

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

Version 11.0



Symantec Data Loss Prevention Network Monitor and Prevent Performance Sizing Guidelines

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version: 11.0

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the *Third-Party License Agreements* document accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrades protection
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our web site at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Managed Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Education Services	Education Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about enterprise services, please visit our web site at the following URL:

www.symantec.com/business/services/

Select your country or language from the site index.

Symantec Data Loss Prevention Network Performance Sizing Guidelines

This document includes the following topics:

- [About network performance tests](#)
- [About network performance sizing guidelines](#)
- [About the Network Monitor test environment](#)
- [About the Network Monitor test methodology](#)
- [Network Monitor test results and sizing guidelines](#)
- [About the Network Prevent \(Email\) test environment](#)
- [About the Network Prevent \(Email\) test methodology](#)
- [Network Prevent \(Email\) test results and sizing guidelines](#)
- [About the Network Prevent \(Web\) test environment](#)
- [About the Network Prevent \(Web\) test methodology](#)
- [Network Prevent \(Web\) test results and sizing guidelines](#)
- [Test policy details for Network Prevent servers](#)

About network performance tests

Network Monitor, Network Prevent (Email), and Network Prevent (Web) are tested to assess their performance under load. Network Prevent (Email) and Network Prevent (Web) are also tested to compare performance between dedicated systems and virtual machine (VM) configurations.

The key objective of these tests is to obtain data on the overall performance and throughput of Network Monitor and Network Prevent and to assist customers with network sizing efforts. These tests are designed to determine the achievable throughput of different system resource configurations, and to estimate how many servers may be needed for a given policy profile and data set.

The test results provide general guidelines that a network administrator or email administrator can use to estimate the number of servers or virtual system resources that are required to support traffic loads on a network. Symantec recommends that you conduct your own testing with more representative traffic profiles and loads to validate that your results are in line with the sizing assumptions used in tests conducted by Symantec.

About network performance sizing guidelines

In general, when using a virtualized environment you should expect some performance degradation as compared to running on a dedicated system with similar system resources. Note that you may be able to minimize performance degradation by optimizing the VMware configuration specific to your environment.

Follow these guidelines when you plan any server deployment:

- The data presented in the Symantec Data Loss Prevention documentation is meant to be used as a reference for estimating deployment requirements. You should validate sizing guidelines in your own test environments before deployment.
- You should test with those policies and configurations that are consistent with expected deployments. For example, EDM-, IDM-, and DCM-based policies, configuration filters, and so on.
- You should evaluate results using a traffic profile that is consistent with your live production environment.

See [“Network Monitor test results and sizing guidelines”](#) on page 12.

See [“Network Prevent \(Email\) test results and sizing guidelines”](#) on page 16.

See [“Network Prevent \(Web\) test results and sizing guidelines”](#) on page 21.

About the Network Monitor test environment

Symantec conducted Network Monitor performance testing in a lab environment that was designed to demonstrate the comparative accuracy of all available capture methods against a replicated offered traffic load. Tests were performed using both a standard hardware configuration and large hardware configuration.

Table 1-1 describes the test hardware environment.


Table 1-1 Network Monitor test hardware

Component	Small system hardware configuration (4 cores)	Medium system hardware configuration (4 cores)	Large system hardware configuration (8 cores)
Processor	1 x Intel Xeon Processor X3220 (2.40 GHz, 1066 MHz FSB)	1 x Intel Xeon Processor X5160 (3.00 Ghz, 1333 MHz FSB)	2 x Intel Xeon Processor E5430 (2.66 GHz, 1333 MHz FSB)
Memory	8 GB RAM	16 GB RAM	16 GB RAM
Ethernet controller for testing native capture	Intel 82546EB Gigabit Ethernet Controller		
Network Monitoring Interface Card (NMIC) for testing high-speed capture	Endace DAG 4.5G2 (PCI-X) and 7.5G2 (PCI-E) cards with DAG v3.3.1 and v3.4.2 drivers, utilities, and runtime libraries.		
Network tap	A multi-port regenerative gigabit Ethernet tap facilitated distribution of the output from a traffic source to target Network Monitor servers.		

