As a result, the SNMP poller cannot reach the device. Use the `MaxRequestSize` value to allow the Data Collector to monitor these devices.

**Example:** If the interface SNMP request has 27 varbinds and `MaxRequestSizeDefault` is set to 15, the outgoing request is split into two smaller requests. One request contains 14 varbinds, and the other contains 13 varbinds.

**Example:** The following example from an SNMP throttle policy shows that the policy ID is "601" for IP Domain "2" with no limit on the number of varbinds:

```
<SnmpThrottlePolicy version="1.0.0">
  <ID>601</ID>
  <MaxOutstandingRequestsDefault>15</MaxOutstandingRequestsDefault>
  <QueueLength>600</QueueLength>
  <TimeoutFailSafeThrottleDefault>15</TimeoutFailSafeThrottleDefault>
  <MaxRequestSizeDefault>0</MaxRequestSizeDefault>
  <IPDomainID>2</IPDomainID>
</SnmpThrottlePolicy>
```

4. Open a REST client editor or HTTP tool that sends requests and gets responses, and set the Content-type to `application/xml`.

5. Open and edit the SNMP throttle policy for your IP Domain by entering the following criteria:

– **URL:** `http://<DA_host:port>/rest/snmpthrottlepolicies/<policyID>`

**Example:** `http://<DA_host:port>/rest/snmpthrottlepolicies/601`

- **DA\_host:port**

Specifies the data aggregator hostname and the port number where you are accessing the REST web services from.

- **policyID**

Specifies a unique identification number that is assigned to the SNMP throttle policy for the IP Domain that contains your sensitive device.

– **HTTP method:** PUT

– Adjust the following values for your IP Range on the Body tab in the HTTP Request pane:

- **<MaxOutstandingRequests>**

SNMP traffic threshold

- **<TimeoutFailSafeThrottle>**

SNMP timeouts threshold

**NOTE**

Both values are required for every IP Range entry. You can disable either parameter by setting the value to "0."

– Remove the following lines:

- `<ID>`

- `<IPDomainID>`

Results are returned in the Body tab of the HTTP Response pane.

**Example:** In the following example, the thresholds are lowered to "10" for device `10.231.41.7` only. For this device, the number of varbinds is limited to 50. The default thresholds and other IP Range thresholds continue using the default value of "15." For devices `10.231.41.1` - `10.231.41.255`, SNMP requests are limited to 30 varbinds:



```

<SnmptThrottlePolicy version="1.0.0">
  <IPRangeList>
    <IPRange>
      <IPRangeText>10.231.41.7</IPRangeText>
      <MaxOutstandingRequests>10</MaxOutstandingRequests>
      <TimeoutFailSafeThrottle>10</TimeoutFailSafeThrottle>
      <MaxRequestSize>50</MaxRequestSize>
    </IPRange>
    <IPRange>
      <IPRangeText>10.231.41.1-10.231.41.255</IPRangeText>
      <MaxOutstandingRequests>15</MaxOutstandingRequests>
      <TimeoutFailSafeThrottle>15</TimeoutFailSafeThrottle>
      <MaxRequestSize>30</MaxRequestSize>
    </IPRange>
  </IPRangeList>
  <MaxRequestSizeDefault>0</MaxRequestSizeDefault>
  <MaxOutstandingRequestsDefault>15</MaxOutstandingRequestsDefault>
  <QueueLength>600</QueueLength>
  <TimeoutFailSafeThrottleDefault>15</TimeoutFailSafeThrottleDefault>
</SnmptThrottlePolicy>

```

**NOTE**

You can adjust the thresholds for a single device or a range of devices. The IP Range definition and the IP Range order determine which threshold applies. The IP Ranges are listed in priority order. That is, the first IP Range that applies to a device determines the threshold value to apply.

6. Always include the `MaxOutstandingRequestsDefault`, `MaxRequestSizeDefault`, `TimeoutFailSafeThrottleDefault`, and `QueueLength` parameters in the update/POST XML at the root level. Include the parameters even if the values do not differ from the default.

**Example:**

This PUT command generates the following policy:

**Update XML: PUT**

**URL:** `http://<DA_host>:8581/rest/snmptthrottlepolicies/21`

```

<SnmptThrottlePolicy version="1.0.0">
  <IPRangeList>
    <IPRange>
      <IPRangeText>130.119.103.8</IPRangeText>
      <MaxOutstandingRequests>10</MaxOutstandingRequests>
      <TimeoutFailSafeThrottle>10</TimeoutFailSafeThrottle>
      <MaxRequestSize>20</MaxRequestSize>
    </IPRange>
  </IPRangeList>
  <MaxRequestSizeDefault>0</MaxRequestSizeDefault>
  <MaxOutstandingRequestsDefault>15</MaxOutstandingRequestsDefault>
  <TimeoutFailSafeThrottleDefault>15</TimeoutFailSafeThrottleDefault>
  <QueueLength>600</QueueLength>
</SnmptThrottlePolicy>

```

This command generates the following policy:

```

<SnmptThrottlePolicy version="1.0.0">

```

```

<ID>21</ID>
<QueueLength>600</QueueLength>
<TimeoutFailSafeThrottleDefault>15</TimeoutFailSafeThrottleDefault>
<IPDomainID>2</IPDomainID>
<IPRangeList>
  <IPRange>
    <IPRangeText>130.119.103.8</IPRangeText>
    <MaxOutstandingRequests>10</MaxOutstandingRequests>
    <TimeoutFailSafeThrottle>10</TimeoutFailSafeThrottle>
    <MaxRequestSize>20</MaxRequestSize>
  </IPRange>
</IPRangeList>
<MaxRequestSize>0</MaxRequestSize>
<MaxOutstandingRequestsDefault>15</MaxOutstandingRequestsDefault>
</SnmpThrottlePolicy>

```

## Schedule Data Purges

You can schedule how often the data aggregator purges data that is older than the specified retention periods. You can modify the start hour, start minute, and the start second. By default, the data aggregator purges data every day at 2:00:00 AM.

### Follow these steps:

1. Access the `globalretentionscheduledefinition` webservice endpoint by entering the following URL in a web browser:  
`http://DA_host:port/rest/globalretentionscheduledefinition`
  - **DA\_host:port**  
Specifies the data aggregator hostname and the port number from where you are accessing the REST web services.
2. Take note of the ID that is assigned to the `globalretentionscheduledefinition` parameter.
3. Look for the elements that have the `StartMinute`, `StartHour`, and `StartSecond` parameters. Use this information to determine whether you want to modify the start hour, start minute, or start second when old data is purged.
4. Open a REST client editor or HTTP tool that sends requests and gets responses, and then set the Content-type to `application/xml`.
5. Enter the following criteria:
  - URL: `http://DA_host:port/rest/globalretentionscheduledefinition/ID`
    - **ID**  
The unique identification number that is assigned to the `globalretentionscheduledefinition` parameter.
  - HTTP method = PUT
  - Enter the time values that you want to change in the **Body** tab of the **HTTP Request** pane.

For example:

```

<GlobalRetentionScheduleDefinition version="1.0.0">
  <StartMinute>28</StartMinute>
  <StartHour>17</StartHour>
  <Enabled>true</Enabled>
  <Status>Scheduled to run everyday at 17:28:00</Status>

```

```
</GlobalRetentionScheduleDefinition>
```

### IMPORTANT

Be sure that there is no white space at the beginning of each of these lines, otherwise the PUT operation fails.

In this example, the start hour has been changed to 17 and the start minute has been changed to 28.

### NOTE

To disable the purge job, set `Enabled` to false. To reenable the purge job, set `Enabled` to true.

Results are returned in the **Body** tab of the **HTTP Response** pane.

For example:

```
<GlobalRetentionScheduleDefinitionList>
<GlobalRetentionScheduleDefinition version="1.0.0">
  <ID>9</ID>
  <StartMinute>28</StartMinute>
  <StartHour>17</StartHour>
  <Enabled>true</Enabled>
  <JobStatus>Has never run</JobStatus>
  <Status>Scheduled to run everyday at 17:28:00</Status>
  <StartSecond>0</StartSecond>
</Item version="1.0.0">
  <CreateTime>Thu Dec 15 15:52:20 EST 2011</CreateTime>
  <Name>Global Retention Schedule Definition</Name>
</Item>
</GlobalRetentionScheduleDefinition>
</GlobalRetentionScheduleDefinitionList>
```

In this example, data that is older than the specified retention periods is purged every day at 17:28:00.

## Schedule Rollup Processing and Baseline Calculations

Rollup processing and baseline calculations are intensive operations. By default, NetOps Portal runs these operations at the bottom of the hour, every hour of the day. But these operations can impact users who generate reports during business hours. Administrators can schedule these operations to occur during off-hours to alleviate the impact during business hours.

### Follow these steps:

1. Enter the following URL in a web browser:  
[http://DA\\_host:port/rest/rollups/config](http://DA_host:port/rest/rollups/config)
  - **DA\_host:port**  
 Specifies the data aggregator host name and the port number.  
**Default port:** 8581
2. Take note of the unique identification number that is assigned to the configuration item.
3. Open a REST client editor or HTTP tool that sends requests and gets responses. Enter the following criteria:
  - **URL:** [http://DA\\_host:port/rest/rollups/config/ID](http://DA_host:port/rest/rollups/config/ID)
    - **DA\_host:port**  
 Specifies the data aggregator host name and the port number.  
**Default port:** 8581

For more information about the ports that are required for DX NetOps Performance Management to work properly, see [Review Installation Requirements and Considerations](#).

- **ID**

The unique ID that is assigned to the configuration item (that you noted in the previous step).

- **HTTP method:** PUT
- Enter the hour of the day when you want rollup processing to begin and end in the **Body** tab of the **HTTP Request** pane.

By default, the following results are returned:

```
<RollupsConfigurationList>
  <RollupsConfiguration version="1.0.0">
    <ID>8</ID>
    <StartHour>0</StartHour>
    <EndHour>23</EndHour>
  </RollupsConfiguration>
</RollupsConfigurationList>
```

- **StartHour**

Defines the hour of the day (in your local timezone) in 24-hour time format when rollup processing will begin.

- **EndHour**

Defines the hour of the day (in your local timezone) in 24-hour time format when rollup processing should end. No new rollups will be kicked off after the end-hour, but any rollups that are in-progress will be allowed to complete.

For more information about these attributes, see [http://DA\\_host:port/rest/rollups/config/documentation](http://DA_host:port/rest/rollups/config/documentation).

### Example:

In this example, you change the schedule so that rollup processing and baseline calculations run only from 20:00 to 7:00.

```
<RollupsConfigurationList>
  <RollupsConfiguration version="1.0.0">
    <ID>8</ID>
    <StartHour>20</StartHour>
    <EndHour>6</EndHour>
  </RollupsConfiguration>
</RollupsConfigurationList>
```

The `<EndHour>` is inclusive. In this example, that means if you specify 6 as the `EndHour`, rollup processing and baseline calculations will be initiated in at the bottom of the hour during the 06:00 hour, but they will not be initiated at the 07:00 hour. Any calculations that are in progress will be allowed to complete.

### WARNING

Any modifications to the default schedule can result in a larger delay in data showing up in reports pertaining to the corresponding resolution. For example, if hourly rollups are delayed, then a reporting showing hourly resolution data will not be current until the hourly rollup has been performed.

## Manage Discovery Using REST

Tenant administrators can manage discovery using the `profiles` Data Aggregator REST web service.

### NOTE

You can also do the following using this web service:

- [Get a list of unreachable devices.](#)
- [Automate device inventory synchronization, including reviewing the discovered devices and instances.](#)

## Automate Creating Discovery Profiles

Discovery profiles specify how discovery operates. They determine the IP domain, IP addresses, IP address ranges, and host names for discovery.

### NOTE

You can also create discovery profiles by way of NetOps Portal. For more information, see [Create Discovery Profiles](#).

### Determine the IDs

Discovery profiles require a tenant ID, an IP domain ID, and an SNMP profile ID.

### TIP

You can pull up the REST-based URLs for the data aggregator for this procedure in a browser tab instead of using a REST API client.

Follow these steps:

1. Find the Default Tenant's ID by issuing a GET request to the `tenants` endpoint for the Data Aggregator REST service:

```
http://<DA_host>:8581/rest/tenants
```

A list of tenants from the data aggregator is shown in the response.

2. Make a note of the Default Tenant's ID (`daTenantId`).

#### Example:

```
<Tenant version="0.0.0">
<ID>1</ID>
<ChangedOn>1568991597</ChangedOn>
<Name>Default Tenant</Name>
<Activated>true</Activated> <Culture>en-US</Culture>
<Description>The default tenant area</Description>
<AccountID/>
<Theme>CA-Blue or CA-White</Theme>
<Flags>0</Flags>
<IsAlso>
<IsA name="SyncableTenant" rootURL="syncabletenant"/>
<IsA name="DataAggregatorTenantMetering" rootURL="datenantmetering"/>
</IsAlso>
</Tenant>
```

In this example, the ID for the `daTenantId` is 1.

3. Find the Default Domain's IP domain ID by issuing a GET request to the `ipdomains` endpoint for the Data Aggregator REST service:

```
http://<DA_host>:8581/rest/ipdomains
```

### TIP

You can also scope the list by tenant.

A list of IP domains is shown in the response.

4. From the list, make a note of the Default Domain's IP domain ID (`aIpDomainId`).

#### Example:

```
<IPDomain version="0.0.0">
<ID>2</ID> <Status>ACTIVE</Status>
```

```

<ChangedOn>1568991597</ChangedOn>
<Flags>0</Flags>
<Name>Default Domain</Name>
<DNSProxyEnabled>false</DNSProxyEnabled>
<Description>The default domain for devices, interfaces, interface addresses and networks.</Description>
<IsAlso>
<IsA name="SyncableIPDomain" rootURL="syncableipdomain"/>
</IsAlso>
</IPDomain>

```

In this example, the ID for the `daIpDomainId` is 2.

