realtimepublishers.com™

The Definitive Guide™ To

Controlling Malware, Spyware, Phishing, and Spam

Dan Sullivan

Copyright Statement

Chapter 8: Implementation Issues in Securing Internet Content

Securing Internet content in an enterprise is now a basic element of the broader information security practices of organizations. The Internet is now woven into the fabric of business much like telephones and shipping services; it is difficult to imagine doing business without it. At the same time, with the benefits of the Internet come the downsides: viruses, spam, phishing messages, potentially unwanted programs (PUPs), and time wasted browsing and downloading offensive material.

Throughout, this guide has examined the responsibilities of organizations to protect the integrity of their information and infrastructure, specific threats to that mission, and technologies for combating those threats. This chapter continues the discussion started in Chapter 7 about technologies for securing Internet content with an examination of the implementation issues associated with applying those technologies.

The particulars topics addressed include:

- Criteria for choosing a secure content mechanism
- Benefits and drawbacks of implementation approaches
- Management issues in securing Internet content
- Best practices in securing Internet content

Let’s begin with a discussion of the core features that a secure content system should support.

Choosing a Secure Content Mechanism

Ideally, a secure content mechanism will posses several characteristics:

- Provide comprehensive coverage against common threats
- Reside in a secure and reliable platform
- Support open standards
- Offer easily customized black lists and white lists
- Align with organizational structures and line of business needs
- Provide adequate reporting
- Prevent simple bypass techniques

Comprehensive Coverage

Network and systems administrators are well acquainted with common threats that use Internet content transmissions as a way into an enterprise: viruses, worms, Trojan horses, keyloggers, video frame grabbers, spam, phishing messages, and PUPs (Figure 8.1).

