

www.securecomputing.com

Secure Computing® has been solving the most difficult network and application security challenges for over 20 years. We help our customers create trusted environments both inside and outside their organizations.

Secure Computing Edge: Protecting the network border

Table of contents

Abstract	2
Introduction	2
Problems associated with email volume	2
Five best practices for managing email volume	3
Complete border security with Secure Computing Edge	3
Powered by TrustedSource intelligence	4
True throughput	4
Failsafe delivery	4
Simple setup and administration	4
Benefits of implementing Secure Computing Edge	4
In conclusion	5

Secure Computing Corporation

Corporate Headquarters

4810 Harwood Road
San Jose, CA 95124 USA
Tel +1.800.379.4944
Tel +1.408.979.6100
Fax +1.408.979.6501

European Headquarters

1, The Arena
Downshire Way
Bracknell
Berkshire, RG12 1PU UK
Tel +44.0.870.460.4766
Fax +44.0.870.460.4767

Asia/Pac Headquarters

1604-5 MLC Tower
248 Queen's East Road
Wan Chai Hong Kong
Tel +852.2520.2422
Fax +852.2587.1333

Japan Headquarters

Shinjuku Mitsui Bldg. 2, 7F
Nishi-Shinjuku 3-2-11
Shinjuku-ku, Tokyo, 160-0023
Japan
Tel +81.3.5339.6310
Fax +81.3.4496.4537

For a complete listing of all our global offices, see www.securecomputing.com/goto/globaloffices

© 2006 Secure Computing Corporation. All Rights Reserved.
CTEdge-WP-Oct06Vf, Bess, enterprise strong, IronMail, MobilePass, PremierAccess, SafeWord, Secure Computing, SecureOS, SecureSupport, Sidewinder G2, SmartFilter, SoftScan, Sinkbook, Type Enforcement, CyberGuard, and Webwasher are trademarks of Secure Computing Corporation, registered in the U.S. Patent and Trademark Office and in other countries. Anti-Virus Multi-Scan, Anti-Virus ProScan, Application Defenses, Edge, G2 Enterprise Manager, Global Command Center, IronMail, IronMail Live Reporting, Message Profiler, MethodMix, On-Box, Outbreak Defender, Power-It-On!, Radar, RemoteAccess, Secure Encryption, SecureWire, SmartReporter, SnapGear, Threat Response, Total Stream Protection, TrustedSource, TrustedSource Portal, and ZAP are trademarks of Secure Computing Corporation. All other trademarks used herein belong to their respective owners.

Abstract

This white paper examines the problem of rising email volume and its effects on today's enterprise email networks. The reader will gain an understanding of effective volume management techniques and will learn how the Secure Computing Edge™ email security appliance solves the volume problem for organizations of any size.

Introduction

The current volume of email sent worldwide is now over 50 billion messages per day. By 2008, this number is expected to rise to a volume of 100 billion per day or more. With the exponential increase in message volume comes a corresponding rise in the threats to corporate email systems.

Secure Computing research has indicated that more than 85% of all email messages are unwanted spam, viruses, denial-of-service (DOS) attacks, Trojans and other malicious threats. Additionally, these threats are constantly evolving to evade detection by traditional security techniques, presenting a major challenge to all organizations, regardless of the type of mail server or message transfer agent installed in the network. How can companies handle the massive increases in email volume without sacrificing accurate detection of threats and without being forced to add hardware to process the additional messages?

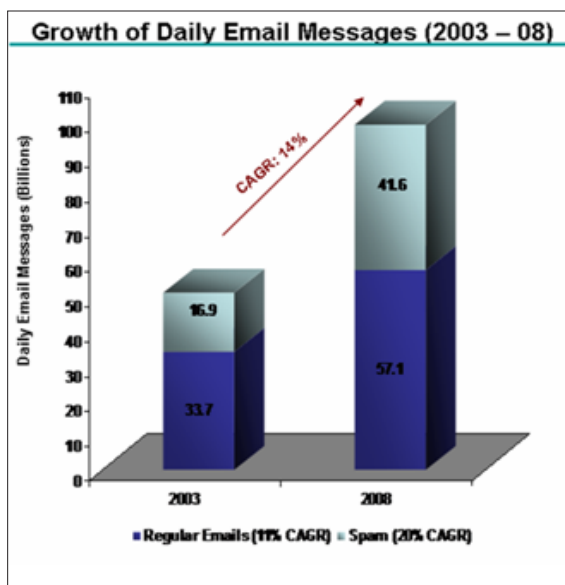


Figure 1: The daily volume of global email is expected to reach well over 100 billion messages by 2008, with spam growing at nearly twice the rate (20% Compound Annual Growth Rate) of legitimate mail (11% CAGR).