5. Find the Default SNMP profile's ID using the one of the following options based on the SNMP profile type:
  - For SNMPv1 and SNMPv2c profiles, issue a GET request to the `snmpv1` endpoint for the Data Aggregator REST service:
 

```
http://<DA_host>:8581/rest/profiles/snmpv1
```

 A list of SNMPv1 and SNMPv2c profiles appears.
  - For SNMPv3 profiles, issue a GET request to the `snmpv3` endpoint for the Data Aggregator REST service:
 

```
http://<DA_host>:8581/rest/profiles/snmpv3
```

 A list of SNMPv3 profiles appears.
6. From the list, make a note of the Default SNMP profile's ID (`defaultSNMPProfileId`).

**Example:**

```

<SNMPv1Profile version="1.0.0">
<ID>1254</ID>
<CommunityName>QUVTOv2rU7kbjNnWycBidE9etgw= </CommunityName>
<PortNumber>161</PortNumber>
<IsAlso>
<IsA name="CommunicationProfile" rootURL="profiles"/>
<IsA name="Syncable" rootURL="syncable"/>
</IsAlso>
<CommunicationProfile version="1.0.0">
<UseForWrite>false</UseForWrite>
<Rank>1</Rank>
<ProfileName>default</ProfileName>
</CommunicationProfile>
<CommunicationFailurePolicy version="1.0.0">
<Retries>2</Retries>
<Timeout>3000</Timeout>
</CommunicationFailurePolicy>
</SNMPv1Profile>

```

In this example, the ID for the `defaultSNMPProfileId` is 1254.

## Create the Discovery Profile

The following example uses the default options.

**Follow these steps:**

1. Create a discovery profile for the Default Tenant by issuing a POST request to the `discoveryprofiles` endpoint for the Data Aggregator REST service:

```
http://<DA_host>:8581/rest/tenant/<daTenantId>/discoveryprofiles
```

For example, using the `daTenantId` reference value of 1:

```
http://<DA_host>:8581/rest/tenant/1/discoveryprofiles
```

**Body (successful result):**

```

<DiscoveryProfile version="1.0.0">
<Item version="1.0.0">
<Name>Paris</Name>
</Item><!-- IPRangesList can be omitted if empty -->
<IPRangesList>
<IPRanges>1.1.1.1-2.2.2.2</IPRanges>
</IPRangesList><!-- IPListList can be omitted if empty -->
<IPListList>
<IPList>1.1.1.1</IPList>
<IPList>2.2.2.2</IPList>
</IPListList><!-- HostNamesList can be omitted if empty -->
<HostNamesList>
<HostNames>HostA</HostNames>
<HostNames>HostB</HostNames>
</HostNamesList><!-- If omitted, use all SNMP Profiles in Tenant -->
<SNMPPProfileIDList><!-- These ID's are DA Item ID's for the profiles -->
<SNMPPProfileID>defaultSNMPPProfileId</SNMPPProfileID>
</SNMPPProfileIDList><!-- ActivationStatus is required (true/false) -->
<ActivationStatus>true</ActivationStatus><!-- IcmpDiscoveryEnabled is optional -->
<IcmpDiscoveryEnabled>true</IcmpDiscoveryEnabled>
<IPDomainMember version="1.0.0"><!-- DA's IP Domain ID -->
<IPDomainID>daIpDomainId</IPDomainID>
</IPDomainMember>
</DiscoveryProfile>

```

### Sample Body

This is a sample Body using the IDs that you retrieved. This example creates a discovery profile called Paris. This discovery profile tries to discover 2 IPs and has **Use ICMP** enabled (the discovered devices can respond to ICMP).

```

<DiscoveryProfile version="1.0.0">
<Item version="1.0.0">
<Name>Paris</Name>
</Item>
<IPListList>
<IPList>1.2.3.4</IPList>
<IPList>5.6.7.8</IPList>
</IPListList>
<SNMPPProfileIDList>
<SNMPPProfileID>1254</SNMPPProfileID>
</SNMPPProfileIDList>
<ActivationStatus>true</ActivationStatus>
<IcmpDiscoveryEnabled>true</IcmpDiscoveryEnabled>
<IPDomainMember version="1.0.0">
<IPDomainID>2</IPDomainID>
</IPDomainMember>
</DiscoveryProfile>

```

The discovery profile is created.

## Automate Running Discovery

Inventory discovery finds devices as specified by a discovery profile. You can automate running discovery by issuing requests to the Data Aggregator web services.

Device discovery finds devices as specified by a discovery profile. You can automate running device discovery by issuing requests to the Data Aggregator REST web services. With this method of running discovery, you can export and modify the results data in formats such as text files or Excel spreadsheets.

In this article:

- [Run Device Discovery Using REST](#)
- [View Device Discovery Profile History Using REST](#)
- [Run Device Rediscovery Using REST](#)

#### NOTE

You can also run device discovery and rediscovery manually or on a schedule using NetOps Portal.

For more information about these other methods of running discovery, see [Run Device Discovery](#) and [Run Device Rediscovery](#).

### Run Device Discovery Using REST

Follow these steps:

1. Get the Default Tenant's discovery profile ID by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/discoveryprofiles
```

#### NOTE

If you are using a multi-tenant configuration, specify the tenant's ID in the call:

```
http://<DA_host>:8581/rest/tenant/<tenantId>/discoveryprofiles
```

- **tenantId**

Specifies the tenant's ID.

If you do not know the Tenant ID, see [determine the ID](#).

A list of discovery profiles is shown in the response.

2. Make a note of the Default Tenant's discovery profile ID (`<daDiscoveryProfileId>`).

#### Example:

```
<DiscoveryProfile version="1.0.0">
  <ID>132614</ID>
  <IcmpDiscoveryEnabled>true</IcmpDiscoveryEnabled>
  <IPListList>
    <IPList>1.2.3.4</IPList>
    <IPList>5.6.7.8</IPList>
  </IPListList>
  <SNMPProfileIDList>
    <SNMPProfileID>1254</SNMPProfileID>
  </SNMPProfileIDList> <CreatePingables>true</CreatePingables>
  <UseListOfSnmpProfiles>false</UseListOfSnmpProfiles>
  <RunStatus>READY</RunStatus>
  <ActivationStatus>true</ActivationStatus>
  <IsAlso>
    <IsA name="IPDomainMember" rootURL="ipdomainmember"/>
  </IsAlso>
  <IPDomainMember version="1.0.0">
    <IPDomainID>2</IPDomainID>
  </IPDomainMember>
  <Item version="1.0.0">
    <Name>Paris</Name>
```



```
<CreateTime>Thu May 28 9:57:13 2020 -0400</CreateTime>
</Item>
</DiscoveryProfile>
```

In this example, the Default Tenant's discovery profile ID is 132614.

3. Start the discovery profile and run it by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **PUT** for the **HTTP Method**, with the following content:

```
http://<DA_host>:8581/rest/tenant/<daTenantId>/discoveryprofiles/<daDiscoveryProfileId>
```

– **discoveryProfileId**

Specifies the Default Tenant's discovery profile ID that you identified.

**Example:**

```
http://<DA_host>:8581/rest/tenant/1/discoveryprofiles/132614
```

**Body:**

```
<DiscoveryProfile version="1.0.0">
<RunStatus>START</RunStatus>
</DiscoveryProfile>
```

You have ran device discovery.

### **View Device Discovery Profile History Using REST**

You can view device discovery profile history (discovery instance) entries using REST.

Use the following process to view device discovery profile history using REST:

1. [Determine the IDs for the discovery profile and \(if you are using a multi-tenant configuration\) the tenant.](#)
2. [View the discovery profile history.](#)

### **Determine the IDs**

**TIP**

You can pull up the REST-based URLs for this procedure in a browser tab instead of using a REST API client.

### **Follow these steps:**

1. (Optional) If you are using a multi-tenant configuration, get the tenant's ID by completing the following steps:
  - a. Find the tenant's ID by entering the following URL for the `tenants` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/tenants
```

A list of tenants from the data aggregator is shown in the response.

- b. Make a note of the tenant's ID (`tenantId`).

**Example:**

```
<Tenant version="0.0.0">
<ID>8</ID>
<ChangedOn>1568991597</ChangedOn>
<Name>Tenant</Name>
<Activated>true</Activated>
<Culture>en-US</Culture>
<Description>The default tenant area</Description>
<AccountID/>
<Theme></Theme>
<Flags>0</Flags>
<IsAlso>
<IsA name="SyncableTenant" rootURL="syncabletenant"/>
<IsA name="DataAggregatorTenantMetering" rootURL="datenantmetering"/>
```

```
</IsAlso>
</Tenant>
```

In this example, the ID is 8.

2. Get the discovery profile's ID by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/discoveryprofiles
```

#### NOTE

If you are using a multi-tenant configuration, specify the tenant's ID in the call:

```
http://<DA_host>:8581/rest/tenant/<tenantId>/discoveryprofiles
```

- **tenantId**

Specifies the tenant's ID.

A list of discovery profiles is shown in the response. Each represents a history entry for the discovery profile that is in NetOps Portal.

3. Make a note of the discovery profile's ID value (`discoveryProfileId`).
4. Test to validate that you selected the correct discovery profile ID by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/discoveryprofiles/<discoveryProfileId>
```

- **discoveryProfileId**

Specifies the discovery profile's ID that you identified.

#### NOTE

If you are using a multi-tenant configuration, specify the tenant's ID in the call:

```
http://<DA_host>:8581/rest/tenant/<tenantId>/discoveryprofiles/<discoveryProfileId>
```

- **tenantId**

Specifies the tenant's ID.

The discovery profile is shown in the response.

## View a Discovery Profile History

### Follow these steps:

1. List the details for the discovery profile by entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/discoveryprofiles/<discoveryProfileId>
```

- **discoveryProfileId**

Specifies the discovery profile's ID that you identified.

#### NOTE

If you are using a multi-tenant configuration, specify the the tenant's ID in the call:

```
http://<DA_host>:8581/rest/tenant/<tenantId>/discoveryprofiles/<discoveryProfileId>
```

- **tenantId**

Specifies the tenant's ID.

The details of the discovery profile are shown in the response. The `DiscoveryInstanceIdList` shows a list of IDs, one for each discovery instance (`discoveryInstancesId`). The higher the number, the newer the entry and the lower the number, the older the entry. `MostRecentInstance` shows the most-recent discovery run results data.

2. (Optional) Get the details for a given discovery instance ID entry by entering the following URL for the `discoveryinstances` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/discoveryinstances/<discoveryInstancesId>
```

- **discoveryInstancesId**

Specifies the discovery instance's ID that you identified.

**NOTE**

If you are using a multi-tenant configuration, specify the tenant's ID in the call:

```
http://<DA_host>:8581/rest/tenant/<tenantId>/discoveryInstances/<discoveryInstancesId>
```

- **tenantId**

Specifies the tenant's ID.

The discovery instance details are shown in the response.

**Sample Output:**

The following is sample REST output for the Default Domain discovery profile that the DX NetOps Spectrum integration often creates using the `discoveryProfileId` ID value (1340):

**NOTE**

This output is from entering the following URL for the `discoveryprofiles` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

```
http://<DA_host>:8581/rest/discoveryprofiles/1340
```

```
<DiscoveryProfile version="1.0.0">
<ID>1340</ID>
<IcmpDiscoveryEnabled>true</IcmpDiscoveryEnabled>
<MostRecentInstance>156692</MostRecentInstance>
<CreatePingables>true</CreatePingables>
<UseListOfSnmpProfiles>false</UseListOfSnmpProfiles>
<RunStatus>SCHEDULED</RunStatus>
<RelatesTo>
<DiscoveryInstanceIDList relatesURL="relatesto/instances" rootURL="discoveryinstances">
<ID>156483</ID>
<ID>156692</ID>
<ID>156690</ID>
<ID>156663</ID>
<ID>155064</ID>
<ID>155112</ID>
<ID>156620</ID>
<ID>155145</ID>
<ID>155088</ID>
<ID>156624</ID>
</DiscoveryInstanceIDList>
<ScheduleIDList relatesURL="relatesto/schedule" rootURL="schedule">
<ID>1341</ID>
</ScheduleIDList>
</RelatesTo>
<IsAlso>
<IsA name="IPDomainMember" rootURL="ipdomainmember"/>
<IsA name="SystemCreatedDiscoveryProfile" rootURL="systemcreateddiscoveryprofiles"/>
</IsAlso>
<IPDomainMember version="1.0.0">
<IPDomainID>2</IPDomainID>
</IPDomainMember>
<Item version="1.0.0">
<Name>Default Domain</Name>
<CreateTime>Wed Oct 2 18:13:04 2019 -0400</CreateTime>
</Item>
</DiscoveryProfile>
```

The following is sample REST output for a discovery instance entry using the `MostRecentInstance` ID value (156692).

