



**The Symantec  
Mail Security™ 8160 Appliance**  
A Technical White Paper on  
Antispam Defenses at the  
TCP/IP Level

# The Symantec Mail Security 8160 Appliance

## A Technical White Paper on Antispam Defenses at the TCP/IP Level

### Contents

<b>Introduction</b> .....	4
<b>The power of layering</b> .....	5
<b>Narrowing the pipe</b> .....	5
TCP/IP level antispam defenses .....	6
SMTP-level antispam defenses .....	7
The Symantec Mail Security 8160 appliance .....	8
<b>Characterizing sources</b> .....	10
Known non-spam sources .....	10
Known spam sources .....	11
Unknown sources .....	11
Suspicious sources .....	12
Gradations in-between .....	12
<b>Upstream effects</b> .....	12
Upstream effect on spam sources .....	12
Upstream effect on valid email sources .....	13
<b>Downstream effects</b> .....	14
<b>Conclusion</b> .....	14
<b>About the author</b> .....	16

## Introduction

Unsolicited bulk email (spam) is now a major concern of IT managers. Having defended users' PCs against email-borne viruses, managers now have to protect user mailboxes from being swamped under a barrage of, at best unwanted, and often offensive (pornographic) or fraudulent (phishing), email messages.

Managers have responded by installing antispam defenses at the message content level, scanning incoming messages in an attempt to distinguish valid email messages, which are then delivered, from spam, which is then discarded or quarantined.

Unfortunately, managers have discovered that, while these antispam defenses significantly reduce the amount of spam arriving in users' mailboxes, three major problems still remain:

1. The volume of incoming spam continues to rise inexorably at a rate determined by the spammers. Thus, the amount of time that users have to spend searching for false positives continues to grow. The number and capacity of servers that must be devoted to message content analysis continues to grow.
2. The amount of disk storage allocated to spam quarantine also continues to grow, and for organizations that are required to retain all incoming mail, the size and complexity of email archives explodes.
3. False positives remain a major headache. This is because there is a message content analysis trade-off between false negatives and false positives. In order to reduce one, a content analyzer has to increase the other. There is also the problem that the onus of dealing with false positives falls on the recipient, who has to scan large numbers of quarantined messages looking for relatively rare false positives.

In the remainder of this paper we will examine approaches that can be employed at other protocol levels to address these three problems.

## **The power of layering**

While installing antispam defenses at the message content level is an obvious first step, there are two other protocol levels at which antispam defenses can operate: the SMTP level and the TCP/IP level.

Suitable antispam measures operating at the TCP/IP level can both reduce the absolute volume of incoming spam requiring message content analysis and, more importantly, flatten its growth.

Suitable antispam measures operating at the SMTP level can also reduce the absolute volume of incoming spam requiring message content analysis, but it is not clear that it will flatten its growth.

Antispam measures operating at the TCP/IP or SMTP level can have two impacts on the false positive problem:

1. Reducing the absolute volume of inbound spam allows the content analysis trade-off between false negatives and false positives to be adjusted in favor of significantly fewer false positives. Stated another way, for the same absolute volume of spam (false negatives) reaching users' mailboxes, significantly fewer valid messages (false positives) will be quarantined.
2. A failure to accept a valid email message at the TCP/IP or SMTP level will result in a non-delivery notification being returned to the sender, who can then take remedial action. For a variety of reasons, non-delivery notifications are not produced when a message is identified (misidentified) as spam at the message content analysis level.


In the remainder of this white paper, we will examine the operation and impact of antispam defenses operating at the TCP/IP and SMTP levels.

## **Narrowing the pipe**

All Internet email entering an organization is transferred over an SMTP transfer channel between a sending Mail Transfer Agent (MTA), acting on behalf of a sending organization, and a receiving MTA, acting on behalf of a receiving organization.

Each SMTP transfer channel is realized atop a TCP/IP transmission channel. Both a sending MTA and a receiving MTA are able to establish and operate a large number of SMTP transfer channels in parallel. Similarly, both a sending and a receiving organization are able to deploy a large number of sending and receiving MTAs. The email transfer capacity between any two organizations is, thus, potentially very large, consisting of many SMTP transfer channels between

many MTAs (see Figure 1). Spammers exploit this large transfer capacity to bombard an organization with large volumes of subtly different spam, hoping that a few will get through message content-level defenses.



**Figure 1. A high-volume SMTP source can establish many TCP/IP channels to an SMTP target**

The objective of antispam defense operating at the TCP/IP or SMTP level is to reduce this transmission/transfer capacity. While there is some overlap, differences between the two protocols mean that the antispam techniques that can be employed to reduce capacity at each level differ.

