Symantec™ Data Loss Prevention Network Monitor and Prevent Performance Sizing Guidelines

Version 15.7

Symantec™
A Division of Broadcom

Symantec™ Data Loss Prevention Network Monitor and Prevent Performance Sizing Guidelines

Documentation version:

Legal Notice

Broadcom
1320 Ridder Park Drive
San Jose, California
95131

https://www.broadcom.com

Symantec Support

All support services will be delivered in accordance with your support agreement and the then-current Enterprise Technical Support policy.

Knowledge Base Articles and Symantec Connect

Before you contact Technical Support, you can find free content in our online Knowledge Base, which includes troubleshooting articles, how-to articles, alerts, and product manuals. In the search box of the following URL, type the name of your product:

https://support.symantec.com

Access our blogs and online forums to engage with other customers, partners, and Symantec employees on a wide range of topics at the following URL:

https://www.symantec.com/connect

Technical Support and Enterprise Customer Support

Symantec Support maintains support centers globally 24 hours a day, 7 days a week. Technical Support's primary role is to respond to specific queries about product features and functionality. Enterprise Customer Support assists with non-technical questions, such as license activation, software version upgrades, product access, and renewals.

For Symantec Support terms, conditions, policies, and other support information, see:

https://entced.symantec.com/default/ent/supportref

To contact Symantec Support, see:

https://support.symantec.com/en_US/contact-support.html

Contents

Chapter 1

About Symantec Data Loss Prevention Network Monitor and Prevent Performance Guidelines

This chapter includes the following topics:

About network performance tests

Symantec tests Network Monitor, Network Prevent for Email, and Network Prevent for Web to assess their performance under load. The objective of these tests is to obtain data on the overall performance and throughput of Symantec Data Loss Prevention Network Monitor and Network Prevent. These tests are designed to determine the achievable throughput of different system resource configurations.

The test results as presented in this document provide general guidelines. Network and email administrators can use these guidelines to estimate the number of servers that are required to support traffic loads on a network. The guidelines can also be used to estimate the required virtual system resources.

Note: Symantec recommends that you conduct your own testing with more representative traffic profiles and loads. Running your own tests validates that your results are in line with the sizing assumptions provided by Symantec.

Symantec conducted tests using both Endace and Napatech high-speed packet capture adapters. Network Monitor test environments for each of these high-speed packet capture adapters are described in separate sections.

See “About network performance sizing guidelines” on page 8.

About network performance sizing guidelines

When you use a virtual environment you should expect some performance degradation (as compared to running on a physical system with similar system resources). You may be able to minimize performance degradation by optimizing the virtual configuration specific to your environment.

Tests were performed using VMware ESX.

Follow these guidelines when planning any server deployment:

- All of the data that is presented in this documentation can be used as a reference for estimating deployment requirements. Validate sizing guidelines in your own test environments before deployment.
- Perform these tests with the policies and the configurations that are consistent with expected deployments. For example, IDM-, EDM-, EMDI-, and DCM-based policies and configuration filters.
- Evaluate the results using a traffic profile that is consistent with your live production environment.

For more details about planning your deployment, see "Deployment planning considerations" and "Minimum system requirements for Symantec Data Loss Prevention servers" in the *Symantec Data Loss Prevention System Requirements and Compatibility Guide* at http://www.symantec.com/docs/doc10602.html.

See “About the Network Monitor performance test environment with Endace cards” on page 9.

See “About the Network Monitor performance test environment with Napatech cards” on page 14.

See “About the Network Prevent for Email performance test environment” on page 18.

See “About the Network Prevent for Web performance test environment” on page 22.

Chapter 2

Network Monitor Performance Guidelines

This chapter includes the following topics:

About the Network Monitor performance test environment with Endace cards

Symantec conducted Network Monitor performance testing in a lab environment. This lab was designed to demonstrate the comparative accuracy of all available capture methods against a replicated traffic load.

Table 2-1 describes the hardware environment that was used to test Network Monitor with Endace cards.

Note: Throughout this document, "core" refers to physical CPU cores, not to hyper-threading CPU cores.