#### NOTE

This output is from entering the following URL for the `discoveryinstances` endpoint for the Data Aggregator RESTful web services API, selecting **GET** for the **HTTP Method**:

`http://<DA_host>:8581/rest/discoveryinstances/156692`

```
<DiscoveryInstance version="1.0.0">
<ID>156692</ID>
<CachedProgressPercentage>100</CachedProgressPercentage>
<TestedCommProfilesList>
<TestedCommProfiles>1254</TestedCommProfiles>
<TestedCommProfiles>1339</TestedCommProfiles>
</TestedCommProfilesList>
<ProgressPercentage>100</ProgressPercentage>
<StartTime>Sun Jul 26 20:00:00 2020 -0400</StartTime>
<CompletionTime>Sun Jul 26 20:00:00 2020 -0400</CompletionTime>
<CompletionStatus>SUCCESS</CompletionStatus>
<ProfileID>1340</ProfileID>
<IPSweepTotalSuccess>0</IPSweepTotalSuccess>
<NewAccessibleOnlyButDeletedDevicesCount>0</NewAccessibleOnlyButDeletedDevicesCount>
<Item version="1.0.0">
<CreateTime>Sun Jul 26 20:00:00 2020 -0400</CreateTime>
</Item>
</DiscoveryInstance>
```

The following is a sample of a REST output for a discovery instance entry (78325) where it found a new device (`NewlyCreatedDevices`):

```
<DiscoveryInstance version="1.0.0">
<ID>78325</ID>
<CachedProgressPercentage>100</CachedProgressPercentage>
<TestedCommProfilesList>
<TestedCommProfiles>1254</TestedCommProfiles>
<TestedCommProfiles>1339</TestedCommProfiles>
</TestedCommProfilesList>
<NewlyCreatedManageableDevicesList>
<NewlyCreatedManageableDevices>78327</NewlyCreatedManageableDevices>
</NewlyCreatedManageableDevicesList>
<ProgressPercentage>100</ProgressPercentage>
<StartTime>Tue Mar 24 12:42:26 2020 -0400</StartTime>
<CompletionTime>Tue Mar 24 12:42:38 2020 -0400</CompletionTime>
<CompletionStatus>SUCCESS</CompletionStatus>
<ProfileID>3986</ProfileID>
<IPSweepTotalSuccess>1</IPSweepTotalSuccess>
<NewlyCreatedDevicesList>
<NewlyCreatedDevices>78327</NewlyCreatedDevices>
</NewlyCreatedDevicesList>
<NewAccessibleOnlyButDeletedDevicesCount>0</NewAccessibleOnlyButDeletedDevicesCount>
<Item version="1.0.0">
<CreateTime>Tue Mar 24 12:42:26 2020 -0400</CreateTime>
</Item>
```

```
</DiscoveryInstance>
```

## Run Device Rediscovery Using REST

You can help lower time spent administering device inventory updates using NetOps Portal by automating or scheduling REST calls for device-by-device rediscovery, instead of running a full discovery profile.

Run rediscovery for a single device by accessing the Data Aggregator REST web services using the following cURL statement:

```
curl -v -X POST -H 'Content-type: application/xml' -d '<root/>' http://<DA_host>:8581/genericWS/devices/<deviceId>/inventorydiscovery
```

## Get a List of Unreachable Devices

If discovery fails for a device, you can get a list of the unreachable devices using the `discoveryinstances` endpoint for the Data Aggregator REST service.

**Prerequisite:** You have specified individual IP addresses or host names in your discovery profile.

### Follow these steps:

1. View a list of the discovery instances by issuing a request to the following `discoveryinstances` endpoint for the Data Aggregator REST service:

```
http://<DA_host>:8581/rest/discoveryinstances
```

2. Find the discovery instance that corresponds to when you ran your discovery.

#### Example:

```
<DiscoveryInstance version="1.0.0">
<ID>6508</ID>
<IPSweepStartTime>Mon Sep 18 15:44:31 2017 -0400</IPSweepStartTime>
```

3. Note the discovery instance ID.
4. View the results for your discovery instance by issuing a request to the following `discoveryinstances` endpoint for the Data Aggregator REST service. Replace `<discoveryInstancesId>` with the discovery instance's ID:

```
http://<DA_host>:8581/rest/discoveryinstances/<discoveryInstancesId>
```

5. View the `<UnreachableDevicesList>` section.

#### Example:

```
<DiscoveryInstance version="1.0.0">
<ID>6508</ID>
<IPSweepStartTime>Mon Sep 18 15:44:31 2017 -0400</IPSweepStartTime>
<TestedCommProfilesList>
<TestedCommProfiles>4657</TestedCommProfiles>
<TestedCommProfiles>4603</TestedCommProfiles>
<TestedCommProfiles>4656</TestedCommProfiles>
</TestedCommProfilesList>
<ProgressPercentage>100</ProgressPercentage>
<PingResponseDeviceCount>0</PingResponseDeviceCount>
<StartTime>Mon Sep 18 15:44:31 2017 -0400</StartTime>
<CompletionTime>Mon Sep 18 15:44:39 2017 -0400</CompletionTime>
<IPSweepCompletionTime>Mon Sep 18 15:44:39 2017 -0400</IPSweepCompletionTime>
<CompletionStatus>SUCCESS</CompletionStatus>
<ProfileID>6507</ProfileID>
<IPSweepTotalSuccess>0</IPSweepTotalSuccess>
```

```

<UnreachableDevicesList>
<UnreachableDevices>10.255.255.255</UnreachableDevices>
<UnreachableDevices>192.168.0.0</UnreachableDevices>
<UnreachableDevices>172.31.255.255</UnreachableDevices>
<UnreachableDevices>192.168.255.255</UnreachableDevices>
<UnreachableDevices>10.0.0.0</UnreachableDevices>
<UnreachableDevices>169.254.0.0</UnreachableDevices>
<UnreachableDevices>169.254.255.255</UnreachableDevices>
<UnreachableDevices>172.16.0.0</UnreachableDevices>
</UnreachableDevicesList>
<NewAccessibleOnlyButDeletedDevicesCount>0</NewAccessibleOnlyButDeletedDevicesCount>
<Item version="1.0.0">
<CreateTime>Mon Sep 18 15:44:31 2017 -0400</CreateTime>
</Item>
</DiscoveryInstance>

```

## Automate Configuring Aggregated Components

You can automate managing (creating, editing, and deleting) aggregated components using a Data Aggregator REST web services.

Administrators can manage aggregated components by way of the `aggregatedcomponents` endpoint for the Data Aggregator REST web service or use this API in your scripts for managing aggregated components or using NetOps Portal. This topic describes how to automate aggregated component management by way of the web service.

For more information about how to manage them using NetOps Portal, see [Manage Aggregated Components](#).

In this article:

- [Create an Aggregated Component](#)
- [Add or Remove Members to and from an Aggregated Component](#)
- [Delete an Aggregated Component](#)
- [Configure the Aggregated Items Processor](#)
- [Configure Aggregated Items ETL](#)

### Create an Aggregated Component

Follow these steps:

1. Set up a REST client with a connection to the Data Aggregator server.
2. Issue the following POST request to the `aggregatedcomponents` endpoint, with the following payload:

```
http://<DA_host>:8581/rest/aggregatedcomponents
```

#### Payload Example 1:

The following payload creates an aggregated component containing two interfaces (interface IDs 1751 and 1752 , with metric family `NormalizedPortInfo`) from different devices.

```

<AggregatedComponent version="1.0.0">
<AggregateNormalizedFacetType>{http://im.ca.com/
normalizer}NormalizedAggregatedPortInfo</AggregateNormalizedFacetType>
<MemberNormalizedFacetType>{http://im.ca.com/normalizer}NormalizedPortInfo</
MemberNormalizedFacetType>
<AggregatedType>interface</AggregatedType>
<MemberItemIDsList>

```

```

    <MemberItemIDs>1751</MemberItemIDs>
    <MemberItemIDs>1752</MemberItemIDs>
  </MemberItemIDsList>
  <Item version="1.0.0">
    <Name>An aggregated interface</Name>
    <Description>An aggregated interface</Description>
  </Item>
</AggregatedComponent>

```

### Payload Example 2:

The following payload creates an aggregated component containing three members (interfaces with IDs 1851 , 1852 , and 1853 , with metric family NormalizedPortInfo ) from the same device (device ID 1489 ).

```

<AggregatedComponent version="1.0.0">
  <AggregateNormalizedFacetType>{http://im.ca.com/
normalizer}NormalizedAggregatedPortInfo</AggregateNormalizedFacetType>
  <MemberNormalizedFacetType>{http://im.ca.com/normalizer}NormalizedPortInfo</
MemberNormalizedFacetType>
  <AggregatedType>interface</AggregatedType>
  <MemberItemIDsList>
    <MemberItemIDs>1851</MemberItemIDs>
    <MemberItemIDs>1852</MemberItemIDs>
    <MemberItemIDs>1853</MemberItemIDs>
  </MemberItemIDsList>
  <Item version="1.0.0">
    <Name>An aggregated interface</Name>
    <Description>An aggregated interface</Description>
  </Item>
  <DeviceItemID>1489</DeviceItemID>
</AggregatedComponent>

```

#### NOTE

The values for AggregateNormalizedFacetType , MemberNormalizedFacetType , AggregatedType , and MemberItemIDs must match the component type that you are creating.

#### – <AggregateNormalizedFacetType>

The type of aggregated metric family for the aggregated component (interface, CPU, or memory).

##### Values:

- NormalizedAggregatedPortInfo: For interface components.
- NormalizedAggregatedCPUInfo: For CPU components.
- NormalizedAggregatedMemoryInfo: For memory components.

#### – <MemberNormalizedFacetType>

The type of metric family for the aggregated component member (interface, CPU, or memory).

##### Values:

- NormalizedPortInfo: For interface components.
- NormalizedCPUInfo: For CPU components.
- NormalizedMemoryInfo: For memory components.

#### – <AggregatedType>

The type of component.

**Values:** interface , cpu , memory

#### – <MemberItemIDs>

The ID of an aggregated component member (interface, CPU, or memory) to include in the aggregated component that you are creating.

– **<DeviceItemID>**

The ID of the device to which the aggregated component member belongs.

**Required:** No

The data aggregator creates the aggregated component and assigns it an ID.

### **Add or Remove Members to and from an Aggregated Component**

Issue the following PUT request to the `aggregatedcomponents` endpoint, specifying the ID of the aggregated component that you want to edit, with the following payload:

```
http://<DA_host>:8581/rest/aggregatedcomponents/<ID_of_aggregated_component>
```

#### **Example:**

The following example edits aggregated component 2000 .

```
http://<DA_host>:8581/rest/aggregatedcomponents/2000
```

- ***ID\_of\_aggregated\_component***

Defines the ID for the aggregated component that you want to edit.

#### **Payload Example:**

The following payload example adds a member (interface with ID 1753 ) to an aggregated component (aggregated component with ID 2000 ) with aggregated component members 1751 , 1752 , and 1753 :

#### **NOTE**

Any existing aggregated component members that you do not list in the payload are removed from the aggregated component.

```
<AggregatedComponent version="1.0.0">
  <MemberItemIDsList>
    <MemberItemIDs>1751</MemberItemIDs>
    <MemberItemIDs>1752</MemberItemIDs>
    <MemberItemIDs>1753</MemberItemIDs>
  </MemberItemIDsList>
</AggregatedComponent>
```

- ***<MemberItemIDs>***

The ID of a member (interface, CPU, or memory) to include in the aggregated component that you are editing.

### **Delete an Aggregated Component**

Issue the following DELETE request to the `aggregatedcomponents` endpoint, specifying the ID of the aggregated component that you want to delete:

```
http://<DA_host>:8581/rest/aggregatedcomponents/<ID_of_aggregated_component>
```

#### **Example:**

The following example deletes aggregated component 2000 .

```
http://<DA_host>:8581/rest/aggregatedcomponents/2000
```

- ***ID\_of\_aggregated\_component***

Defines the ID for the aggregated component that you want to delete.



## Configure the Aggregated Items Processor

The aggregated items processor aggregates metric data, by default, every 5 minutes. If data for all aggregated component members is not available, the aggregated items processor reschedules the aggregation for this member, by default, 2 times. After the last attempt to aggregate data, the aggregated items processor aggregates the data even if the data for all members is not available, and calculates the value of the `DataCompleteness` metric. This metric shows data completeness.

### NOTE

The `NormalizedAggregatedCPUInfo`, `NormalizedAggregatedPortInfo`, and `NormalizedAggregatedMemoryInfo` aggregated metric families include the `DataCompleteness` metric.

### Follow these steps:

1. Set up a REST client with a connection to the data aggregator server.
2. Find the ID for the aggregated items processor by issuing a GET request to the `aggregateditemsprocessor/config` endpoint:

```
http://<DA_host>:8581/rest/aggregateditemsprocessor/config
```

The configuration for the aggregated items processor is shown in the response.

### Example response:

```
<AggregatedItemsProcessorConfigurationList>
  <AggregatedItemsProcessorConfiguration version="1.0.0">
    <ID>1392</ID>
    <MaxNumberOfAggregationRescheduling>2</MaxNumberOfAggregationRescheduling>
    <MinAggregationDelay>1</MinAggregationDelay>
    <MaxNumberOfConcurrentAggregations>5</MaxNumberOfConcurrentAggregations>
    <MinAggregationRunInterval>5</MinAggregationRunInterval>
  </AggregatedItemsProcessorConfiguration>
</codeblockAggregatedItemsProcessorConfigurationList>
```

In this example, the ID for the aggregated items processor is 1392.

3. Make note of the ID.
4. Issue the following PUT request, specifying the ID for the aggregated items processor, with the following payload:

```
http://<DA_host>:8581/rest/aggregateditemsprocessor/config/<ID_of_aggregated_items_processor>
```

### Example:

```
http://DA_host:8581/rest/aggregateditemsprocessor/config/1392
```

### – `ID_of_aggregated_component`

Defines the ID for the aggregated component that you want to configure.