### TCP/IP-level antispam defenses

The Transmission Control Protocol (TCP) is layered atop the Internet Protocol (IP), and is referred to in composite as TCP/IP.

The IP is concerned solely with the layout of data packets, and the structure of source and target address information within those data packets. Target address information in each data packet is employed, in conjunction with routing data, to select an output channel. This process allows a transfer to be effected from an IP packet source, through 0 or more intermediates (routers), to an IP packet target. There are no error controls that operate at the IP level. Individual IP packets can be garbled in transmission, time-out, or otherwise go missing. It is, thus, the responsibility of higher-level protocols to detect and deal with errors. TCP/IP is one of these higher-level protocols.

The TCP/IP establishes a TCP/IP session between an initiator and a target. It then operates an error-free, full-duplex, rate-limited transmission channel between these two TCP/IP end points.

To achieve this, the TCP/IP protocol:

- Provides controls to establish and terminate a TCP/IP session
- Adds a checksum and a sequence number to each IP packet (allowing it to discard corrupted packets and re-order packets that arrive out of order)
- Provides a means of dynamically signaling both the successful (error-free) receipt of a complete sequence of IP packets, and the number of IP packets (the window) that may be transmitted in advance of receiving a signal of receipt

Antispam measures that operate at the TCP/IP level can be driven by source IP address and observed SMTP command traffic, and operate by:

- Reducing the transmission capacity of a TCP/IP channel
- Limiting the number of email messages sent per TCP/IP session
- Limiting the establishment of parallel TCP/IP sessions
- Limiting the rate at which TCP/IP sessions can be established

### **SMTP-level antispam defenses**

The Simple Mail Transfer Protocol (SMTP) is a half-duplex command/response protocol operated over a full-duplex TCP/IP transmission channel. An SMTP session is established by a MTA—referred to as an SMTP client—that wishes to transfer a sequence of one or more email messages to an MTA acting on behalf of one or more target domains—referred to as an SMTP server. (The terms SMTP client and SMTP server refer to their roles in the SMTP protocol, and not to the type of system being employed.)

Before the TCP/IP layer acting on behalf of an application accepts a TCP/IP session, it asks the application via an API whether it wishes to accept or reject that session (from a specified IP source address). Under these circumstances, an SMTP server can reject TCP/IP sessions from a putative spam source, in the same way the TCP/IP layer can.


An SMTP server cannot reduce the transfer capacity of an SMTP transfer channel in the same way that the transmission capacity of a TCP/IP transmission channel can be reduced. This is because the only control mechanisms that SMTP servers have are either to delay an SMTP response to a received SMTP command or to delay receiving SMTP data over a TCP/IP API.

Delaying an SMTP response has no impact on a spam source, which usually ignores responses in any case; and delaying the reading of data over a TCP/IP API, if it can be done, will only have an effect when the TCP/IP window is exceeded. For email messages that fit within the TCP/IP window, it will have no effect. This means that reducing the capacity of a TCP/IP transmission channel, and therefore an SMTP transfer channel, can only be achieved at the TCP/IP level.

What antispam defenses operating at the SMTP level can achieve, which defenses operating at the TCP/IP level cannot, are those that are triggered by the content of SMTP commands themselves. For example, an SMTP server can use Sender ID–defined mechanisms to validate the SMTP MAIL FROM address or SUBMITTER parameter as being compatible with the IP address from which the TCP/IP session was established, and terminate an SMTP session with a failure response if it is not.<sup>1</sup>

### The Symantec Mail Security 8160 appliance

The Symantec Mail Security 8160 appliance implements a comprehensive set of antispam defenses at the TCP/IP level. Their objective is to significantly reduce the transfer capacity available to spammers, while continuing to maintain it for legitimate sources of email. This is achieved by interposing a Symantec Mail Security 8160 appliance between an organization's inbound SMTP port firewall and its inbound SMTP MTAs (see Figure 2).



**Figure 2. Symantec Mail Security 8160 appliance interposed between firewall and MTAs**

<sup>1</sup>Sender ID is the name given to a unification of the Sender Policy Framework (SPF) and Microsoft's Caller ID proposal. This unification was originally performed under the auspices of the Internet Engineering Task Force (IETF) MTA Authorization Records in DNS (marid) working group. As the marid working group was unable to reach rough consensus, largely due to concerns about Microsoft intellectual property licensing terms, Sender ID is now progressing as a private effort by Microsoft and others. Details can be found at <http://spf.pobox.com> and [www.microsoft.com/mscorp/twc/privacy/spam/senderid/default.mspx](http://www.microsoft.com/mscorp/twc/privacy/spam/senderid/default.mspx).