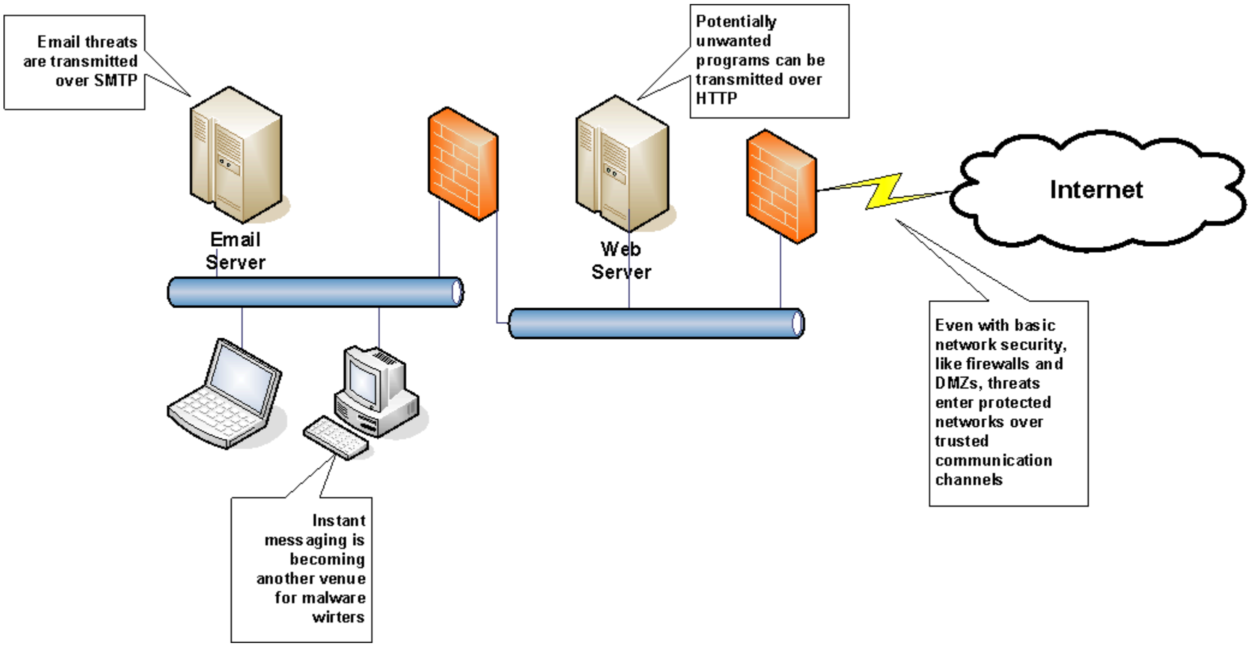

Email threats are transmitted over SMTP

Potentially unwanted programs can be transmitted over HTTP

Even with basic network security, like firewalls and DMZs, threats enter protected networks over trusted communication channels

Instant messaging is becoming another venue for malware wirters

Figure 8.1: Content-based threats can endanger a secure network over trusted communications channels.

The problem of dealing with this breadth of threats is compounded by the fact that the threats can use multiple vectors, or means of entering a private network, including:

- Email messages
- HTTP communications
- Instant messaging protocols
- File transfers

Threats in Email Messages

Viruses and worms are probably the threats most commonly associated with email. Well-known email viruses and worms—such as Melissa, SoBig, and MyDoom—are carried in email messages over the Simple Mail Transport Protocol (SMTP). In addition to malicious programs, emails can carry banned content, such as MP3 files, video clips, and other large, non-business–related information into an organization.

For details about viruses, worms, and other malicious programs, see the McAfee Virus Information Library at http://vil.nai.com.

Threats in HTTP Traffic

The complexity of communications over HTTP is growing as applications become more sophisticated. With the growing complexity of applications come new ways to transmit unwanted content:

- Spyware can be downloaded without a user's knowledge or explicit permission from Web sites when a user browses to that site (see Figure 8.2).
- Users may unintentionally install a Trojan horse program when installing what they assume to be a legitimate application.
- Even basic browser functionality, such as cookies, is now used in elaborate mechanisms for tracking users' behavior across multiple sites.

Because both SMTP and HTTP are fundamental protocols to core Internet applications, network administrators cannot simply block them at the firewall. Rather, the content that is carried on these protocols must be analyzed to prevent unintended uses. Other protocols, such as those used for instant messaging, are not as commonly required as HTTP and SMTP; however, in cases in which instant messaging is a required network service, secure content filters must work with those protocols as well.

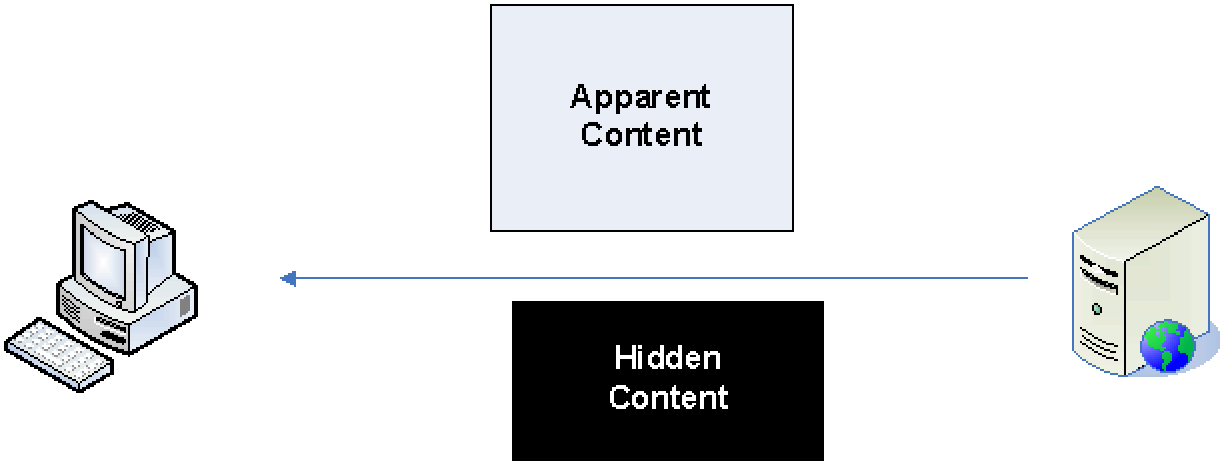

Figure 8.2: In addition to the apparent content that is downloaded when browsing the Web, other content, such as tracking cookies, may be downloaded without a user's knowledge or permission.