Threats to corporate email security can be grouped into four primary categories:

1. **Spam** is broadly defined as any message that is unsolicited and unwanted, or "junk mail." Most spam messages attempt to sell products or services; a large percentage of these messages are pornographic or otherwise offensive.
2. **Phishing** is a scam in which fraudsters "fish" for personal information by pretending to be a legitimate company. These attempts often claim that your account is in jeopardy and ask you to validate, confirm or update information such as credit card numbers.
3. **Viruses** come in many forms. Some are intended merely to cause a nuisance and block network traffic temporarily, while others are designed to steal vital information and relay it to an external server controlled by the hacker who created the virus.
4. **Zombies** are the newest threat to enterprise network security. A zombie PC is one that has been taken over by a remote hacker through the use of Trojans, which are files that appear to be legitimate but instead are viruses that hijack a PC and use it to send spam, viruses, DoS attacks and phishing scams. These zombie machines are networked and used in conjunction with each other to send thousands of messages each, often targeting specific entities.

Problems associated with email volume

Every message that crosses the network gateway uses valuable bandwidth, which is already in short supply in most organizations. By allowing all messages to reach their mail security gateways or mail servers regardless of their quality, mail administrators are both wasting their available bandwidth and, as the volume of inbound email continues to skyrocket, bandwidth to handle mail flow is becoming a critical issue for many companies.

IT departments are being forced to add additional mail security gateways and mail servers to their infrastructure as the volume of mail outstrips the capacity of their existing machines. Considering that the inbound mail volume at many companies is doubling every three to four months, mainly due to bad emails, it's easy to see that IT departments have a significant challenge on their hands trying to purchase, test and install the components of their rapidly growing email infrastructure.

In addition, IT department resources are forced to manage these increasingly complex infrastructures, requiring valuable man-hours. Multiple machines must be maintained, and potentially several different vendors must be managed. In addition, administrators must learn the intricacies of several different programs and control ever-expanding racks full of servers — an expensive proposition for many organizations.

Simply adding hardware is a reactive approach. To take a more proactive approach, many administrators are starting to use products or services that look at the sender's reputation. They hope that by doing so they can work to eliminate bad email at the connection (network or TCP/IP) level. While the intent is laudable the issues with many of these reputation services are numerous.

By deploying an email gateway MTA such as Sendmail™, Postfix™ or any of a number of other alternatives, administrators attempt to cut down the number of messages passing through. Unfortunately, each of these solutions requires additional levels of security in order to accurately and effectively reduce message volume to a tolerable level. For example, to cut down on spam volume, a Sendmail environment may rely on Spam Assassin™ to reduce spam, Panda Perimeter Scan™ for anti-virus protection, and several other products to address other individual threats. The obvious weakness in this approach is that each of these products is designed as a stand-alone application; few, if any, are designed to interact with applications from other vendors, leaving a gaping hole in the correlative intelligence-gathering process necessary for effective overall security. In addition, each application loaded onto a box requires additional processing power, and must query multiple outside sources to obtain up-to-date information each time a sender tries to connect to the network.

Five best practices for managing email volume

1. The first layer of defense against bad emails should occur at the network layer, or at the time of a connection request from an outside mail server.
2. Check the sender's IP address to determine the sender's past behavior.

3. Create rules for handling messages from various types of senders. Connection requests from senders with good behavioral reputations should be allowed and the mail accepted; connection requests from senders with behavioral reputations that show them to be bad senders should have connection requests denied; and those senders classified as suspicious should be subject to traffic shaping techniques to verify their legitimacy.
4. The initial layer of security should be offloaded from your email security gateway or mail servers, so a machine capable of handling a large volume of email messages should be deployed at your network border. This requires deploying a hardened, purpose-built machine to maintain the integrity of your network security.
5. Do not sacrifice your current email infrastructure, end-user experience or network environment when implementing volume control measures. Securing your network does not need to be an invasive operation, nor should it require reconstruction of your existing network infrastructure.

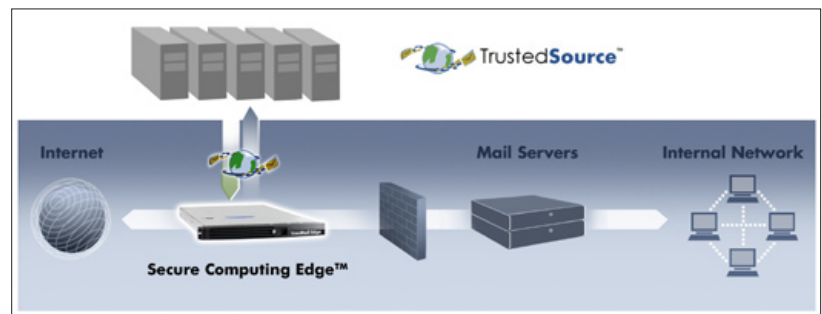


Figure 2: The Edge appliance is installed outside the corporate firewall, providing the first layer of protection against unwanted email. Constant communication with the TrustedSource database ensures accurate spam detection.

