



realtimepublishers.com<sup>tm</sup>

# *Tips and Tricks* *Guide<sup>tm</sup> To*

# Secure Content Appliances

**McAfee**<sup>®</sup>  
Proven Security<sup>™</sup>

*Dan Sullivan*

---

## Introduction to Realtimepublishers

by Sean Daily, Series Editor

The book you are about to enjoy represents an entirely new modality of publishing and a major first in the industry. The founding concept behind [Realtimepublishers.com](http://Realtimepublishers.com) is the idea of providing readers with high-quality books about today's most critical technology topics—at no cost to the reader. Although this feat may sound difficult to achieve, it is made possible through the vision and generosity of a corporate sponsor who agrees to bear the book's production expenses and host the book on its Web site for the benefit of its Web site visitors.

It should be pointed out that the free nature of these publications does not in any way diminish their quality. Without reservation, I can tell you that the book that you're now reading is the equivalent of any similar printed book you might find at your local bookstore—with the notable exception that it won't cost you \$30 to \$80. The Realtimepublishers publishing model also provides other significant benefits. For example, the electronic nature of this book makes activities such as chapter updates and additions or the release of a new edition possible in a far shorter timeframe than is the case with conventional printed books. Because we publish our titles in “real-time”—that is, as chapters are written or revised by the author—you benefit from receiving the information immediately rather than having to wait months or years to receive a complete product.

Finally, I'd like to note that our books are by no means paid advertisements for the sponsor. Realtimepublishers is an independent publishing company and maintains, by written agreement with the sponsor, 100 percent editorial control over the content of our titles. It is my opinion that this system of content delivery not only is of immeasurable value to readers but also will hold a significant place in the future of publishing.

As the founder of Realtimepublishers, my *raison d'être* is to create “dream team” projects—that is, to locate and work only with the industry's leading authors and sponsors, and publish books that help readers do their everyday jobs. To that end, I encourage and welcome your feedback on this or any other book in the Realtimepublishers.com series. If you would like to submit a comment, question, or suggestion, please send an email to [feedback@realtimepublishers.com](mailto:feedback@realtimepublishers.com), leave feedback on our Web site at <http://www.realtimepublishers.com>, or call us at 800-509-0532 ext. 110.

Thanks for reading, and enjoy!

Sean Daily  
Founder & Series Editor  
Realtimepublishers.com, Inc.

**Note to Reader:** This book presents tips and tricks for four topics related to secure content appliances and their role in enterprise security. For ease of use, the questions and their solutions are divided into topics, and each question is numbered based on the topic, including

- Topic 1: Business Justification for Secure Content Appliances
- Topic 2: Policies and Procedures for Secure Content Management
- Topic 3: System Architecture and Secure Content Management
- Topic 4: Secure Content Appliance Performance

Introduction to Realtimerepublishers.....	i
Topic 1: Business Justification for Secure Content Appliances .....	1
Q 1.1: Why does an organization need another security device for securing content? .....	1
Q 1.2: How does a secure content device complement other security devices?.....	4
Desktop Antivirus .....	4
Firewalls.....	4
Intrusion Prevention Devices.....	5
Complementing Secure Content .....	6
Policy Administration .....	6
Q 1.3: What is the ROI for secure content appliances? .....	7
ROI and Related Calculations.....	7
Calculating ROI on Secure Content Devices.....	8
The Cost of Spam.....	9
The Cost of Viruses .....	11
The Cost of Lost Productivity and Non-Business–Related Web Activity.....	12
Q 1.4: How will secure content management aid in regulatory compliance? .....	12
Privacy Regulations .....	12
Protecting Personal Medical Information .....	13
Preventing Identity Theft .....	14
Data Integrity Regulations .....	15
Q 1.5: What are best practices for educating users about spyware, spam, and phishing? .....	17
Preventing Spyware Infections .....	17
How to Reduce Spam.....	19
Phishing Facts Every Email User Should Know .....	20
Topic 2: Policies and Procedures for Secure Content Management.....	22
Q 2.1: What topics should be addressed in secure content policies?.....	22

Policy Types.....	22
Content Policies .....	23
Q 2.2: How can a systems administrator monitor the effectiveness of current settings? .....	25
Establishing a Baseline .....	25
Reporting on Appliance Performance.....	25
Browser-Based Appliance Reports .....	26
Third-Party Tool Reporting .....	27
Email, SNMP, and Syslogging .....	28
Monitoring Tasks .....	29
Task 1: Establish a Monitoring Policy .....	29
Task 2: Establish a Baseline .....	29
Task 3: Analyze Reports.....	29
Task 4: Verify Accuracy of Content Filtering .....	30
Q 2.3: How can administrators tune secure content policies using global and user-specific rules?30	
Global Policies .....	31
Non-Global Policies.....	31
Policy Inheritance .....	31
Q 2.4: How can administrators use quarantine and deferred mail management to secure content?32	
Quarantining Content.....	33
Isolating Virus-Infected Messages.....	33
Isolating Spam .....	34
Deferred Email Management.....	35
Controlling Content Distribution .....	35
Q 2.5: Some spam passes through the filters; how can the filtering be improved?.....	35
Identifying Spam with General Rules.....	36
Erroneous Categorizing .....	38
False Positives.....	38
False Negatives .....	38
Additional Filtering Mechanisms .....	38
Staying Up to Date.....	39
When All Else Fails .....	39
Topic 3: System Architecture and Secure Content Management .....	40
Q 3.1: Where should a secure content appliance be placed? .....	40

Secure Content Device Operational Modes .....	41
Q 3.2: Why are desktop antivirus software and personal firewalls still needed? .....	43
Layered Security .....	44
Example: Firewall Rules.....	44
Example: Antivirus Protection.....	45
Securing Mobile Devices.....	47
Q 3.3: How does a secure content appliance work with Web servers, caching servers, and application servers?.....	47
Network Protocols .....	48
Positioning the Secure Content Appliance .....	48
Explicit Proxy Mode.....	48
Transparent Router Mode .....	49
Transparent Bridge Mode .....	49
Configuring for Performance and Functionality.....	50
Q 3.4: Can a secure content appliance be attacked? .....	51
Security and Operating System Architecture.....	52
Hardening an OS.....	53
Shutting Down Unnecessary Services and Removing Unneeded Programs .....	54
Patching the OS and Services .....	54
Configuring Services to Reduce Vulnerabilities.....	55
Q 3.5: How do appliances stay up to date on the latest threats?.....	55
Tracking Updates .....	55
Updating Antivirus Applications .....	56
Updating the Anti-Spam Application .....	56
Topic 4: Secure Content Appliance Performance.....	57
Q 4.1: What are threats to content and information assets must organizations address? .....	57
Viruses, Worms, and Other Malware .....	57
Spam .....	59
Phishing scams.....	59
Spyware.....	59
Q 4.2: How can an organization protect against spyware?.....	60
Keeping Spyware Out.....	60
Define a Spyware Policy.....	60
Scanning Multiple Protocols.....	61

Monitoring Spyware Detection.....61

Educating Users .....61

Q 4.3 How can an organization protect against phishing?.....62

    Technical Controls for Phishing .....62

    User Education About Phishing.....63

Q 4.4: How can an organization minimize spam? .....63

    Educate Users.....64

    Do Not Contribute to the Problem .....64

        Do Not Become a Zombie .....65

Q 4.5: How can an organization implement better access controls to Internet content? .....65

## Copyright Statement

© 2006 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimepublishers.com and the Realtimepublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimepublishers.com, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at <http://nexus.realtimepublishers.com>.]

## **Topic 1: Business Justification for Secure Content Appliances**

### **Q 1.1: Why does an organization need another security device for securing content?**

**A:** Security professionals are not at a loss for tools and applications to thwart threats to their IT infrastructure. So why do we need to consider yet another kind of application, this time for securing content?

The simple answer is that security tools are tuned for particular problems and no one tool will ever address all security needs. As new threats emerge, so will tools that address those threats. Consider some types of security tools and applications commonly found in enterprise IT environments:

- Network monitoring tools—These programs capture network traffic and allow network administrators to analyze operations on their networks. These are useful for detecting port scans, identifying the source of unusually large volumes of traffic (for example during a Denial of Service—DoS—attack), or probes, such as CGI attacks.
- Firewalls—This ubiquitous security device is a basic tool for controlling the flow of traffic into and out of a network. Firewalls have evolved from being simply a perimeter device to a desktop application that provides protection for individual computers.
- Vulnerability assessment—One of the earliest and best known vulnerability assessment tools, SATAN, garnered mixed responses from security professionals. It provided systems administrators (and hackers) with a single tool to probe computers for a large number of known vulnerabilities, such as unpatched applications. Today, vulnerability assessment tools are considered another essential tool in the security analyst's toolbox.
- Intrusion Prevention Systems (IPS)—These systems began as Intrusion Detection Systems (IDSs) that used sophisticated rules and patterns of network traffic to detect attacks on a network or server. Instead of just detecting attacks, IPS devices can stop them as well by closing down sessions, blocking traffic from specific Internet addresses, and other methods. IPSs can work on either a network or host level.

- Antivirus software—As the name implies, antivirus software identifies and removes malicious code that attaches itself to other programs. Like other security applications, antivirus software has evolved to counter related threats including worms, Trojan horses and blended threats (this type of malware includes multiple malicious programs, such as viruses, worms, keyloggers, file transfer programs, Internet chat clients, and so on).
- Identity management systems—Identity management is the practice of tracking and controlling access to information assets based on a person's—or in some cases, an application's—privileges. Identity management systems combine the features of authentication systems, such as Single Sign-On (SSO) applications, with authorization systems that track a person's roles and privileges.
- Encryption applications—Encryption programs encode data so that only those authorized to see the data may have access to it. This type of application is a basic security technology used in Web protocols, such as SSL; authentication systems; and end-user programs, such as Pretty Good Privacy (PGP). As more and more data is stored on mobile devices and shared storage arrays, such as SANs, encrypting data on disks is becoming more common.

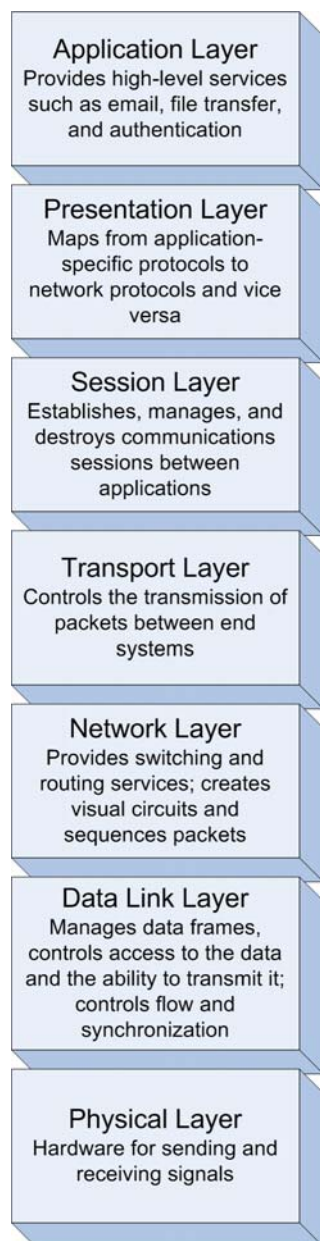
There are other, more specialized tools and applications for security professionals but this list gives a sense of the breadth of security devices already deployed within enterprise environments. Do you really need another security tool thrown into the mix? Couldn't the need be filled by an existing tool or combination of tools?

The answers are, respectively, yes and no. Yes, there is a need for another type of tool to meet threats that are conveyed through content, such as email messages and Web content. The tools mentioned in the previous list are designed and configured to address a narrow range of problems, such as blocking access to ports. With the exception of antivirus software, the tools mentioned do not target the high-level content that moves in and out of a network.

To secure content, you need an application specifically designed for that problem. The characteristics of an appropriate tool include:

- Operates at the application layer of the network—As Figure 1.1 shows, network operations are divided into seven logical layers. Security tools typically function best at one layer; for example, packet sniffers work at the data link and network layers while identity management systems work at the application layer. To effectively analyze the threat embedded in content, the tool must function at the application layer.
- Analyzes all content entering or leaving the network—Email, instant messages, and Web content are the most likely means of bringing malware into an enterprise network. Scanning and removing viruses, worms, Trojan horses, and other threats after they reach the server or desktop is one approach; a better method is to prevent it from entering the network at all.
- High-performance analysis—Securing content should not slow email or Web services. Secure content applications should analyze content and allow permissible content to reach its destination with virtually no impact on transmission times.

- Highly accurate analysis—The process for identifying threatening content should have low rates of false positives (categorizing allowed content as not allowed) and false negatives (categorizing non-allowed content as allowed).
- Tamper-proof—The system that is analyzing content should not itself become compromised by malware or attacks.



**Figure 1.1:** The OSI network model includes seven logical layers; security measures must address threats at every layer.

The best method for achieving all these objectives is to deploy a secure content application, service, or appliance.

## Q 1.2: How does a secure content device complement other security devices?

**A:** A secure content device compliments several of the security devices mentioned in question 1.1, including:

- Desktop antivirus
- Firewalls
- Intrusion prevention systems (IPSs)
- Policy administration

No single security device can address all security threats; in addition, some degree of overlap provides supplementary protection to an enterprise's information infrastructure.

### ***Desktop Antivirus***

Both secure content devices and desktop antivirus software scan for malware—why use both? Ideally, viruses, worms, Trojan Horses, and other malware would never enter an enterprise network. Secure content devices scan network traffic that is likely to carry malware payloads—especially email, file transfer, and Web-based traffic. Secure content devices can identify and block malware at the perimeter of the network; however, there are other means for malware to enter the organization.

Mobile devices, especially laptops, are not continuously protected by secure content devices. Employees take laptops home to insecure home networks. Sales staff travels with laptops, connecting to the Internet from client sites, airports, coffee shops, and other points beyond the enterprise's control. These mobile devices must be protected while they are disconnected from the enterprise network by desktop antivirus software.

Also, although the majority of malware threats propagate through networks, they can still be transferred through shared storage devices. The earliest viruses traveled on floppy disks that were passed between PC users. Today, flash memory devices have largely replaced floppy disks as the preferred storage device for transferring data, but the problem remains the same: infected programs and documents can easily move from one computer to another. Desktop antivirus software can readily scan flash drives and eliminate known malware before it can infect another device.

### ***Firewalls***

Firewalls are the standard means for controlling the type of network traffic that enters and leaves an enterprise network. Firewalls are configured to allow traffic on necessary ports—for example, TCP port 80 for HTTP, 21 for ftp, and 23 for Telnet. When a service provided by a port is not needed, the port is blocked. Firewalls provide a first line, course-grained line of defense; finer-grained security is required in addition.

The basic limit of firewalls is that they work with the structure of network traffic, not its content. For example, an infected file can be transferred into a network using ftp as long as the sender authenticates with the ftp server (assuming anonymous logins are not allowed) and the sender does not violate other basic constraints, such as exceeding storage limits. Similarly, a malicious application could use HTTP tunneling to transfer malware or use otherwise blocked protocols to communicate with malware already infecting a local device. This situation is especially problematic because of the high volume of HTTP traffic in enterprises. In one study of a large institution, more than 40 percent of all incoming and 90 percent of all outgoing traffic used HTTP.

 For more information about HTTP tunneling and detection methods, see “Detecting HTTP Tunneling Activities” at <http://www.ll.mit.edu/IST/pubs/Pack-IEEE2002.pdf>.