Threats in Instant Messages

As more mail traffic is scanned for viruses and similar malicious software, some malware writers are turning to instant messaging systems as a means for deploying their code. Some of the past threats to instant messaging include:

- The Kelvir worm struck the MSN Messenger tool in 2005. Kelvir displayed a Web site link in instant messaging chat sessions that when clicked would download a Trojan horse from the Web site.
- A similar worm, known as Lamo, worm struck the AOL instant messaging service in early 2006. Again, users were enticed to click a link that would have infected their machines.
- The Ocabot-F, similar to the others, enticed users to click a link, which then installed malware. In this case, once installed, the malware would send messages to all parties listed in the victim's instant messaging address book.

One option for administrators is to ban instant messaging from the enterprise network. This method would deny a valuable communication tool legitimately used in some organizations. In addition, some instant messaging clients allow users to use TCP port 80, used for HTTP traffic, bypassing administrator's attempts to prevent IM traffic.

Threats from File Transfers

Peer-to-peer, often referred to as P2P, file sharing has many advantages for groups collaborating or sharing resources, but the tools themselves may be problematic. For example, BitTorrent is a popular open source application for sharing programs as well as audio and video files. Because it is an open source program with freely available code and a liberal distribution license, BitTorrent has been repackaged—some distributors include spyware and possibly other PUPs in the repackaged version.

> 📖 For other file sharing programs that may include PUPs, see "Clean and Infected File Sharing Programs" at http://www.spywareinfo.com/articles/p2p/.

Besides the file sharing tools, the content that is transmitted may contain unwanted content. For example, sharing and storing MP3 files on company servers and desktops can consume large volumes of storage as well as network bandwidth.

A comprehensive, secure content filter should be able to detect malicious and other unwanted programs in emails, HTTP traffic, across peer-to-peer file sharing, and within instant messaging traffic. As the list indicates, there is no single "vulnerable" protocol that must be protected; any protocol and its associated applications are potentially vulnerable to misuse. The secure content mechanism, in addition to being comprehensive, must also be protected from tampering.

Secure and Reliable Platforms for Content Security

Secure content devices must be hardened to minimize the chances that the system is not compromised. Like other servers, a secure content device will run an operating system (OS) and multiple services. Each of these is a source of potential vulnerabilities that must be addressed. Common issues that must be considered to maintain the security of a secure content device include:

- Patching the OS
- Patching system applications
- Shutting down unnecessary services
- Protecting access control information, such as password files
- Reviewing audit logs

One of the advantages of using an appliance for secure content management is that the system is configured securely when shipped and maintained through automatic updates and patches. For those running a secure content application on a device managed internally, see the sidebar "Hardening Servers: Using Bastille to Lockdown Linux."

Hardening Servers: Using Bastille to Lockdown Linux

Locking down, or hardening, an OS requires in-depth knowledge of system services, vulnerabilities, and attack methods. Fortunately, freely available tools, such as the Bastille Hardening program, can assist systems administrators with this task. The tool works interactively with systems administrators by prompting with questions about the servers use and the administrator's needs. The tool then builds policies based on those needs. Bastille also supports an assessment mode that generates a report about available security configurations and changed settings.

Bastille analyzes and configures:

- Patches
- File permissions
- Account security
- Network and other daemons
- Sendmail
- Domain names
- Printing

The value of tools such as Bastille is twofold: first, it tightens the security of a server; second, it educates systems administrators about key OS security issues.

