MailMarshal 6.X Policy Implementation Sizing Guide

Contents

This document has been prepared to provide an insight into the impact that different policy configuration and third party scanners has on mail flow in MailMarshal SMTP 6.X.

The information in this document is generally applicable to all releases of MailMarshal SMTP 6.X. The testing and research was performed with version 6.2.

OVERVIEW

This paper has been prepared to provide an insight into the impact that different policy configuration and third party scanners have on mail flow in MailMarshal SMTP 6.X.

The paper is designed to be used in conjunction with the following Whitepapers also available on the M86 Security website:

- MailMarshal SMTP 6.X Sizing Guide
- MailMarshal SMTP 6.X Performance Benchmarking

By reading these papers together, the reader will be able to gain an understanding of the factors affecting MailMarshal performance and suggestions for maximizing performance.

Information on sizing and hardware requirements for each release is available in the *User Guide* and Release Notes.

POLICY IMPACT ON PERFORMANCE

MailMarshal SMTP 6.X is a flexible software-based email content security product. It is capable of processing sustained mail loads into the millions of messages per day. However, as the volume of mail increases, it is important to consider the impact of adding mail processing policies, to determine whether additional mail processing nodes may be required.

This paper focuses only on the MailMarshal Engine, where an incoming message spends most of its life. The Receiver can also apply rules to messages, and some of the information below is applicable to it. However, the performance bottlenecks for both the Receiver and Sender services are typically related to network performance – that is, how fast a message can be received from or sent to a remote host.

Individual policies within MailMarshal SMTP 6.X are made up of three parts: user matching criteria, rule conditions, and rule actions. These components are combined to create complex, flexible policies to meet the business needs of customers.

Each policy that is added necessarily impacts the speed at which MailMarshal SMTP 6.X can process messages. However, the performance impact varies considerably, depending on which items are selected within the policy. The tables below describe the expected performance impact of selecting each of these items.

User Matching

All user matching criteria in MailMarshal SMTP 6.X have a negligible performance impact on rule processing. The number of members contained in the various MailMarshal user groups can impact performance somewhat, but typical installations with user groups containing a few thousand members do not require any special provisions to minimize impact.

Engine Rule Conditions

Condition Name	Performance Impact	Explanation
Where message attachment is of type file types	Very Low	MailMarshal SMTP 6.X identifies file types as it unpacks each message into its constituent components. Therefore, little additional work is required.
Where attachment fingerprint is/is not known	Low (variable)	This condition's performance impact is directly related to the size of the fingerprint database. In most companies, the database consists of a few hundred images, so the impact is minimal.
Where message size is message size	Very Low	MailMarshal SMTP 6.X sees the file size as it unpacks each message. Therefore, little additional work is required.
Where the estimated bandwidth required to deliver this message is bandwidth	Very Low	This condition requires a simple multiplication of message size by unique recipient domains.
Where message contains attachments named file name	Very Low	MailMarshal SMTP 6.X identifies file names as it unpacks each message. Therefore, little additional work is required.
Where message triggers text censor script(s) scripts	Low (variable)	The MailMarshal SMTP 6.X TextCensor feature is very fast, but its speed relates to the number of tokens in the TextCensor script. Typically, TextCensor scripts consist of a few hundred tokens.
Where the result of a virus scan is scanner result	High to Very High	The impact of a virus scanner is always large, but DLL-based virus scanners integrated with MailMarshal SMTP 6.X are orders of magnitude faster than command-line counterparts.
Where the external command command is triggered	Medium to Very High (variable)	External commands require significant overhead for the Engine to execute, but the actual performance impact is mostly determined by the content of the command.
Where the attachment parent is of type parent types	Very Low	MailMarshal SMTP 6.X identifies attachments and their parents as it unpacks each message into its constituent components. Therefore, little additional work is required.
Where message attachment size is file size	Very Low	MailMarshal SMTP 6.X identifies attachments sizes as it unpacks each message. Therefore, little additional work is required.
Where number of recipients is count	Very Low	MailMarshal SMTP 6.X determines the number of recipients during each message's transmission. Therefore, little additional work is required.
Where message contains one or more headers header match	Very Low to Low	Header matches can use multiple regular expressions. The complexity and number of regular expressions can affect the performance impact.
Where number of attachments is count	Very Low	MailMarshal SMTP 6.X determines the attachment count during the unpacking process for each message. Therefore, little additional work is required.
Where message is categorized as 'Spam'	High	SpamCensor is a complicated category script that contains thousands of heuristic tests to determine the probability that a particular message is spam.
Where message is categorized as 'URLCensor'	Very High	The URLCensor category script performs DNS lookups to a DNS-based blacklist. Rule performance depends on the speed with which the DNS server replies to these queries.
Where message is categorized as 'Spamhaus'	Very High	The Spamhaus category script performs DNS lookups to a DNS-based blacklist. Rule performance depends on the speed with which the DNS server replies to these queries.