Firewalls provide essential security functions but alone they are not enough. Secure content devices examine incoming traffic once it has passed the firewall (or before it reaches the firewall in the case of outgoing traffic). Rather than examine just the structure of the traffic (for example, “this is an SMTP packet”), it examines the content (“Lose weight while you work at home”) allowing the secure content device to identify spam and other unwanted content. IPSs are closely related to firewalls.

### ***Intrusion Prevention Devices***

There are two types of intrusion prevention: host-based and network based. Host-based IPSs protect individual servers and workstations from attacks that cannot, or at least are not, stopped by perimeter defenses, such as firewalls. Host-based IPSs detect anomalous behaviors on servers as well as truly suspect actions, such as an attempt to write a file from a Web browser or the escalation of local privileges.

Host-based IPSs can do some things that other network-based approaches cannot. For example, a host IPS can analyze the content of an encrypted message after it is decoded; a secure content device that monitors network traffic does not have access to the decrypted traffic.

Network-based IPSs use signatures, or patterns of traffic, to detect anomalies in network activity. As with host-based IPSs, there are some attacks that are difficult at best to detect with other methods. One of these attacks is known as Address Resolution Protocol (ARP) poisoning.

ARP is used to map from IP addresses to MAC addresses, the unique physical address on a network interface. ARP, like other Internet protocols, is quite trusting. Devices do not need to authenticate to send an ARP message to another device; any device (or attacker) can send a message telling a server that IP address A maps to physical address B. The server will store that information in an ARP table and use the physical address when addressing messages to IP address A. With ARP poisoning, an attacker can effectively re-route traffic away from a legitimate device to another, compromised machine. Network-based intrusion detection can detect this type of attack in the lower levels of the OSI network model.

## Complementing Secure Content

Both host-based and network-based intrusion detection provide defenses against particular types of attacks. Both protect information infrastructure, such as the integrity of network routing and the operating system (OS) access controls. These complement secure content devices that analyze the content that depends on that information infrastructure.

### ***Policy Administration***

The foundation of a secure infrastructure is a set of well-defined policies governing several aspects of information security, including:

- Authentication
- Authorization
- Vulnerability scanning
- Database access
- Remote access
- Password
- Wireless networking

Information security policy administration tools are relatively new but are emerging to address the difficulties in managing silos of security. One of the key reasons to use policy administration is to centralize management of policies and reporting. This complements secure content devices by providing the means to report on events and defined policies within the secure content device.

 For more information about security policies, see the SANS Security Policy Project at <http://www.sans.org/resources/policies/>.

Deploying multiple defensive layers is a standard practice in information security. Some countermeasures, such as network intrusion detection and firewalls, protect the transmission of network traffic. Host-based intrusion prevention and desktop antivirus software protect the integrity of OSs, applications, and data. Secure content devices protect against the introduction of malware, spyware, spam, and phishing attacks from entering a network. Together these and other tools provide a security infrastructure that can provide a layered defense and address a multitude of threats.

### Q 1.3: What is the ROI for secure content appliances?

**A:** The Return on Investment (ROI) for a secure content appliance is based on a range of factors, including:

- Losing productivity of employees who have to deal with spam in their email accounts
- Additional hardware and software licensing costs to maintain adequate resources to process high volumes of email, including spam
- Losing staff time to eradicating and repairing damage caused by malware infections
- Avoiding fines for failure to comply with regulations
- Avoiding lost intellectual property when proprietary documents are transferred out the of the organization

Clearly, some factors are easily quantified, such as the amount of storage that is taken up with spam. Others, such as lost employee productivity, can at least be roughly estimated. However, some of the largest factors, such as regulatory fines and lost intellectual property are difficult to assess. Nonetheless, organizations can perform some basic ROI analysis on secure content appliances.

#### ***ROI and Related Calculations***

ROI is one of a number of capital expenditure analysis calculations. Depending upon your needs, one or more of these calculations may be used to determine whether deploying a secure content device makes financial sense. The most commonly used calculations are:

- **Present value**—Present value is a calculation that takes into account the value of money over time. For example, if a company saves \$10,000 in 1 year from an investment in IT infrastructure, and one can reasonably expect to earn 6 percent if that savings were invested (that percentage is known as the discount rate), then the present value of \$10,000 a year from now is  $\$10,000/1.06$  (1 plus the discount rate), or approximately \$9434.
- **Net present value**—The net present value is similar to present value but takes into account initial costs.
- **Payback period**—The payback period is the time period in which the total savings from an investment equals the amount of the investment. For example, a \$25,000 investment that saves \$10,000 per year has a payback period of 2.5 years.


- ROI—ROI takes into account both the net present value of money and the net benefits realized by the investment. Net benefits are defined as:

Savings + Increased Revenue - Recurring Costs

Without going into the details of why the calculation is defined as it is, here is the basic formula for calculating ROI over a 3-year period:

$$\frac{[\text{Net Benefit for Year 1} / (1 + \text{Discount Rate})] + \text{Net Benefit for Year 2} / (1 + \text{Discount Rate})^2 + \text{Net Benefit for Year 3} / (1 + \text{Discount Rate})^3]}{\text{Initial Costs}}$$

- Internal rate of return—The internal rate of return (IRR) calculation is the most complex of the group listed here. IRR is often used to compare the benefits of different projects and choose among them. As IRR is expressed as a percentage, it is easy to compare projects of different financial and time scales. IRR calculates the discount rate at which the present value of the net benefit of an investment equals zero. (Microsoft's XIRR function and Star Office and Open Office's IRR function can be used to calculate IRR.)

 For more information about IRR, see and "Internal Rate of Return Revisited" at [http://members.tripod.com/~Ray\\_Martin/DCF/nr7aa003.html](http://members.tripod.com/~Ray_Martin/DCF/nr7aa003.html).

### **Calculating ROI on Secure Content Devices**

To assess the investment value of a secure content device, you need to include in calculations the costs of spam, viruses, phishing attacks, lost productivity that result from non-business-related Web activity, violations of regulations, and loss of intellectual property. The last two items are difficult to estimate. The following example will ignore those values, as they are very environment dependent—thus, the results of the calculations may underestimate the true value of the investment.

## The Cost of Spam

First, let's examine the cost of spam. There are basically three types of costs: lost productivity, additional hardware costs, and additional administrative costs. To calculate lost productivity, start with the number of email users, the average number of spam messages, and the time required to read and delete those messages. The basic formula for calculating lost productivity is:

$$\text{Number of email users} * \text{number of spam messages per day PER USER} * \\ \text{time in minutes to read/delete spam message} * (\text{average hourly} \\ \text{rate} / 60) * \text{number of work days per year}$$

To calculate storage costs, start with the number of email users, the average number of spam messages, the average spam message size, the number of days the message resides on the server, and the average cost of storage. The basic formula for calculating spam storage costs is:

$$\text{Number of email users} * \text{number of spam messages per day per user} * \\ * \text{average spam size} * \text{average cost of 1MB storage per year} / 365 * \\ \text{number of days message stored}$$

To calculate additional administrative costs, you need to estimate the number of minutes per day email administrators manage spam problems and the hours per month, on average, email administrators and systems administrator spend addressing storage and network traffic-related problems as a result of spam.

With these three factors—lost productivity, additional hardware costs, and additional administrative costs—you can estimate the cost of spam to an enterprise. Figure 1.2 shows an example calculation of a payback period.

User Inputs		Results	
Number of email users	500	<b>Software Costs</b>	
Number of years to perform the analysis	1	Lost user productivity due to incoming spam	\$ 100000
<b>Assumptions</b>		<b>subtotal</b>	<b>\$ 100000</b>
Daily number of junk emails received per user	10	<b>Hardware Costs</b>	
Time in minutes to read an email message	0.1	Spam mail storage space costs	\$ 16569
Average hourly employee rate (fully burdened)	\$ 50.00	Cost of managing spam and junk mail infestations	\$ 3000
International factor	1	Email server downtime recovery costs due to large message stores	\$ 900
Annual number of work days	240	<b>subtotal</b>	<b>\$ 20469</b>
Average message size in bytes	17000	<b>Total Spam Costs</b>	<b>\$ 120469</b>
Average cost of storage per MB	\$ 0.28	<b>Cost of McAfee SpamKiller</b>	<b>\$ 7185</b>
Average storage length in days of spam messages	2	<b>Your Saving with McAfee SpamKiller</b>	<b>\$ 113284</b>
Daily time in minutes spent by IT managing spam mail problem	15		
Potential email downtime in hours per month from large message stores or excessive message traffic	1.5	Time to recoup investment (in days)	23

\*SpamKiller pricing is based on protection using SpamKiller for WebShield appliances.

**Figure 1.2: Example savings and pay-back period on one component of a secure content device—McAfee SpamKiller anti-spam software.**

To calculate the ROI of this investment, simply take the net benefit and divide it by 1 plus the discount rate. Assume the net benefit is the savings on spam costs calculated in Figure 1.2, \$120,469, the initial costs are \$7185, and a discount rate of 6 percent. The ROI for this investment is:

$$ROI = (\$120,469 / 1.06) / \$7,185 = 1581\%$$

For every dollar spent on spam protection, \$15.81 is saved. The cost of other threats to secure content are calculated in a similar manner.

## The Cost of Viruses

Spam is constant, so users are always having to deal with it. Viruses, while still prevalent, do not present the same constant and sustained level of success in reaching targets. Therefore, one of the key factors in estimating the value of antivirus devices is understanding the probability that a virus will successfully infect an unprotected device. The other factors, which Figure 1.3 shows, are comparable to those used in the spam calculation.

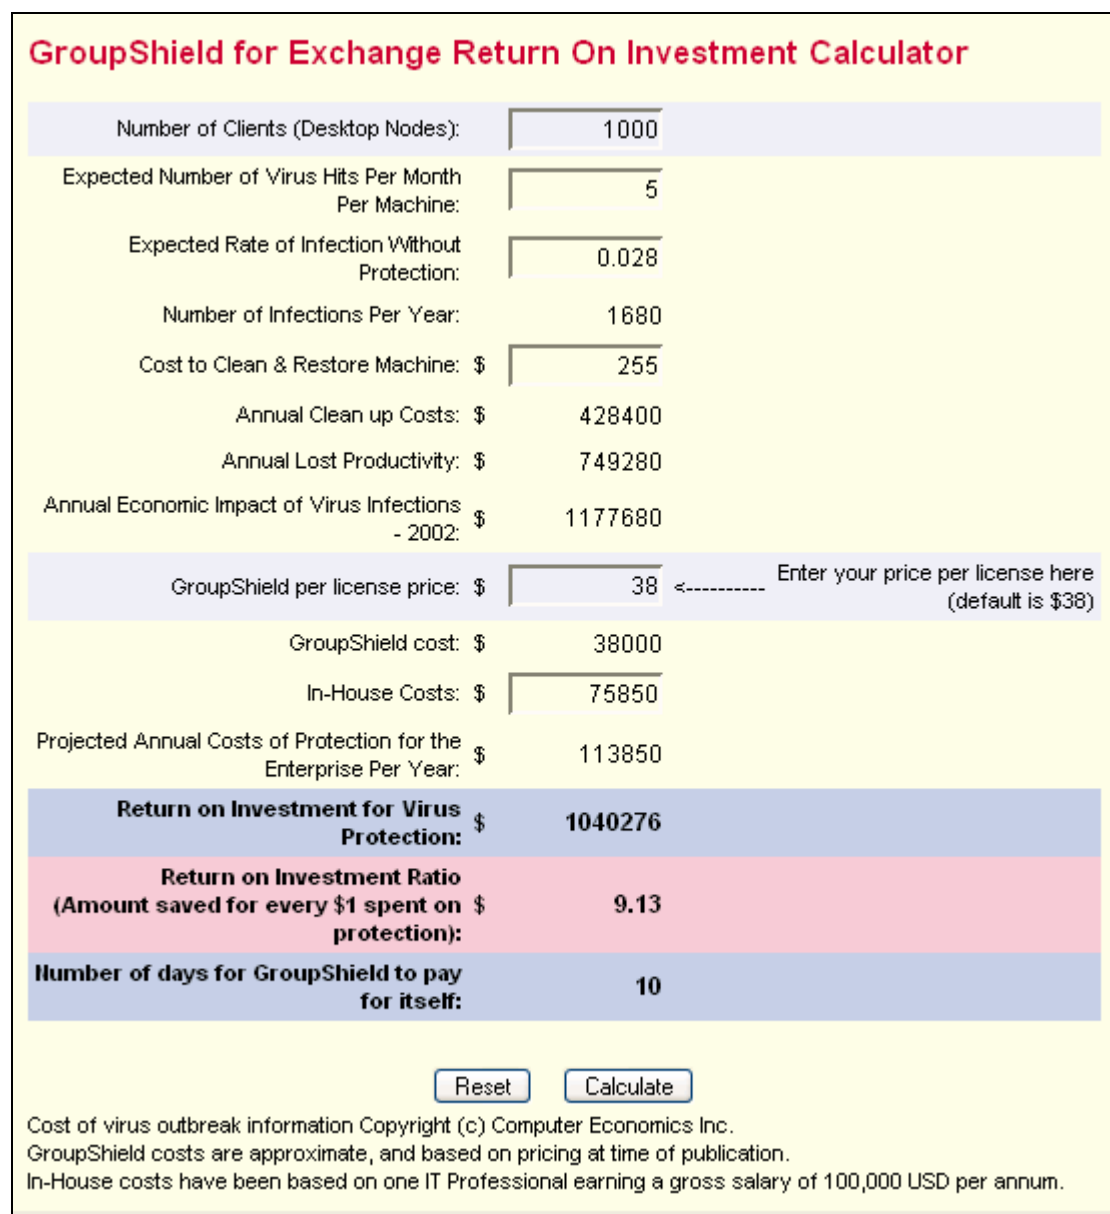


Figure 1.3: Example savings and ROI ratio for antivirus protection.

## The Cost of Lost Productivity and Non-Business–Related Web Activity

Checking personal email, shopping online, browsing online casinos, and other non-business–related activities can put a drain on productivity. Secure content devices can prevent not only malware and unwanted content from entering an enterprise network but also users from browsing sites unrelated to business operations. Even a cursory examination of Web logs can give some indication of the level of this problem within an organization. How many users are visiting time-wasting sites? What is the duration of time spent at those sites? With estimates of those two measures, you can calculate the expected savings in productivity by blocking those sites.

However, blocking sites does not guarantee that the time will be used 100 percent productively. When calculating productivity savings, consider using an adjustment factor to account for this fact.

The ROI from a secure content device is substantial even when considering only easily quantifiable measures, such as savings due to spam and virus protection. “Soft” benefits—such as avoiding regulatory fines and preventing the disclosure of proprietary and trade secret information—provide additional, but difficult-to-quantify incentives for investing in secure content devices.

### Q 1.4: How will secure content management aid in regulatory compliance?

**A:** For the past several years, governments have been actively changing the regulatory environment with respect to personal privacy and the integrity of business information.