Table 1-1 Network Monitor test hardware (*continued*)

Component	Small system hardware configuration (4 cores)	Medium system hardware configuration (4 cores)	Large system hardware configuration (8 cores)
Operating system configurations	Windows Server 2003 (32-bit) with native packet capture. Windows Server 2008 R2 (64-bit) with native packet capture. Red Hat Enterprise Linux 5 (32-bit) with native packet capture. Red Hat Enterprise Linux 5 (64-bit) with native packet capture.	Red Hat Enterprise Linux 5 (64-bit) with Endace NMIC.	Windows Server 2003 (32-bit) with Endace NMIC. Red Hat Enterprise Linux 5 (32-bit) with Endace NMIC.

Figure 1-1 shows the relationship of test computers to the network traffic generator and regenerative Ethernet tap.

Figure 1-1 Network Monitor test configuration

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

- Windows systems were configured to run with a 3 GB/1 GB user/kernel address split (/3 GB /userva=3030).
 - Red Hat Enterprise Linux 5 Servers were configured in both Endace and native tests. Linux native capture was tuned to use a 128 MB ring buffer.
 - Network Monitor advanced settings were tuned as follows:
 - BUFFER_POOL_PACKETS ranged from .85 million to 1.2 million packets.
 - SMALL_POOL_PACKETS ranged from .75 million to 1 million packets.
 - The KERNEL_BUFFER_SIZE was left at the default value (16 MB for Windows 32-bit, and 64 MB for all other systems).
- For 64-bit platforms, given a kernel buffer that is large enough to handle the NIC driver's processing capability, 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 test methodology”](#) on page 12.

See [“Network Monitor test results and sizing guidelines”](#) on page 12.

About the Network Monitor test methodology

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 delivered at rates ranging from 15,000 to over 200,000 packets per second. The resulting sustained offered 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. A given Monitor under test was considered to successfully handle the offered load if it correctly generated an incident for all 20 iterations of the target file at a 100% match rate. The point at which a given capture method was no longer able to deliver total match accuracy was considered to be the limit of its performance capabilities.

See [“Network Monitor test results and sizing guidelines”](#) on page 12.

Network Monitor test results and sizing guidelines

Network Monitor servers were tested with different capture methods that accommodate different levels of network traffic. Based on this performance testing, Symantec rates the tested configurations as shown in [Table 1-2](#).

Table 1-2 Supported pre-filter performance for Network Monitor capture methods

Server configuration	Operating system	Bandwidth (Mbps)
4 cores, native packet capture	Windows Server 2003 (32-bit)	100
	Windows Server 2008 (64-bit)	300
	Red Hat Enterprise Linux (32-bit)	650
	Red Hat Enterprise Linux (64-bit)	650
4 cores, Endace NMIC	Red Hat Enterprise Linux (64-bit)	900

Table 1-2 Supported pre-filter performance for Network Monitor capture methods (*continued*)

Server configuration	Operating system	Bandwidth (Mbps)
8 cores, Endace NMIC	Windows Server 2003 (32-bit)	900
	Red Hat Enterprise Linux (32-bit)	900

The test results for your network environment may be different due to variations in the protocol composition, protocol configuration, and policy load in a production deployment. Symantec recommends testing in advance against live or recorded feeds from your production infrastructure and your target protocol/policy configuration to assess capability to meet the demands of your deployment. Note that a wide divergence of your performance numbers from those presented in this document may indicate a configuration issue 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 Monitor begins to decline for each capture method. 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. For example, assuming that your test results of a single Network Monitor Server were similar to those presented here, [Table 1-3](#) shows the number of Network Monitor servers required for different traffic levels.

Table 1-3 Estimating the number of Network Monitor servers

Network traffic (Mbps)	Windows native, 4 core servers		Linux native, 4 core servers (32-bit or 64-bit)	Endace NMIC (any tested configuration)
	32-bit	64-bit		
50	1	1	1	1
100	1	1	1	1
500	5	2	1	1
750	8	3	2	1
1,000	10	4	2	2

The above estimates assume that:

- There is equal load distribution across all servers
- There is no redundancy

See “[About the Network Monitor test environment](#)” on page 9.