### Payload Example:

The following example sets the `MaxNumberOfAggregationRescheduling` to 1.

```
<AggregatedItemsProcessorConfiguration version="1.0.0">
  <MaxNumberOfAggregationRescheduling>1</MaxNumberOfAggregationRescheduling>
  <MinAggregationDelay>1</MinAggregationDelay>
  <MaxNumberOfConcurrentAggregations>5</MaxNumberOfConcurrentAggregations>
  <MinAggregationRunInterval>5</MinAggregationRunInterval>
</AggregatedItemsProcessorConfiguration>
```

### – `MaxNumberOfAggregationRescheduling`

Defines the maximum number of times the aggregated items processor attempts to aggregate data when data for all aggregated component members is not available during the `MinAggregationRunInterval`. For example, you have an aggregated component with two CPU members, and you have set this parameter to 1, if data for one of the CPU members is not available, the aggregated items processor waits one `MinAggregationRunInterval` before it aggregates data.

**Default: 2**

– **MinAggregationDelay**

Defines the amount of time before the aggregated items processor attempts to aggregate metric data. The amount of time is calculated using the following formula:

$\text{MinAggregationRunInterval} * \text{MinAggregationDelay}$

**Default: 1**

**Values:**

- **0:** The aggregated items processor attempts to aggregate data immediately after it receives notification of End of Cycle (EOC).

**IMPORTANT**

Setting the value to 0 can increase the load on the data aggregator, as data for all aggregated component members might not be available.

- **1:** The aggregated items processor waits for the `MinAggregationDelay` after the EOC before attempting to aggregate the data.
- **2:** The aggregated items processor waits for  $2 * \text{MinAggregationDelay}$  after the EOC before attempting to aggregate the data.

For example, if `MinAggregationRunInterval` is set to 5, and `MinAggregationDelay` is set to 2, then the aggregated items processor waits 10 minutes.

– **MaxNumberOfConcurrentAggregations**

Defines the maximum number of aggregation processes that the aggregated items processor can run concurrently in the data aggregator, usually one aggregation process per specified aggregated metric family.

**Default: 5**

– **MinAggregationRunInterval**

Defines the minimum rate (in minutes) at which the aggregated items processor aggregates metric data. This rate is the same for all aggregated metric families.

**Default: 5**

## **Configure Aggregated Items ETL**

The aggregated items ETL creates the `v_dim_aggregate_item` view that contains aggregated items. If you are not using aggregated components, you can turn the aggregated items ETL off (disabled).

### **Follow these steps:**

1. Find the ID for the aggregated items ETL by issuing a GET request to the `dimaggregateditems` endpoint:

`http://<DA_host>:8581/rest/batch/dimaggregateditems/config`

The configuration for the aggregated items processor is shown in the response.

**Example response:**

```
<DimAggregatedItemsBatchConfigurationList>
  <DimAggregatedItemsBatchConfiguration version="1.0.0">
    <ID>1401</ID>
    <Enabled>true</Enabled>
    <Interval>5</Interval>
  </DimAggregatedItemsBatchConfiguration>
</DimAggregatedItemsBatchConfigurationList>
```

In this example, the ID for the aggregated items ETL is 1401 .

2. Make note of the ID.
3. Issue the following PUT request, specifying the ID for the aggregated items ETL, with the following payload:

`http://<DA_host>:8581/rest/batch/dimaggregateditems/config/1401`

**Payload Example:**

```
<DimAggregatedItemsBatchConfiguration version="1.0.0">
  <Enabled>true</Enabled>
  <Interval>5</Interval>
</DimAggregatedItemsBatchConfiguration>
```

- **Enabled**  
Specifies if the aggregated items ETL is enabled (true) or disabled (false).  
**Default:** true
- **Interval**  
Sets how often (in minutes) the aggregated items ETL should run, similar as any other ETLs.  
**Default:** 5

## OpenAPI

The OpenAPI is a flexible tool that you can use to extract data from the DX NetOps Performance Management database. The OpenAPI enables integration between DX NetOps Performance Management data and external applications.

This public API enables integration between DX NetOps Performance Management data and external applications. It uses the QueryBuilder GUI. You can extract and explore performance data by creating custom Query URLs using the QueryBuilder. These URLs return customized data in the specified format. You can view the data in a browser or process the data in a custom web application.

The OpenAPI uses the OData 2.0 industry standard.

For more information about this standard, see [the OData 2.0 documentation](#).

## Use the OpenAPI QueryBuilder

The OpenAPI QueryBuilder creates query URLs that export configuration and polled data.

All users who have NetOps Portal credentials can use QueryBuilder and the OpenAPI. The OpenAPI uses Single Sign-On (SSO) for credential authentication.

### NOTE

OpenAPI SSO authentication requires access to NetOps Portal and SSO ports. If the ports are blocked for the listed host name or IP address, update the URL host for the device manager service.

For more information, see [Configure the Web Service Security Settings Using the SSO Configuration Tool](#).

In this article:

- [Access the QueryBuilder](#)
- [Create an OpenAPI Query](#)
- [OpenAPI Controls](#)
- [Find Metric Family Names for QueryBuilder](#)
- [Best Practices for OpenAPI Performance](#)

### Access the QueryBuilder

**Follow these steps:**

1. Go to the following URL:

```
http://da_hostname:8581/odataquery
```

In a fault-tolerant data aggregator environment, specify the proxy server as the *da\_hostname*.

2. Log in using your NetOps Portal credentials.

**NOTE**

You cannot run OpenAPI queries outside the QueryBuilder as a user with Security Assertion Markup Language (SAML) 2.0 authentication. SAML authentication works in the QueryBuilder UI. LDAP authentication and NetOps Portal authentication support direct OpenAPI queries.

The schema XML description provides detailed information about the items and relationships in your system. To review the schema XML description and metadata, go to the following URL:

`http://da_hostname:8581/odata/api/$metadata`

**Create an OpenAPI Query**

Extract customized data sets from the data source using the QueryBuilder. Generate OpenAPI queries.

**Follow these steps:**

1. Click the **Query Expression** field and start a query.  
The tokens appears.

**NOTE**

The Web browser address bar updates to show the tokens that you add to the **Query Expression** field. To continue editing the query, copy and save this URL.

2. In the **Query Expression** field, do the following:
  - Define what the item type that you want included the data set, such as device, by adding the `for` token.
  - In the **Query Expression** field, define the output data by adding the `select`, `expand metrics`, and `expand` tokens, and then refine the results by adding more tokens.

**WARNING**

Queries that return large sets of results can negatively affect your system. Ensure that they return only the results that are relevant to your needs.

For more information about the tokens that you can add, see [OpenAPI Controls](#).

QueryBuilder creates the OData URL.

3. Do one of the following:
  - To run the query in the QueryBuilder browser window, click **Run**.
  - To run the query from a Web browser, REST tool, or custom application, copy the OData URL.

**OpenAPI Controls**

The QueryBuilder uses tokens that represent logical elements in the OpenAPI query syntax. The tokens define the type of attribute, filter, metric family, metric, or time range for the query. Adding tokens updates the Query URL.

Use the following tokens to build the query:

- **for**

This token determines the type of item that the query retrieves data about and includes in the data set. This token appears only once in each query.

Most options for this token include selections that set up an automatic filter. For example, when you select **interface** in the token, you can select **within Groups that...** This option adds a filter token with the group filter selected as the first criterion.

**NOTE**

If a newly-added metric family does not appear in the QueryBuilder, refresh the Web browser.

- **select**

This token determines the item properties to include in the data set.

**NOTE**

Property order is not supported in the results.

- **expand metrics**

This token determines the performance metrics to include in the data set.

**NOTE**

Property order is not supported in the results.

- **filter**

This token adds custom filters that are based on logical functions using the **AND** and **OR** operations. Select whether the filter is case-sensitive for all logical functions.

When OData evaluates a filter expression, the **Any** operator is used to determine whether the Boolean expression is True or False for a collection of items. The following examples return True:

```
odata/api/devices?$select=ID,Name&$filter=((startswith(Name, 'cisco') eq true))
```

```
odata/api/cpus?$top=10&$filter=cpumfs/im_Utilization gt 80 where cpumfs/im_Utilization is the collection [60,65,81,68,61]
```

**NOTE**

In large systems, some queries might time out if the filters are applied inefficiently. OpenAPI evaluates filter expressions in the order they appear in the query. Use the filter that limits the data more first. For example, if you want all groups where interface utilization is great than 50 percent in North America, place the North America filter first.

- **filter metrics**

This token adds custom filters that are based on the selected metrics.

- **time range**

This token limits the time range of the results for performance metrics and sets the granularity of the data.

**NOTE**

If the time granularity is set to Day and the Time Zone is set to any value other than the default, the **Start Time** and **End Time** fields are disabled.

- **expand**

This token adds data that is related to the base item to the returned data set. For example, if you selected data for interface, use **expand** to get data about the related devices.

In the **expand** token, select which columns to add to the data set. For example, for device, you might select name, primary IP address, and location.

Each expanded attribute appears in the HTML table in the same format as a metric column. For example, to display interface information for a specific device, you can use the following query:

```
select=device/Name&$expand=device
```

- **group/aggregate**

This token defines grouping and aggregation for the returned data set. For example, aggregate CPU utilization to the device level, and average the utilization for the last hour.

**NOTE**

You can apply only the **groupby** and **aggregate** OData transformations. If the starting entity is a metric family, then you can apply the **groupby** transformation to **ID**, **DeviceItemID**, and **Timestamp**. If the starting entity is a config entity and aggregation is for metric family data, then you can apply the **groupby** transformation to **ID** and **DeviceItemID**. You can apply the **groupby** transformation on the group level using group ID. You can apply the **aggregate** transformation with up to five expressions.

- **sort**

This token controls the sorting of the query output.

You can sort using the following query:

```
/cpus?$orderby=Name
```

You can sort on the properties of the target entity of your query:

```
http://hostname:8581/odata/api/TargetEntity
```

For example, sorting of metric columns works best when the target entity of your query is the metric family.

- **limit (top)**

This token specifies the maximum number of rows or expanded rows in the query output.

- **Maximum number of rows**

The number of rows in the Results table.

- **Maximum number of expanded rows**

The number of rows in the Expand window.

This number only applies to the number of rows when the Expand token is used.

**NOTE**

To modify the default values or the maximum number of rows, see [Configure OpenAPI Defaults and Limits](#).

- **Skip**

The number of leading rows to omit from the query output

- **format**

This token determines the format of the returned data set. The OpenAPI supports the following formats:

- **HTML Table**

If the query does not include this token, HTML Table is the default format.

**NOTE**

The HTML table format is supported only in the OpenAPI QueryBuilder. Direct OpenAPI queries support JSON, XML, Atom, and CSV.

- **JSON**

**NOTE**

The JSON format removes the metadata from the resulting data set. Doing so reduces the network transfer time, the payload size up to 50-75 percent, and the client processing time. To add metadata to JSON output manually, use the following HTTP header:

```
Accept: application/json;odata=verbose
```

You can also add metadata to JSON output. To do so, add the \$format query option as the following text to the OData URL:

```
$format=application/json;odata=verbose
```

- **XML**

- **Atom**

- **CSV**

To specify CSV as a format manually, append the following text to the end of the OData URL:

```
$format=text/csv
```

- **custom parameter**

This token adds custom parameters in OData syntax to the OpenAPI query.

## **Find Metric Family Names for QueryBuilder**

The metric family names in the OpenAPI are shortened versions of the internal metric family names. The OpenAPI metric family names remove the prefix 'Normalized' and the suffix 'Info'. The capitalization is removed and the suffix 'mf' is added to the string. For example, the internal name of the Interface metric family is 'NormalizedPortInfo'. The OpenAPI metric family name is 'portmf'.

To see the internal metric family name, look up the name in NetOps Portal.

For more information, see [Manage Metric Families](#).

## Best Practices for OpenAPI Performance

The OpenAPI enables quick and flexible extraction of data from the database. For queries that produce large results, use the following best practices to improve performance:

- When you apply multiple filters to a query, ensure that all filters rules against the same object are adjacent. Splitting these filters reduces the efficiency of the query.
- Avoid filtering by strings. For example, to filter on a group, look up the group ID instead of filtering on the group name.
- Combine group expressions. Using an OR operator between a group expression and another expression is unsupported. For example, the following expression returns an error:  
`odata/api/interfaces?&$select=ID,Description&$filter=(groups/Name eq 'Manageable Devices') or (substringof('Gigabit',Description) eq true)`
- Aggregation functions against large data sets take a long time.
- OpenAPI is not a bulk data export tool. Use OpenAPI to extract only the data that you need.
- Use the top and skip values to apply paging for large data sets.
- Use the largest granularity that is relevant for your data set. For example, if you want data that is aggregated by sum, the hourly or daily values provide the same information as rate data.

## OpenAPI QueryBuilder Examples

The following examples highlight the flexibility of the OpenAPI. Use these examples as a model to create your own OpenAPI queries:

### Export Inventory and Configuration Details

The following examples show how to extract lists of items from inventory with specific configuration details.

#### Get a List of All Devices with Cisco in the Name

You want to get a list of Cisco devices and IP addresses of those devices. For these devices, you know that the device name includes "Cisco".

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name, PrimaryIPAddress
- **filter:** name contains cisco

#### OpenAPI URL:

```
http://da_hostname:8581/odata/api/devices?$select=ID,Name,PrimaryIPAddress&
$filter=((groups/Name eq 'cisco'))
```

#### Get a List of Devices and Location Information

You want a list of Cisco devices and the location information (location, latitude, longitude, location description, and elevation) for those devices.