#### *Privacy Regulations*

There has been growing concern over the use of private information for unauthorized business purposes. For example, should a pharmaceutical company know of a patient’s congenital heart condition so that they can market a new cardiovascular drug? Should banks be allowed to share account information with business partners so that their partners can sell personal financial planning services? The consensus answer to these and similar questions is no. The widespread adoption of privacy protections has been rapid in the United States, the European Union, Canada and Australia. Some well-known regulations governing personal privacy include:

- State of California, United States passed SB 1386, a law directing companies and government agencies to inform California residents of any unauthorized disclosure of personal information.
- The United States passed the Health Insurance Portability and Accountability Act (HIPAA), which dictates how personal medical information is used and shared.
- The Australian Federal Privacy Act defines principals for the collection and use of personal information of Australian citizens.
- The European Union Directive 95/46/EC defines regulations about how personal data about European citizens is collected, stored, shared, and updated.
- Canada has passed the Personal Information Protection and Electronic Documents Act (PIPEDA) defining standards for protecting personal data.


The details of these regulations vary, but the objectives and requirements are similar. First, organizations must exercise due care when collecting and storing personal information. In some cases, regulations define the circumstances in which information may be shared. For example, under HIPAA, physicians can share information about a common patient but not with drug company sales persons. To limit a company's exposure with regard to privacy protection, most will implement access controls.

Access controls are physical and technical safeguards used to protect the integrity and confidentiality of data. Typical controls include access control lists (ACLs) and file protections that define which users may read and change data. Although these are essential controls, they are not always sufficient. Consider the following two examples.

### **Protecting Personal Medical Information**

A hospital administrator receives a request from an executive steering a committee working on long range plans for hospital expansion. The committee needs aggregate information about the geographic distribution of patients and the types of medical services provided to patients from various areas. The administrator is pressed for time and cannot assign anyone to summarize the raw data; instead, he or she sends a database extract with detailed patient information, including personally identifying information. The steering committee is not making medical evaluations of those patients, so their detailed, personal information should not be shared. An access control system will not prevent this violation because the administrator has legitimate access to the data on a day-to-day basis.

What is needed in this case is a content-based control such as a content filter that can be configured to detect patterns indicative of personal medical records. For example, if data is frequently shared between systems in the hospital and with resident doctors' offices, there may be a standard program for extracting a patient record. This program may use a well-defined XML scheme with labels such as Patient-First-Name, Patient-Last-Name, and Primary-Diagnosis. These labels can be detected as data is transmitted across the network and, depending upon other conditions in the filter rules, the transmission can be blocked (see Figure 1.4).



**Figure 1.4:** Extracted data analyzed by a content filtering mechanism can prevent the transmission of protected data.


In this case, there may be no intent to violate the regulation, the busy administrator just did not know that the extract with personally identifying information should not have been copied to others outside of the hospital's group of medical professionals. Not all violations are so benign.

## Preventing Identity Theft

In the past, criminals robbed banks because that is where the money was. Now, stealing identities can lead to the money. Several high-profile security breaches of credit card processing and financial institutions are raising awareness of the threat of identity theft that results from poor security measures. Perhaps the most telling example to date is the exposure of as many as 40 million credit card accounts due to a breach at CardSystems, a one-time transaction processor for MasterCard, Visa, and American Express.

Businesses, governments, and other organizations with responsibility for protecting financial and personal data will often use several security mechanisms including access controls, firewalls, and intrusion detection systems (IDSs). Even with these safeguards in place, users inside the organization with knowledge of systems, patch levels, and application vulnerabilities can avoid security countermeasures and access confidential data. However, when that information is transmitted, it is subject to analysis by content filtering safeguards—preventing unauthorized transmission of protected data.

Filters could be constructed, for example, to detect patterns indicating credit card information being sent outside the organization—for example, a 16-digit number (credit card number) followed by a 4-digit number with the first two digits representing a number between 1 and 12 (the expiration date) being routed to an address outside the network (see Figure 1.5).



**Figure 1.5:** Depending on the content and the location of the recipient, a content filter can prevent protected data from being transmitted outside an intranet.

Security professionals have long known that no single security safeguard will eliminate threats to information systems and their data. Multiple countermeasures are required to reduce the wide variety of threats that are present today. Content filtering is one layer of a multi-layered defense against privacy violation as well as other compliance violations.

### **Data Integrity Regulations**


Names such as Enron, Tyco, and WorldCom once elicited images of successful companies that set standards for performance in the market. Now they are more likely to conjure images of executives entering federal courthouses and stories of lost investments. Governments, especially the United States federal government, has reacted to these and other corporate scandals with laws designed to preserve the integrity of information provided to investors and other stakeholders in public companies.

The best-known regulation governing the integrity of business information is the United States' Sarbanes-Oxley Act. For IT professionals, Sarbanes-Oxley creates new demands for ensuring integrity of financial reports, for establishing internal procedures appropriate to ensure data integrity, and for reporting material changes in a company's operations. Other well-known regulations target particular industries, such as the Gramm-Leach-Bliley Act which applies to banks, and Title 21, Code of Federal Regulations, Part 11 (21 CFR Part 11), which applies to the pharmaceutical industry.

These regulations cover a broad range of topics but can be distilled to a set of core principals with respect to the due care that is required to protect information. Business must be able to

- Protect personal information with which they are entrusted
- Ensure appropriate security measures are in place so that data is not tampered with
- Establish well-defined controls and procedures for managing data
- Audit and report on changes to data under their control

Secure content devices can contribute substantially to the multiple layers of security measures that must be in place to meet these regulations.

 For more information about best practices on compliance and IT governance, see the Information Systems Audit and Control Association Web site. Especially useful is the Control Objectives for Information and related Technology (COBIT) framework., available at: <http://www.isaca.org/Template.cfm?Section=Downloads10&Template=/TaggedPage/TaggedPageDisplay.cfm&TPLID=63&ContentID=13742 - COBIT.>

The primary defenses for protecting privacy and integrity are access controls. Users should be granted permission to view and change data based on their role in an organization. However, primary defenses are not enough:

- Someone with legitimate access to confidential data might disclose his or her password to another employee, either accidentally or because of a social engineering scam (such as someone calling the user and pretending to be the Help desk).
- A vulnerability in a server operating system (OS) allows an attacker to gain control of the system and copy files with customer data or overwrite data such as financial projections.
- Spyware could include keyloggers, which record keystrokes (including usernames and passwords) and send the captured data to servers controlled by the perpetrators.
- A traveling executive might download information to an unmanaged device, such as a desktop computer in a hotel business center, which is locally cached. The data may be left for other users to retrieve without the executive's knowledge.

Secure content devices provide additional levels of protection in these cases. For example, in the case of the disclosed password, the unauthorized user may use the stolen credentials to log in remotely to a server. The unauthorized user may then attempt to download a file with customer names and credit card numbers. If the secure content device is configured to detect the proper patterns (for example, file headers, credit card number, key phrases, and so on), the file transfer will be blocked. The device could also block the transfer of data when an attacker exploits a vulnerability on a compromised server and attempts to copy sensitive information.

Spyware and keyloggers are especially menacing because large numbers of these threats can be deployed to automatically transmit significant amounts of information. Once the information is transferred to servers controlled by the perpetrator, text processing tools can be used to scan large amounts of data looking for valuable information, such as Social Security numbers, bank account numbers, and credit card numbers. Secure content devices can detect those same pieces of valuable information and prevent them from being transmitted in the first place.

## Q 1.5: What are best practices for educating users about spyware, spam, and phishing?

**A:** The first step in educating users about spyware, spam, and phishing is to explain the nature of the problem and the consequences of unwanted emails and programs.

### Preventing Spyware Infections

In addition to using spyware blocking and removal tools, there are several steps users can take to reduce the chances of a spyware infection. These include:

- Properly configuring your browser—For Internet Explorer (IE), proper configuration includes disabling the Enable Install On Demand options (see Figure 1.6 for an example).

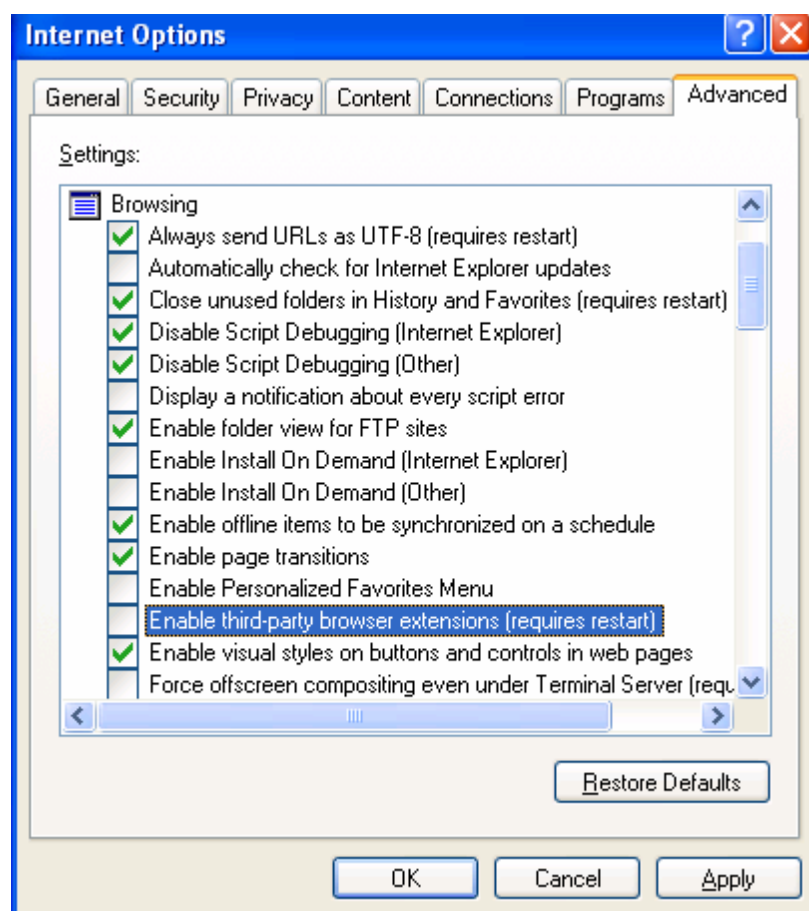
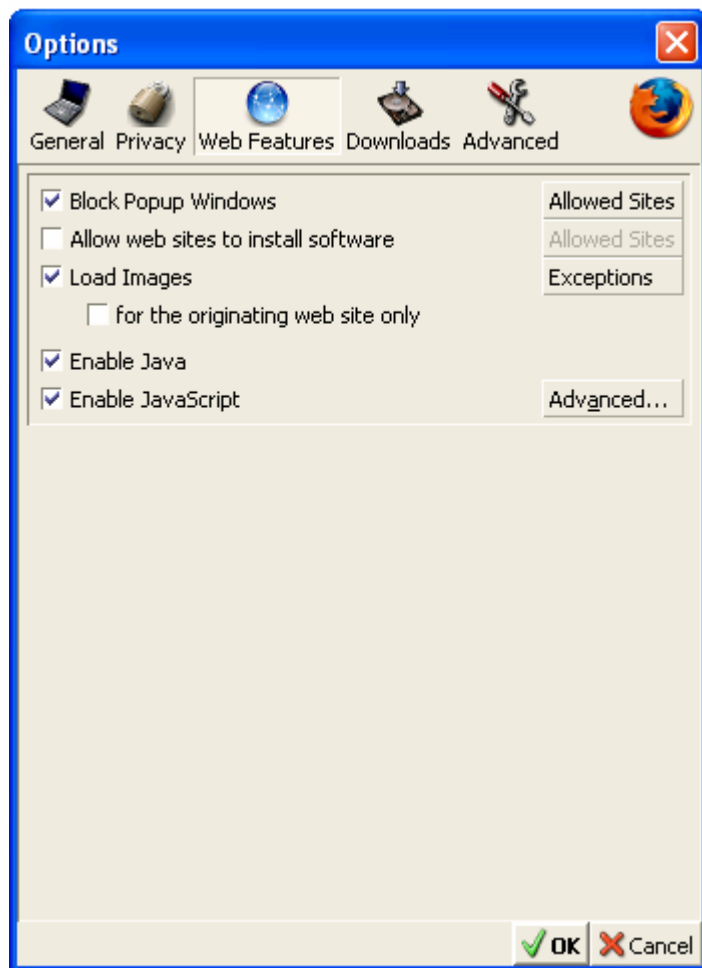


Figure 1.6: Configuring browser security settings is one step to reducing the chances of a spyware infection.

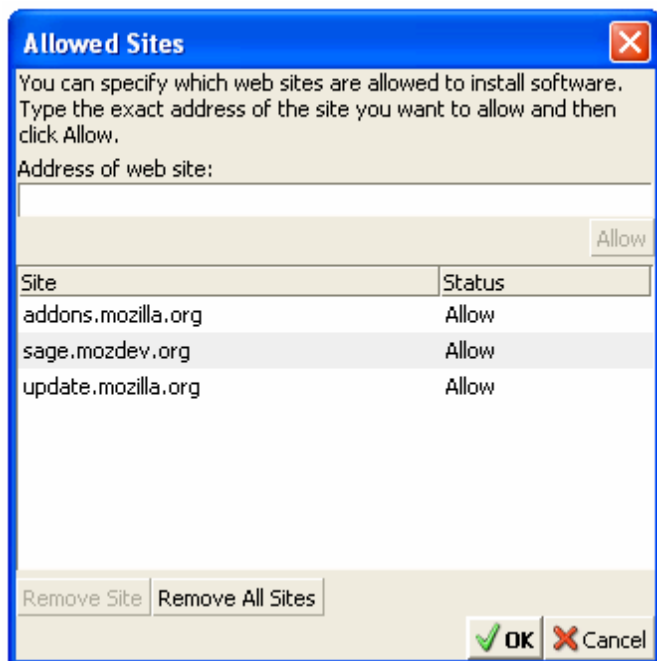
For more information about controlling IE security, see “Working with Internet Explorer 6 Security Settings” at <http://www.microsoft.com/windows/ie/using/howto/security/settings.mspx>.

Mozilla Firefox, another popular browser, has fewer known vulnerabilities than IE and is thus a good option for security-conscious users. As with IE, Firefox users can disable the automatic installation of software, as Figure 1.7 shows.



**Figure 1.7:** Firefox allows users to disable the automatic installation of software.

If the “Allow web sites to install software” option is selected, users can specify which sites are allowed to install software (see Figure 1.8).



**Figure 1.8:** Firefox allows users to list trusted sites for the purpose of automatically installing software.

- Browse trusted sites—File sharing sites, such as music “sharing” sites, are likely spots for picking up spyware and other malware.
- Keep your browser and operating systems (OSs) up to date with patches.
- Use a firewall—Doing so can help prevent spyware from transmitting information from your computer.

### **How to Reduce Spam**

It’s safe to assume spam will always be with us. Although spam filtering can be quite successful, reaching into the 90+ percent success rates for blocking spam, it’s better to never receive spam in the first place. The following list highlights options for minimizing the amount of spam you receive:


- Do not add your email address to newsletters, opt-in offers, or other lists without understanding how your email address will be used. Reputable businesses and organizations publish and adhere to privacy policies. Make sure you understand an organization’s privacy policy before giving them your address.
- Do not post your email on newsgroups or Web-based message boards. These sites are culled by spammers for email addresses.
- Do not respond to spam. A reputable mass emailer may take you off a list as requested but don’t assume you are working with a reputable business if you don’t know them.
- Use a disposable email address when you must give an email address. If this disposable address begins to attract too much spam, drop the address.

Phishing attacks are a more troubling type of spam that warrant their own set of guidelines.