📖 For more information about the Bastille Hardening tool, see http://www.bastille-linux.org/index.html.

It is often said of IT that the only constant is change. With that in mind, a secure content system should support a wide variety of open communications standards to ensure that as users' needs and usage patterns change, content remains protected.

Support Open Standards

The Internet is built on a number of open standards. To understand the complexities of securing content on the Internet, it helps to reference a standard model for internetworking services known as the Open Standards Interconnection (OSI) model. The model consists of seven layers, as Figure 8.3 shows.

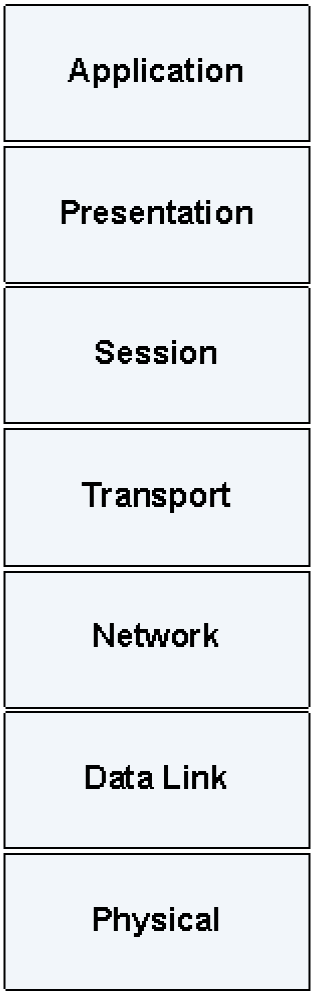

Figure 8.3: The OSI model highlights how higher-level networking services build on lower-level services.

The bottom layer, the physical layer, is responsible for converting bits into electrical signals and controlling certain physical aspects of the data. The data link converts messages into bits, controls how a computer accesses the network (for example, with collision detection for Ethernet), and converts frames between different types of networks. The networking layer provides logical internetworking services, routing, and addressing services. For the most part, secure content applications can ignore much of the complexity at these levels; information about operations at these levels is hidden by the protocols that run at the upper levels.

The Internet Protocol (IP) operates at the networking layer and underlies most of the protocols that are used by secure content applications. The transport layer manages transmission, message segmentation, integrity checking, flow control, and, if used in the protocol, sequencing. Both TCP and UDP function at this level. The session layer manages connections between applications. The presentation layer manages the translation of data into standard formats such as ASCII, EBCDIC, JPEG, and MPEG.

The top division, the application layer, provides the services that most users think of when they think of networking. Some of the common applications at this level include:

- SMTP
- File transfer protocol (FTP)
- HTTP
- Simple Network Management Protocol (SNMP)

Much of the focus of secure content management is on the application layer.

✎ Instant messaging also falls into the application category but there is not yet a single protocol that is commonly used across instant messaging services; a number have been proposed—including Session Initiation Protocol (SIP), SIP for Instant Messaging and Presence Leverage (SIMPLE), and Application Exchange (APEX)—but have not been widely adopted. Thus, to work across different protocols, users need an instant messaging client that supports multiple protocols, as is the case with iChat, Gaim, and Trillian, or they must use a server, such as Jabber, that uses software bridges to provide support for multiple instant messaging protocols. The Jabber protocol, known as the Extensible Messaging and Presence Protocol (XMPP) is now an Internet standard.

📖 For more information about Extensible Messaging and Presence Protocol (XMPP) as an Internet standard, see http://www.ietf.org/rfc/rfc3920.txt.

The services provided by the lower layers of the networking model are used in a variety of content transmission applications. Although email, hypertext, and file transfers have been popular tools for a long time (at least in terms of Internet history), and instant messaging has emerged as a valuable business tool, it is likely that other applications and new protocols, such as XMPP, will become commonly used communications tools. When that occurs, secure content applications that have worked well with other applications should readily adopt to protect the new means of communication. In addition to broadly advantageous characteristics such as comprehensive coverage, secure platform, and extensibility, secure content systems should have characteristics that support customizable functions that meet specific organizational needs.