#### NOTE

The `Location` attribute is tied to the MIB `sysLocation` value. The `Latitude`, `Longitude`, `LocationDesc`, and `Elevation` attributes are user-defined values. An administrator must set these values on the data aggregator. If these values are not set, they return a null result.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name, PrimaryIPAddress, Location, Latitude, Longitude, LocationDesc, Elevation
- **filter:** name contains cisco

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?
$select=ID,Name,PrimaryIPAddress,Location,Latitude,Longitude,LocationDesc,Elevation&
$filter=((groups/Name eq 'cisco'))
```

**Get a List of Devices and Components If the Device IP Address Begins with 10.251**

You want a list of all devices where the IP address begins with 10.251. You also want a list of the components for those devices.

**TIP**

Use the **expand** token to get items that are related to the main entity of the query.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name, PrimaryIPAddress
- **filter:** PrimaryIPAddress starts with 10.251
- **expand:** components, DisplayName

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$expand=components&
$select=ID,Name,PrimaryIPAddress,components/DisplayName&
$filter=((startswith(PrimaryIPAddress, '10.251') eq true))
```

**Get a List of Devices and Components in the Boston Group**

You want a list of all devices in the Boston group. You also want a list of the components for those devices.

**TIP**

When you selected a group filter, OpenAPI always includes subgroups of the selected group.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name, PrimaryIPAddress
- **filter:** groups/Name equal Boston
- **expand:** components, DisplayName

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$expand=components&
$select=ID,Name,PrimaryIPAddress,components/DisplayName&$filter=((groups/Name eq
'Boston'))
```

**Get a List of Metric Families and Vendor Certifications for a Specific Device**

You want to know which metric families and vendor certifications are used to monitor specific devices. From this list, you can derive the available metrics for a device. The name of the device is MyDevice-123.

Build the query by adding the following tokens:



- **for:** device
- **select:** ID, Name
- **filter:** DisplayName equal MyDevice-123
- **expand:** metricfamilyhistory, MetricFamilyID, VendorCertDisplayName

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$expand=metricfamilyhistories&
$select=ID,Name,metricfamilyhistories/MetricFamilyID,metricfamilyhistories/
VendorCertDisplayName&$filter=((DisplayName eq 'MyDevice-123'))
```

**Get a List of Supported Metric Families**

You want to know which metrics are available for reporting. From a list of metric families, you can determine which metric families are supported by which monitored devices.

Build the query by adding the following tokens:

- **for:** metricfamilyhistory
- **select:** Status
- **filter:** Status/Status equal SUPPORTED
- **expand:** device/ID, device/Name, metricfamilydef/ID, metricfamilydef/DisplayName, metricfamilydef/DisplayDescription, metricdefs/ID, metricdefs/DisplayName, metricdefs/DisplayDescription

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/metricfamilyhistories?
$expand=metricfamilydef,metricdefs,device&$select=device/ID,device/
Name,Status,metricfamilydef/ID,metricfamilydef/DisplayName,metricfamilydef/
DisplayDescription,metricdefs/ID,metricdefs/DisplayName,metricdefs/DisplayDescription&
$filter=((Status eq 'SUPPORTED'))
```

**Export Inventory and Monitoring Data**

The following examples show how to extract monitoring data for items with specific configuration details or monitoring status.

**List the Interfaces in the North America Group That Are Not Being Polled**

You want to know which interfaces in the North America group are not currently being monitored.

Build the query by adding the following tokens:

- **for:** interface
- **select:** ID, Name
- **filter:** groups/Name equal North America and IsPolled False

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/interfaces?$select=ID,Name&$filter=((groups/Name eq
'North America') and (IsPolled eq false))
```

### **Get the Polling Statistics for Devices in a Group for the Last 24 Hours**

You want to know the success and failure rate of devices in the Boston group for the last 24 hours. The database does not store the number of unsuccessful polls. To derive this value, subtract the number of successful polls from the number of polls sent using the **group/aggregate** token.

#### **TIP**

Time selections are absolute. When you select the time for the query, that time is fixed. To create a query with a flexible time, use a script to replace the time and data in the query.

Build the query by adding the following tokens:

- **for:** device
- **select:** Name
- **expand metrics:** devicepollingstatisticsmfs: im\_NumSuccessfulPolls
- **group/aggregate:**
  - **Group By:** devicepollingstatisticsmfs, ID
  - **Add an Aggregation Function:** Sum Of im\_NumPollsSent sub im\_NumSuccessfulPolls
- **filter:** groups/Name equal Boston
- **time range:** Last24Hours

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$apply=groupby(devicepollingstatisticsmfs/ID, aggregate(devicepollingstatisticsmfs(im_NumPollsSent sub im_NumSuccessfulPolls with sum as Value)))&period=1d&resolution=RATE&$expand=devicepollingstatisticsmfs&$select=Name,devicepollingstatisticsmfs/im_NumSuccessfulPolls&$filter=((groups/Name eq 'Boston'))
```

### **Get a List of Devices with More Than 25 Percent Polling Failures in the Last 30 Days**

You want to identify devices that respond inconsistently to poll requests. For this query, use the aggregation function to calculate the poll failure rate, then set a limit against that value.

#### **TIP**

You can derive values from metrics in the system using the **group/aggregate** token. In the results, the aggregated data is reported as `Value`. You can use `Value` as a filter criteria for the filter metrics token.

Build the query by adding the following tokens:

- **for:** device
- **select:** Name
- **group/aggregate:**
  - **Group By:** devicepollingstatisticsmfs, ID
  - **Add an Aggregation Function:** Sum Of ((im\_NumPollsSent sub im\_NumSuccessfulPolls) div im\_NumPollsSent) mul 100
- **filter metrics:** devicepollingstatisticsmfs/Value greater 25
- **time range:** Last30Days

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$apply=groupby(devicepollingstatisticsmfs/ID, aggregate(devicepollingstatisticsmfs(((im_NumPollsSent sub im_NumSuccessfulPolls) div im_NumPollsSent) mul 100 with sum as Value)))&period=1m&resolution=RATE&$select=Name&$filter=((devicepollingstatisticsmfs/Value gt 25))
```

## **Export Device Status**

The following examples show how to extract details regarding device statuses:

### **Get a List of All Devices from the Last 24 Hours with Management Issues**

You want to identify recent devices with management issues.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** availabilitymfs, DeviceItemID
  - **Add an Aggregation Function:** Count Distinct Of (Timestamp)
- **filter metrics:** availabilitymfs/Value less 1
- **time range:** Last24Hours/As polled
- **limit(top):** top=2000, skip=0, expandtop=0

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$apply=groupby(availabilitymfs/DeviceItemID, aggregate(availabilitymfs(Timestamp with countdistinct as Value)))&resolution=RATE&period=1d&$top=2000&$skip=0&top=0&$select=ID,Name&$filter=((availabilitymfs/Value lt 1))
```

### **Get a List of All Devices from the Last 24 Hours with Reachability Issues**

You want to identify recent devices with reachability issues.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** reachabilitymfs, DeviceItemID
  - **Add an Aggregation Function:** Count Distinct Of (Timestamp)
- **filter metrics:** reachabilitymfs/Value less 1
- **time range:** Last24Hours/As polled
- **limit(top):** top=2000, skip=0, expandtop=0

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$apply=groupby(reachabilitymfs/DeviceItemID, aggregate(reachabilitymfs(Timestamp with countdistinct as Value)))&resolution=RATE&period=1d&$top=2000&$skip=0&top=0&$select=ID,Name&$filter=((reachabilitymfs/Value lt 1))
```

### **Get a List of All Devices from the Last 24 Hours that are Up without a Single CPU Polled**

You want to identify devices that need CPU metric polling restarted.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** cpumfs, DeviceItemID
  - **Add an Aggregation Function:** Count Distinct Of (Timestamp)
- **filter metrics:** cpumfs/Value less 1
- **time range:** Last24Hours/As polled
- **limit(top):** top=2000, skip=0, expandtop=0

**OpenAPI URL:**

```
http://da_hostname:8581//odata/api/devices?$apply=groupby(cpumfs/DeviceItemID,
  aggregate(cpumfs(Timestamp with countdistinct as Value)))&resolution=RATE&period=1d&
  $top=2000&$skip=0&top=0&$select=ID,Name&$filter=((cpumfs/Value lt 1))
```

### **Get a List of All Devices from the Last 24 Hours that are Up without a Single Interface Polled**

You want to identify that the devices that need Interface metric polling restarted.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** portmfs, DeviceItemID
  - **Add an Aggregation Function:** Count Distinct Of (Timestamp)
- **filter metrics:** portmfs/Value less 1
- **time range:** Last24Hours/As polled
- **limit(top):** top=2000, skip=0, expandtop=0

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$apply=groupby(portmfs/DeviceItemID,
  aggregate(portmfs(Timestamp with countdistinct as Value)))&resolution=RATE&period=1d&
  $top=2000&$skip=0&top=0&$select=ID,Name&$filter=((portmfs/Value lt 1))
```

### **Extract Inventory and Time-Series Data**

The following examples show how to extract trend and time-series data for items with specific configuration details:

#### **Export As-Polled Bits-In and Bits-Out for Interfaces in a Group**

You want to extract key interface monitoring data from the Boston group to use in a new visualization. You want to extract one week of as polled data.

**TIP**

To get the as polled data for a week, set the maximum number of expanded rows to the correct value. For this example, the poll interval is 5 minutes, so 2016 data point is one week of data.

Build the query by adding the following tokens:

- **for:** interface
- **select:** Name
- **expand metrics:** portmfs: im\_BitsIn, im\_BitsOut
- **filter:** groups/Name equal Boston
- **time range:** Last7Days/As polled
- **limit:** Maximum number of expanded rows: 2016

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/interfaces?period=1w&resolution=RATE&$top=50&
$skip=0&top=2016&$expand=portmfs&$select=Name,portmfs/im_BitsIn,portmfs/im_BitsOut&
$filter=((groups/Name eq 'Boston'))
```

**Get Total Hourly Bits-In and Bits-Out for Interfaces Associated to Devices by IP Address Substring**

You want to extract key throughput information for interfaces where the parent device IP address begins with 10.251 .  
You want the hourly throughput totals for the month.

**TIP**

For 30 days of hourly data, set the maximum number of expanded rows to 720.

Build the query by adding the following tokens:

- **for:** interface
- **select:** Name
- **expand metrics:** portmfs: im\_BitsIn, im\_BitsOut
- **filter:** device/PrimaryIPAddress starts with 10.251
- **time range:** Last30Days/Hour
- **limit:** Maximum number of expanded rows: 720

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/interfaces?period=1m&resolution=HOURL&$top=50&
$skip=0&top=720&$expand=portmfs&$select=Name,portmfs/im_BitsIn,portmfs/im_BitsOut&
$filter=((startswith(PrimaryIPAddress, '10.251') eq true))
```

**Get Total Daily Bits-In & Bits-Out for Interfaces That Have Less Than 1-GB Daily Throughput for the Last 30 Days**

You want to identify which interfaces are underutilized.

**TIP**

Because bits in and bits out aggregate by sum, daily resolution provides the same value and the query runs faster.

Build the query by adding the following tokens:

- **for:** interface
- **select:** Name
- **expand metrics:** portmfs: im\_BitsIn, im\_BitsOut
- **filter metrics:** portmfs/im\_Bits less 1000000000
- **time range:** Last30Days/Day

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/interfaces?period=1m&resolution=DAY&$expand=portmfs&
$select=Name,portmfs/im_BitsIn,portmfs/im_BitsOut&$filter=((portmfs/im_Bits lt
10000000000))
```

### **Extract Inventory and Aggregated Values**

The following examples show how to extract aggregated data for items with specific configuration details.

#### **TIP**

To apply customizable run-time analytics to the data set, use the **group/aggregate** token.

#### **NOTE**

You cannot mix percentile aggregations with other aggregation types.

### **Show Me the 95th Percentile Utilization Value for Each Interface in a Group for the Last Seven Days**

You want the 95th percentile of utilization for each interface in the Boston group.

Build the query by adding the following tokens:

- **for:** interface
- **select:** Name
- **expand metrics:** portmfs: im\_Utilization
- **group/aggregate:**
  - **Group By:** portmfs, ID
  - **Add an Aggregation Function:** Percentile95 im\_Utilization
- **filter:** groups/Name equal Boston
- **time range:** Last7Days/As polled

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/interfaces?$apply=groupby(portmfs/
ID, aggregate(portmfs(im_Utilization with percentile95 as
Value)))&resolution=RATE&period=1w&$expand=portmfs&$select=Name,portmfs/ID&
$filter=((groups/Name eq 'Boston'))
```

### **Show Me the Average CPU Utilization for Each Device in a Group for the Last 30 Days**

You want to the average CPU utilization for routers in the Boston group

Build the query by adding the following tokens:

- **for:** device
- **select:** Name
- **group/aggregate:**
  - **Group By:** cpumfs, ID
  - **Add an Aggregation Function:** AverageOf im\_Utilization
- **filter:** groups/Name equal Boston
- **time range:** Last30Days/As polled

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$apply=groupby(cpumfs/ID,
aggregate(cpumfs(im_Utilization with average as Value)))&period=1m&resolution=RATE&
$select=Name&$filter=((groups/Name eq 'Boston'))
```

### **Show Me the Interfaces in a Group That Have Averaged Less Than 40 Percent Hourly Utilization**

You want to see which interfaces in your network were underutilized during the last quarter.