Complete border security with Secure Computing Edge

The Edge email security appliance was designed specifically to address the issue of rising email volume. Secure Computing Edge is positioned at the perimeter of the mail system, controlling traffic at the network border, rather than at the mail server or desktop. Secure Computing has designed Edge with a hardened operating system, a proprietary MTA and a Mail IDS (intrusion detection system). What this means to an organization is that Secure Computing Edge can safely be placed at the network edge to perform its role. Secure Computing Edge does not rely on any commonly used MTA software, as many of those are known to have vulnerabilities. Additionally, Edge will block hacker attacks that use methods such as denial-of-service attacks (syn flood), Telnet or ping attacks, and buffer overflow attacks.

Powered by TrustedSource intelligence

Secure Computing Edge relies on TrustedSource™, Secure Computing's revolutionary reputation system, for information about every sender that attempts to connect to the protected enterprise's mail servers. TrustedSource is the first and only reputation system to combine traffic data, whitelists, blacklists and outbreak detection with the unparalleled strength of Secure Computing's global customer network of more than 1600 customers in 40 countries, including over one-third of the Fortune 500. It is also the only reputation system available that is able to provide numerical scoring for every IP address across the Internet.

When the Edge appliance receives an SMTP connection request, the box will hold the response to the sender until the sender reputation is understood. Secure Computing Edge utilizes the intelligence provided by TrustedSource to make high-speed decisions about whether messages should be rejected or allowed, based on a quick IP lookup operation. Secure Computing Edge maximizes speed and efficiency by caching the TrustedSource data locally, with regular updates streamed from the central TrustedSource server. Using the TrustedSource data, Edge can take any of the following actions:

- When Secure Computing Edge receives a connection request from a known bad sender, such as a spammer or hacker, it rejects the connection immediately without accepting any data into corporate network. The sender receives an error code telling them not to retry the connection, as it will only lead to another rejection.
- When a sender receives a score from TrustedSource that falls into the "suspicious" range, Secure Computing Edge will again reject the connection, but will ask the suspicious sender to retry. This traffic shaping, or throttling, is very effective in slowing down the volume of bad email. Legitimate senders will receive the request and resend the message, which will then be accepted. Conversely, spammers, phishers and the like typically will not retry; resending mass quantities of messages is expensive, and their mass-mailing programs are not written to include "retry" logic.
- Messages from "good" senders will pass through the Edge box to the mail server without any processing. Secure Computing Edge will not acknowledge to the sender that the message has been received until the mail gateway confirms it has received the message.

True throughput

Secure Computing Edge is scalable to meet the needs of organizations of any size, with true throughput capabilities of three million email messages per box, per hour. The measure of true throughput includes receiving the connection request, performing a TrustedSource query, and handling the message as discussed above.

Failsafe delivery

Secure Computing Edge is interoperable with all existing mail security gateways and mail servers. Additional functionality provides round-robin load balancing and fail-over between next hops for all mail gateways. If all mail gateways are unavailable, Edge will write the messages to disk, and will reattempt delivery intermittently until a mail gateway is available.

Simple setup and administration

Secure Computing Edge is easy to install. The SmartStart™ wizard-driven set-up is intuitive, guiding the administrator through a step-by-step process with a point-and-click interface. Administrators are required only to determine the threshold score range that classifies senders as good, suspicious or bad. To simplify this process, Secure Computing provides best practices data that advises administrators of exactly which settings will work best in their environment. The configuration of the appliance is handled automatically at installation; when the Edge appliance connects to the Internet, Secure Computing will deliver a configuration package to the machine that reflects the current best practices in use by IronMail® appliances around the world.

Benefits of implementing Secure Computing Edge

The Edge adds value to corporate email security efforts in several areas:

- **Decreased server load** – By identifying and blocking known "bad" IP addresses, Secure Computing Edge can reduce the intake of messages into the network by more than 50%. This is critical in handling the constantly increasing load of mail.
- **Infrastructure growth control** – By significantly decreasing the load on mail servers, Secure Computing Edge allows organizations to handle the additional mail volume, often without the expense of adding expensive mail gateways or servers.

- **Increased effectiveness** – If a known spammer tries to use a new technique to evade detection, Secure Computing Edge will still recognize the origin of the message, causing it to be blocked. Email security solutions that do not employ IP-based reputation will be unable to maintain their effectiveness against new threats.

In conclusion

The exponential rise in email volume over the past several years is showing no signs of relenting, and will likely continue into the future as email continues to expand its role as the primary method of business communication. In order to effectively manage this rising volume, enterprises must take an approach to securing their networks that begins at the outer edge. This need for border security is what compelled Secure Computing to create the Edge email security appliance. Secure Computing Edge is the only email volume control appliance on the market to receive the global intelligence of Secure Computing's TrustedSource reputation system, providing the most accurate and effective detection of unwanted mail, blocking more than half of all email messages before they can reach the mail gateway or mail servers. TrustedSource reputation system, providing the most accurate and effective detection of unwanted mail, blocking more than half of all email messages before they can reach the mail gateway or mail servers.