## The Symantec Mail Security 8160 Appliance

A Symantec Mail Security 8160 appliance reduces the transfer capacity available to both known and putative spammers in four dimensions.

- It reduces the bandwidth of a spammer's individual TCP/IP transmission channels and, therefore, his or her individual SMTP transfer channels. It achieves this by two means. It reduces the TCP/IP window size—the number of IP packets that a TCP/IP receiving system indicates to a sending system that it may transmit in advance of receiving an acknowledgement for an earlier IP packet. It also reduces the TCP/IP acceptance rate—the number of IP packets that a TCP/IP receiving system acknowledges per unit time. A primary role of the second form of reduction is to enforce the window size even when a sending system violates the protocol and ignores it.
- It limits the number of email messages that a spammer may send per TCP/IP session. It achieves this by tearing down a TCP/IP session following a prescribed number of completed SMTP DATA commands. This further reduces the capacity of an individual TCP/IP transmission channel.
- It limits the number of individual TCP/IP sessions that a spammer can establish per unit time. It achieves this by tracking the TCP/IP bind frequency from a single source, and rejecting those that exceed a prescribed threshold.
- It limits the number of parallel TCP/IP sessions that a spammer can establish. It achieves this by tracking TCP/IP session concurrency from a single source, and rejecting those that exceed a prescribed threshold.

These four actions have a profound effect on the amount of spam reaching an organization's inbound SMTP MTAs—both in the short and the medium to longer term. The effect in the short term is obvious: By narrowing both the message transfer capacity and the number of SMTP transmission channels, the volume of spam from a single spam source (one or more MTAs) will be very significantly reduced. But this in turn produces a positive effect. To maintain their efficiency, spammers need to transfer large numbers of messages per unit time; when they fail to do so, they need to take remedial action. Their only choice is to eliminate problematic target domains from their target address lists. They give up trying to send to them. So in the medium to longer term, incoming spam volumes also decrease. This effect is very noticeable when one tracks message volumes through a Symantec Mail Security 8160 appliance. Rather than continuing on a never-ending upward trend, email volumes have been observed to flatten out and remain flat (see Figure 3).




Figure 3. Symantec Mail Security 8160 appliance inbound email volume over time

### Characterizing sources

In the preceding section, we described how a Symantec Mail Security 8160 appliance narrows the width and number of TCP/IP transmission channels employed by SMTP as a means of reducing spam volumes. The degree to which the number and capacity of TCP/IP transmission channels are narrowed is dependent on identifying the source of an incoming TCP/IP session. Is it from an unknown source, a suspicious source, a known spam source, a known non-spam source, or some gradation in between?

### Known non-spam sources

Following installation, a Symantec Mail Security 8160 appliance rapidly builds and maintains a list of known non-spam sources and, more crucially, known high-volume non-spam sources. This is based on an ongoing combination of protocol analysis (see “Suspicious sources,” below), subsequent content analysis, and reputation (provided by the Symantec Sender Reputation Service). High-volume, non-spam sources are allocated large TCP/IP windows, can transfer many email messages per TCP/IP session, and are allowed many parallel TCP/IP transmission channels. Low-volume, non-spam sources are allocated large TCP/IP windows, can transfer many email messages per TCP/IP session, and are allowed a small number of TCP/IP transmission channels.

This approach differs from a conventional white list in that it operates at the IP address range level, and non-spam sources are both added and removed dynamically. A Symantec Mail Security 8160 appliance is not dependent on subscription updates for its identification of a non-spam source.

### **Known spam sources**

Following installation, a Symantec Mail Security 8160 appliance rapidly builds and maintains a list of known spam sources. This is based initially on their reputation (provided by the Symantec Sender Reputation Service), and subsequently on ongoing protocol (see “Suspicious sources,” below) and content analysis. Known spam sources are allocated small TCP/IP transmission windows, are only allowed to transmit one email message per TCP/IP session, and can only establish a very small number of TCP/IP sessions per unit time. If a spammer continues to transmit under these constraints (they may well not), then this allows a small number of email messages to be received and analyzed in order to maintain an accurate characterization of the source. It is assumed that downstream SMTP MTAs and content analysis filters will appropriately characterize and block the small volume of putative spam received from known spam sources.