Table 2-1 Network Monitor test hardware with Endace cards

Component	System hardware configuration
Processor	1 x Intel Xeon E5620 processor (quad-core) (2.4 GHz, 1066 MHz FSB)
Memory	8 GB RAM with Intel Gigabit Internet Controller or 16 GB RAM with Endace DAC Controller
Ethernet controller for native capture testing	Intel 82571 Gigabit Ethernet controller
High-speed packet capture card	Endace 7.5 G2/G4 (PCIe) cards with DAG v5.7.1 drivers, utilities, and run-time libraries.
Network tap	A multi-port regenerative gigabit Ethernet tap facilitated distribution of the output from a TCP replay computer to the target Network Monitor servers.

Tests were performed using the following operating-system configurations:

- Red Hat Enterprise Linux 7.5 (64-bit) with native packet capture.
- Red Hat Enterprise Linux 7.5 (64-bit) with Endace high-speed packet capture card.

Figure 2-1 shows the relationship of test computers to the network traffic generator.

Figure 2-1 Network Monitor performance test environment with Endace cards

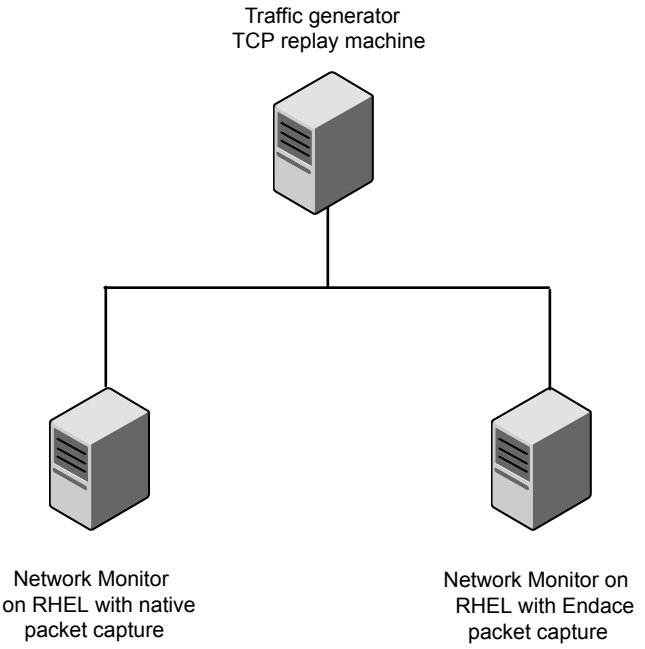

The Network Monitor servers were tested with a standard one quad-core processor configuration using both native capture and Endace capture methods. The servers were tested on Linux platforms. They were configured as follows:

- Red Hat Enterprise Linux 7.5 servers were configured for both Endace cards and for native tests.
- Network Monitor advanced settings were tuned as follows:

Network Monitor advanced setting	Linux 64-bit systems
PacketCapture.NUMBER_BUFFER_POOL_PACKETS	1200000
PacketCapture.NUMBER_SMALL_POOL_PACKETS	1000000
PacketCapture.KERNEL_BUFFER_SIZE (native capture)	128M
PacketCapture.KERNEL_BUFFER_SIZE (Napatech and Endace)	64M
PacketCapture.RING_CAPTURE_LENGTH	18000
PacketCapture.NUMBER_JUMBO_POOL_PACKETS	10000
PacketCapture.SIZE_JUMBO_POOL_PACKETS	18000

For 64-bit systems with a kernel buffer large enough to handle the processing capability of the NIC driver, increasing the buffer further showed no substantial increase in performance.

- All standard protocols were active, in addition to custom protocol definitions for Telnet, SSH, and SSL.

See "About the Network Monitor performance test methodology for an environment with Endace cards" on page 12.

About the Network Monitor performance test methodology for an environment with Endace cards

A single IDM policy was enabled that covered a target 20 MB document.

Sizing guidelines were derived from a background load of real-world traffic samples. These samples were delivered at rates ranging from 15,000 to over 200,000 packets per second. The resulting sustained background load ranged from 70 Mbps to near gigabit-level saturation.

At each background load interval, 20 copies of the target file were played at a constant rate of 3000 packets per second. Monitors that correctly generated an incident for all 20 iterations of the target file at a 100% match rate were considered a success. That is, these monitors successfully handled the offered load. When a given capture method was no longer able to deliver total match accuracy, or when packet capture discards occurred, it was considered to have reached the limit of its performance capabilities.

See "Network Monitor performance test results and sizing guidelines for environments with Endace cards" on page 12.