See “[About the Network Monitor test methodology](#)” on page 12.

About the Network Prevent (Email) test environment

Using load generators and sample content, both standard hardware and 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 configurations described in this section.

Network Prevent (Email) servers were tested on the Symantec recommended hardware specifications for dedicated systems, and on two VM configurations with different virtual CPU resources.

[Table 1-4](#) shows the hardware and operating system configuration used for the dedicated server computer and for the virtual machine host computer.

Table 1-4 Network Prevent (Email) test hardware

Component	Dedicated server hardware configuration	Virtual machine host server hardware configuration
Processor	2 dual-core 3.0 GHz CPUs	2 dual-core 3.0 GHz CPUs
Memory	8 GB RAM	12 GB RAM
Disk space	140 GB Ultra-fast SCSI	140 GB Ultra-fast SCSI
Operating system	Microsoft Windows Server 2003 Enterprise Edition (32-bit) Microsoft Windows Server 2008 R2, Enterprise Edition (64-bit)	Red Hat Enterprise Linux 5 (32-bit and 64-bit)
NIC	Copper 1 Gb/100 Mb Ethernet NIC	Copper 1 Gb/100 Mb Ethernet NIC

Network Prevent (Email) was tested on the virtual machine host using two different VM configurations running on the following platform:

- VMware: ESX Server 3.5.0

- 8 GB VM container

Two VM configurations with a different number of virtual CPUs were tested:

- 2 CPU VM container
- 4 CPU VM container

Note that hyper-threading was not enabled for the test VM configurations.

See [“About the Network Prevent \(Email\) test methodology”](#) on page 15.

See [“Network Prevent \(Email\) test results and sizing guidelines”](#) on page 16.


About the Network Prevent (Email) test methodology

Network Prevent (Email) servers were tested using a representative set of ten policies. These policies included a variety of detection types.

See [“Test policy details for Network Prevent servers”](#) on page 24.

To simulate an email environment, an auto-load generation tool was used to send email traffic in forwarding mode between a client and server with a Network Prevent (Email) Server between them.

Figure 1-2 Network Prevent (Email) test configuration



The Network Prevent (Email) servers were tested using the same corpus 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 = 16,032
- Number of attachments = 15,807
- Attachment size = from 17 bytes to almost 3 MB
- Average size of attachments = 115.4 KB per message

- Attachments were a mixture of asp, cpp, doc, gif, gz, h, html, js, pdf, ppt, rtf, txt, vbs, xls, and zip file types

22.1% of these message attachments contained content that violated one or more of the test policies.

See “[Network Prevent \(Email\) test results and sizing guidelines](#)” on page 16.

Network Prevent (Email) test results and sizing guidelines

The following table presents benchmark results for throughput, message volume, and average transfer time that can be expected from a single Network Prevent (Email) Server.

Note: The results in [Table 1-5](#) do not include any redundancy, failover, or TLS processing requirements. Symantec tested Network Prevent (Email) TLS support in a configuration with 12 concurrent, open TCP connections between the upstream MTA and Network Prevent. In this configuration, enabling TLS processing caused a 20% reduction in throughput compared to the value shown in [Table 1-5](#). If your MTA does not optimize TLS connection setup and reuse, throughput may be reduced further due to the increased processing necessary to establish secure connections. Consult your MTA documentation and perform additional testing to evaluate TLS performance in your environment.

Table 1-5 Network Prevent (Email) test results

System configuration	Throughput (Mbps)	Message volume (messages/second)	Average transfer time (in seconds)
Standard Dedicated System (2 dual-core CPU, 8 GB RAM)	40	30	0.3
VM Container (2 CPU, 8 GB RAM)	14	10	1.2
VM Container (4 CPU, 8 GB RAM)	22	17	0.7

Prevent servers scale linearly to handle volumes in excess of the figures shown here. Most MTAs can distribute load to the corresponding Prevent servers as necessary. It is common for Prevent servers to be paired with MTAs in an N:N redundant, load-balanced configuration.