## Phishing Facts Every Email User Should Know

Email users should understand the following about phishing attacks:

- Phishing is a form of social engineering, otherwise known as a con. The phisher gains the confidence of their target and then elicits useful information.
- Phishers gain confidence of victims by *appearing* to be a trusted and legitimate entity, such as a bank or other business. It is relatively easy to steal logos and even entire Web pages from Web sites, so users should not be fooled by the well-known logo or look and feel of their bank.
- Rather than judge an email by appearance, examine the content closely. For example, is your bank emailing you that they need to verify your account number, Social Security number, or online banking password? If so, the message is most likely a phishing scam. When in doubt, ask yourself, has the bank ever called or sent postal mail with these questions? Chances are they haven't. Banks and businesses have other ways to verify their records.
- Beware of emails with upsetting or exciting messages. Phishers count on emotion to overtake rational assessment. If something serious has occurred with your personal finances, would a business email you about it? Not likely. When credit card companies suspect fraud, they call the home phone number of the card holder.
- If you do decide to follow a link in an email, always type the URL into the browser. Doing so will help avoid any tricks used to exploit vulnerabilities in browsers. It will also make it more likely you will catch a minor difference in a legitimate Web address that links to a bogus site.
- Watch for promises of prizes, get-rich-quick schemes, and similar appeals to greed coupled with a request for personal information. (Spam sometimes makes similar promises but requires the recipient to purchase something in return).
- Spammers blanket thousands of email addresses, pretending to be EBay, Pay Pal, banks, and other institutions. Don't be surprised if you receive a phishing email from a business with which you have an account; phishers count on getting at least some customers from those businesses.
- Do not put personal or financial information in an email or Web site in response to an email. When in doubt, contact the business either by phone or through the business's Web site (by typing in the Web address or getting it from a search engine; don't use the URL in the email).

 For more information to educate users about phishing scams, see the Anti-Phishing Working Group's "Consumer Advice: How to Avoid Phishing Scams" at [http://www.antiphishing.org/consumer\\_rec.html](http://www.antiphishing.org/consumer_rec.html), the United States Federal Trade Commission's "How Not to Get Hooked by a Phishing Scam" at <http://www.ftc.gov/bcp/online/pubs/alerts/phishingalrt.htm>, and the United States Department of Justice's "Special Report on Phishing" available at [http://www.antiphishing.org/DOJ\\_Special\\_Report\\_On\\_Phishing\\_Mar04.pdf](http://www.antiphishing.org/DOJ_Special_Report_On_Phishing_Mar04.pdf).

As with other security threats, phishing is a process of evolving threats and countermeasures. As email users become more educated about the nature and consequences of phishing, phishers have to hone their appeals. The latest evolution of phishing has earned the name “spear-phishing” because it uses low volumes but highly targeted phishing emails. This method makes these types of attacks more difficult to detect by businesses and more likely to fool the intended victim.

 See Computer World’s “Training Needed to Halt Spear-Phishing Attacks” at <http://www.computerworld.com/securitytopics/security/holes/story/0,10801,104087,00.html?source=x584> for more information about spear-phishing, including the results of a test conducted by the State of New York on email users at five state agencies.

---

## **Topic 2: Policies and Procedures for Secure Content Management**

### **Q 2.1: What topics should be addressed in secure content policies?**

**A:** Content policies can be organized around two dimensions: first, services provided on the network, including:

- SMTP email
- POP3 email
- HTTP
- FTP

Second, based on threats, such as

- Spam
- Viruses and other malware
- Disclosure of private or confidential information
- Banned content
- Use of time-wasting Web sites

There is clearly overlap, for example, between how spam is handled in an SMTP email system and a POP3 email system. At the same time, different protocols or services have different vulnerabilities and require different types of monitoring. For example, private or confidential information can be transmitted via email or FTP; however, FTP's long history of vulnerabilities warrants attention to those conditions.

### **Policy Types**

Policies are rules applied to groups of users. Policies can be either all users, known as global policies, or non-global policies, which apply to groups of users. It is preferable to place as many content rules in global policies as possible. Non-global policies should be used for exceptions to the rules. For example, as a global policy, you might limit the size of attachments to 10MB but want to allow attorneys and others in the legal department, who tend to work with large volumes of documents, to accept attachments as large as 20MB.

Non-global policies are assigned to groups of users. In general, the groups should be logically related by their organizational function rather than common characteristics of the policies that are applied to them. For example, both the legal and marketing departments might be allowed to receive attachments as large as 20MB. However, that similar requirement is basically a coincidence; tomorrow the requirement of one of the departments may change. Thus, an administrator would want to establish separate legal and marketing groups rather than a single 20MB Attachment group.

## Content Policies

Content policies need to address a range of topics, including scanning, encryption and digital signatures, disclaimers, and content size. Scanning is a broad set of activities geared toward ensuring unwanted content is not allowed into an organization and protected content is not allowed out. Scanning policies should define rules for

- Antivirus
- Anti-spam
- Banned words and phrases
- Private, proprietary, and trade secret information

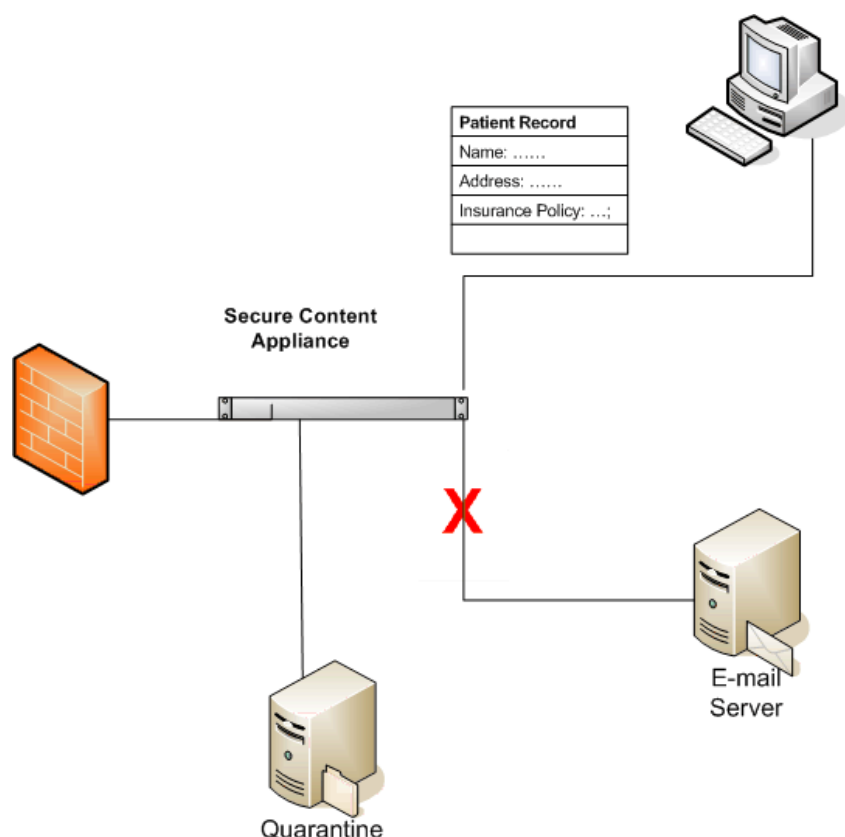
Antivirus scanning will use both signatures (binary patterns indicating a virus or other malware) and heuristic analysis (general patterns that indicate the existence of malware, such as an attempt to modify files with prior user command) to detect malware. Antivirus policies should define how an infected file is handled. It could be deleted, quarantined, or cleaned and passed through. Often the best option is to clean and pass through; the recipient gets the message but not the malware. One drawback to this approach is that the message is changed and so any associated digital signature will no longer match when a new signature is calculated upon arrival. Antivirus policies should include frequent updates of virus signature files as well as patches and upgrades to the virus detection engine.

Anti-spam scanning policies need to balance the need for comprehensive rules that capture most spam while not restricting access to legitimate email with false positives. To ensure the best possible spam protection, keep spam files up to date. Also, use the white list feature, Permit Sender, to list senders allowed to bypass spam scanning, which prevents the chance of false positives (legitimate email classified as spam) from trusted senders. Similarly, blacklists are used to prevent known spammers from sending messages. This setup prevents potential problems, as spammers routinely hijack email accounts so that legitimate emailers may have their messages blocked.

Policies should define how spam should be handled once it is identified: it can be refused, deleted, or forwarded to a special recipient who is monitoring spam activity on the network. Another option is to add a message to the email indicating the message is potential spam and let the recipient decide how to deal with the message.

Both incoming and outgoing messages should be scanned for banned words and phrases. Content rules should include checks to words or phrases in email messages that might be considered as contributing to an offensive or hostile work environment. Outgoing messages should be checked for private, proprietary, or other confidential information. This setup can require careful crafting of rules.

Consider a healthcare provider that uses email to transmit patient information between doctors. In the United States, the Health Insurance Portability and Accountability Act (HIPAA) regulation places restrictions on how protected health information is shared among doctors, insurances companies, and others involved in healthcare. Sending an email with an attached patient record might violate HIPAA unless the recipient is allowed by the patient to receive that information. Of course, the email system or content filter will not have access to databases or paper files that identify legitimate recipients of protected healthcare information. In such cases, a provider may require that all patient records are sent from a single email account that is used only by patient records personnel. A set of content rules could then be defined to block the transmission of the message with indicators of protected patient information, and forward it to an email account for further review (see Figure 2.1).



**Figure 2.1:** Quarantine areas can be established to retain sensitive information so that it can be reviewed before sending it outside the organization.

Proprietary and trade secret information is protected in a similar way. Policies should include rules for blocking or quarantining confidential information before it leaves the email system. This setup will require the custom definition of a dictionary of terms and phrases, such as project and process names, that are kept confidential.

Another area that should be addressed in content policies is time-wasting Web sites. Although many organizations have no interest in blocking occasional visits to news sites, those same organizations have no need for gambling or adult sites. Policies should be defined that block access to known time-wasting, non-work–related sites. Content policies cover a broad range of topics but are essentially rules for filtering what you would want coming into or going out of your organization.

## **Q 2.2: How can a systems administrator monitor the effectiveness of current settings?**

**A:** To monitor the effectiveness of a secure content appliance’s settings, administrators must establish a baseline of activity and then regularly examine the volumes and types of events on the network.

### ***Establishing a Baseline***

A baseline should be established when the appliance is first installed. The objective is to understand what constitutes “normal” activity on the network, including the number of emails sent and received, volumes of HTTP and FTP traffic, the number of spam messages, and malware applications detected. The baseline should include both absolute measures on protocols, such as the volume of HTTP traffic per day, and percentages, such as the percent of emails classified as spam.

In addition to these measures, the administrator should assess the accuracy of the content filtering to determine the rate of false positives, the number of messages incorrectly categorized as spam, attachments incorrectly identified as malware, and the rate of false negatives (that is, banned content that was not detected by the appliance). This work requires manual review of quarantined messages and careful tracking of malware infections.

Once the baseline is established, administrators should monitor the same measures over time using the same reports and analysis used to establish the baseline.

### ***Reporting on Appliance Performance***

The basic reporting tools at the administrator’s disposal are:

- Browser-based reports available through the appliance
- Centralized reporting through ePolicy Orchestrator
- Email, Simple Network Management Protocol (SNMP), and Syslogging

Each type of reporting has its advantages and requires varying levels of configuration and integration.

## Browser-Based Appliance Reports

The secure content appliance includes several reports providing summary and detailed information about traffic volumes and significant events. Although the details of each of these reports are more thoroughly described in the appliance documentation, let's take an overview look at example reports provided.

For example, a secure content appliance can provide information such as system status that includes details about protocols, hardware, load sharing, and general status information. The protocol status displays counts of the amount of traffic scanned, the number of viruses detected, emails deferred, spam messages blocked, and volumes of HTTP, FTP, and email traffic. You can also gather information about the status of each protocol and the workload processed. The hardware, load, and general status information display low-level details, such as the RAID status of hard drives and MAC addresses of the appliance's NICs (see Figure 2.2).

Technical Support | Submit a sample | Virus Information Library | About WebShield | Resources | Help Topics

McAfee WebShield appliance v3.0

192.168.100.50 **System Status** [Reset Counters](#) [Settings](#) [Refresh](#)

**Monitor**  
[Status](#)  
[Performance](#)  
[Logs](#)  
[Charts](#)  
[Updates](#)  
[Resources](#)

**Policy**

**Configure**

**Update**

**E-Mail**

**System**

**Network**

**Troubleshoot**

[Home](#)

[Show Quick Help](#)

**Protocol status (Counters Reset Jan 01 2005 00:00:01)**

<a href="#">SMTP Viruses Detected</a>	From Inside 3 : From outside 40	<a href="#">SMTP E-Mails Received</a>	3457 / Hour (12432)
<a href="#">HTTP Viruses Detected</a>	From Inside 50: From outside 64	<a href="#">HTTP Traffic</a>	187.3 MB / Hour (1236425678)
<a href="#">FTP Viruses Detected</a>	From Inside 3 : From outside 58	<a href="#">FTP Traffic</a>	12.4 MB / Hour (145425843)
<a href="#">POP3 Viruses Detected</a>	From Inside 0 : From outside 12	<a href="#">POP3 E-Mails Received</a>	45 / Hour (234)
<a href="#">Total Viruses Detected</a>	230	<a href="#">SMTP Spam Detected</a>	343
<a href="#">E-Mails Deferred</a>	5	<a href="#">SMTP Spam Blocked</a>	189
<a href="#">Viruses Quarantined</a>	82	<a href="#">SMTP Spam Blocked by RBL</a>	23
<a href="#">Content Quarantined</a>	2		

**Dashboard**

<a href="#">SMTP Health</a>		<a href="#">Scanning Partition Used</a>	10%
<a href="#">HTTP Health</a>		<a href="#">Logging Partition Used</a>	14%
<a href="#">FTP Health</a>		<a href="#">Quarantine Partition Used</a>	18%
<a href="#">POP3 Health</a>		<a href="#">Deferred Partition Used</a>	44%
<a href="#">Memory Swap Rate</a>	12 /sec	<a href="#">Processors Used</a>	66%
<a href="#">Load Average</a>	110		

**General Status**

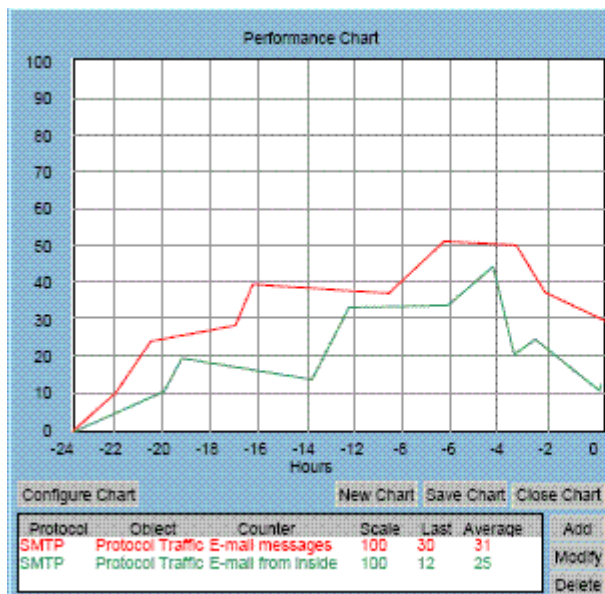
**Hardware Status**

**Load Sharing Status**

Copyright (c) 2004 Network Associates Technology, Inc. All rights reserved.