Network Monitor performance test results and sizing guidelines for environments with Endace cards

Network Monitor servers were tested with different capture methods accommodating different levels of network traffic. All systems were tested using Red Hat Enterprise Linux 7.5 (64-bit). Based on this performance testing, Symantec rates the tested configurations as shown in Table 2-2.

Table 2-2 Supported pre-filter performance for Network Monitor capture methodsfor environments with Endace cards

Server configuration	Bandwidth (Mbps)
Native packet capture	650
Endace card packet capture (driver version 5.7.1)	900

The test results for your network environment may be different. Variations in the protocol composition, protocol configuration, and policy load in a production deployment make a difference. Symantec recommends testing in advance against live or recorded feeds from your production infrastructure and your target protocol and policy configuration. This advance testing enables you to assess the capability to meet the demands of your deployment. Note that you may have a configuration issue if there is a wide divergence of your performance numbers from those presented in this document. The issue may be with your network architecture, tap or span configuration, network card, or capture settings.

See "About network performance sizing guidelines" on page 8.

The Network Monitor tests were designed to determine at what level of overall network traffic the detection capability of a Network Monitor Server begins to decline. As traffic rates increase, additional servers should be added to balance the total load so that no individual server's load exceeds the target level.Table 2-3 shows the estimated number of Network Monitor servers that are required for different traffic levels. This estimation assumes that test results of a single Network Monitor Server are similar to those presented here.

Table 2-3 Estimating the number of Network Monitor servers for testing with Endace cards

Network traffic (Mbps)	Linux native packet capture	Endace card packet capture
50	1	1
100	1	1
500	1	1
750	2	1
900	2	1

The traffic estimates shown assume:

- Equal load distribution across all servers
- No redundancy

See "About the Network Monitor performance test environment with Napatech cards" on page 14.

About the Network Monitor performance test environment with Napatech cards

Symantec conducted Network Monitor performance testing in a lab environment. These tests were designed to demonstrate the comparative accuracy of all available capture methods against a replicated offered traffic load.

Table 2-4 describes the hardware environment that was used to test the performance of Network Monitor in an environment with Napatech packet capture cards.

Note: Throughout this document, "core" refers to physical cores, not to hyper-threading cores.

Table 2-4 Network Monitor performance test hardware for environments with Napatech cards

Component	System hardware configuration
Processor	2 Quad-Core Xeon E5620 CPU (2.4 GHz, 1066 MHz FSB)
Memory	16 GB RAM
Ethernet controller used to test native capture	Intel 82571 Gigabit Ethernet Controller
High-speed packet capture adapters	Napatech NT4E (non-STD), or NT40A01 adapter with Napatech 3G driver package 8.0.3 for Windows 2016. Napatech NT4E (non-STD), or NT40A01 adapter with Napatech 3G driver package 8.1.0 for RHEL 7.5.

Tests were performed using the following operating-system configurations:

- Windows Server 2016 (64-bit) with native packet capture.
- Red Hat Enterprise Linux 7.5 with native packet capture.
- Windows Server 2016 (64-bit) with Napatech adapters.
- Red Hat Enterprise Linux 7.5 (64-bit) with Napatech adapters.

Figure 2-2 shows the relationship of test computers to the network traffic generator.

Figure 2-2 Network Monitor performance test environment using Napatech cards

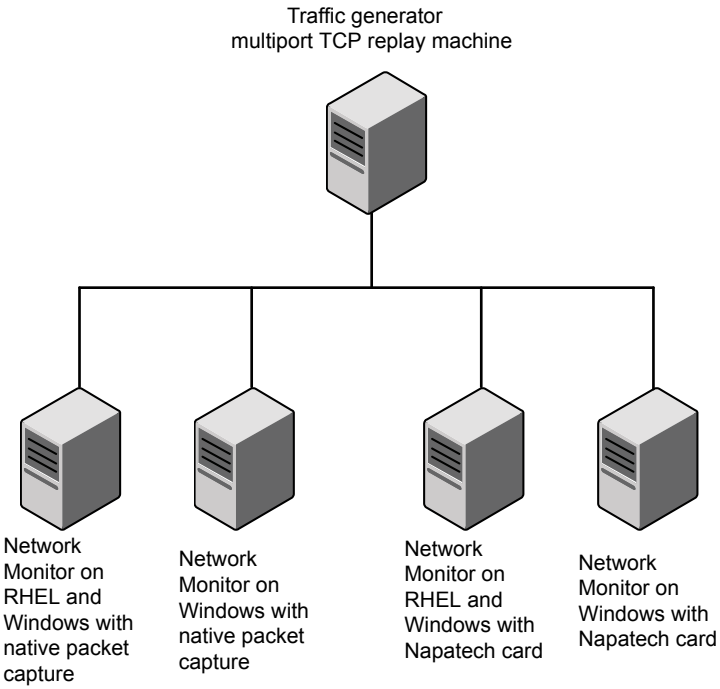

The Network Monitor servers were tested on both a standard hardware configuration and a large system hardware configuration. Both native capture and Napatech capture methods on Linux and Windows platforms were used. The systems were configured as follows:

- Network Monitor advanced settings were tuned as follows:

Network Monitor advanced setting	All 64-bit systems
PacketCapture.NUMBER_BUFFER_POOL_PACKETS	1200000
PacketCapture.NUMBER_SMALL_POOL_PACKETS	1000000
PacketCapture.KERNEL_BUFFER_SIZE	64M

  Increasing the buffer further showed no substantial increase in performance for 64-bit systems. This assumes that the kernel buffer was large enough to handle the processing capability of the high-speed packet capture adapter.

- All standard protocols were active, in addition to custom protocol definitions for Telnet, SSH, and SSL.

See “About the Network Monitor performance test methodology for an environment with Napatech cards” on page 16.

About the Network Monitor performance test methodology for an environment with Napatech cards

A single IDM policy was enabled that covered a target 20 MB document.

Sizing guidelines were derived from a background load of real-world traffic samples. The samples were delivered at rates ranging from 15,000 to over 200,000 packets per second. The resulting sustained background load ranged from 70 Mbps to near gigabit-level saturation.

At each background load interval, 20 copies of the target file were played at a constant rate of 3000 packets per second. If a Network Monitor Server correctly generated an incident for all 20 iterations of the target file at a 100% match rate, it was considered a success. When a given capture method was no longer able to deliver total match accuracy, it had reached the limit of its performance capabilities.

See “Network Monitor performance test results and sizing guidelines for environments with Napatech cards” on page 16.

Network Monitor performance test results and sizing guidelines for environments with Napatech cards

Network Monitor servers were tested with different capture methods. The capture methods accommodated different levels of network traffic. Tests were performed with Windows Server 2016 (64-bit) and Red Hat Enterprise Linux 7.5 (64-bit). The results of this performance testing are shown in Table 2-5.

Table 2-5 Supported pre-filter performance for Network Monitor capture methods in environments with Napatech cards

Server configuration	Operating system	Bandwidth (Mbps)
Native packet capture	Windows Server 2016	300
	Red Hat Enterprise Linux 7.5	650
Napatech cards: NT4E NT40A01	Windows Server 2016	900 x 4 capture interfaces
	Red Hat Enterprise Linux 7.5	900 x 4 capture interfaces

Variations in the protocol composition, protocol configuration, and policy load in a production deployment may produce different test results for your network environment. Symantec recommends testing in advance against live or recorded feeds from your production infrastructure and your target protocol and policy configuration. This way, you can assess the capability of your setup to meet the demands of your deployment.

If your performance numbers diverge from those presented in this document, that might indicate a configuration issue. Possible issues include problems with your network architecture, tap or span configuration, network card, or capture settings.

See "About network performance sizing guidelines" on page 8.

The Network Monitor tests determine what level of overall network traffic causes the detection capability to decline for each capture method. As traffic rates increase, add more servers to balance the total load. Then, no individual server's load will exceed the target level. Table 2-6 shows the number of Network Monitor Servers that are required for different traffic levels. This assumes that test results of a single Network Monitor Server are similar to those presented here.

Table 2-6 Estimating the number of Network Monitor Servers for testing in an environment with Napatech cards

Network traffic (Mbps)	Windows native packet capture	Linux native packet capture	Napatech card packet capture
50	1	1	1
100	1	1	1
500	2	1	1
750	3	2	1
900	4	2	1
900 x 4 capture interfaces	N/A	N/A	1

These estimates assume:

- Equal load distribution across all servers
- No redundancy

Chapter 3

Network Prevent for Email Performance Guidelines

This chapter includes the following topics:

About the Network Prevent for Email performance test environment

Using load generators and sample content, virtual machine (VM) configurations were tested to simulate different customer environments. These test results provide a point-in-time measurement that was generated using the specific variables and the configurations that are described in this section.