Where message is categorized as 'CountryCensor'	Low	The CountryCensor category script checks IP addresses contained within each message header against a fast, custom database of IP address ranges.
Where message is categorized as category	Very Low to Very High	Category script performance is completely dependent on the script's content. Typical custom category scripts are quite small and specialized, with few tests.
Where message spoofing analysis is based on criteria	Very Low	Spoofing is determined by comparing the source IP address and email domain to a list of IP addresses held in memory.
Where the sender is/is not in the recipient's safe senders list	Low	MailMarshal SMTP 6.X keeps safe senders lists in memory.
Where the sender is/is not in the recipient's blocked senders list	Low	MailMarshal SMTP 6.X keeps blocked senders lists in memory.
Where the attached image is/is not/maybe inappropriate	Medium	MailMarshal SMTP 6.X uses Image Analyzer to determine an image's appropriateness. This deep image scanning can be costly, but the condition is only performed on messages with images. Therefore, the total impact is medium.

Engine Rule Actions

Action Name	Performance Impact	Explanation
Copy the message to folder	Low	MailMarshal SMTP 6.X copies the message to a nominated quarantine folder.
BCC a copy of the message to Email address	Low	A duplicate of the message is created and sent to a particular email address. This introduces extra sending overhead, but doesn't take much effort for the MailMarshal Engine.
Run the external command	Medium to Very High	External commands require significant overhead for the Engine to execute, but the actual performance impact is mostly determined by the content of the command.
Send a mail template notification message	Low	A message is generated from a template and placed into the queue. This introduces extra sending overhead, but doesn't take much effort for the MailMarshal Engine.
Strip attachment	Medium	Because this action changes the content of the message, MailMarshal is forced to repack the message.
Write log message(s) with classifications	Low	A classification is sent from the Controller to the Array Manager. This introduces some additional database load, but takes little effort for the Engine to generate.
Stamp message with message stamp	Medium	Because this action changes the content of the message, MailMarshal is forced to repack the message.
Rewrite message headers using expressions	Medium	Header rewrites can use multiple regular expressions. The complexity and number of regular expressions can affect the performance impact. Because this changes the content of the message, MailMarshal is forced to repack the message.
Add attachments to valid fingerprints list	Low	Selected attachments are copied to the fingerprints database.
Set message routing to host	Very Low	Message is re-routed to a particular IP address by adding a field to the MailMarshal envelope.
Add message users into group	Low	MailMarshal SMTP adds additional users into a group.

Move the message to folder	Low	MailMarshal SMTP 6.X moves the message to a nominated quarantine folder.
Park the message in folder	Low	MailMarshal SMTP 6.X moves the message to a nominated parking folder.
Delete the message	Low	MailMarshal SMTP 6.X deletes the message.
Pass message to rule	Very Low	MailMarshal SMTP 6.X moves forward to a particular portion of the email policy.

TESTING THE IMPACT

MailMarshal SMTP 6.X has many configuration options that can impact its performance. M86 Security tested many of these options to measure the approximate performance changes in a real-world environment.

The results of each test are broken into three figures: the average message per second throughput per node, the average message per second change per node, and the average time impact percentage per node.

M86 Security has included specific figures in this paper on the elapsed time of each test and the average number of messages processed per second. However, these specific figures can vary greatly depending on the environment where MailMarshal SMTP is installed. Important factors that will change these figures are the speed of the processor, the amount of memory installed on the server, the speed of the disk subsystem, the responsiveness of DNS servers, the content of the mail flowing through MailMarshal SMTP, and other processes running on the MailMarshal SMTP server.

Therefore, the reader is encouraged to look at their own baseline message per second rate, and use the cited percentages to approximate the expected impact, rather than applying the gain or loss in messages per second values directly.