Figure 2.2: An example report provided by the secure content appliance.

Additional information provided by a secure content appliance includes Web pages that let administrators see the history of selected counters over a 24-hour period. For example, an administrator can track the number of email messages originating inside and outside the network (see Figure 2.3).



**Figure 2.3:** Examining a historical view of select counters.

In addition to storing performance detail on the appliance, administrators have the option of centralizing reporting with a third-party tool or basic network event monitoring tool such as SNMP traps and syslog. Events are filtered and sent to reporting tools based on three criteria:

- Protocol—Protocol filtering allows administrators to monitor events by traffic type; for example, all email traffic might be consolidated in a third-party tool-provided report, while FTP reporting is consolidated in a legacy syslog reporting application.
- Severity—Severity settings are used to control the volume of events by limiting reporting to only events deemed severe enough to warrant an administrator’s attention.
- Event type—Event type classifies events into antivirus, anti-spam, content filter, and system events.

### Third-Party Tool Reporting

A third-party tool, such as McAfee’s ePolicy Orchestrator, centrally manages security policies and procedures across a network. These applications passively monitor activity on a network, prevent changes to system configurations, and ensure that workstations and servers remain in compliance with security policies. These tools provide consolidated reporting and library of predefined reports.

For example, when an ePolicy agent is installed on a secure content appliance, events are sent to the ePolicy Orchestrator and included in that application's reports. Reports specifically designed for the WebShield secure content appliance include:

- Content Filter Report by Rule
- Content Filter Report by Rule and Time
- Content Filter Report Rules Triggered
- URLs Blocked
- Viruses Detected
- Spam Detections by Appliance
- Top Ten Spammers
- Infection History

Other reports provide statistics on throughput and more detailed information on viruses and spam prevention. In addition to the predefined reports, administrators have the option of defining custom reports using Crystal Reports, an industry-leading reporting application.

### **Email, SNMP, and Syslogging**

Many organizations may already have existing event reporting systems or practices in place based on email alerts, SNMP messages, and syslog. Third-party tools usually support email notification of an administrator when an event occurs.

SNMP is a protocol designed for sending messages from a managed device to a network management system. An agent resides on the managed device—in this case, the secure content appliance—and sends messages to a management device that logs the message. Syslog is an application that allows distributed systems to centralize their logging information.

 For more information about SNMP, see <http://www.snmpLink.org/>. For more information about syslog, see the Syslog RFC at <http://www.faqs.org/rfcs/rfc3164.html>.

Regardless of how administrators choose to log events and report on performance, the overall monitoring process is essentially the same.

## Monitoring Tasks

There are several steps to maintaining an effective monitoring procedure.

### Task 1: Establish a Monitoring Policy

The first step is to establish a monitoring policy that defines which measures will be tracked, how often measures will be taken, who is responsible for collecting measures, who is responsible for reviewing measures, who is responsible for acting on particular information, and the conditions that warrant the creation of a new baseline set of measures.

### Task 2: Establish a Baseline

The baseline should be documented along with secure content policies and current appliance configurations. The baseline document should include:

- Volumes of traffic by protocol. Regular variations from the average, such as peak periods of FTP traffic at the end of the week caused by backup files copied to an offsite location, should also be noted.
- Number of email messages sent and received.
- The number and percent of total messages of viruses and spam and phishing messages detected over a period of time, such as a day or week. Track these by protocol as well.
- The number of misclassified spam and phishing messages, both false positives and false negatives.
- Number of URLs blocked.
- Number of reports of mistaken URL blocks.
- Number of reports of missed viruses and other malware.

The negative measures, such as missed viruses and mistaken URL blocks, will be relatively infrequent, so those should be measured over long time periods to get reasonably accurate measures.

### Task 3: Analyze Reports

Once the baseline is established, administrators should review reports, event logs, and other information collected to monitor variations from the baseline:

- Substantial increase in virus detections
- Detection of new, high-volume spamming sources
- Users with unusually high number of blocked URLs
- Blocked sites with unusually large number of attempts to access
- Unexpected changes in traffic by protocol

#### Task 4: Verify Accuracy of Content Filtering

Checking the accuracy of content filtering is a time-consuming task and can rarely be performed as frequently as analyzing reports. As time permits, administrators should review message quarantines and reports of virus or other malware infections to determine the number of errors in the filtering process. Depending on the type of error, one of several actions may be warranted:

- Notifying the appliance's vendor about missed spam or emailing the miscategorized message to customer support
- Adding or removing names from white lists and black lists
- Educating users about browsing non-business-related Web sites
- Verifying the version and patch level of secure content appliance software and libraries to ensure the latest versions are used

In addition to the regular four tasks, administrators may need to redefine their baseline under some circumstances. Significant changes in network infrastructure and number of users, the introduction of new enterprise applications, or major organizational changes, such as a merger, typically justify creating a new baseline measure for secure content monitoring.

#### Q 2.3: How can administrators tune secure content policies using global and user-specific rules?

**A:** The first step in tuning content filtering rules is to understand how they are organized and how they are applied. Content filtering is dependent on a wide range of rules that control what content can enter and leave a network. These rules are grouped into related sets of rules known as *policies*. Policies are generally defined for a particular function or protocol and apply to either inbound or outbound traffic. Policies are typically created for:

- Simple Mail Transfer Protocol (SMTP)
- Post Office Protocol 3 (POP3)
- File Transfer Protocol (FTP)
- Virus scanning
- Hyper Text Transfer Protocol (HTTP)

Policies include rules for different aspects of a protocol. An SMTP policy, for example, typically includes rules governing:

- Data size limits, the number of characters per line, and the maximum number of received lines
- Denial of Service (DoS) protection by checking for trivial commands and the length of SMTP conversation
- The number of MX records received after a DNS query for a mail exchanger

There are two types of policies: global and non-global. Global rules apply to everyone. A typical global rule performs a medium-level virus scan on all incoming SMTP traffic. Often, groups of users within an organization will have slightly different requirements and require a non-global policy.

For example, the data warehouse group may download a large data file every Friday night. Due to processing constraints, the file must be downloaded in 1 hour or less; using a medium virus scan can lead to a download time of more than 1 hour. An exception is made for this download and only a low-level virus scan is performed on this file to ensure the data warehouse process meets its service level agreements (SLAs).

### **Global Policies**

Global policies are designed to cover a wide range of threats. Ideally, all users would be covered by a single set of global policies, but that is not always possible. Still, from a management perspective, it is best to minimize non-global policies and keep global and non-global policies closely coupled.

In general, try to keep rules as general as possible. This method allows rules to be applied to the maximum number of cases, which in turn can help minimize the number of rules required. A corollary to this guideline is to not use more conditions than required in a rule. This allows the rule to be applied to the broadest number of events possible. When exceptions to the rule are discovered, non-global policies can be defined to address the exceptions.


### **Non-Global Policies**

Expect exceptions to at least some policies. For example, as a general rule, trade secret information is not allowed to leave the organization. A research and development group may have executed a joint development agreement with a business partner and now need to share information about a narrow range of trade secret processes the company uses. In this case, a new, non-global policy is defined and applied to members of the research and development group working on this project. The policy will allow transmission of email messages that contain phrases associated with the proprietary process only when the email recipient is included in a list of registered researchers at the partner company. In the case of global policy, the rules were as general as possible, in the case of non-global policies, the rules should be as specific as possible.

### **Policy Inheritance**

Non-global policies inherit rules from global policies. This setup helps to ensure that the default behavior of global policies is carried over to non-global policies. Administrators can change only the minimal number of rule attributes to implement the exception they need. For example, a POP3 policy may allow for email attachments as large as 5MB, perform a medium-level virus scan, check for banned words and phrases, and check for known spammers in the send address.

Now engineers in the product design department may have to exchange large computer-aided design (CAD) drawings. The administrator could define a non-global policy that inherits the global policy and then override the 5MB limit on attachments and replace it with a more appropriate limit, such as 30MB. The administrator does not need to change any other part of the inherited policy for it to also apply to engineers. In addition, if in the future, the global policy changes, those changes will be reflected in the non-global policy (except, of course, if the change applies to a rule that is overridden by the non-global policy).

 Two principals of rule design are worth calling out for emphasis:

When designing a global rule, it should be as general as possible. Global policies implement the global behavior of the content filtering system, so rules in global policies should apply to as much content as possible.

Non-global rules are designed for exceptions and should apply to as few instances as possible. These are exceptions to the default behavior of content filtering and should apply to as few instances as possible. This will minimize the number of false negatives (that is, improper content that is not detected).

## Q 2.4: How can administrators use quarantine and deferred mail management to secure content?

**A:** When there is a problem with a message or a document, how can it be addressed? There are several options:

- Delete the message or document
- If the problem is a threat, such as a virus infection, clean the message or document and send it on
- If the message cannot be cleaned, store the message for further review
- If the content of the transmission contains banned words or phrases, store the message for further review
- If the message cannot be delivered, store it for future delivery attempts

The first option works well with known spam or messages that are sent containing banned language. The second option is used to allow messages through after the threat has been eliminated and there is no concern for harm to network resources. The remaining options all require the ability to isolate and store content before acting on them further. A secure content appliance should do so by providing quarantine areas and deferred mail storage as Figure 2.4 shows.




Figure 2.4: At any time, a systems administrator can find the number of quarantined and deferred messages stored in the appliance.


### Quarantining Content

Quarantining content isolates a threat to a storage area on the appliance so that the threat cannot harm network resources. A message should be quarantined if one of two things occurs:

- The message is infected with a virus that cannot be removed and the governing policy dictates that, as a secondary action, the message must be quarantined
- The message is categorized as spam


### Isolating Virus-Infected Messages

Once quarantined, administrators should review the messages and take appropriate action (see Figure 2.5).



**Figure 2.5:** Quarantined viruses are kept in a secure area on the appliance until they are acted upon by the administrator.


Administrators have several options for dealing with these isolated messages. The messages can be viewed, deleted, or forwarded.

 Unless the message is being sent to an antivirus vendor as a sample for analysis, use forwarding with care. The message will still be infected.

Quarantining a virus-infected message keeps the malware from reaching the desktop. Although the desktop is likely protected by local antivirus software, it is better to keep the virus from reaching its destination. The desktop, for example, may not have the latest antivirus signature files or scanning engine. Some malware is now designed to disable desktop antivirus programs. Another technique changes the local host file so that automatic updates programs cannot reach vendor's sites to download updates. Quarantining at the appliance is another example of layered protection that compensates for vulnerabilities and threats to individual components of the security system.

## Isolating Spam

Identifying spam is not an exact science. Some messages that may appear to be spam are legitimate emails and vice versa.

 Question 2.5 explores how spam filtering depends upon rules that score messages as to the likelihood of them being spam.

Some spam is easily identified and systems administrators can be confident that it is actually spam. Borderline cases are more problematic. Administrators do not want to raise the threshold too much on what is considered spam or else spam that should be blocked will make it to recipients' inboxes. At the same time, email administrators do not want to delete a legitimate email that is mistakenly categorized as spam. Quarantining provides a middle ground. Messages that are considered spam can be stored in the email quarantine area until they are reviewed by the administrator and dispatched accordingly.

## Deferred Email Management

Information security is often described as providing CIA: confidentiality, integrity, and availability. Deferred email management contributes to integrity and availability by providing the ability to defer the delivery of messages if there is a problem relaying a message. Ideally, a high-level dashboard display, as Figure 2.6 shows, will include the status of deferred messages.

Protocol status (Counters Reset Jan 01 2005 00:00:01)			
<a href="#">SMTP Viruses Detected</a>	From inside 3 : From outside 40	<a href="#">SMTP E-Mails Received</a>	3457 / Hour (12432)
<a href="#">HTTP Viruses Detected</a>	From inside 50: From outside 64	<a href="#">HTTP Traffic</a>	187.3 MB / Hour (1236425678)
<a href="#">FTP Viruses Detected</a>	From inside 3 : From outside 58	<a href="#">FTP Traffic</a>	12.4 MB / Hour (145425843)
<a href="#">POP3 Viruses Detected</a>	From inside 0 : From outside 12	<a href="#">POP3 E-Mails Received</a>	45 / Hour (234)
<a href="#">Total Viruses Detected</a>	230	<a href="#">SMTP Spam Detected</a>	343
<a href="#">E-Mails Deferred</a>	5	<a href="#">SMTP Spam Blocked</a>	189
<a href="#">Viruses Quarantined</a>	82	<a href="#">SMTP Spam Blocked by RBL</a>	23
<a href="#">Content Quarantined</a>	2		

Figure 2.6: A secure content appliance interface should provide summary information about the number of quarantined and deferred items.

## Controlling Content Distribution

Quarantining is also used with content filtering to ensure that controlled content—such as proprietary information, personal documents, and other confidential material—is not sent outside the organization inappropriately. Like spam, the suspect content may be stored on the appliance and held for review by the application administrator.

Quarantining and deferring are two common methods for creating middle grounds. In the case of quarantining, infected messages, borderline spam, and suspect content can be held and reviewed before letting it pass or deleting it completely. In the case of deferred email, messages can be stored and forwarded at a later time rather than discarding a message after initial attempts fail.

### Q 2.5: Some spam passes through the filters; how can the filtering be improved?

**A:** Spam filters depend upon rules that are designed to identify messages that are truly spam without mistakenly categorizing legitimate email as spam. These rules are created by examining large numbers of spam messages to identify characteristics common to spam. Typically, these rules take into account phrases commonly found in spam and their location within the structure of an email.

### Identifying Spam with General Rules

Spam messages will, of course, vary, but there are common characteristics that anti-spam designers can use to identify the most obvious spam:

- Promises of easily earned money
- Free or discounted items
- Apparent messages from a customer service department
- Surveys
- Fear-inducing messages, such as claims that the reader's PC is unprotected

Marketers have long known that short, well-phrased pitches can get a user's attention. Spammers use the same principal, and, fortunately for the rest of us, this is the Achilles heel of spam. Table 2.1 lists several phrases that are good indicators of spam along with scores, or weights, indicating the relative confidence that this phrase indicates a piece of spam.

Spam Indicator	Score
Get paid for your opinion	2.0
On sale	1.0
Limited time	0.8
Unbelievable prices	0.8
From: Antivirus Administrator	1.2
Dear Friend	1.8
Congratulations you are a winner	2.0

**Table 2.1: Example spam phrases and scores of the likelihood that they are spam.**

Scores are essential to measuring the likelihood of spam. All of the phrases listed in Table 2.1 can be used in legitimate emails. However, if enough of them are used, even if they only weakly indicate spam (for example, through the use of "limited time"), there is a good chance the message is actually spam. Similarly, if only two or three phrases are used but are strongly correlated with spam (for example, "Get paid for your opinion"), chances are good that the message is spam.

The filtering rules add the score associated with each matching spam phrase to find the total score for a message. If the score exceeds a threshold, the message is considered spam.

To illustrate how this works, consider two example emails. The first is a legitimate message in response to a sales call.

Dear Frank,

Thanks for taking the time to meet with me yesterday about our new line of office furniture *on sale* through the end of the month. I'm sure you'll agree that some of our specials are at *unbelievable prices* but we are only offering these to select customers and for a *limited time*.

I've attached a formal proposal for your review. Please feel free to contact me with any questions; otherwise I will call you Friday to follow up.