No significant performance differences were noted between 32-bit and 64-bit operating systems on dedicated hardware.

When the policy set and size of the message set is known, server requirements can be roughly estimated by extrapolating from the testing numbers shown in this document. Understanding your organization's current email traffic will help with determining how many Network Prevent (Email) servers are needed to stay within the throughput and response time limits shown here. For example, the SMTP traffic that needs to be processed in a Network deployment can be obtained 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, including the number of CPUs and amount of dedicated RAM, as well as resource reservations for CPU cycles and RAM. The virtualization/guest operating system overhead can lead to a modest performance degradation in messaging throughput compared to a standard dedicated 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 (Email) will run as its own VM image. If the MTA is also virtualized, such as is the case with Brightmail Virtual edition, then both Network Prevent (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 (Email) requirements. For example, assuming that your test results of a single Network Prevent (Email) Server were similar to those presented in [Table 1-6](#).

Note: The recommendations in [Table 1-6](#) do not account for redundancy or failover, or TLS processing requirements. TLS processing performance has not been tested in VM configurations.

Table 1-6 Estimating the number of Network Prevent Servers (Email)

Traffic volume	Number of dedicated servers needed	Number of 4-CPU VM containers needed	Number of 2-CPU VM containers needed
70 Mbps	2	4	5

Table 1-6 Estimating the number of Network Prevent Servers (Email)
(continued)

Traffic volume	Number of dedicated servers needed	Number of 4-CPU VM containers needed	Number of 2-CPU VM containers needed
80 Messages/Second	3	5	8

The estimates shown in [Table 1-6](#) assume that:

- There is equal load distribution across all servers
- There is no redundancy

Your test results for your network environment may be different. Note however, that a wide divergence of your performance numbers from the results shown above may indicate a configuration issue between your email system and the Network Prevent system.

See [“About network performance sizing guidelines”](#) on page 8.

About the Network Prevent (Web) test environment

Network Prevent (Web) servers were tested on the Symantec recommended hardware specifications for dedicated systems, and on two VM configurations with different virtual CPU resources.

[Table 1-7](#) shows the hardware and operating system configuration used for the dedicated server computer and for the virtual machine host computer.

Table 1-7 Network Prevent (Web) test hardware

Component	Dedicated server hardware configuration	Virtual machine host server hardware configuration
Processor	2 dual-core 3.0 GHz CPUs	2 dual-core 3.0 GHz CPUs
Memory	8 GB RAM	16 GB RAM
Disk space	140 GB Ultra-fast SCSI	140 GB Ultra-fast SCSI

Table 1-7 Network Prevent (Web) test hardware (continued)

Component	Dedicated server hardware configuration	Virtual machine host server hardware configuration
Operating systems tested	<p>Microsoft Windows Server 2003 Enterprise Edition (32-bit)</p> <p>Microsoft Windows Server 2008 R2, Enterprise Edition (64-bit)</p> <p>Red Hat Enterprise Linux 5 (64-bit)</p>	<p>Microsoft Windows 2003 Enterprise Edition (32-bit)</p> <p>Microsoft Windows Server 2008 R2, Enterprise Edition (64-bit)</p>
NIC	Copper 1 Gb/100 Mb Ethernet NIC	Copper 1 Gb/100 Mb Ethernet NIC

Network Prevent (Web) was tested on the virtual machine host using two different VM configurations running on the following platform:

- VMware: ESX Server 3.5.0
- 8 GB VM container

Two VM configurations with a different number of virtual CPUs were tested:

- 2 CPU VM container
- 4 CPU VM container

Note that hyper-threading was not enabled for the test VM configurations.

See [“About the Network Prevent \(Web\) test methodology”](#) on page 19.

See [“Network Prevent \(Web\) test results and sizing guidelines”](#) on page 21.

About the Network Prevent (Web) test methodology


Using load generators and sample content, both standard hardware and 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 configurations described in this section.