Testing Methodology

Each MailMarshal SMTP email processing node was fed 100,000 email messages collected from a real world corporate environment. These messages were submitted to the MailMarshal servers at the maximum rate that the servers could accept them.

After each test, the MailMarshal performance counters and MailMarshal text logs were analyzed and the processing times for each rule were collected and averaged.

Hardware and Software Environment

The testing environment consisted of a single Windows Server 2003 R2 machine running the MailMarshal Array Manager and Microsoft SQL 2005 Server environment. This Array Manager managed three Windows Server 2003 R2 Servers running MailMarshal Nodes.

The mail corpus

80,0000 incoming spam messages

12,0000 incoming ham messages

8,000 outgoing ham messages

The Core Hardware for the servers

MailMarshal Array Manager (ARRAYMANAGER)

HP ProLiant DL380

2GB RAM

2 Intel P4 2.8GHz CPU's

MailMarshal email processing nodes (NODE1, NODE2, NODE3)

IBM System x3550 Server

2GB RAM

1 Intel Dual Core P4 2.33GHz CPU

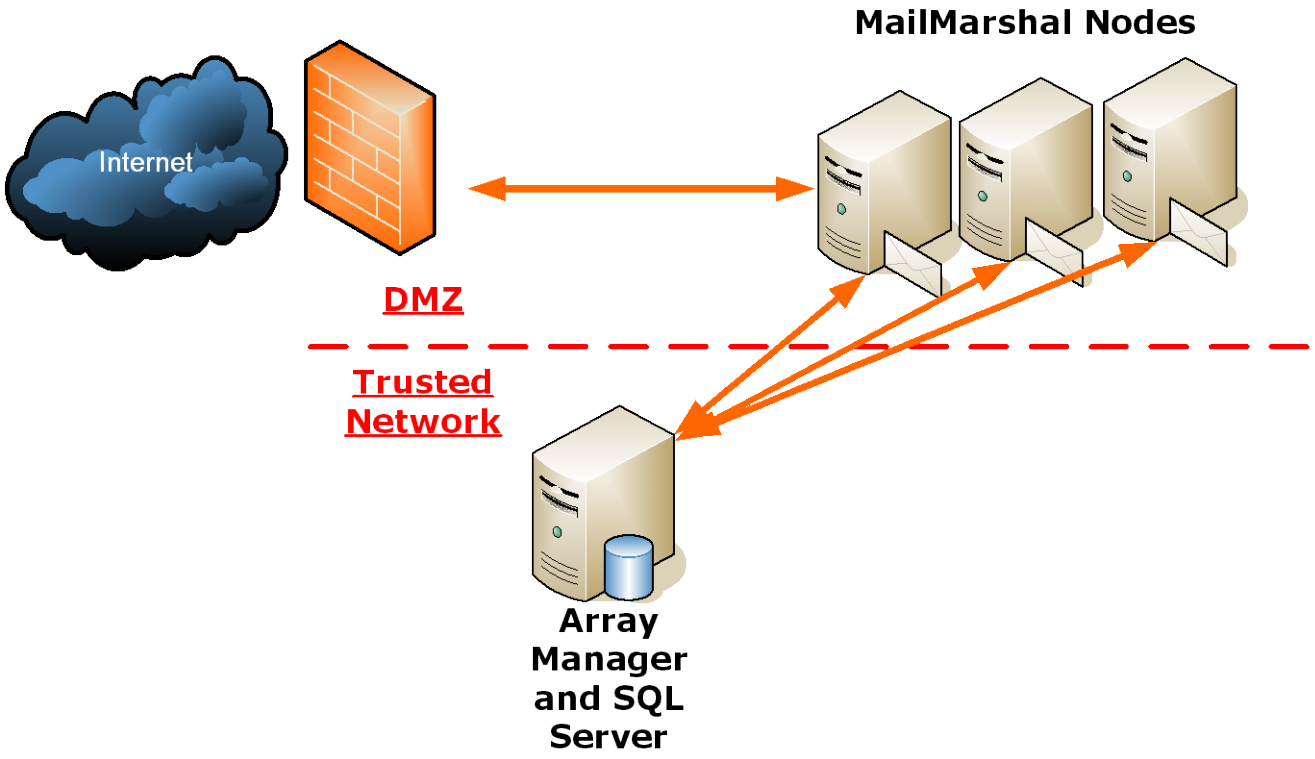

TEST RESULTS

Default Rules

The first test M86 Security performed was running the test corpus through an unmodified ruleset. This ruleset includes message archiving, anti-spam scanning using SpamCensor, TextCensor scripts, and DNS blacklists, attachment blocking, and profanity blocking.