Regards,  
Mary Jones  
Acme Office Furniture

This message has three phrases commonly found in spam (indicated by bold italics). The total score is calculated as:

On Sale	1.0
Unbelievable prices	0.8
Limited time	<u>0.8</u>
	2.6

Assuming a threshold of 4 (a low tolerance for spam), this message is not considered spam and would pass through the filter. The following example is a fictional but representative spam:

*Dear Friend,*

*Congratulations you are a winner!* For a *limited time*, you can claim your prize from Grand Award Sweepstakes! Just click the link below, provide us with your name and address and the bank account number where you would like the funds deposited.

Again, congratulations,  
Yours truly,  
Grand Award Sweepstakes Prize Committee

This message also has three phrases commonly found in spam (indicated by bold italics). The total score for this message is:

Dear Friend	1.8
Congratulations you are a winner	2.0
Limited time	<u>0.8</u>
	4.6

As the total is greater than the 4.0 threshold, this message would be correctly categorized as spam.

This technique generally works well because it's fast (there is no complex analysis, just string matching and simple arithmetic), and with well-crafted rules, correctly categorizes most spam. However, occasionally, things do not work as planned.

## ***Erroneous Categorizing***

Spam filter rules are not perfect. They apply rules derived from examining large samples of spam and non-spam messages. Using statistical methods, rule designers can generalize rules from large samples to find the best indicators of spam. As with the use of statistical methods in other applications, there is a margin of error. These errors come in two forms: false positives and false negatives.

### **False Positives**

A false positive mistake categorizes a legitimate email as a piece of spam. This mistake occurs when phrases commonly found in spam are used in the email. If false positives are occurring at an unacceptable rate, the threshold for classifying a message as spam may be raised. Doing so will cause fewer messages to fall into the spam category and reduce the chances of a false positive because the message has some, but not many, characteristics in common with spam. For example, the first sample message would have been categorized as spam if a lower threshold, such as 2.5, had been set. Although lowering the threshold decreases the chance of false positives, it increases the chances of a false negative.

### **False Negatives**

False negatives are mistakes that allow spam to pass through as legitimate email. In the ideal spammer world, spammers would be able to maximize their use of marketing phrases that catch readers attention while still “flying under the radar” of the anti-spam filters. As spammers learn the phrases that cause their messages to be filtered, they will vary the content of the message to avoid trigger matches with spam rules. If they can avoid triggering enough rules, their message scores will fall below the threshold and the spam will make its way to the recipient. Such messages are false negatives.

Clearly, there is not a definitive set of rules that will correctly identify all spam while avoiding false positives. Even if you could compile an ideal set of rules for all known spam, it would not necessarily work as well for new spam created by spammers with those very rules in mind. Filtering spam is a cat-and-mouse game. Spammers are constantly trying to avoid detection and will continuously vary their content.

### ***Additional Filtering Mechanisms***

Besides filtering rules, spam can be controlled through the use of white lists and black lists. A white list contains email addresses and domains trusted not to send spam. Business partners, clients, patients, government agencies, public companies, and other organizations that do business with a company may be added to the white list. Any messages that are sent from those addresses are not subject to filtering by spam rules.

The white list is useful for two reasons. First, as the messages are not scanned for spam phrases, the anti-spam application can operate more efficiently. This benefit is especially useful when a small number of domains send a large proportion of all email to a business or organization.

Black lists contain a list of addresses and domains of known spammers. Any message from an address on the black list is categorized as spam and not allowed through. Black lists complement filtering rules based on content phrases. Rather than crafting rules to cover all the spam that may come from known spammers, the black list effectively shuts down traffic from those addresses.

## Staying Up to Date

It is also important to stay up to date on spam filtering rules and the anti-spam engine. Anti-spam designers build new rule sets to address the changing patterns of spam as they emerge. Appliance administrators can use the Update | Anti-Spam option in the appliance to download and install the latest rules and anti-spam engine (the application that executes the rules).

## When All Else Fails

There may be times when the up-to-date spam filters, black and white lists, and threshold adjustments are not effectively blocking spam. When that occurs, contact your vendor and submit samples of the spam that is slipping through. Anti-spam designers will then be able to study the spam and develop appropriate countermeasures.

### Where to Send Spam Examples?

Most anti-spam vendors accept examples of spam that are not blocked by their products. The following list provides some of the most popular vendors along with instructions for submitting spam:

- McAfee customers can send examples of spam to [customer+false-positive@clicknet.com](mailto:customer+false-positive@clicknet.com) and [customer+missed-spam@clicknet.com](mailto:customer+missed-spam@clicknet.com)
- TrendMicro users can send their examples to [spam@support.trendmicro.com](mailto:spam@support.trendmicro.com)
- Symantec users can send spam examples by following the instructions posted at [http://service1.symantec.com/SUPPORT/ent-brightmailkb.nsf/0588786bcfa7fb9888256f72007c8a4b/5d0964f0afa6403f88256f93008006cb?OpenDocument&src=bar\\_sch\\_nam&seg=sb](http://service1.symantec.com/SUPPORT/ent-brightmailkb.nsf/0588786bcfa7fb9888256f72007c8a4b/5d0964f0afa6403f88256f93008006cb?OpenDocument&src=bar_sch_nam&seg=sb)

Finally, when there is a sudden outbreak of a particular type of spam, vendors may develop specialized rules, known as extra rules in the technical documentation, to combat the outbreak.

To summarize, the following steps should be followed to improve the spam detection rate of a secure content appliance:

- Update anti-spam filtering rules
- Update the anti-spam engine
- Adjust spam score threshold
- Add entries to the black list and white list
- Submit samples of missed spam to your anti-spam vendor
- Download extra, specialized rules as needed

Following the first two steps will help to ensure the appliance is configured to filter the latest and broadest range of spam. Adjusting thresholds and configuring black and white lists can fine tune the appliance's performance. In special circumstances, submitting a sample of missed spam and downloading specialized rules may be the correct course of action.


## Topic 3: System Architecture and Secure Content Management

### Q 3.1: Where should a secure content appliance be placed?

**A:** Secure content appliances are used to control what is allowed to enter and leave an organization's network. It follows logically that the device should be located on the perimeter of the network. Perimeters can use a single layer of defense with a single level of firewalls that block ports and filter network traffic at the lower levels of the OSI network model (see question 1.1 for more information about the OSI network model). A common configuration creates a multi-level perimeter known as a DMZ (de-militarized zone).

DMZs use multiple network segments to create three zones: the external zone, which includes the Internet; the internal zone, which includes an organization's network, servers, desktops, and other devices accessible to the internal network; and the DMZ, which lies between the internal and external zone.

Typically, the Web servers are located in the DMZ. These servers need to be accessible from the Internet but also protected from it. Applications running on the Web server invoke applications and services running on servers on the internal network. For example, an application running on the Web server may query a customer database, located in the internal zone. Internet users have no direct access to the database server; all interaction is through a proxy application on the Web server. This configuration balances the need for accessibility with the need for controlled access to the mission-critical servers (see Figure 3.1).



**Figure 3.1:** DMZs provide additional protection of mission-critical servers by limiting access to those servers to trusted proxies in the DMZ.

## Secure Content Device Operational Modes

In addition to the existing configuration of a network, a secure content administrator needs to consider how the device will be configured. There are three options:

- Explicit proxy mode
- Transparent router mode
- Transparent bridge mode

The choice of mode determines whether other devices are aware of its use, how the secure appliance is configured and connected to the network, and how much configuration work is required during installation. The focus here is on how the device is connected to the network.

### **Explicit Proxy Mode**

In explicit proxy mode, servers that communicate with the secure content device are configured to communicate directly with the device. For example, incoming mail is routed from the Internet to the secure content device and then passed to the internal email server. When in explicit proxy mode, only HTTP, SMTP, POP3, and FTP traffic should be routed to the secure content device; all other protocols are refused by the device.

In explicit proxy mode, administrators have great flexibility in placing a secure content server. The one configuration rule is that the secure content appliance must be located behind a firewall; other than that, the device can be placed anywhere in a DMZ or internal network.

The reason for the flexibility in positioning the devices is that other devices that use the secure content device are configured to explicitly send traffic to and receive traffic from the device. For example, in a switched network, the secure content device can be connected to any router or switch in the network. Although, the device can be placed anywhere in the network, some positions will still be better than others.

Consider the flow of traffic. As all incoming and outgoing email and Web traffic will pass through the device, it should be located in a segment that minimizes additional traffic. For example, if an email server and Web server are in the same network segment, it makes sense to position the secure content device there because traffic will eventually flow into that segment to reach the email and Web servers.

Also consider the bandwidth utilization on a segment. If network traffic is near capacity on a segment, adding a secure content device to that segment will only increase the network load.


### **Transparent Router Mode**

In transparent router mode the secure content device acts as both a content filter and a router. As in explicit proxy mode, the secure content device should be behind the firewall.

In transparent router mode, clients are not reconfigured to send and receive traffic to and from the secure content device. Only the firewall and other routers must be reconfigured to direct traffic to the secure content device. This setup eliminates the need to change client configurations but it does limit options for placing the device.

It is recommended that a secure content device in transparent router mode be placed between the firewall and a router. In cases in which the secure content device is the only router, it should be placed immediately after the firewall.


Transparent router mode is recommended for networks that use firewall rules to control traffic. In explicit proxy mode, packets are redirected to the secure content device making the client's IP address unavailable to the firewall rules engine (see Figure 3.2).



**Figure 3.2:** In transparent router mode, the secure content device functions as both a content filter and a router.

### **Transparent Bridge Mode**

Transparent bridge mode is similar to transparent router mode, but simpler. As Figure 3.3 shows, the secure content device does not route traffic, it simply passes traffic between two network segments. No clients, firewalls or routers must be reconfigured. In transparent bridge mode, the secure content appliance should be placed between the firewall and the router.



**Figure 3.3:** In transparent bridge mode, the device joins two segments of a network and passes traffic between the two without routing.

Secure content appliances should be placed inside a well-configured firewall. Administrators should also determine which mode the appliance will use. Transparent bridge mode is the easiest to configure; however, all traffic will pass through and be filtered. In high-traffic networks, you may find better performance by configuring the appliance in explicit proxy mode and sending only relevant traffic (for example, HTTP, FTP, SMTP and POP3). Of course, if the routing features of the appliance are used, the appliance should be placed at the junction of two or more network segments.

### Q 3.2: Why are desktop antivirus software and personal firewalls still needed?

**A:** There are two primary reasons for deploying desktop antivirus software and personal firewalls on networks that contain secure content devices:

- No security device can address all potential security threats; multiple layers of security are required to maintain network and system integrity
- Mobile devices are not protected by a secure content device when they are not connected to the network

All organizations are subject to the limits of any single security technique or tool and a growing number are faced with the challenge of managing and securing mobile devices.


## **Layered Security**

The principal of layered security dictates that multiple forms of countermeasures are required to ensure the integrity of a network and related infrastructure. This approach is also known as defense-in-depth—security measures are placed throughout the network on the perimeter, servers, workstations, and mobile devices. Defenses also operate at different levels of the network; for example, network firewalls can filter based on protocols and ports and application firewalls can examine XML message content to identify invalid or unauthorized messages.

### **Example: Firewall Rules**

Take a simple example. A company's network includes an IBM mainframe, a Microsoft SQL Server system, and a number of desktop and laptop Microsoft Windows clients. Let's assume the mainframe and SQL Server support business partners in a supply chain by providing access to parts and inventory information. Many users within the organization do not need access to those applications but their desktops and laptops are connected to the same network so they are exposed to the same traffic. By installing personal firewalls on those devices, the network administrators can block traffic on ports that must be open on the network but are not used locally. This action provides an extra layer of defense against threats that are propagated on those protocols or ports (see Figure 3.4).

Similarly, if a device were to become infected with a blended threat that included a keylogging program, a personal firewall could prevent the transmission of information over, for example, an Internet Relay Chat (IRC) protocol. There may be legitimate reasons for others to use that protocol, so the network firewalls would allow the traffic through. Using a personal firewall, one can configure finer-grained rules and better protect information flow.



*Figure 3.4: Defense-depth allows for both course-grained security for the overall network and fine-grained filtering localized to systems that require it.*


### Example: Antivirus Protection

Another example of the benefits of layered security involves the use of encryption. To ensure the privacy and integrity of messages, email users can encrypt a message and apply a digital signature to an email. Encryption scrambles the message so that it may not be read in transit to its destination. A digital signature is a string of characters appended to a message that is generated using a hash algorithm and a code known as the sender's private key. Upon receiving the message, the receiver uses another code, called the sender's public key, to decrypt the message and recalculate the digital signature. If the calculated signature matches the one in the message, the recipient knows the message is authentic and has not been tampered with.

Digital signatures and encryption are well-designed to meet the needs of privacy and integrity, but what happens when a virus is attached to a protected message? Secure content managers have a few options:

- The secure content appliance can reject the email and not deliver it to the recipient.
- The secure content device can remove the virus and send the rest of the message to the recipient. In that case, the content of the message has changed, so the recipient will not calculate the same digital signature; he or she will not know if any change other than removing the virus has occurred.
- The administrator can leave the virus embedded in the message and have a desktop antivirus program detect and remove the malware.

None of these options is ideal. Administrators must choose between denying a service to a user (either reliable delivery of email or message integrity checks) or allow a known piece of malware into the network (see Figure 3.5).



**Figure 3.5: Encrypted, digitally-signed emails are sometimes better handled by desktop antivirus than by network-level scans.**

By combining countermeasures at different points in the network and using multiple types of security tools, network and security administrators have more options for configuring security mechanisms that allow them to accommodate the needs of users and applications within the network. This type of layered defense also provides multiple points of protection should a single point become compromised. However, even within a well-secured network, mobile devices present security challenges.

## Securing Mobile Devices

One of the advantages of using a network appliance for securing content is that any content entering or leaving the network is protected. Unfortunately, network devices do not always stay put. Laptops and PDAs with network access pose a particular threat to network security because they are allowed to physically disconnect and reconnect to the network, often at will.

Consider how easy it is to circumvent perimeter defenses with mobile devices. Imagine an employee who is blocked from browsing his favorite music sharing site during his lunch hour, so he decides to disconnect his laptop from the internal network, walk across the street to the local coffee shop with a WiFi hotspot, download music files (and unknowingly, some spyware), then return to the office to reconnect to the network. Spyware, that would have been blocked by the antispyware mechanism in the secure content appliance had the appliance not blocked the URL (another example of a layered defense) is now within the organization's network.

A single point of detection and prevention is not sufficient with mobile devices. As mobile devices will not always have the security services of the network available to them, these devices must have local versions of antivirus, anti-spyware, and personal firewalls.



For additional protection of mobile devices, consider a third-party tool such as McAfee ePolicy Orchestrator, which can ensure devices remain in compliance with security policies. If a device is changed and no longer in compliance, a third-party tool can enforce compliance and updates as well as notify administrators when threats or rogue systems attach to the network.

Layered security is considered a best practice among security professionals and should be practiced to levels appropriate to an organization's needs and capabilities. A secure content appliance adds a layer of protection and in fact uses multiple layers within itself (for example, malicious content missed by URL blocking can be caught by antivirus software). Mobile devices by their nature circumvent perimeter defenses such as firewalls and secure content devices. They must have their own localized security tools in addition to those available on the network.

### Q 3.3: How does a secure content appliance work with Web servers, caching servers, and application servers?

**A:** There are many secure content appliances that provide content filtering services. Like other servers in a distributed networking environment, the solutions use common protocols to communicate with other servers. Key topics to consider when introducing a secure content appliance are:

- What protocols are used on the network to be protected?
- Where should the appliance be positioned for maximum protection?
- How will the secure content appliance affect overall system performance and functionality?