Build the query by adding the following tokens:

- **for:** interface
- **select:** Name
- **expand metrics:** portmfs: im\_utilization
- **group/aggregate**
  - **Group By:** portmfs, ID
  - **Add an Aggregation Function:** AverageOf im\_utilization
- **filter metrics:** portmfs/Value less 40
- **time range:** Last3Months/Hour

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/interfaces?$apply=groupby(portmfs/ID,
  aggregate(portmfs(im_Utilization with average as Value)))&period=3m&resolution=HOURL&
  $expand=portmfs&$select=Name,portmfs/im_Utilization&$filter=((portmfs/Value lt 40))
```

### **Show Me the Devices with Average CPU and Memory Utilization More Than 75 Percent Based on IP Address Substring**

You want to see devices in the subnet starting with 10.251 that have high CPU and memory utilization.

Build the query by adding the following tokens:

- **for:** device
- **select:** Name
- **filter:** PrimaryIPAddress starts with 10.251
- **expand metrics:** cpumfs: im\_MemoryUtilization, im\_Utilization
- **filter metrics:** cpumfs/im\_MemoryUtilization greater 75 OR cpumfs/im\_Utilization greater 75

#### **OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$expand=cpumfs&$select=Name,cpumfs/
  im_MemoryUtilization,cpumfs/im_Utilization&$filter=((cpumfs/im_MemoryUtilization gt
  75) or (cpumfs/im_Utilization gt 75)) and ((startswith(PrimaryIPAddress, '10.251') eq
  true))
```

### **Extract Metrics Aggregated at the Group Level**

The following examples show how to extract metrics aggregated at the group level. Use these queries to understand capacity that is related to the logical groupings of resources. For example, by customers, geography, services, or other.

#### **NOTE**

In large systems, some queries might time out if the filters are applied inefficiently. OpenAPI evaluates filter expressions in the order they appear in the query. Use the filter that limits the data more first. For example, if you want all groups where interface utilization is greater than 50 percent in North America, place the North America filter first.

### **Show Me the Top Ten Interfaces in Group ID 10 with the Highest Bandwidth Utilization**

You want to see the top ten interfaces in Group ID 10 with the highest bandwidth utilization.

Build the query by adding the following tokens:

- **for:** portmf
- **select:** ID, im\_utilization
- **filter:** groups/ID equal 10
- **expand:** interface/Name
- **sort:** im\_Utilization (DESC)
- **limit(top):** top=10, skip=0, expandtop=10

**OpenAPI URL:**

```
http://da_hostname:8581 /odata/api/portmfs?$orderby=im_Utilization desc&$top=10&
$skip=0&top=10&$expand=interface&$select=ID,im_Utilization,interface/Name&
$filter=((groups/ID eq 10))
```

### **Show Me the Average Interface Utilization of Each Interface in Group ID 10**

You want to know the average interface utilization of each interface in Group ID 10.

**NOTE**

The following query works best for small groups. Aggregated values are not paged.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name, Description

**NOTE**

Value is selected by default on the server side.

- **filter:** ID equal 10
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portmfs(im\_utilization with average as Value))

**OpenAPI URL:**

```
http://da_hostname:8581 /odata/api/groups?$apply=groupby(portmfs/ID,
aggregate(portmfs(im_Utilization with average as Value)))&$expand=interfaces&
$select=ID,Name,Description,interfaces/ID,interfaces/Name&$filter=((ID eq 10))
```

### **Show Me the Average Interface Utilization for All Interfaces Within Group ID 10 and Group ID 11**

You want to know the average interface utilization for all interfaces within Group ID 10 and Group ID 11.

**NOTE**

The following query works best for small groups. Aggregated values are not paged.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name
- **filter:** (ID equal 10) or (ID equal 11)
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portmfs(im\_Utilization with average as Value))

**OpenAPI URL:**



```
http://da_hostname:8581 /odata/api/groups?$apply=groupby(portmfs/ID,
  aggregate(portmfs(im_Utilization with average as Value)))&$select=ID,Name&$filter=((ID
  eq 10) or (ID eq 11))
```

### **Show Me the Maximum Interface Utilization Value for Each Child Group Within the Boston Group**

You want to know the maximum interface utilization value for each child group within the Boston group.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name
- **filter:** (GroupPathLocation contains Boston)
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portmfs(im\_Utilization with max as Value))

#### **OpenAPI URL:**

```
http://da_hostname:8581 /odata/api/groups?$apply=groupby(portmfs/ID,
  aggregate(portmfs(im_Utilization with max as Value)))&$select=ID,Name&
$filter=((substringof('Boston:', GroupPathLocation) eq true))
```

### **Show Me the 95th Percentile of Interface Utilization for Each Group That Has the Word "Boston" in the Description**

You want to know the 95th percentile of interface utilization for each group with the word "Boston" in the Description.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name
- **filter:** (Name contains Boston)
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portmfs(im\_Utilization with percentile95 as Value))

#### **OpenAPI URL:**

```
http://da_hostname:8581 /odata/api/groups?$apply=groupby(ID,
  aggregate(portmfs(im_Utilization with percentile95 as Value)))&$select=ID,Name&
$filter=((substringof('Boston',Name) eq true))
```

### **Show Me the Top Five Group IDs with the Highest Group Interface Utilization for All Child Groups in the Boston Group**

You want to see the top five groups with the highest group interface utilization for all child groups in the Boston group.

Build the query by adding the following tokens:

- **for:** portmf
- **select:** ID, Value
- **filter:** (groups/GroupPathLocation contains Boston) and (Value greater than or equal 0)
- **group/aggregate:**

- **Group By:** groups/ID
- **Add an Aggregation Function:** aggregate(im\_Utilization with average as Value)
- **sort:** Ordering, Value (DESC)

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/portmfs?$apply=groupby(groups/ID,
  aggregate(im_Utilization with average as Value))&$orderby=Value desc&$select=ID,Value&
$filter=((substringof('Boston:',groups/GroupPathLocation) eq true) and (Value ge 0))
```

**Extract Metrics within Specified Business Hours**

The following examples show how to extract metrics within specified business hours. For more information see, [Advanced OpenAPI Query Examples](#).

**Show Me Data for Monday through Friday 9:00 AM to 5:00 PM Eastern Standard Time (EST)**

You want to see data within your business hours, which are Monday through Friday 9:00 AM to 5:00 PM EST.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portmfs(im\_Utilization with average as Value))
- **time range:** Last 7 Days/As polled/-05/Mon,Tue,Wed,Thu,Fri 9:00-17:00

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/groups?$apply=groupby(portmfs/ID,
  aggregate(portmfs(im_Utilization with average as
  Value)))&resolution=RATE&period=1w&bh=Mon-Fri 9:00-17:00&tz=-05&$select=ID,Name
```

**Show Me Data for Monday, Tuesday, and Friday 9:00 AM to 12:00 PM and 1:00 PM to 6:00 PM**

You want to see data within your business hours, which are Monday, Tuesday, and Friday 9:00 AM to 12:00 PM and 1:00 PM to 6:00 PM.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portmfs(im\_Utilization with average as Value))
- **time range:** Last 7 Days/As polled/Mon,Tue,Fri 9:00-12:00 13:00-18:00

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/groups?$apply=groupby(portmfs/ID,
  aggregate(portmfs(im_Utilization with average as
  Value)))&resolution=RATE&period=1w&bh=Mon-Tue 9:00-12:00 13:00-18:00,Fri 9:00-12:00
13:00-18:00&$select=ID,Name
```

**Show Me Data for Monday through Friday 8:00 PM to 4:00 AM**

You want to see data within your business hours, which are Monday through Friday 8:00 PM to 4:00 AM.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portms(im\_Utilization with average as Value))
- **time range:** Last 7 Days/As polled/Mon,Tue,Wed,Thur,Fri 20:00-4:00

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/groups?$apply=groupby(portmfs/ID,
  aggregate(portmfs(im_Utilization with average as
  Value)))&resolution=RATE&period=1w&bh=Mon-Fri 20:00-4:00&$select=ID,Name
```

**Show Me Data for Monday through Friday 9:00 AM to 5:00 PM in January**

You want to see data for January within your business hours, which are Monday through Friday 9:00 AM to 5:00 PM.

Build the query by adding the following tokens:

- **for:** group
- **select:** ID, Name
- **group/aggregate:**
  - **Group By:** portmfs/ID
  - **Add an Aggregation Function:** aggregate(portms(im\_Utilization with average as Value))
- **time range:** 2017-01-02T09:00~2017-01-31T17:00/As polled/Mon,Tue,Wed,Thu,Fri 9:00-17:00

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/groups?$apply=groupby(portmfs/ID,
  aggregate(portmfs(im_Utilization with average as
  Value)))&resolution=RATE&starttime=1483365600&endtime=1485900000&bh=Mon-Fri 9:00-17:00&
  $select=ID,Name
```

**Extract SD-WAN Data**

The following examples show how to extract SD-WAN data.

**Show Me All Tunnels and Attributes in a Group**

You want to see data for your SD-WAN tunnels.

Build the query by adding the following tokens:

- **for:** sdntunnel
- **filter:** groups/Name contains Site 1

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/sdntunnels?$filter=((substringof('Site 1', groups/Name) eq true))
```

**Show Me the Latency, Loss, and Jitter Values for All Tunnels in a Group**

You want to see the latency, loss, and jitter values for your SD-WAN tunnels.

Build the query by adding the following tokens:

- **for:** sdntunnel
- **filter:** groups/Name contains Site 1
- **expand metrics:** sdntunnelmfs(im\_Jitter,im\_Latency,im\_PacketLossPercentage)

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/sdntunnels?$expand=sdntunnelmfs&$select=sdntunnelmfs/
std_im_Jitter,sdntunnelmfs/std_im_Latency,sdntunnelmfs/std_im_PacketLossPercentage&
$filter=((substringof('Site 1', groups/Name) eq true))
```

**Show Me the System Metrics for the Source and Destination Devices for All Tunnels in a Group**

You want to see the system metrics for your source and destination devices.

Build the query by adding the following tokens:

- **for:** sdntunnel
- **filter:** groups/Name contains Site 1
- **expand metrics:** sdntunnelmfs(std\_im\_Jitter,std\_im\_Latency,std\_im\_PacketLossPercentage)

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/sdntunnels?$expand=sdntunnelmfs&$select=sdntunnelmfs/
std_im_Jitter,sdntunnelmfs/std_im_Latency,sdntunnelmfs/std_im_PacketLossPercentage&
$filter=((substringof('Site 1', groups/Name) eq true))
```

**Show Me the Virtual Interface Metrics for the Source and Destination Interfaces in a Group**

You want to see the virtual interface metrics for your source and destination interfaces.

Build the query by adding the following tokens:

- **for:** sdnvirtualinterface
- **filter:** groups/Name contains Site 1
- **expand metrics:** stdvitalinterface(std\_im\_BitsIn,std\_im\_BitsOut)

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/adnvirtualinterfaces?$expand=virtualinterfacemfs&
$select=virtualinterfacemfs/std_im_BitsIn,virtualinterfacemfs/std_im_BitsOut&
$filter=((substringof('Site 1', groups/Name) eq true))
```

**Retrieve a Device Life Cycle State**

You want to retrieve a device's life cycle state. Use the `devicelifecycle` entity with the expand operator.

Build the query by adding the following tokens:

- **for:** device
- **select:** ID, Name
- **expand:** devicelifecycle, State

**OpenAPI URL:**

```
http://da_hostname:8581/odata/api/devices?$expand=deviceLifecycle&
$select=ID,Name,deviceLifecycle/State
```

## Advanced OpenAPI Query Examples

To take further advantage of the flexibility of the OpenAPI, use advanced OpenAPI queries. You can use advanced OpenAPI queries to build your own OpenAPI applications.

The OpenAPI service declares its structure in the Metadata Document, so that you can view the following details:

- The requests that can be executed
- The structure of the results
- How the service can be navigated

You can invoke the Metadata Document using the following URI:

```
http://da_host:8581/odata/api/$metadata
```

### **Special Aggregation Functions**

Special aggregation functions can be used to:

- Calculate projection values
- Determine the projected time to a threshold
- Get the properties of the linear data model such as the slope and intercept that are used for these calculations

The following special aggregation functions are available for these calculations:

- **intercept**  
Where the projection line crosses the Y-axis
- **slope**  
The steepness of the projection line
- **projection**  
A line predicting the future performance of a metric
- **datasetcount**  
The number of data points that are used to calculate the linear data model
- **timetothreshold**  
The number of seconds required for the projection line to cross a threshold  
This value is calculated from the end of the time interval that is defined by the `starttime` and `endtime` parameters

#### **NOTE**

Note: If `starttime` and `endtime` are unspecified, the default time interval is used.

### **Threshold Projection Parameters**

To extract projections and threshold projections, include the following parameters in advanced OpenAPI queries:

- **prjoffset**  
The amount of time in seconds to calculate out the projection
- **threshold**  
The threshold value for the query

## **Extract Projections and Threshold Projections**

Projections predict future performance. Threshold projections indicate when a metric is predicted to cross a threshold. Use this functionality to manage the capacity of your infrastructure and reprovision or acquire hardware as needed.

The following examples show how to extract projections and threshold projections.

### **Show Me the Projection for a Single Metric**

You want to see the projected CPU utilization for the next 24 hours starting from Wednesday, Aug 10, 2016 at 14:05:00 GMT.

To get this data, use the following OpenAPI URL:

```
http://da_host
:8581
/odata/api/cpumfs?apply=groupby(ID,aggregate(im_Utilization with projection as Value))&
$select=ID,Value&starttime=1470837900&endtime=1470927900&resolution=DAY&prjoffset=86400
```

### **Show Me the Threshold Projection for a Single Metric**

You want to see when CPU utilization is predicted to cross 40 percent. The default resolution for calculating the data is as polled. The threshold projection is returned in seconds.

To get this data, use the following OpenAPI URL:

```
http://da_host
:8581
/odata/api/cpumfs?$apply=groupby(ID,aggregate(im_Utilization with timetothreshold as Value))&
$select=ID,Value&threshold=40
```

### **Show Me the Threshold Projection for a Specified Resolution**

You want to see when the CPU utilization is predicted to cross 40 percent. Your desired resolution for calculating the data is daily. The threshold projection is returned in seconds.

To get this data, use the following OpenAPI URL:

```
http://da_host
:8581
/odata/api/cpumfs?$apply=groupby(ID,aggregate(im_Utilization with timetothreshold as Value))&
$select=ID,Value&resolution=DAY&threshold=40
```

### **Show Me Multiple Projections from a Single Query**

You want to see the projected throughput (bits per second) and discard rate for both inbound and outbound traffic in 60 days. You want to use hourly data from the past three months.

To get this data, use the following OpenAPI URL:

```
http://da_host:8581/odata/api/portmfs?apply=groupby(ID,aggregate((im_BitsPerSecondIn
with projection as im_BitsPerSecondIn),(im_BitsPerSecondOut with
projection as im_BitsPerSecondOut),(im_DiscardsIn with projection as
im_DiscardsIn),(im_DiscardsOut with projection as im_DiscardsOut)))&
$select=ID,Value&starttime=1463320821&endtime=1471269621&resolution=HOURL&prjoffset=5184000
```

## Business Hours Parameters

To extract data based on business hour time ranges, include the following parameters in advanced OpenAPI queries:

- **duration**  
Machine-based time interval in time units (**s** second, **h** hour, **d** day, **w** week)  
Machine-based time models a quantity or amount of time in terms of seconds. Machine-based time can be accessed using other duration-based units, such as hours, days, and weeks. In addition, the days unit is treated as exactly equal to 24 hours, thus ignoring daylight savings effects. Use this parameter when you want to ignore daylight savings effects. You can use with `starttime` and `endtime` parameters to specify the start point or endpoint of the duration.  
**Example:** `duration=1w`
- **period**  
Human-based time interval in time units (**s** second, **h** hour, **d** day, **w** week, **m** month, **y** year)  
Human-based time models a quantity or amount of time in terms of years, months, and days. Human-based time includes daylight savings and other effects. Use this parameter when you want to account for daylight savings effects. You can use with `starttime` and `endtime` parameters to specify the start point or endpoint of the period.  
**Example:** `period=1w`
- **tz**  
Time zone offset from Greenwich Mean Time (GMT) applied to the start and end times of the query and any specified business hours. If no `tz` parameter is used, the default time zone of the server is used.  
**Example:** `period=20d&tz=-4:00`

### NOTE

If you use a daily or weekly resolution in your query, the `tz` parameter is ignored.

- **bh**  
Business hours  
**Examples:**  
`bh=Mon-Fri`  
`bh=Mon,Tue,Fri`  
`bh=Mon 8:30am-5:30pm`  
`bh=Mon 8:30-17:30`  
`bh=Mon 8:30-12:30 13:30-17:30`

You can use `starttime` and `endtime` parameters with the business hours parameters as shown in the following examples:

- `starttime=1493145475&endtime=1493145775&bh=Mon-Fri 8:30am-5:30pm`
- `starttime=2017-03-15T18:00:00&period=20d&bh=Mon-Fri 8:30am-5:30pm&tz=-8:00`
- `endtime=2017-04-15T18:00:00&period=20d&bh=Mon-Fri 8:30am-5:30pm&tz=-8:00`
- `starttime=2017-03-15T18:00:00&period=20d&bh=Mon-Fri 8:30am-5:30pm&tz=-8:00`
- `starttime=2017-03-15T18:00:00&endtime=2017-04-15T18:00:00&bh=Mon-Fri 8:30am-5:30pm`

## Custom Functions

The following custom functions are available:

- **getSchemaVersion**  
See the schema version information.  
**Example:** `http://da_host:8581/odata/api/getSchemaVersion`
- **getDataCollectors**  
See a list of your Data Collectors and their details including hostname, IP address, and status.  
**Example:** `http://da_host:8581/odata/api/getDataCollectors`
- **getGroupMetricFamilies**

See a list of metric families by group.**Example:** `http://da_host:8581/odata/api/getGroupMetricFamilies?&ID=GroupID&format=json&timeout=30`

– **GroupID**

Specify the Group ID from the Data Aggregator.

• **getDataAggregators-**

See a list of Data Aggregators and their details including hostname, IP address, and status.**Example:**

`http://da_host:8581/odata/api/getDataAggregators`

• **getResultLimiters**

See a list of result limiters.**Example:** `http://da_host:8581/odata/api/getResultLimiters`

## Configure OpenAPI Defaults and Limits

To perform operations that heavily affect the system during off-hours, override the QueryBuilder properties.

You can override several parameters in the QueryBuilder. The default parameter values that are defined in the OData limiters configuration file protect the system from OpenAPI queries that negatively affect the overall system performance. These parameters either limit the returned set of results, or define the timeout threshold for potentially large operations. You can customize and override the default parameter values, which vary according to the scale and capability of the system.

### Override an OpenAPI Parameter

To override a parameter, append the following code to the OpenAPI URL:

`&<Override_parameter>=<override_value>`

For example, to override the number of rows that are returned for a device query, use the following URL:

`http://da_host:8581/odata/api/devices?$expand=cpumfs&$select=Name,PrimaryIPAddress,cpumfs/im_MemoryUtilization,cpumfs/im_Utilization &$top=200`

The following parameters can be overridden:

Parameter	Default Value	Override	Description
defaultTopLimit	50	\$top	Number of rows to return
defaultExpandTopLimit	100	top	Custom parameter for number of expanded rows to return
defaultQueryTimeoutSecs	30 (sec)	timeout	Custom parameter for overall time query can execute before timeout exception

### Configure Web Service Parameter

The OpenAPI web service parameters limit concurrent use and limit the processing time for requests beyond that limit.

To override a web service parameter, edit the value in the following configuration file:

`<installation_directory>/apache-karaf/etc/com.ca.im.odata.filters.OpenAPIRequestLimiterFilter.cfg`

• **installation\_directory**

The default installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

You can override the following web service parameters.



- **maxRequests**  
Number of simultaneous OpenAPI requests. Other requests are suspended until one of the current queries finishes.  
**Default: 4**
- **suspendMs**  
The length of time, in milliseconds, that each additional request is suspended when the OpenAPI reaches the maximum number of simultaneous requests. For example, specifying 20000 suspends the request for 20 seconds, after which time the request is rejected. A rejected request generates the 503 Unavailable error, and the OpenAPI does not run the query again automatically.  
**Default: -1** (Use the value of defaultQueryTimeoutSecs)
- **waitMS>**  
The length of time, in milliseconds, to wait before trying to accept a new request. This parameter is used when the **maxRequests** limit is reached. Set this parameter to the same value as defaultQueryTimeoutSecs in the `<installation_directory>/apache-karaf/etc/com.ca.im.odata.beans.ODataLimiters.cfg` file.

### Configure Maximum Results

To modify the defaults or increase the limit, configure the values for the number of rows that OpenAPI queries return.

#### **Follow these steps:**

1. Log on to the data aggregator host.
2. Locate and edit the `<installation_directory>/apache-karaf/etc/com.ca.im.odata.beans.ODataLimiters.cfg` file.
  - **installation\_directory**  
The default installation directory for the data aggregator.  
**Default:** `/opt/IMDataAggregator`
3. Modify the file to set the limits and defaults. The following example shows the default values for this file. The bold attributes control the limits and defaults:
 

```

defaultTopLimit=50
defaultExpandTopLimit=100
maxTopLimit=20000
maxSubQueryLimit=2000000
defaultRateTimeIntervalSecs=3600
defaultHourlyTimeIntervalHours=168
defaultDailyTimeIntervalDays=30
defaultWeeklyTimeIntervalWeeks=52
defaultQueryTimeoutSecs=30
maxQueryTimeoutSecs=120
      
```

  - **defaultTopLimit**  
Defines the default value for the maximum number of rows in the output.  
**Default: 50**
  - **defaultExpandTopLimit** Defines the default value for the maximum number of expanded rows in the output.  
**Default: 100**
  - **maxTopLimit** Defines the limit for the value in the maximum number of rows in the output.  
**Default: 20000**
  - **maxSubQueryLimit**  
Defines the maximum number of expanded rows that an OpenAPI query can return. In QueryBuilder, the limit for the maximum number of expanded rows is calculated by dividing this value by the specified value for the maximum number of rows.

**Default:** 2000000

– **defaultRateTimeIntervalSecs**

Defines the default value for the query time interval when the resolution equals rate. If the time interval is unspecified, the query time is for 3,600 seconds (one day).

**Default:** 3600

– **defaultHourlyTimeIntervalHours**

Defines the default value for the query time interval when the resolution equals hour. If the time interval is unspecified, the query time is for 168 hours (7 days).

**Default:** 168

– **defaultDailyTimeIntervalDays**

Defines the default value for the query time interval when the resolution equals day. If the time interval is unspecified, the query time is for 30 days.

**Default:** 30

– **defaultWeeklyTimeIntervalWeeks**

Defines the default value for the query time interval when the resolution equals week. If the time interval is unspecified, the query time is for 52 weeks.

**Default:** 52

– **defaultQueryTimeoutSecs**

Defines the standard timeout for all queries.

**Default:** 30

– **maxQueryTimeoutSecs**

Defines the maximum timeout, which can be specified in the URL with the `timeout` parameter.

**Default:** 120

4. Save the changes.

The new values apply when you load or reload QueryBuilder in the browser.

## OpenAPI Apps

You can serve custom content to OpenAPI app views on dashboards and context pages using OpenAPI apps. They deliver and present data in a customizable way using the flexibility of OpenAPI queries.

You can build your own apps or you can select from the various available apps. You can then deploy and display these apps in NetOps Portal.

The following video demonstrates how to download, deploy, and add OpenAPI apps to NetOps Portal:

In this article:

- [OpenAPI App Development](#)
- [Download an App](#)
- [Deploy an App](#)
- [Display an App in NetOps Portal](#)

### OpenAPI App Development

OpenAPI apps are made up of HTML, JavaScript, CSS, and metadata contained in a zip file. Generally, apps use NetOps Portal web services and OpenAPI.

Use the following process to guide you through the OpenAPI app development and testing process:

1. If one exists, find a similar OpenAPI app to use as the basis for your new app.

2. Set up a local folder (a top-level single directory) to house the files for the app, such as the `appConfig.properties` file. Use a single directory that contains all the required directories and files for the app.
3. [Create the `appConfig.properties` file in the local folder that you created to house the files for the app.](#) This properties file is required for deployment in NetOps Portal and for the app view.
4. Use OpenAPI QueryBuilder to create sample OData queries for your app. Save the output from these queries as to a static file. Start developing your code with the static data.  
After development, static data is useful for debugging to ensure that your parsing code works correctly.
5. Build a demo page. Ensure that the page can load the static data. Also, ensure that it can save the data to a variable attached to the window object. To enable browser debugging tools, add the debugger statement in the JavaScript code. Use the browser debugging tool to look at the data that came over, and verify that part of the app worked.
6. Start parsing parameters. Include the URL you used to get the original data, and use a parameter to switch into debug mode. Use parameters from the URL to the build OpenAPI query. Use the `console.log` function to log any URL to the console, such as `console.log("url: " + url)`.
7. Pick the JavaScript libraries that you want to use for your app. Look at existing apps for examples and copy the libraries into your app folder. Remove unnecessary files.
8. Start writing the app. If the data includes multiple items, start with one item first. At this stage, do not focus on the app appearance. However, *do* include CSS classes in the resulting HTML to make it easier. Put your labels in a single place in your code, in case you localize the app.
9. When you have something working, deploy the app to your NetOps Portal test system. Verify that it works on a remote server. If you refactor your code to get it to work remotely, there is less code to refactor at this stage.
10. Zip the local folder that you created to house the files for the app, including the `appConfig.properties` file, and then [deploy the app through NetOps Portal](#).
11. Test the app in NetOps Portal by completing the following:
  - a. Create a dashboard.
  - b. Add two app views to the dashboard.
  - c. Edit one of the app views, and select your app from the drop-down list.  
If the app does not appear in the drop-down list, something in the `appConfig.properties` file is incorrect.
  - d. After the app appears in the first view, edit the other app view, and add the same app.  
This test ensures that the app can support multiple instances on a single page.
12. Continue to work locally, and periodically redeploy the app to the test system. If you update the URL in the metadata, edit the app view. Reselect the app to get the updated URL.
13. Use a more restricted user account to verify that the app still works correctly. If your app is context or time range aware, change the associated controls and verify that your app responds correctly.
14. If appropriate, use the redirector to add context page links to your app. The redirector makes the app feel like a native part of the system.
15. When the app is complete, consider contributing your app to the GitHub repository for others. Before you add the app to the repository, remove any personal information.

### **The Format for the `appConfig.properties` File**

Use the following format for the `appConfig.properties` file:

```
appName=App Name
description=Optional app description
url=Required_URL_parameters
height=height in pixels
supportedContext=context_code
```

- **App Name**

- A unique name for the app.
- (Optional) **Optional app description**  
A description that appears in the App View when you select the app.
- **Required URL parameters**  
The app URL and URL parameters.  
For information about URL parameters in NetOps Portal, see [Browser Views](#).
- (Optional) **height in pixels**  
The height of the app view in pixels.  
**Default:** 250
- (Optional) **context\_code**  
The context where the app appears in the app view. You can use this parameter to restrict the available apps to specific contexts. To designate multiple contexts, separate the values with a comma.  
**Values:**
  - **d:** Device
  - **i:** Interface - The app is available only in app views on interface context pages and the app is unavailable at the dashboard level.
  - **s:** Server
  - **r:** Router
  - **g:** Group
  - **nc:** None - The app appears in all contexts.**Default:** nc

**Example:**