Network Prevent (Web) servers were tested using a representative set of ten policies. These policies included a variety of detection types.

See [“Test policy details for Network Prevent servers”](#) on page 24.

To simulate Web traffic, an auto-load generation tool was used to send HTTP POSTs through a Web proxy server (Bluecoat Model: 510-D running proxy SGOS 5.2.4.8) to a Web server acting as an HTTP sink with a Network Prevent (Web) Server in place. All traffic consisted of HTTP POSTs, with no FTP or HTTPS traffic. Each HTTP POST request consisted of a small body of text that contained no policy violation and a file attachment selected from the data set. Multiple runs for each data set were executed with each test run lasting for five minutes.

Figure 1-3 Network Prevent (Web) test configuration



Three data sets with different characteristics were used to simulate HTTP traffic.

Table 1-8 Network Prevent (Web) test data sets

	Small data set	Medium data set	Large data set
File sizes	Between 1 and 4 KB	Between 100 and 150 KB	Between 1 and 3 MB
Total number of files	440	300	140
Number of incidents	380	146	83

Table 1-8 Network Prevent (Web) test data sets (continued)

	Small data set	Medium data set	Large data set
File types (by extension)	asc, asp, bat, cfm, cpp, doc, eml, gif, h, htm, html, java, js, lnk, pdf, rtf, shala, shtml, txt, vbs, xml, zip	C, Doc, h, htm, html, js, mht, mpp, pdf, ppt, rtf, xls, zip	doc, gz, h, htm, inf, jpg, log, pdf, ppt, rtf, txt, xls, zip

The number of incidents shown for each data set specifies how many incidents were created by a single run of the data set against the test policy set.

See “[Network Prevent \(Web\) test results and sizing guidelines](#)” on page 21.

Network Prevent (Web) test results and sizing guidelines

Without Network Prevent (Web) in place, approximate throughput is shown in the following table:

Table 1-9 Proxy throughput (without Network Prevent Web)

Data set	Proxy throughput (Mbps)
Large	130
Medium	110
Small	20

With Network Prevent (Web) in place, performance data was determined by logging request size and request processing time. These two data points were used to determine the goodput and incremental delay.

Table 1-10 Network Prevent (Web) throughput and incremental delay test data for dedicated hardware

Test server	Operating system	Large data set		Medium data set		Small data set	
		Average processing time per request (milliseconds)	Throughput (Mbps)	Average processing time per request (milliseconds)	Throughput (Mbps)	Average processing time per request (milliseconds)	Throughput (Mbps)
Standard Dedicated System	Windows (32-bit)	1.14	62.5	.16	51.1	.05	2.8
	Windows (64-bit)	1.54	84.1	.14	51.1	.04	2.9
(2 dual-core CPU, 8 GB RAM)	Linux (64-bit)	1.71	81.7	.11	51.1	.01	2.9

Table 1-11 Network Prevent (Web) throughput for virtual machines

Test server	Operating system	Large data set throughput (Mbps)	Medium data set throughput (Mbps)	Small data set throughput (Mbps)
VM Container (2 CPU, 8 GB RAM)	Windows (32-bit)	29.4	12.3	1.6
	Windows (64-bit)	48	16.2	2.3
VM Container (4 CPU, 8 GB RAM)	Windows (32-bit)	58	22.3	2.3
	Windows (64-bit)	69.2	25.3	2.6

Note: Virtual machine testing for Network Prevent (Web) showed an average processing time per request ranging from 0.1 seconds to 2.1 seconds on 64-bit Windows with a 4 CPU test server. You should 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 (Web) takes to receive the HTTP transaction (encapsulated in ICAP) from the proxy, perform a DLP inspection, and send the inspected transaction back to the proxy. It does not include the time the proxy takes to intercept the HTTP transaction, transform

it to an ICAP transaction, and re-transform the ICAP response from Network Prevent (Web) to an HTTP transaction.