Network Prevent for Email Servers were tested on the two AWS virtual machine configurations with different virtual CPU resources: m4.2x large and AWS m4.4x large.

For more information on AWS instances, see the documentation on Amazon EC2 Instance Types at https://aws.amazon.com/ec2/instance-types.

Although memory configurations are listed, be aware that Network Prevent for Email performance is primarily CPU-bound, not memory-bound.

Table 3-1 shows the hardware configurations that are used for the AWS host computers. AWS configurations were tested using Red Hat Linux Server release 6.10.

Table 3-1 Network Prevent for Email performance test environments

Component	AWS m4.2x large configuration
Processor	8 vCPU 2.3 GHz Intel Xeon ES-2686 v4 (Broadwell)
Memory	32 GB RAM

See "About the Network Prevent for Email performance test methodology" on page 19.

About the Network Prevent for Email performance test methodology

Network Prevent for Email Servers were tested using a set of 139 policies. These policies included a variety of detection types and were a representative sampling of policies used by three large Symantec Data Loss Prevention customers. An auto-load email generation tool was used to simulate an email environment. The tool sent email traffic in forwarding mode between a client and server with a Network Prevent for Email Server between them.

Figure 3-1 shows the Network Prevent for Email test configuration used to produce the data in the following tables.

Figure 3-1 Network Prevent for Email performance test environment

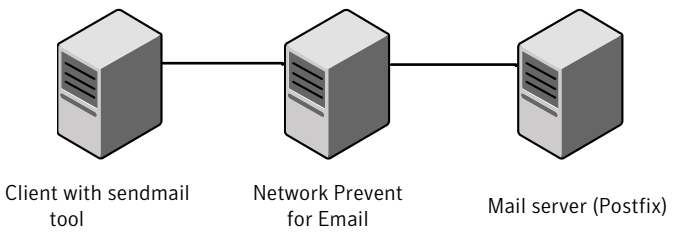

The Network Prevent for Email Servers were tested using the same set of email message attachments. For test purposes, message attachments were used to control message size and volume and to generate incidents. The test messages contained minimal body text with no content that violated policies.

- Number of email messages = 10,000
- Number of attachments = 811
- Average attachment size = 57 KB

- Average size of emails = 70 KB per message
- Attachments = a mixture of `doc`, `html`, `jpg`, `pdf`, `png`, `ppt`, `txt`, `xls`, and `zip` file types

Approximately 5% of these message attachments contained content that violated one or more of the test policies.

See "Network Prevent for Email performance test results and sizing guidelines" on page 20.

Network Prevent for Email performance test results and sizing guidelines

The following table presents benchmark results. It indicates the message volume (throughput) and latency (average processing time) that you can expect from a single Network Prevent for Email Server for deployments using very large 5 GB EDM indexes.

Symantec tested Network Prevent for Email TLS support between MTAs and Network Prevent for Email Servers to determine the number of concurrent SMTP and TCP connections. Your throughput may be reduced if your MTA does not optimize TLS connection setup and reuse. This reduction happens because of the increased processing overhead necessary to establish secure connections. Consult your MTA documentation and perform additional testing to evaluate TLS performance in your environment.

The results in Table 3-2 do not include any redundancy or failover, but do include TLS processing requirements.

Table 3-2 Network Prevent for Email performance test results with TLS for very large (5 GB) deployments

System configuration	Message volume (messages per second)	Latency (in seconds)
VM container 8-core VM, 8 message chains	88	0.26

Network Prevent for Email Servers scale linearly to handle volumes in excess of the figures that are shown here. Most MTAs can distribute load to the corresponding Network Prevent for Email Servers as necessary. Network Prevent for Email Servers are commonly paired with MTAs in an N:N redundant, load-balanced configuration.

You can estimate server requirements by extrapolating from the testing numbers that are shown here. You need to know the policy set and size of the message set. Understanding your organization's current email traffic helps you determine how many Network Prevent for

Email servers are needed to stay within the throughput and the response time limits shown. For example, the SMTP traffic that needs to be processed in a network deployment can be obtained from two sources. It can come from the MTA itself or a general sizing guideline of X outbound messages per user may be estimated.

A variety of factors influence performance of the virtual configurations. These factors include the number of CPUs and amount of physical RAM, as well as resource reservations for CPU cycles and RAM. Overhead from virtualization and guest operating systems can lead to a performance degradation in messaging throughput. This performance is compared to a standard physical system running on the same hardware. You may want to run multiple virtual instances on the same hardware to extract maximum performance and take full advantage of system resources.