```
appName=Percentile Trend App
description=This is a Percentile Trend App
url=index.html?
id={ItemIdDA}&startTime={TimeStartUTC}&endTime={TimeEndUTC}&metric=im_UtilizationIn
height=750
supportedContext=i,d
```

**App Development Best Practices**

Use the following best practices for developing OpenAPI apps:

- Use only internally-sourced JavaScript libraries. If you need libraries for your app, include the library in the app folder. Also include any files that describe the license terms of the library.
- To avoid the app from breaking, do not use any NetOps Portal JavaScript, CSS, or images.
- Test the app to verify that it works at scale and with different configurations. Many apps depend on system resources, such as web services. Verify that the app is not generating unnecessary load.
- Build test modes into the app to help debug problems. Add an optional parameter that switches the app from using a web service API to a canned set of data included in the app.
- If you distribute the app to other users, verify that no private information is included. Sanitize the static data files that are included for debugging.
- Use relative paths and never include the full URL. If the system is behind a firewall, or changes DNS names, using full URLs breaks the app. Using relative paths helps the transition between working locally and then deploying the app.
- Verify that the app works in an iframe. Both the Browser View and App View use iframes to isolate the apps from the rest of the page.
- The OpenAPI uses Data Aggregator item IDs. NetOps Portal IDs are not recognized. Use the Data Sources web services to convert between the Data Aggregator item ID and the NetOps Portal ID.

For more information, see [Data Sources Web Service](#).

- To access OpenAPI data, the request must come from an app on the NetOps Portal host. The app deployment places the app folder on the NetOps Portal host.
- For direct OData queries in the apps, use the relative path for the NetOps Portal system to use the OpenAPI proxy: `/pc/odata/api/...`

## **Download an App**

Various apps are available on GitHub. These applications are supported through GitHub open-source collaboration:

Copyright (c) 2018 CA Technologies

The MIT License

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

### **Follow these steps:**

1. Go to <https://github.com/CA-PM>.  
A list of available apps with descriptions appear.
2. Click the app that you want to download.  
The contents of the app, a sample visualization, and readme instructions appear.
3. Click **Clone or download**, and then **Download ZIP**.
4. Save the files to your preferred directory.

The app is downloaded.

## **Deploy an App**

To deploy OpenAPI apps, load the app through NetOps Portal. You can deploy apps while the system is running. App deployment does not require a restart. For successful deployment, the app must meet the following requirements:

- The app files are contained in a single folder.
- The folder includes the `appConfig.properties` file.
- The app folder is zipped in a ZIP file. The maximum ZIP file size is 100 MB.

This procedure requires the Administrator role.

### **Follow these steps:**

1. Hover over **Administration, Configuration Settings**, and then click **App Deployment**.  
The **App Deployment** page appears.

2. In the **App** field, browse for the file, select the ZIP file, and then click **Open**.
3. Click **Add**.

The app is copied to the user app directory.

### **Display an App in NetOps Portal**

Use the app view in NetOps Portal to display an app.

#### **NOTE**

PDF and CSV printing and emailing options are unsupported for app views.

#### **Follow these steps:**

1. Add or edit a dashboard or page for the intended item context.
2. Add an app view for the app.
3. Select the app for the view.
4. Update the parameters and height as desired.
5. Click **Save**.

## **Audit OpenAPI Usage**

To track the performance of the OpenAPI QueryBuilder, see the usage statistics.

To view the usage statistics of the OpenAPI QueryBuilder, review the OpenAPI log file or create an OpenAPI query. Auditing OpenAPI usage helps to diagnose system performance issues.

### **OpenAPI Usage Log File**

The log file provides detailed information about each query that is executed, and whether the query succeeded or failed. Each entry for a single query in the log file documents the start and end of the request. The log file is located in the `<installation_directory>/apache-karaf/data/log/odata-services.impl.log` directory.

- ***installation\_directory***

The default installation directory for the data aggregator.

**Default:** `/opt/IMDataAggregator`

The following query details are documented in the log file:

- Host that made the API call
- Tenant
- User
- URL
- Length of time of the request

### **Example**

The following entry is an example of a successful query request:

```
INFO | qtp1695827227-67 | 2015-06-01 12:17:55,207 | ODataRequestProcessor |
odata.impl.ODataRequestProcessor 910 || Start request (Thread id 67): host:
'10.42.200.100' tenant: 'Default Tenant' user: 'admin' url: 'http://da.ca.com:8581/
odata/api/components?$expand=device&$select=ID,Name,device/ID,device/Name&$format=json&
$top=20000' - Run budget 30000 ms (30 seconds).
```

```
INFO | qtp1695827227-67 | 2015-06-01 12:17:55,482 | ODataRequestProcessor |
odata.impl.ODataRequestProcessor 948 || Done request (Thread id 67): host:
'10.42.200.100' tenant: 'Default Tenant' user: 'admin' url: 'http://da.ca.com:8581/
odata/api/components?$expand=device&$select=ID,Name,device/ID,device/Name&$format=json&
$top=20000' - Execution time 273 ms.
```

The following entry is an example of a failed query request:

```
INFO | p1695827227-1474 | 2015-06-01 12:33:30,892 | ODataRequestProcessor |
odata.impl.ODataRequestProcessor 910 || Start request (Thread id 1474): host:
'10.42.200.100' tenant: 'Default Tenant' user: 'admin' url: 'http://da.ca.com:8581/
odata/api/components?$expand=device&$select=ID,Name,device/ID,device/Name&$format=json&
$top=500000' - Run budget 30000 ms (30 seconds).
ERROR | p1695827227-1474 | 2015-06-01 12:33:30,894 | ODataRequestProcessor |
odata.impl.ODataRequestProcessor 1047 || Failed to complete request (Thread id
1474): host: '10.42.200.100' tenant: 'Default Tenant' user: 'admin' url: 'http://
da.ca.com:8581/odata/api/components?$expand=device&$select=ID,Name,device/ID,device/
Name&$format=json&$top=500000' - Execution time 0 ms. Error: The top value must be in
the range (0 .. 20000).
```

### Track OpenAPI Usage

You can track the overall OpenAPI usage by running an OpenAPI query. This query can show the number of queries during a specific time period or the average processing time. The following statistics are available:

- **Number of Queries**  
The number of queries that occurred since the last read
- **Number of Successful Queries**  
The number of successful queries that occurred since the last read
- **Number of Failed Queries**  
The number of failed queries that occurred since the last read
- **Time to Process Queries**  
The average time to process queries


#### NOTE

This data is collected every 5 minutes.


### Follow these steps:

1. Select **Click to start query** in the Query Expression field.
2. Select **metric family, openapiquerymf** from the Query Expression field.
3. Select **select** from the Query expression field.
4. Select the statistics that you want to view. For example, select the following metrics:
  - ID
  - im\_TimeToProcessQuery
  - im\_NumberOfQueries
5. Click **Run**.  
The OpenAPI usage statistics appear, as shown in the following example:

Query Expression

 for


openapiquerymf

 select

ID, im\_TimeToProcessQuery, im\_NumberOfQueries

▼ OData URL

[http://vik-smoke-da1:8581/odata/api/openapiquerymf?\\$select=ID,im\\_TimeToProcessQuery,im\\_NumberOfQueries](http://vik-smoke-da1:8581/odata/api/openapiquerymf?$select=ID,im_TimeToProcessQuery,im_NumberOfQueries)

 Copy to clipboard

Run

Table - Results

ID	im_TimeToProcessQuery	im_NumberOfQueries
4828	1398.12068965517	58.0
4828	997.307692307692	39.0
4828	0.0	0.0



---

## Product Accessibility Features

---

Broadcom is committed to making sure that all customers, regardless of ability, can successfully use its products and supporting documentation to accomplish vital business tasks. This section outlines the accessibility features of DX NetOps Performance Management.

### NOTE

If you use a screen reader, use a browser other than Internet Explorer to access DX NetOps Performance Management.

For more information about known screen reader limitations, see [Known Issues](#).

### Product Enhancements

DX NetOps Performance Management offers accessibility enhancements in the following areas:

- Display
- Sound
- Keyboard
- Mouse

### NOTE

The following information applies to Windows-based and Macintosh-based applications. Java applications run on many host operating systems, some of which already have assistive technologies available to them. For these existing assistive technologies to provide access to programs written in JPL, they need a bridge between themselves in their native environments and the Java Accessibility support that is available from within the Java virtual machine (or Java VM). This bridge has one end in the Java VM and the other on the native platform, so it will be slightly different for each platform it bridges to. Sun is currently developing both the JPL and the Win32 sides of this bridge.

### Display

To increase visibility on your computer display, you can adjust the following options:

- **Font style, color, and size of items**

Defines font color, size, and other visual combinations.

- **Screen resolution**

Defines the pixel count to enlarge objects on the screen.

- **Cursor width and blink rate**

Defines the cursor width or blink rate, which makes the cursor easier to find or minimize its blinking.

- **Icon size**

Defines the size of icons. You can make icons larger for visibility or smaller for increased screen space.

- **High contrast schemes**

Defines color combinations. You can select colors that are easier to see.

### Sound

Use sound as a visual alternative or to make computer sounds easier to hear or distinguish by adjusting the following options:

- **Volume**

Sets the computer sound up or down.

- **Text-to-Speech**

Sets the computer's hear command options and text read aloud.

- **Warnings**

Defines visual warnings.

- **Notices**

Defines the aural or visual cues when accessibility features are turned on or off.

- **Schemes**

Associates computer sounds with specific system events.

- **Captions**

Displays captions for speech and sounds.

## **Keyboard**

You can make the following keyboard adjustments:

- **Repeat Rate**

Defines how quickly a character repeats when a key is struck.

- **Tones**

Defines tones when pressing certain keys.

- **Sticky Keys**

Defines the modifier key, such as Shift, Ctrl, Alt, or the Windows Logo key, for shortcut key combinations. Sticky keys remain active until another key is pressed.

## **Mouse**

You can use the following options to make your mouse faster and easier to use:

- **Click Speed**

Defines how fast to click the mouse button to make a selection.

- **Click Lock**

Sets the mouse to highlight or drag without holding down the mouse button.

- **Reverse Action**

Sets the reverse function controlled by the left and right mouse keys.

- **Blink Rate**

Defines how fast the cursor blinks or if it blinks at all.

- **Pointer Options**

Let you do the following:

- Hide the pointer while typing
- Show the location of the pointer
- Set the speed that the pointer moves on the screen
- Choose the pointer's size and color for increased visibility
- Move the pointer to a default location in a dialog box

## Keyboard Shortcuts

The following table lists the keyboard shortcuts that *DX NetOps Performance Management* supports:

Keyboard	Description
Ctrl+X	Cut
Ctrl+C	Copy
Ctrl+K	Find Next
Ctrl+F	Find and Replace
Ctrl+V	Paste
Ctrl+S	Save
Ctrl+Shift+S	Save All
Ctrl+D	Delete Line
Ctrl+Right	Next Word
Ctrl+Down	Scroll Line Down
End	Line End

## Keyboard Navigation

Keyboard	Description
Tab	<ul style="list-style-type: none"> <li>• Advance the focus to the next field or control in the page.</li> </ul>
Shift+Tab	<ul style="list-style-type: none"> <li>• Move the focus to the previous field or control in the page.</li> </ul>
Enter	<ul style="list-style-type: none"> <li>• Activate the control that has focus (for example, click a button).</li> </ul>
Shift+Enter	<ul style="list-style-type: none"> <li>• Make a selection in a menu or tree.</li> </ul>
Space	<ul style="list-style-type: none"> <li>• Change selection of a checkbox that has focus.</li> </ul>
Down Arrow	<ul style="list-style-type: none"> <li>• Move the focus into a menu or into a tree.</li> <li>• Move the focus to the next item in a menu, tree, radio-button group, or chart.</li> </ul>
Up Arrow	<ul style="list-style-type: none"> <li>• Move the focus to the previous item in a menu, tree, or radio-button group, or chart.</li> </ul>
Left Arrow	<ul style="list-style-type: none"> <li>• Open a closed item in a menu or tree.</li> <li>• Move to the next element of a radio button group and select it.</li> </ul>
Right Arrow	<ul style="list-style-type: none"> <li>• Close an open item in a menu or tree.</li> <li>• Move to the previous element of a radio button group and select it.</li> </ul>
Esc	<ul style="list-style-type: none"> <li>• Close an open menu or dialog box.</li> </ul>
Ctrl+Left Arrow	<ul style="list-style-type: none"> <li>• When the focus is on a grid column header, make the column narrower.</li> </ul>

Ctrl+Right Arrow	<ul style="list-style-type: none"><li>When the focus is on a grid column header, make the column wider.</li></ul>
Ctrl+Up Arrow	<ul style="list-style-type: none"><li>When the focus is on a grid column header, move the column to the left.</li></ul>
Ctrl+Down Arrow	<ul style="list-style-type: none"><li>When the focus is on a grid column header, move the column to the right.</li></ul>

## Product References and Abbreviations

---

This documentation references the following products and abbreviations:

- CA Application Delivery Analysis (ADA)
- DX NetOps Mediation Manager (DX NetOps MM)
- DX NetOps Network Flow Analysis (NFA)
- DX NetOps Performance Management (PM)
- DX NetOps Spectrum (Spectrum)
- CA Unified Communications Monitor (CA UCM)
- Unified Infrastructure Management (UIM)
- DX Application Performance Management (APM)

## Documentation Legal Notice

---

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by Broadcom at any time. This Documentation is proprietary information of Broadcom and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Broadcom.

If you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all Broadcom copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to Broadcom that all copies and partial copies of the Documentation have been returned to Broadcom or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, BROADCOM PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL BROADCOM BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF BROADCOM IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The check mark in a Circle design is the registered trademark of NortonLifeLock Inc. and is used under license therefrom.

The manufacturer of this Documentation is Broadcom Inc.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b) (3), as applicable, or their successors.

Copyright © 2005–2024 Broadcom. All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