This approach differs from a conventional black list in that it operates at the IP address range level, and spam sources are both added and removed dynamically based on their actual behavior. A Symantec Mail Security 8160 appliance is not dependent on subscription updates for its identification of a spam source.

### **Unknown sources**

When a TCP/IP bind is received from a new source, a Symantec Mail Security 8160 appliance employs its reputation (provided by the Symantec Sender Reputation Service) to determine the TCP/IP transmission window size and transmission channels it will be initially allocated. In the absence of a positive reputation, TCP/IP sessions from unknown sources will be moderately restricted, until their behavior and the email messages received from them can be analyzed in order to more accurately characterize the source. It is assumed that downstream SMTP MTAs and content analysis filters will appropriately characterize and block the small volume of putative spam received from previously unknown sources.

## **Suspicious sources**

Suspicious sources are sources that manifest illegal or inappropriate TCP/IP or SMTP behavior. For example, they employ forged IP addresses, ignore the window size, don't wait for SMTP responses, and so on. In the absence of other positive indications based on reputation and analyzed message content, suspicious sources are treated as spam sources.

## **Gradations in-between**

A key feature of the Symantec Mail Security 8160 appliance is that it does not operate an "all or nothing" approach by either accepting or rejecting all TCP/IP binds. Instead it dynamically grows and shrinks the aggregate TCP/IP transmission capacity offered to an SMTP source based on the observed behavior of that source.

## **Upstream effects**

TCP/IP sessions are established on behalf of an upstream SMTP MTA acting on behalf of an organization that wishes to transfer email to one or more SMTP MTAs acting on behalf of one or more target Internet domains. We have already described how a Symantec Mail Security 8160 appliance restricts the aggregate capacity of TCP/IP transmission channels from SMTP MTAs acting on behalf of putative spammers. In this section we will consider the effects of this restriction on both spammers and valid email sources.

### **Upstream effect on spam sources**

When a Symantec Mail Security 8160 appliance restricts the aggregate capacity of TCP/IP transmission channels from SMTP MTAs (spam cannons) acting on behalf of a spam source, it significantly reduces the amount of incoming spam that they will be able to transmit to a target domain. This has an immediate benefit to the protected domains.

In the longer term, as the effectiveness of a spam cannon is measured by the number of spam messages they successfully transmit in a unit of time, any slowdown in message transmission capacity to a target domain has a negative impact on their effectiveness. Spammers react to this reduction in effectiveness by automatically or manually blocking transmission to MTAs that produce this reduction in effectiveness.

### **Upstream effect on valid email sources**

There will be circumstances in which a Symantec Mail Security 8160 appliance restricts the aggregate capacity of TCP/IP transmission channels from SMTP MTAs acting on behalf of a valid source. In the case of a new low-volume source without any reputation, this restriction will have little noticeable impact. Messages will take longer to transmit, but that's it. At low volume, this is almost undetectable.

What about a new high-volume source without a reputation? There are two issues here. First, what form might such a beast take? High-volume sources of valid email don't just appear without warning one day. Yes, an organization may deploy a new outbound MTA on a new IP address, but the Symantec Mail Security 8160 appliance will detect this new IP address as belonging to an existing high-volume valid source and not restrict it. And in any case, if one does, subsequent analysis of received content by the Symantec Mail Security 8160 appliance will cause it to reclassify the source as valid and over time increase the capacity, duration, frequency, and concurrency of the TCP/IP transmission channels it can establish.

Finally, valid email is transmitted by MTAs acting on behalf of valid users in an organization or subscribers to an ISP. Organizational MTAs serve closed domains that are presumed to employ valid SMTP MAIL FROM addresses. ISP MTAs are able to validate SMTP MAIL FROM addresses as being assigned an authenticated source and/or allocated IP addresses. This allows them to reject email from invalid sources and with invalid SMTP MAIL FROM addresses. When these MTAs fail to transfer an email message to a target domain within a specified time, they can return a non-delivery notification to a valid original sender. This sender can then take remedial action, for example, telephoning the intended recipient.

Valid email that is rejected as spam (a false positive) at the content analysis level does not usually result in a non-delivery notification being returned to the original sender. This is because there is no way to determine, at this point, whether the Return-Path address sourced from the SMTP MAIL FROM address has been forged or not—in the case of spam, it almost always has been. Since non-delivery notifications sent to forged addresses will, at best, fail to be delivered, and at worst, fuel “bounce storms” (a form of email denial-of-service attack), they are usually not sent. There is, therefore, no event beyond lack of response to that email that would allow a sender of a false positive to take remedial action.