This test was used as a baseline, to determine the speed at which a basic installation processed messages through its Engine on the test hardware.

Test Name	Average Msgs/Sec per Node	Average Msg/Sec Change per Node	Average Time Impact Percentage
Default Rules	31.191	N/A	N/A

Virus Scanners

The McAfee for Marshal virus scanner add-on was added to the default ruleset. Inbound and outbound virus scanning was enabled.

Of all the policy elements available in MailMarshal SMTP, virus scanning rules have the largest performance impact.

The virus scanners available for the MailMarshal SMTP product come in two basic types: DLL-based and command line scanners. The command line scanners, while not tested for this paper, have been found to be several orders of magnitude slower than DLL-based virus scanners like McAfee for Marshal.

Rules Enabled:

Virus & Threats (Inbound)/Block Virus

Virus & Threats (Inbound)/Virus Scanner Errors

Virus & Threats (Outbound)/Block Virus

Virus & Threats (Outbound)/Virus Scanner Errors

Test Name	Average Msgs/Sec per Node	Average Msg/Sec Change per Node	Average Time Impact Percentage
Default Rules	31.191	N/A	N/A
Virus Scanner (McAfee)	20.947	-10.244	-32.842%

Image Analyzer

Image Analyzer is an add-on available for MailMarshal SMTP 6.X that scans attached images for suspected pornographic content. It is relatively lightweight, but its total performance impact is largely dependent on the quantity of images found within the mail stream. For this test, both inbound and outbound image analysis was enabled.

Rules Enabled:

Attachment Management (Inbound)/Block Suspected Pornographic Images

Attachment Management (Outbound)/Block Suspected Pornographic Images

Test Name	Average Msgs/Sec per Node	Average Msg/Sec Change per Node	Average Time Impact Percentage
Default Rules	31.191	N/A	N/A
Image Analyzer	29.492	-1.699	-5.447%

Spyware Scanners

Spyware scanners work similarly to virus scanners, but attempt to detect unwanted or malicious code from unethical companies. Because they have smaller signature databases, they tend to be faster than virus scanners, but the performance impact is still significant.

MailMarshal SMTP 6.X currently offers two spyware scanners – PestPatrol and Counterspy. Both were tested on inbound and outbound email.

Rules Enabled:

Virus & Threats (Inbound)/Block Spyware

Virus & Threats (Outbound)/Block Spyware

Test Name	Average Msgs/Sec per Node	Average Msg/Sec Change per Node	Average Time Impact Percentage
Default Rules	31.191	N/A	N/A
Spyware Scanner (PestPatrol)	29.241	-1.950	-6.252%
Spyware Scanner (Counterspy)	30.353	-0.838	-2.687%

TextCensor and Category Scripts

MailMarshal SMTP 6.X is able to scan email for SEC or SOX compliance related material, racist language, social security numbers, and other items largely through its TextCensor feature or its category script feature.

Both TextCensor scripts and category scripts are lightweight, and typically have a fairly negligible performance impact. In this case, five TextCensor scripts and two category scripts were enabled with a combined total of over five hundred individual tests, with slightly over a 1% performance impact.

Rules Enabled:

Policy Management (Inbound)/Social Security Number detection

Policy Management (Inbound)/SEC Compliance Rule

Policy Management (Inbound)/Credit Card Number detection

Policy Management (Inbound)/Sarbanes-Oxley Compliance Rule

Policy Management (Inbound)/Racist and Hate content

Policy Management (Inbound)/Weapons related content

Policy Management (Outbound)/Block Common & Mild Profanity

Policy Management (Outbound)/Social Security Number detection

Policy Management (Outbound)/SEC Compliance Rule

Policy Management (Outbound)/Credit Card Number detection

Policy Management (Inbound)/Sarbanes-Oxley Compliance Rule

Policy Management (Inbound)/Racist and Hate content

Test Name	Average Msg/Sec per Node	Average Msg/Sec Change per Node	Average Time Impact Percentage
Default Rules	31.191	N/A	N/A
SOX, SEC, and Racism Checks	30.858	-0.333	-1.068%

File Type Identification

Many MailMarshal SMTP 6.X rules cause negligible performance impact because much of the work is done during the unpacking process. For example, a policy can be based on the file type condition. During the unpacking process, MailMarshal SMTP recursively identifies and unpacks message components. Once these components are fully unpacked, the Engine begins processing rules against the message contents.