Easily Customized Black Lists and White Lists

White lists and black lists are used to allow content from certain sources into an organization and to block content from other organizations.

Managing Black Lists

Secure content applications generally ship with defined black lists of known sites that provide online gambling, contain sexually explicit content, incite or promote illegal activities, contain illegal material, and promote hate speech or other patently offensive behavior. There may be occasions in which sites may need to be added or removed from a black list.

What is generally inappropriate for most business activity will have legitimate uses in some cases. For example, journalists and academic researchers may need access to sites professing hate speech, while the staff of a creative marketing firm might need access to lingerie catalogs online when developing a campaign for a fashion industry client. A “one-size-fits-all” approach rarely works in IT and secure content management is no exception.

Web sites are easily changed, so keeping up with all the potentially inappropriate sites is virtually impossible. Automated tools for collecting the URLs of sites containing offensive content and user-submitted sites to Internet repositories of black lists help, but they are not guaranteed to find all problematic sites. Network administrators should have the ability to easily add URLs to the black list as needed. In addition to keeping track of which sites should be blocked, administrators do not want to block legitimate communications.

Managing White Lists

White lists are used to define sources of content that are allowed to transmit information to the organization. Manually constructing white lists can be time consuming, but application-specific techniques can speed the process.

Lists of contacts maintained in email clients are obvious sources for white list content. If a person is listed in an employee’s contact list, it is reasonable to assume that the person has legitimate business with the organization and should have his or her email sent to the recipient without blocking.

Web server log files can be another source of information for building URL white lists. For example, by analyzing Web server logs, administrators may be able to determine that the majority of Web traffic at their site is generated by a few hundred sites. By reviewing the list of most popular sites, an administrator can construct a white list of known popular and verified sites.

Black lists and white lists are well-established tools in the area of secure content management. The ease of customization, along with the size of the database of black listed URLs, is an important consideration when choosing a secure content system.

> 📖 For more information about white lists and black lists, see Chapter 7.

Just as organizations will have differing needs with regard to black lists and white lists, they will have varying needs in terms of their organizational structure.

Align with Organizational Structures and Line of Business Needs

In this chapter and throughout this guide, the terms *system manager* and *network manager* have been used to refer to IT personnel responsible for network services and application management. There has been little, if any, distinction regarding the scope of these managers' responsibilities.

In practice, particularly in large organizations, there are many systems and network managers. Some are responsible for managing email servers, others maintain Web servers, and still others install, configure, and monitor database servers. Each of these has a role in securing content. As Figure 8.4 shows, administrators responsible for maintaining email systems may belong to different parts of the IT organization than those that control the Web servers. Securing email traffic (SMTP content) and Web traffic (HTTP content) should fall under the control of their respective departments or groups.