## **Network Protocols**

The Internet uses several protocols, or standards for communication, but not all of them are relevant to securing content. For example, the low-level Open Shortest Path First (OSPF) protocol used by routers is not subject to content filtering. The most important protocols from a content filtering perspective are:

- Hyper Text Transfer Protocol (HTTP) used by the Web
- File Transfer Protocol (FTP)
- Simple Mail Transfer Protocol (SMTP) used to send email messages between servers
- Post Office Protocol 3 (POP3) used to retrieve email from servers

Firewalls generally allow traffic using these protocols to pass in and out of a protected network. Therefore, a secure content appliance would have to be configured with policies defined for each protocol to ensure maximum protection. In addition to defining policies, the level of protection is also dependent on how the secure content appliance is positioned in the network and what traffic is analyzed.

In some cases, a systems administrator may want to scan all HTTP, FTP, SMTP, and POP3 traffic as soon as it passes through the firewall. In other cases, there may be high volumes of traffic to a program running on an application server that need not be analyzed. For example, the traffic may be an XML data exchange between two managed servers, so the content is well understood and filtering it would just put an additional, unnecessary load on the appliance.

## **Positioning the Secure Content Appliance**

The position of the secure content appliance will depend, in part, on which operational mode is used. There are three operational modes:

- Explicit proxy mode
- Transparent router mode
- Transparent bridge mode


### **Explicit Proxy Mode**

In explicit proxy mode, network devices are configured to send traffic directly to the secure content appliance. The appliance in this mode acts as a proxy processing traffic for other network devices.

This mode is used, for example, if only SMTP traffic is filtered. In that case, external mail servers would be configured to send mail messages to the appliance, which would filter the traffic and then forward it on to the internal mail server. Similarly, explicit proxy mode could be used to scan HTTP traffic by using the secure content appliance as a proxy for Web servers.

The advantage of this mode is that systems administrators can target a subset of all network traffic for filtering and avoid unnecessary processing by the appliance. One relative disadvantage of this approach is that it requires additional work on the part of the systems administrators to configure servers to explicitly send traffic to the secure content appliance.

As Figure 3.6 shows, in explicit proxy mode, the position of the appliance is determined more by traffic patterns across network segments than the need to have the appliance in a particular position. As devices are configured to send traffic to the appliance, it can be positioned virtually anywhere on the network. Of course, the secure content appliance should still be positioned behind a firewall.




**Figure 3.6:** In explicit proxy mode, the appliance can be placed anywhere in the network because traffic is routed as needed to the appliance.

### Transparent Router Mode

In transparent router mode, the appliance acts as a router as well as a content filter. Other network devices do not need to be configured to explicitly send traffic to the device unless it is also acting as your default gateway. The appliance should be placed just inside the firewall so that all traffic entering or leaving the network is scanned. The content filtering is done transparently to the devices generating traffic. When in transparent network mode, the appliance routes traffic between two networks.

### Transparent Bridge Mode


In transparent bridge mode, the secure content appliance joins two physical networks, allowing them to be treated as a single network. No routing is performed. This setup is a simpler configuration than transparent routing model and requires less configuration. Like transparent router mode, in transparent bridge mode, the appliance should be positioned between the firewall and other network devices (see Figure 3.7).



**Figure 3.7:** In both transparent bridge mode and transparent router mode, the secure content appliance should be placed just inside the firewall.


### Configuring for Performance and Functionality

Maintaining acceptable performance levels is a major concern in many network environments. Tools, such as caches and load-balancing hardware, are often introduced to compensate for increasing demands on network devices. Caches improve performance by locally storing frequently used content so that the same content is not constantly retrieved from its source Web site or database. When processing loads on applications servers increase, load-balancing hardware is sometimes used to divide the workload between multiple applications servers. In both of these cases, a secure content appliance can still easily fit into the network to provide the content scanning functionality needed. As Figure 3.8 shows, the secure content appliance can be positioned between the firewall and Web cache so that any content stored in the cache has been filtered.



**Figure 3.8:** The secure content appliance is placed between the intranet firewall and the Web cache to ensure content that reaches the cache has been appropriately filtered.

In situations in which load-balancing hardware is used with application servers, the secure content appliance is placed prior to the load-balancing device in the traffic flow.

 The secure content device is designed for load sharing. In cases in which traffic and performance demands are so high that a single appliance cannot meet performance needs, multiple appliances can be configured in a load-sharing arrangement. In this configuration, a master appliance receives all traffic and then passes it to load-sharing devices, which perform single functions, such as virus and spam scanning. This division of labor allows for a more simplified configuration than had traditional load-balancing techniques been used with multiple appliances.

There are multiple ways to configure the secure content appliance to work with other network devices. By considering the protocols to filter, the devices which require filtered traffic, the routing services required, and the performance demands on the network, systems administrators can find an optimal configuration of the secure content appliance with other network devices.


### Q 3.4: Can a secure content appliance be attacked?

**A:** A widely held belief in the security community is that any device can be compromised if a group of skilled perpetrators has the time, resources, and desire to break in. However, some security countermeasures, such as encryption with very long keys, can take years and massive computing resources to break. As a result, most security measures are designed to keep attackers and others at bay for a long enough time for the attempts to be discovered or to raise the cost of breaking in to a level so high that the value of the information stolen is no longer worth the cost of retrieving it.

There are certainly incentives to attacking a secure content appliance. For example, if an attacker were able to compromise the secure content appliance and change the virus scanning policy for the Hyper Text Transfer Protocol (HTTP), the attacker could deliver to a device spyware that includes a keylogger. If the antivirus scanning level of a global policy could be reduced, a blended threat could be transmitted, which could exploit vulnerabilities in database applications and steal private or company confidential information. Like firewalls, a secure content appliance is a front-line safeguard; unlike firewalls, though, secure content appliances perform complex analysis well beyond the abilities of a firewall. Breaking through a secure content device can be much more advantageous to an attacker than compromising a firewall. Clearly, there is no shortage of motive to attack a secure content device.

### **Security and Operating System Architecture**

Sound security begins with sound design. The Linux operating system (OS) used in some third-party secure content appliances is based on a ringed architecture, as Figure 3.9 shows.




**Figure 3.9:** The Linux OS uses four major subsystems that provide a ringed architecture.

The purpose of this type of architecture is to isolate critical functions, such as process scheduling and memory management, from user programs that may contain errors or malicious code. Each subsystem is designed to perform specific tasks. The Linux kernel manages five main tasks:

- Process scheduling
- Memory management
- Virtual file system
- Network interface
- Interprocess communications

The kernel depends upon hardware controllers to provide some services and in turn provide services to the next layer, OS services. Only the kernel has access to hardware features related to memory and processing. Users may not change code in the kernel. OS services provide file system and window management services, which are used by user applications, such as databases, Web servers, and other applications.

 For more information about the Linux kernel, see Ivan Bowman's "Conceptual Overview of the Linux Architecture" at <http://plg.uwaterloo.ca/~itbowman/CS746G/a1/>.

By separating duties between levels, the system is protected from malicious code while still allowing programmers to invoke OS services as needed. For example, an application can make a request to write a block of data to a file system and may even specify exactly where in a file the data block is to be written, but the application may not specify a location using disk geometry (such as the track, cylinder, and sector of a disk). The kernel hides those details behind the virtual file system that is further abstracted by the file system in the OS services layer. The benefits of this type of protection become clear when you consider the potential impact if it were missing.

An early form of computer virus was the boot sector virus. These viruses write to specific areas of disks, known as Master Boot Records (MBRs), which contained OS files and code. By changing critical code used to manage the disk, a virus writer could control the behavior of the disk.

A ring architecture, such as used in Linux, provides a well-established and effective mechanism for preventing disruption of critical OS functions from most malware. Although users and attackers can add programs, even ones with malicious code, isolating these programs minimizes the risks that malicious code can disrupt core services.

### ***Hardening an OS***

Hardening an OS consists of several steps:

- Shutting down unnecessary services
- Patching the OS and services
- Configuring services to reduce vulnerabilities

Like other areas of security, no one of these steps is enough to protect a server, but in concert these steps can significantly reduce the risk of exposure to a security breach.

## Shutting Down Unnecessary Services and Removing Unneeded Programs

Linux distributions provide a dizzying array of applications and utilities including compilers, graphical interfaces, databases, Web servers, communications programs, multimedia systems, personal productivity packages, file transfer programs, windows managers, and more. Very few of these are necessary to perform content filtering. In the best case scenario, installing these services simply consumes disk storage; in the worst case scenario, they introduce vulnerabilities.

Take a compiler for example. The source code for Linux and other open source systems is readily available on the Internet. If an attacker could introduce a piece of code onto the server and then recompile the program, the attacker could compromise a server regardless of the hardware platform. Simply removing the compiler in this case would ensure the server could not be compromised in this way.

In other cases, an attacker does not even have to introduce a vulnerability—it exists already. For example, programs that do not perform range checking are subject to buffer overflow attacks. During these attacks, the overflow either disrupts the functions of a program or can facilitate the introduction of new code during the program execution. The new code changes the behavior of the program to perform some malicious action, such as acquire root access and copy the password file to an ftp site controlled by the attacker. Needless to say, if the program with the vulnerability is not running, the attacker cannot exploit it.

## Patching the OS and Services


Another method of protecting a secure content device is to apply patches to services and OSs. Many seasoned IT administrators have mixed feelings about patching. On the one hand, it is comforting to know that developers are continually correcting vulnerabilities, improving performance, and making other enhancements to their systems. On the other hand, many systems administrators have learned the hard way about dependencies between components. A critical business application may break after a service pack is applied because, in addition to patching a known vulnerability, the service pack might include dozens of other changes to code. Such is not the case with network appliances.

One of the benefits of network appliances is that they are strictly controlled by vendors. Every piece of software, every service that runs, and every dependency between modules is known and tested by the vendor before the appliance ships. The reduced flexibility to systems administrators is actually a benefit: the vendor only needs to support a small number of possible configurations.

## Configuring Services to Reduce Vulnerabilities

The Bastille Hardening program has been used in secure content appliances, such as those provided by McAfee. This program analyzes a configuration and guides administrators through the hardening process. Bastille is a well-known and widely used hardening program for Linux and HP-UX and is recommended by the Center for Internet Security. Following Bastille recommendations can help reduce exposure to vulnerabilities in a number of areas including:

- Patches
- File permissions
- Account security
- Miscellaneous daemons
- Sendmail
- DNS
- Printing

 For more information about the Bastille Hardening program, see <http://www.bastille-linux.org>. In addition, the Center for Internet Security is an excellent resource for security benchmarks (<http://www.cisecurity.org/>).

### Q 3.5: How do appliances stay up to date on the latest threats?

**A:** Vendors typically provide frequent updates for secure content appliances. Virus developers and spammers change their techniques and content to avoid detection, but vendors keep abreast of these changes and update both signature files (virus definition and spam definition files) and the scanning engines that use those signature files. Fortunately, keeping the appliance up to date on the latest threats is a relatively simple matter because the appliance automates virtually all of the work.

#### Tracking Updates

The option Monitor | Updates tool allows administrators to set the schedule for updating both signature files and scanning engines.



**Figure 3.10:** The automatic update facility allows for separate scheduling of signature file (rules) and scanner (engine) updates.

The option Monitor | Updates displays information about the status of antivirus and anti-spam automatically scheduled updates. The display includes

- Names of scheduled updates
- Current status of each scheduled update
- Date and time of last update

Antivirus and anti-spam updates are configured separately.

### **Updating Antivirus Applications**


Antivirus applications should be updated at least once per week but more frequently is preferable. Most antivirus designers and developers are regularly updating virus definition files to counter emerging threats. Although the application uses heuristic, or “rule of thumb,” filters to catch a class of viruses as well as virus-specific rules to detect specific viruses, it is best to keep the virus definition files up to date to ensure new threats are consistently detected. When a virus spreads suddenly or can inflict significant damage, antivirus vendors will create extra definition files and release them as soon as possible to combat the spread of the virus.

In addition to keeping the virus definition file up to date, the appliance keeps the antivirus engine up to date. The engine is the program that reads the virus definition file and uses those definitions to identify viruses and clean infected files. The engine is updated to detect new types of viruses that do not have the same characteristics as older viruses. In general, the engine is updated every few months.

Update files can be downloaded to the appliance from three sources:

- An authorized vendor FTP site
- A proxy on the local area network (LAN), if the appliance is not configured to access external FTP sites directly
- A server within the network that has already downloaded the files

To keep up to date, use the Monitor | Status | General Information or the Monitor | Updates option to determine the current revision level of the antivirus engine and virus definition file.

 Virus definition files and antivirus engine updates are available at <ftp://ftp.nai.com/virusdefs/4.x/>.

### **Updating the Anti-Spam Application**

Like the antivirus application, the anti-spam components of a secure content appliance are updated frequently by vendors. Both spam definition files and the anti-spam engine are kept up to date to counter changes in spamming practices.

New anti-spam definition files will contain rules to identify spam that may have slipped past earlier filters. In addition, extra rules are created for sudden widespread instances of spam that are otherwise not caught by existing filters. By keeping up to date on the latest virus and spam threats, secure content appliances are able to deploy appropriate countermeasures as new threats emerge.

---

## **Topic 4: Secure Content Appliance Performance**

### **Q 4.1: What are threats to content and information assets must organizations address?**

**A:** The major threats to information assets include:

- Viruses, worms, and other malware
- Spam
- Phishing scams
- Spyware

Left unchecked, these threats can leave organizations with compromised computers, security breaches, loss of information, identity theft victims, and reduced ROI on information technology (IT) investments because resources are consumed with non-business related content.

### **Viruses, Worms, and Other Malware**

Malicious programs have evolved from small, machine-language programs propagated by sharing floppy disks to sophisticated collections of programs that can gain control of systems, steal personal information, and replicate rapidly. This malicious software, or malware, falls into several broad categories:

- Viruses
- Worms
- Trojan horses
- Keyloggers
- Backdoors
- Rootkits

#### **Viruses**

Viruses are malicious programs that attach themselves to other programs to execute and propagate. Viruses can run either as executable programs or as macro-viruses embedded in applications, such as Microsoft Word. Viruses consist of two basic parts, a replication mechanism and a payload, the destructive part of the virus.


In the early days of antivirus protection, vendors could discover identifying patterns within a virus that uniquely identify that virus. This identification allowed researchers to create libraries of signatures to detect a virus that could then be removed or at least quarantined. Virus writers responded with encryption to hide the tell-tale signs of a virus, then with the development of mutating viruses, which change in structure but retain the same functionality.

Mutating viruses require a radically different detection approach: rather than look for the same pattern, antivirus researchers must look at the behavior of a program to determine whether it is malicious. Some indicators are commands to change a file without first being commanded by a user and writing to particular memory locations used for low-level system tasks.

### **Worms**

Worms are similar to viruses in that they are malicious programs that self-propagate. Unlike viruses, worms do not depend upon other programs. Worms exploit vulnerabilities in systems and move, sometimes quite rapidly, from one system to another.

One of the most famous worms is SQL Slammer, which flooded large sections of the Internet within 15 minutes of its release. SQL Slammer exploited a vulnerability in Microsoft's database, SQL Server, forcing unpatched servers to generate database server requests sent to random IP addresses. The worm was not sophisticated in how it targeted other victims; instead it depended upon flooding the Internet with packets knowing at least some of the requests would target a SQL Server database.