For this test, both inbound and outbound image blocking was enabled.

Rules Enabled:

Attachment Management (Inbound)/Block IMAGE Files

Attachment Management (Outbound)/Block IMAGE Files

Test Name	Average Msg/Sec per Node	Average Msg/Sec Change per Node	Average Time Impact Percentage
Default Rules	31.191	N/A	N/A
Block Images	31.167	-0.024	-0.769%

MINIMIZING PERFORMANCE IMPACT

MailMarshal SMTP 6.X Engine performance can be improved by adjusting both the policy and the environment in which MailMarshal operates.

Careful Policy Placement

Attempt to place higher-impact rules later than lower-impact rules. Any mail that is quarantined early in the email policy will not have the costly rules applied against it.

One example of judicious policy placement is to run the SpamCensor rule before the URLCensor or Spamhaus rules. DNS lookups take a longer amount of time, on average, than checking the message with the SpamCensor.

On the test systems used to research this paper, the SpamCensor rule takes on average 18ms to run, while the Spamhaus rule takes anywhere from 150ms to 300ms. At the time of this writing, SpamCensor catches roughly 96% of the total spam quarantined by MailMarshal. By ensuring that the SpamCensor is run first, only around 4% of spam messages require that an expensive DNS lookup be performed.

Reject Messages at the Receiver

MailMarshal SMTP 6.X Receiver rules can be used to block messages outright, ensuring that they are never processed by time-consuming Engine rules. If, for example, the Receiver blocks messages using the Spamhaus DNS blacklist, obvious spam will never be run against virus scanners or SpamCensor. Relying solely on the Spamhaus Engine rule to perform DNS blacklist checks incurs a significant penalty.

Use Local DNS Servers

Many DNS blacklist providers, such as Spamhaus, offer a blacklist feed to companies for a nominal fee. By mirroring an oft-used DNS blacklist locally, a MailMarshal administrator ensures that DNS network latency is minimal.

Experiment with Engine Threads

MailMarshal SMTP 6.X uses *number of processors* + *1* Engine threads by default. For example, on a dual-processor machine, 3 Engine threads will be used.

However, this value can be modified by a MailMarshal administrator to improve performance in some cases. Some tasks that MailMarshal SMTP 6.X performs, such as DNS lookups or calls to virus scanners, can take quite some time. The Engine must idle while waiting for a response.

By increasing the available Engine threads, MailMarshal SMTP can use this idle time to process additional messages. However, if the thread count is too high, the MailMarshal Engine will spend much of its time context-switching, and performance can dramatically decrease. Most administrators will have to experiment with this value to find one that suits the environment.

ABOUT M86 SECURITY

M86 Security is a global provider of Web and messaging security products, delivering comprehensive protection to more than 20,000 customers and over 16 million users worldwide. As one of the largest independent internet security companies, we have the expertise, product breadth and technology to protect organizations from both current and emerging threats. Our appliance, software and cloud-based solutions leverage real-time threat data to proactively secure customers' networks from malware and spam; protect their sensitive information; and maintain employee productivity. The company is based in Orange, California with international headquarters in London and offices worldwide. For more information about M86 Security, please visit www.m86security.com.

TRY BEFORE YOU BUY

M86 Security offers free product trials and evaluations. Simply contact us or visit www.m86security.com/downloads

Corporate Headquarters
828 West Taft Avenue
Orange, CA 92865
United States

Phone: +1 (714) 282-6111
Fax: +1 (714) 282-6116

International Headquarters
Renaissance 2200
Basing View, Basingstoke
Hampshire RG21 4EQ
United Kingdom
Phone: +44 (0) 1256 848080
Fax: +44 (0) 1256 848060

Asia-Pacific
Millennium Centre, Bldg C, Level 1
600 Great South Road
Ellerslie, Auckland, 1051
New Zealand
Phone: +64 (0) 9 984 5700
Fax: +64 (0) 9 984 5720

Version 25.06.10