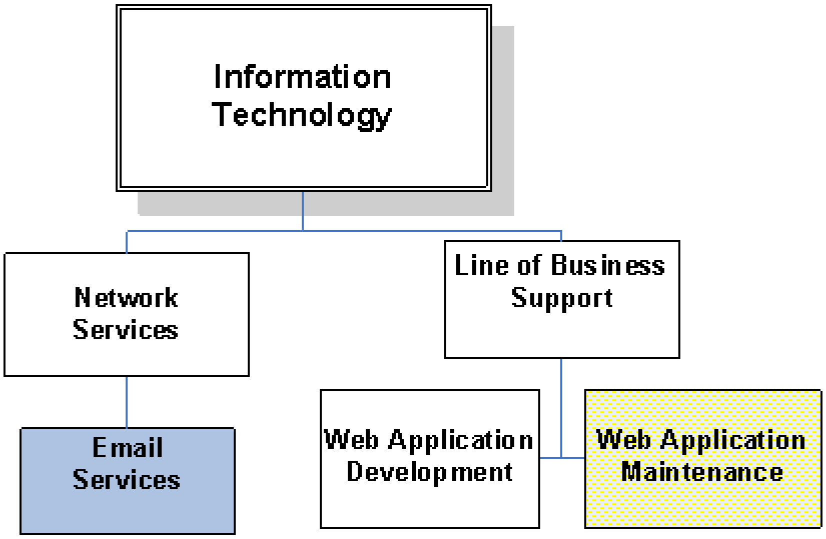

Figure 8.4: Secure content systems should be managed along the same organizational lines as IT responsibilities.

When choosing a secure content system, it is useful to partition the responsibilities of the secure content application along organizational lines. For example, one set of administrators may be responsible for mail servers, another is responsible for Web servers, and yet another group manages firewall and router configurations. The email managers may want full control over the secure content application that analyzes SMTP traffic, while the Web server managers expect total control over HTTP traffic; they both need to coordinate to some degree with firewall and router managers.

Provide Adequate Reporting

It is not enough for a secure content system to block targeted material; it must also be able to report to systems administrators about the status of the application as well as measure the state of the network traffic. At the least, secure content systems should report on:

- Application status
- System performance
- Significant event logs
- Configuration and update status

These are the core attributes systems administrators will need to monitor.

Application Status

The application status is essentially a snapshot of the state of content processing (see Figure 8.5). The overall status of the system consists of measures about each protocol that is analyzed by the application. Protocol measures examples include:

- Number of viruses detected
- Number of email messages deferred
- Number of email messages quarantined
- Reason for quarantine, such as virus detected or spam detected
- Volume of traffic processed in a given period of time

These measures are closely related to another set of key measures, system performance.