## Downstream effects

What effects will a Symantec Mail Security 8160 appliance have on the SMTP MTAs that it protects and the content analysis antispam defenses that an organization has already deployed? The answer, beyond having to deal with significantly reduced message volumes, is none. But a significantly reduced message volume is important for several reasons:

1. Organizations can reduce the computing resources that have to be allocated to SMTP MTAs and to the servers hosting content analysis antispam defenses.
2. In addition, organizations won't have to add more resources later, as a Symantec Mail Security 8160 appliance flattens inbound spam volume growth.
3. Some freed-up computing resources can be deployed to more stringent, but also more compute-intensive, content analysis antispam filters than was previously possible.
4. The balance between false positives and false negatives can be altered so as to significantly reduce the former with no absolute increase in the latter.

## Conclusion

In this white paper, we have demonstrated how antispam defenses operating at the TCP/IP level, such as those provided by a Symantec Mail Security 8160 appliance, can significantly improve both the effectiveness and the efficiency of an organization's composite antispam defenses.

It does so in several ways:

- It reduces the number of TCP/IP sessions from spammers that have to be accepted by an organization's inbound SMTP MTAs.
- It reduces the volume of spam that a spammer is able to transmit to an inbound MTA, by reducing both the transmission capacity of individual TCP/IP transmission channels and the number of email messages that may be transmitted per session.

This, in turn:

- Reduces the computing resources that have to be devoted to inbound SMTP MTAs and message content-based antispam defenses
- Reduces the amount of disk storage that has to be devoted to spam quarantine

## The Symantec Mail Security 8160 Appliance

- Reduces the size and complexity of retained email archives for organizations that are required to retain all email accepted by their inbound MTAs
- Reduces the volume of quarantined spam that has to be scanned by end users looking for false positives
- Largely eliminates the otherwise inexorable growth in spam volumes
- Allows an organization to deploy more stringent content filters
- Allows an organization to redefine the balance between false positives and false negatives, resulting in fewer false positives for the same volume of false negatives
- It reacts rapidly to actual behavior, allowing new spam sources to be rapidly throttled and valid high-volume sources to have their transmission capacity increased.
- In the case of valid inbound messages that fail to be transferred to inbound MTAs, non-delivery notifications will be returned to the sender, allowing her or him to take remedial action.

In conclusion, comprehensive antispam defenses at the TCP level, such as those provided by a Symantec Mail Security 8160 appliance, can have a significant impact on both the cost and effectiveness of an organization's other antispam measures. Reductions in the computing resources that have to be devoted to antispam defenses at the message content level, in inbound Internet bandwidth devoted to email, and in the size of email quarantines/archives mean that a Symantec Mail Security 8160 appliance yields a rapid return on investment (ROI), even before intangibles such as the reduction and improved handling of false positives are considered.

### **About the author**

Nick Shelness is currently an Independent Technology Consultant and a Senior Analyst at Ferris Research. He was previously the Chief Technology Officer (1998–2001) of Lotus, an IBM company. Prior to that he served as Chief Messaging Architect of Lotus (1994–1998), as Chief Scientist of Soft-Switch Inc., and as a tenured member of the faculty of the University of Edinburgh's Department of Computer Science. For his contributions to the development of email, Shelness was awarded the Electronic Messaging Association's Distinguished Service Award in 1993, and was named as both a Lotus and an IBM Fellow. Shelness lives in rural Scotland, where his other interests include gardening and walking.

## About Symantec

Symantec is the global leader in information security, providing a broad range of software, appliances, and services designed to help individuals, small and mid-sized businesses, and large enterprises secure and manage their IT infrastructure.

Symantec's Norton™ brand of products is the worldwide leader in consumer security and problem-solving solutions.

Headquartered in Cupertino, California, Symantec has operations in 35 countries.

More information is available at [www.symantec.com](http://www.symantec.com).

Symantec has worldwide operations in 35 countries. For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 800 745 6054.

Symantec Corporation  
World Headquarters  
20330 Stevens Creek Boulevard  
Cupertino, CA 95014 USA  
408 517 8000  
800 721 3934  
[www.symantec.com](http://www.symantec.com)

Symantec and the Symantec logo are U.S. registered trademarks of Symantec Corporation. Symantec Mail Security and Symantec Mail Security 8160 are trademarks of Symantec Corporation. All other brand and product names are trademarks of their respective holder(s). Copyright © 2005 Symantec Corporation. All rights reserved. Printed in the USA. 4/05 10408721