The above results assume that:

- Network Prevent (Web) is configured to inspect all requests larger than 1 KB in size (the default setting is 4 KB).
- The Web proxy is set to operate in transparent mode.

Tests in your network environment may have different results. However, a wide divergence of your performance numbers from those presented in this document may indicate an issue in how your network and Prevent server are configured.

A variety of factors influence performance of the virtual configurations, including the number of CPUs and amount of dedicated RAM, as well as resource reservations for CPU cycles and RAM. The virtualization/guest operating system overhead can lead to a modest performance degradation in Web throughput of large datasets compared to a standard dedicated 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.

See [“About network performance sizing guidelines”](#) on page 8.

Your own test results should be used as a basis for sizing your Network Prevent (Web) server requirements. For example, assuming that your test results of a single Network Prevent (Web) Server were similar to the Medium data set results shown in [Table 1-10](#), you should expect the following:

Table 1-12 Estimating the number of Network Prevent Servers (Web)

HTTP traffic volume	Number of dedicated servers needed	Number of 4-CPU VM containers needed	Number of 2-CPU VM containers needed
25 Mbps	1	2	2
50 Mbps	1	3	4
100 Mbps	2	5	8
200 Mbps	4	10	16

The above estimates assume that:

- Traffic flows are comparable to the Medium data set, with a size of between 100 KB to 150 KB
- There is equal load distribution across all servers
- There is no redundancy

Test policy details for Network Prevent servers

All tests of Network Prevent (Email) and Network Prevent (Web) were run using the set of Symantec Data Loss Prevention polices shown in the following table.

Table 1-13 Network Prevent test policies

Policy	Type	Comments
Credit Card Numbers	Data Identifiers (DI)	
U.S. Social Security Numbers	Data Identifiers (DI)	
State Data Privacy	Data Identifiers (DI)	
OMB Memo 06-16/FIPS 199	Keywords	
NERC	Keywords	
Encrypted Data	Keywords and metadata	
Source code	Regular expressions and metadata	
GLBA EDM Policy	EDM	1 million rows, 14.25 MB EDM. Incident created on 3 or more matches.
Fake Customer Policy	EDM	1,040,001 rows, 5.3 MB EDM. Incident created on 3 or more matches.
Longevity IDM Policy	IDM	1600 documents

See [“About the Network Prevent \(Email\) test methodology”](#) on page 15.

See [“About the Network Prevent \(Web\) test methodology”](#) on page 19.

Index

B

background loads 12
Bluecoat 20

C

capture methods 13
cores 10, 14, 16, 18

D

data sets 21

E

Endace DAG 9
environment 9
ethernet controllers 9

G

guidelines 8

H

hardware 10, 14, 19
HTTP POSTs 20
HTTP transactions 23
hyper-threading 15

I

ICAP 23
IDM policies 12
incremental delay 22

M

memory 9, 14, 18
MTAs 17

N

network cards 9, 14, 19
Network Monitor
 sizing guidelines for 13

Network Monitor (*continued*)

 test configuration for 11
 test environment for 10
 test methodology for 12
 test results for 12

Network Prevent (Email)

 sizing guidelines for 18
 test environment for 14
 test methodology for 15
 test policy for 24
 test results for 16

Network Prevent (Web)

 data sets for testing 21
 sizing guidelines for 23
 test environment for 18
 test methodology for 19
 test policy for 24
 test results for 21
 throughput results for 22

network taps 9

NIC 14, 19

NMIC 9

O

outbound messages 17

P

packet capture methods 13
performance guidelines 23
performance tests 8
policies 15, 24
processors 9, 14
protocols 11

R

RAM. *See* memory
redundancy 18, 23
request sizes 23

S

- sizing guidelines 8
- SMTP traffic 17
- SSH 11
- SSL 11

T

- telnet 11
- test environment 9, 14, 18
- test hardware 10
- test methodology 12, 15, 19
- test objectives 8
- test policies 24
- test results 12, 16, 21
- throughput 21–22
- TLS authentication 16
- transparent mode 23

V

- virtual machines 8, 14, 16–17, 19

W

- Web proxies 20, 23