 For a description of the spread of the worm from the perspective of an Internet operations center, see Paul Boutin's "Slammed!: An inside view of the worm that crashed the Internet in 15 minutes" at <http://www.wired.com/wired/archive/11.07/slammer.html>.

SQL Slammer demonstrated the need for both patching and content filtering. Once a malicious piece of software is released on the Internet there may be little time to craft a custom response.

### **Trojan Horses**

Trojan horses are programs that appear to serve one purpose and actually perform another. A program that promises to synchronize your desktop computer's clock with a highly accurate atomic clock but also collects personal information about your surfing habits is a Trojan horse.

Trojan horses, unlike other malware, may be installed intentionally on a system. A user may not realize that a peer-to-peer (P2P) file sharing application he or she downloads to share music files also contains a program to capture usernames and passwords that are then transmitted to an attacker's server. Again, content filtering can help identify malicious programs that are brought into a network intentionally, albeit, under false pretenses.

### **Keyloggers, Backdoors, and Rootkits**

Keyloggers, backdoors and rootkits are some of the most dangerous forms of malware. Keyloggers simply record keystrokes and send the captured information back to an attacker for analysis. Text scanning programs can quickly analyze those files for personal information, such as Social Security numbers, bank account numbers, credit card numbers, as well as usernames and passwords.

Backdoors are changes to a system's configuration and create a way for an attacker to gain control of a system. Creating an administrator or root account controlled by the attacker is one example of a backdoor.

Rootkits allow attackers to gain control of a system but also hide the attacker's tracks, making detection especially difficult. As rootkits can gain control over any aspect of an operating system (OS), the only way to ensure the malware is eliminated is to format all drives and restore the system from a known uninfected backup.

## Spam

Spam is unwanted, unsolicited email. Spam not only wastes the time of end users but also taxes system resources, such as storage and bandwidth. In addition, it places demands on email administrators who have to manage the additional volume of email.

The key to controlling spam is to identify it as it comes into the network and deleting or quarantining it. This ability assumes that the spam detection software is highly accurate: it does not identify legitimate mail as spam (known as a false positive) or miss identifying spam (a false negative).

## Phishing scams

Phishing scams are cons that use email to lure victims into divulging personal information of sending money to a bogus charity or get rich scheme. Phishing scammers masquerade as legitimate businesses (banks, eBay, and PayPal are favorites of phishing scammers) by sending emails with official logos and urgent messages about the need to update account information or verify personally identifying information.

Once they have the victims' attention, scammers lead victims to Web sites that appear legitimate but are actually phony versions set up to capture information such as bank account numbers, usernames, and passwords. Phishing scams are difficult to detect and user education is one of the best defenses for this threat.

 To learn more about phishing, see the Anti-Phishing Working at <http://www.antiphishing.org>.

## Spyware

Spyware, sometimes called adware, is malicious code that captures information about users and their online activities without their knowledge. In addition to violating users' privacy, spyware can negatively impact system performance. Consequences of spyware include:

- Loss of privacy and identity theft
- Decreased system performance
- Disabling security software, such as antivirus and firewalls, which leaves systems vulnerable to other malware infections
- Displaying unwanted pop-up advertisements
- Changing host files and other networking files causing users to unintentionally navigate to spyware-promoted sites.

As with other malware, spyware can be removed, but it is better to prevent its introduction in the first place by filtering content and blocking spyware.

## Q 4.2: How can an organization protect against spyware?

**A:** Spyware, and its slightly more benign variation, adware, are programs that track user's activities and gather information without the user's knowledge. Unlike viruses and worms, spyware is not intended to cause direct and immediate damage to IT infrastructure. Instead, these programs are designed to collect information about users' identities, including account numbers, drivers' license numbers, usernames, passwords, Social Security numbers, and other personal details. This information is transmitted back to those who deployed the spyware.

Regardless of the original intent, spyware often leads to poor system performance and unstable systems. Multiple infections result in numerous processes consuming system resources and interacting in unpredictable ways. Some spyware changes system configurations, for example, turning off firewalls to ensure the spyware can function as expected. This behavior leaves infected systems open to further damage from other malware. Spyware is used to distribute advertisements, aid in identity theft, and perform affiliate fraud to steal fees from legitimate referring sites.

### ***Keeping Spyware Out***

The best way to deal with spyware is to keep it off your network. Perimeter defenses, including a secure content appliance, can block spyware as it enters the network. Spyware, like viruses and spam, can be detected using signature matching engines. This detection requires a library of up-to-date spyware signatures and a high-performance pattern-matching engine that can scan incoming traffic.

### **Define a Spyware Policy**

Using the Internet inherently requires a balancing of risks and benefits. Spyware is one of those risks, and organizations should articulate the tradeoffs they are willing to make in balancing the utility of Internet services. A policy should include:

- A statement of acceptable use with regard to Internet sites. Spyware is often downloaded from popular entertainment sites and peer-to-peer networks.
- A description of the protocols that will be scanned, how spyware will be disposed of, and acceptable impact on network performance by spyware-scanning tools.
- A general approach to monitoring and event response. Procedures should be defined with detailed descriptions of how to respond to particular events; for example, if more than  $x$  pieces of spyware are detected from a site, the URL for that site will be added to the content filter blacklist.

Once a spyware policy is in place, its objectives should be implemented using a secure content appliance along with user education and maintenance procedures.

## Scanning Multiple Protocols

Spyware can travel over multiple protocols. Someone browsing a peer-to-peer site opens up his or her system to downloading spyware at the same time. A blended threat piece of malware attached to an email can include keylogging and cookie tracking programs that record a user's online activity. Spyware can also come in a Trojan Horse, such as a utility downloaded via FTP that supposedly keeps your computer's clock synchronized with an atomic clock. It is essential to scan HTTP, FTP, SMTP, and POP-3 traffic as it enters the network.

## Monitoring Spyware Detection

In addition to understanding the importance of blocking spyware, security administrators must grasp the volume of spyware reaching the network perimeter. Sudden spikes in spyware detection may indicate:

- An increase in spyware deployment on the Internet in general
- Better detection of spyware by the secure content appliance
- An increase in user browsing and downloading at sites from which spyware is launched
- An increase in spyware components included in blended threat viruses

We will likely have to live with the ever-increasing spyware deployments on the Internet for the foreseeable future; there is not much we can do, from a technical perspective, to slow that trend. Legislative and legal means might ameliorate that problem some. but we should assume spyware, like viruses, are a threat that cannot be eliminated but can be controlled. Better detection methods will also lead to increases in spyware detection and of course are welcome.

The third cause of increased spyware detection is best addressed by user education and clearly defined policies on legitimate use of IT infrastructure. Analysis of logs can identify the sites at which spyware is entering the network and the URLs can be blocked using content filtering functions of the secure content appliance.

As companies and other organizations get better at patching operating systems (OSs) and locking down networks, attackers will find other vulnerabilities to exploit in the never-ending cat-and-mouse game of creating new threats in response to countermeasures deployed by security professionals. Keylogging and URL hijacking are just two ways attackers can collect useful information—both can be done with spyware.

Understanding changes in the patterns of spyware detection can help administrators better pinpoint the specific threats to their organizations.

## Educating Users

Like so many other areas of information security, user education is a key element of a successful strategy. Users must understand how spyware works, how it gets onto computers, and how to minimize the chance of getting stuck with it. When users understand this malware can steal personally identifying information, reduce system performance, and lead to unstable systems, they will have plenty of incentive to help address the problem.

### Q 4.3 How can an organization protect against phishing?

**A:** Phishing is the practice of tricking individuals into disclosing private information, especially financial and identifying information. Organizations should implement both educational and technical measures to protect against phishing.

#### ***Technical Controls for Phishing***

Phishing attacks, like spam, have distinguishing characteristics that make automatic detection possible. Content filtering, such as that provided by a secure content appliance, can identify and block many phishing attacks. Some of the tell-tale signs of a phishing email are:

- Requests for account numbers, Social Security numbers, or other identifying numbers
- Requests for funds in return for exceptional returns later
- Links to Web sites with names similar to legitimate financial institutions' Web sites

Phishing perpetrators are adapting their techniques to avoid detection. Researchers at the Anti-Phishing Working Group found that although the number of phished brands remains about the same, phishers are now targeting smaller brands, making the pool of potential targets much larger. Nonetheless, technical measures can effectively reduce the threats from phishing attacks.


For example, antivirus techniques can counter phishing attacks that deploy keylogging malware to capture passwords and account numbers. An emerging threat is the use of Trojan programs that change users' host files in order to redirect users from legitimate bank sites to phishing sites. Again, antivirus-type scanning can detect and prevent this type of malware from reaching users' desktops. Another technique for controlling phishing is URL filtering. As phishing sites are discovered, they can be include on URL blacklists to prevent users from inadvertently reaching those sites and disclosing usernames, passwords, and account numbers.


Technical countermeasures that include content filtering, antivirus scanning, and URL blocking will contribute to reducing the risk of phishing, but no technical solution will be 100% effective. Phishers are constantly changing and adapting techniques in response to these technical countermeasures. Educating users will continue to be another essential component in the battle with phishers.

### User Education About Phishing

When users understand the threat of phishing and are made aware of the techniques used by phishers, they have a better chance of avoiding the scams. The Anti-Phishing Working Group has compiled several suggestions for avoiding a phishing scam:

- Use caution with emails or other messages asking for financial account information, especially if the message is not personalized.
- Do not trust messages from financial institutions that are not digitally signed, they may be spoofed messages.
- Avoid filling out online forms that prompt for personal information.
- Check URLs that should be secured; the URL of such sites begin with https:, not http:
- Patch your browser

 In addition to these measures, consider the recommendation from the United States Computer Emergency Readiness Team (US-CERT) issued in 2004 that recommended Microsoft Internet Explorer (IE) users switch browsers because of security flaws in the domain/zone security model, the DHTML object model, MIME-type detection, and ActiveX.

 For more information about how to avoid becoming the victim of a phishing attack, see the Anti-Phishing Working Group recommendations at [http://www.antiphishing.org/consumer\\_recs.html](http://www.antiphishing.org/consumer_recs.html).

### Q 4.4: How can an organization minimize spam?

**A:** Spam will not be eliminated in the foreseeable future, but there are several measures organizations can take to minimize spam:

- Educate users
- Filter content to prevent spam from entering the network
- Prevent inadvertent relaying of spam
- Prevent Trojan horse programs from distributing spam through your computers

These measures use a combination of human and technical countermeasures to decrease the likelihood of spam.

## Educate Users

Getting a spam message in front of a reader is a key goal of spammers; it is also a critical point for actions that can affect the level of spam this person receives. If spam does make it to a user's inbox, the user should flag the message as spam or junk mail if their email client has a built-in filter. In addition, the user should *not*

- Reply to the message
- Unsubscribe to the message
- Buy anything solicited by the email
- Click on a link embedded in the email


These actions will inform the spammers they have found an active email address. In the case of clicking on a link, making a purchase, or even just loading an HTML-based email, which sends back tracking information, these actions can improve the effectiveness of spam. The cost of spamming is so low (especially when the spammers use computers and bandwidth belonging to someone else) that even a small number of responses can make the whole spamming operation worth their efforts.

In addition to properly handling spam, users should provide their email addresses only to trusted parties. When signing up for online services, users should read the privacy agreement and understand for what purposes their email addresses will be utilized and whether their addresses will be sold to a third party. Do not post a personal email address online. Also, do not forward email from an unknown sender; doing so may lead to another user replying to the original spam, purchasing something from the spammer, or otherwise contributing to the success of the spam operation.

## Do Not Contribute to the Problem

Also, organizations should make sure they are not contributing to the problem. First, ensure that email servers do not provide for third-party relay. This service is provided by email servers that allow external users to send messages through the server without checking that the sender is a legitimate user. This feature of email servers allows spammers to use the resources of other organizations to send their junk email. To help minimize spam and reduce the threat of impacting the efficiency and productiveness of your own servers, configure email servers to prevent third-party relay.

The secure content appliances have to act as relay agents, so it is important to configure the appliances to relay only locally originating mail. Adding at least one entry to the local domains list on the appliance will enable anti-relaying functions and protect against spammers appropriating your mail services. This entry should be added when the appliance is installed. An open relay can easily be discovered on the Internet, often in a matter of just several hours.

 If you are not sure whether third-party relay is enabled on your email servers, test for it using the mail relay testing service from the Network Abuse Clearinghouse. The service is available at <http://www.abuse.net/relay.html>.

## Do Not Become a Zombie

Zombies are computers that have been compromised to the point where an attacker can control the functions of the computer. A number of well-known blended threats (malware that includes multiple pieces of malicious code, such as a virus, worm, Trojan horse, keylogger, and video frame grabber) include code to gain some control over the infected computer. Once in place, the malware opens a communication channel with a chat room or private server where it finds additional instructions, updated code, or new code to execute on the compromised machine. Networks of these zombie computers may be used by spammers to conduct their mass mailings.

To prevent this type of breach, use layered defenses, including desktop antivirus software, network and personal firewalls, and content filtering on the network. A secure content appliance provides effective countermeasures to several different types of threats and provides a first line of defense to keep servers, desktops, and other network devices from becoming unwitting participants in a spammer's efforts.

### Q 4.5: How can an organization implement better access controls to Internet content?

**A:** Controlling access to Internet content is a challenge. Browsers are ubiquitous, there are many sites accessible to users, and the number and nature of sites are changing constantly. Rather than defining access controls on all objects, as is commonly done with operating systems (OSs) and applications, systems administrators are better able to manage Internet content by applying filters as the content is accessed. This method changes the typical access control model of "User A is allowed to perform operations 1, 2, and 3 on file X" to "Users are not allowed to download content with attributes A, B, and C."

Filtering is generally performed using two techniques:

- Content filtering rules prevent content from entering the local network by detecting patterns indicative of inappropriate or malicious content.
- Black lists and white lists block or allow content, respectively, without further processing

The advantage of content filtering based on patterns is that administrators do not have to block all potentially problematic sites. Any content that has characteristic keywords or phrases can be blocked. At the same time, administrators may know of specific sites that should not be accessed from business systems. Black lists are used to block all content sites regardless of the patterns found in the content. Similarly, white lists are used to ensure that sites with legitimate use in business operations—such as business partners' Web sites, professional references, and general information sites—are always accessible.

Often, these two types of content controls are used in conjunction with one another. For example, a policy can be defined to prevent users from accessing gambling sites from business computers. This limitation can be accomplished with black lists that identify specific sites that are banned. Additionally, content filter rules can be defined to block Web pages with words such as “Blackjack,” “poker,” “roulette,” and “slots.”

As network scanning appliances are typically placed just inside the firewall, they are ideally positioned to scan content as a form of Internet access control:

- Content with banned words is blocked just inside the firewall; it never reaches the users
- Web pages with potential malware in the form of ActiveX components or JavaScripts can be scanned and blocked before they have a chance to execute on a client machine
- Spam and phishing messages are detected at the perimeter and are therefore prevented from putting undo load on email servers or client devices

An additional benefit of an appliance-based approach to Internet access control is that antivirus and anti-spam services are available as well. If a user manages to access an untrusted site, such as a peer-to-peer file-sharing network used for downloading shareware, and downloads an infected file, the antivirus scanner will be able to stop the virus before it can infect local devices.

## Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.