Note that when virtualized, Network Prevent for Email runs as its own VM image. If the MTA is also virtualized, then both Network Prevent for Email and the MTA can run on the same physical server within a given virtual container. A dedicated network interface should be used for each VM container.

Your own test results should be used as a basis for sizing your Network Prevent for Email requirements.

Note: The recommendations in Table 3-3 do not account for redundancy, failover, or TLS processing requirements. The recommendations are based on using tuned settings.

Table 3-3 Estimating the number of Network Prevent for Email servers

Traffic volume	Number of 8-CPU physical servers needed	Number of 8-CPU VM containers needed	Number of 16-CPU VM containers needed
70 messages per second	1	1	1
80 messages per second	1	1	1

The traffic estimates shown in Table 3-3 assume:

- Equal load distribution across all servers
- No redundancy

Your test results for your network environment may be different. If there is a wide divergence of your performance numbers from the results presented, there may be a configuration issue between your email system and Network Prevent for Email.

Symantec recommends that you keep your CPU usage at 90% or lower for optimum performance. Many very large enterprises with thousands of employees running hundreds of policies have higher processing requirements than those listed here, and should plan accordingly.

Chapter 4

Network Prevent for Web Performance Guidelines

This chapter includes the following topics:

About the Network Prevent for Web performance test environment

Network Prevent for Web servers were tested on the AWS (Amazon Web Services) virtual configurations.

Table 4-1 shows the system configuration that was used for the physical server computer. Tests were performed using Red Hat Linux Server release 6.10.

Table 4-1 Network Prevent for Web virtual test hardware configuration

Component	AWS m4.2x large configuration
Processor (hyper-threading enabled)	8 vCPU 2.3 GHz Intel Xeon ES-2686 v4 (Broadwell)
Memory	32 GB

See “About the test methodology” on page 23.

About the test methodology

Standard hardware configurations were used to simulate a typical environment. The tested environment that is described in this section was used to provide the published results in this document.

A Server was tested using a standard set of policy types.

An auto-load traffic generator was used to simulate network traffic from iPads to a network with a Server in place. All of the data that was generated consisted of HTTP POSTs, with no FTP or HTTPS traffic. The maximum individual packet size that was used in this environment was 1095 bytes to simulate 3G data packets.

The data stream is routed through a 3G latency simulator which delays the transmission of each packet. The delay is used to simulate the network latency that is associated with 3G data communication over the Internet.

The data stream is sent from the 3G latency simulator to a proxy server. The proxy server routes the data stream to the Server for detection. After the data has been processed, the data stream is routed back to the proxy server. Then, the proxy server forwarded the data stream to a Web server acting as an HTTP sink.

Figure 4-1 test configuration

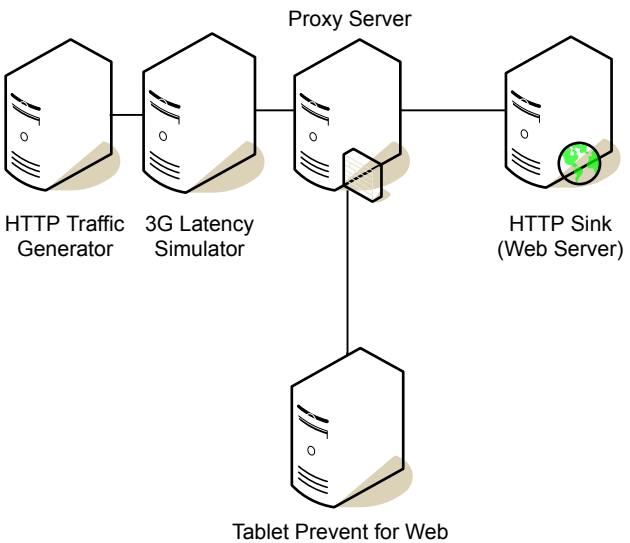

The following file types were used during testing:

Table 4-2 test data sets

Data set	File size
Small data set	Between 1 KB and 4 KB
Medium data set	Between 100 KB and 150 KB
Large data set	Between 1 MB and 3 MB
File types (by extension)	`asp, bat, C, cfm, cpp, doc, eml, gif, h, htm, html, inf, java, js, log, pdf, ppt, rtf, sha1a, shtml, txt, vbs, xls, xml, zip`

Network Prevent for Web performance test results and sizing guidelines

With Network Prevent for Web in place, performance data was determined by logging request size and request processing time. These two data points were used to determine the throughput and incremental delay. Three different data set sizes were used: large, medium, and small. The following points about the testing and the test results represented in the following tables are important to note.

- When Network Prevent for Web is tuned to the defaults presented here, it runs at about 90% capacity. Adding a spare CPU can increase the performance of your Network Prevent for Web servers. Since Network Prevent for Web is CPU bound. It is not RAM bound; adding more RAM does not necessarily improve performance.
- Each test was done with one virtual core to one message chain.
- All tests were done on AWS virtual servers.
- The test set comprised 139 policies, which together used 9 GB on the disk.
- In the following tables, latency refers to average processing time.

Table 4-3 details Network Prevent for Web performance with virtual hardware and with one virtual core to one message chain.

Table 4-3 Network Prevent for Web throughput and latency (incremental delay) test data for virtual hardware – 1 CPU to 1 message chain

Test server 1 vCPU to 1 message chain	Small data set		Medium data set		Large data set	
	Latency (milliseconds)	Throughput (Mbps)	Latency (milliseconds)	Throughput (Mbps)	Latency (milliseconds)	Throughput (Mbps)
AWS 8 vCPU 16 connections	83	10	628	47	1167	207

The results that are shown in the previous tables assume that Network Prevent for Web is configured to inspect all requests larger than 1 KB in size. The default setting is 4 KB.

With Network Prevent for Web in place, performance data was determined by logging request size and request processing time on virtual hardware. These two data points were used to determine the throughput and incremental delay. Three different data set sizes were tested: small, medium, and large.

Note: Virtual machine testing on the AWS test server for Network Prevent for Web showed average processing times per request as shown in the table for small, medium, and large data sets. Perform in-house testing with your chosen hardware, virtual machine, and operating system configuration to validate performance results before deployment.

The average processing time includes the time that Network Prevent for Web takes to receive the HTTP POST transaction (encapsulated in ICAP) from the tool. It also includes the time to perform a Data Loss Prevention inspection and send the inspected transaction back to the tool.

Your test results for your network environment may be different. There may be an issue with your network and your Network Prevent for Web Server configuration if your results diverge widely from the results that are presented here.

A variety of factors influence performance of the virtual configurations. These factors include

- Number of CPUs
- Amount of RAM
- Resource reservations for CPU cycles and RAM

The virtualization and the guest operating system overhead can lead to a modest performance degradation (5%) in web throughput of large data sets. This degradation is compared to a standard physical system running on the same hardware. Run multiple virtual instances on the same hardware to extract maximum performance and take full advantage of system resources.

See "About network performance sizing guidelines" on page 8.

Your test results should be used as a basis for sizing your Network Prevent for Web Server requirements. If your test results of a single Network Prevent for Web Server are similar to the results for the medium data set shown in Table 4-3, you should expect the results that are shown in Table 4-4.

Note: Symantec recommends that you keep your CPU usage at 90% or lower for optimum performance. Many very large enterprises with thousands of employees running hundreds of policies have higher processing requirements than those listed here, and should plan accordingly.

Table 4-4 Estimating the number of Network Prevent for Web Servers

HTTP traffic volume	Number of virtual CPUs with 1:1 ratio: 1 vCPU to 1 message chain	Number of 2 quad-core physical servers with 1:1 ratio: 1 CPU to message chain
64 Mbps - 8 vCPU	1	1
128 Mbps - 2 boxes x 8 vCPU	1	1
256 Mbps - 4 boxes x 8 vCPU	2 *	1

* Assuming 2 times the throughput for the 16 vCPU instance.

Note: When the CPU to message chain ratio is set to 1:2, CPU utilization is near 95%. Ensure that you have no other processors running and have a backup redundant Network Prevent for Web Server as CPU utilization is so high. Monitor CPU utilization closely to understand the system health under load.

All the data used and numbers obtained were for comparison with the numbers from the previous release. No new test scenarios were considered for performance. Your performance improvements may vary depending on your use cases and mix of customer data. The tuning ratio between the CPU and message chain does not change compared to the last release.

The estimates that are shown assume that:

- Traffic flows are comparable to the Medium data set, with an average file size of 220 KB.
- Equal load distribution across all servers.
- No redundancy.

Index

R

S

T

V