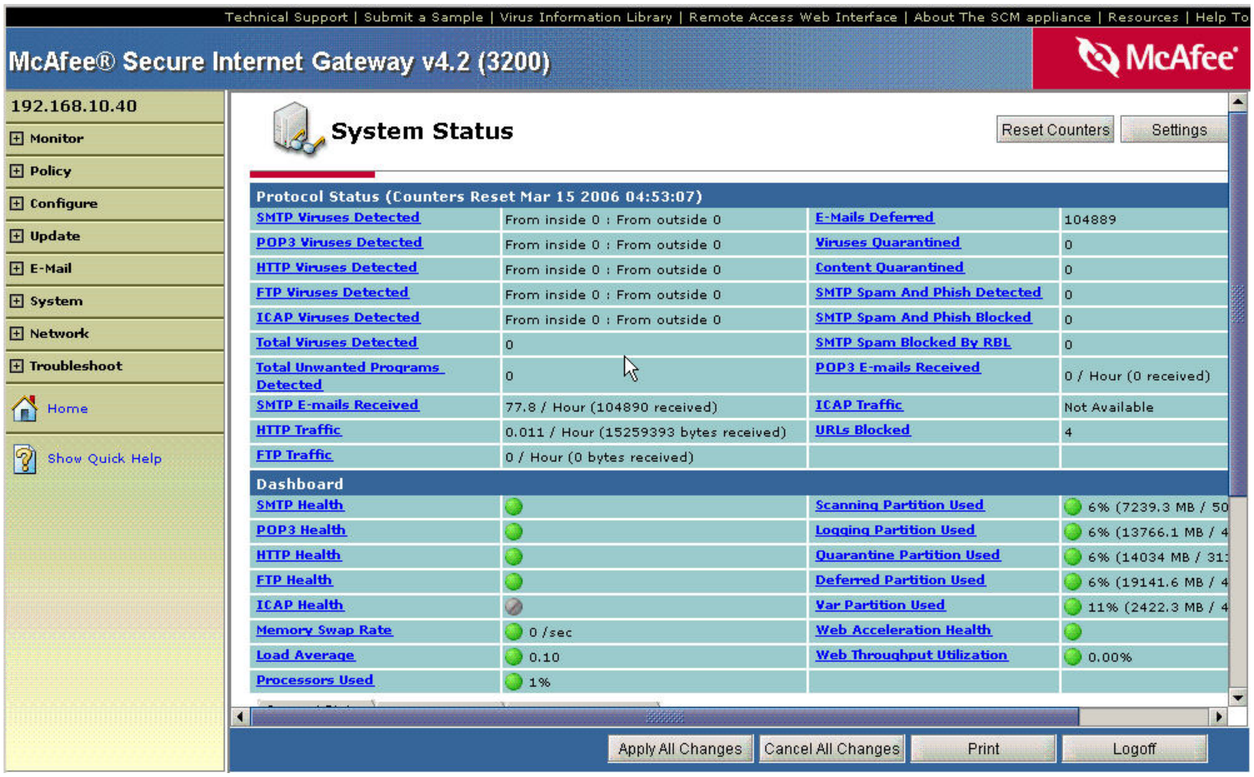

Protocol Status (Counters Reset Mar 15 2006 04:53:07)

SMTP Viruses Detected	From inside 0 : From outside 0	E-Mails Deferred	104889
POP3 Viruses Detected	From inside 0 : From outside 0	Viruses Quarantined	0
HTTP Viruses Detected	From inside 0 : From outside 0	Content Quarantined	0
FTP Viruses Detected	From inside 0 : From outside 0	SMTP Spam And Phish Detected	0
ICAP Viruses Detected	From inside 0 : From outside 0	SMTP Spam And Phish Blocked	0
Total Viruses Detected	0	SMTP Spam Blocked By RBL	0
Total Unwanted Programs Detected	0	POP3 E-mails Received	0 / Hour (0 received)
SMTP E-mails Received	77.8 / Hour (104890 received)	ICAP Traffic	Not Available
HTTP Traffic	0.011 / Hour (15259393 bytes received)	URLs Blocked	4
FTP Traffic	0 / Hour (0 bytes received)		

Dashboard

SMTP Health	●	Scanning Partition Used	● 6% (7239.3 MB / 50
POP3 Health	●	Logging Partition Used	● 6% (13766.1 MB / 4
HTTP Health	●	Quarantine Partition Used	● 6% (14034 MB / 31
FTP Health	●	Deferred Partition Used	● 6% (19141.6 MB / 4
ICAP Health	⊘	Var Partition Used	● 11% (2422.3 MB / 4
Memory Swap Rate	● 0 /sec	Web Acceleration Health	●
Load Average	● 0.10	Web Throughput Utilization	● 0.00%
Processors Used	● 1%		

Figure 8.5: A secure content application should provide aggregate reporting on system status to allow administrators to quickly identify problem areas.

System Performance

System performance monitoring should allow administrators to select measures of interest, such as number of email messages analyzed, number of spam messages blocked, and number of viruses detected. Ideally, administrators will be able to specify how often the measures are updated and how they are displayed (for example, in table or chart form).

Significant Event Logs

Secure content systems analyze multiple protocols and, in the course of those functions, create events that systems administrators should be aware of:

- Events triggered by management events, such as logging in to the appliance
- Software update events
- Failures to communicate between clients and servers
- Events created when email messages are deferred, quarantined, or deleted
- Events created when a virus, spam, spyware, or other PUP is detected
- URL blocking events

Reports on significant events should be available on an on-demand basis or reported automatically, for example, as an email message to an administrator. Systems administrators should have the ability to configure the reporting mechanism to best suit their needs.

Configuration and Update Status

Another area of reporting focuses on system configuration. At any time, systems administrators should be able to review key configuration parameters, such as the date and time of the last software update. Changing system parameters should also trigger events that are tracked in the system event log. Configuration information, such as port settings (see Figure 8.6), should also be readily available.

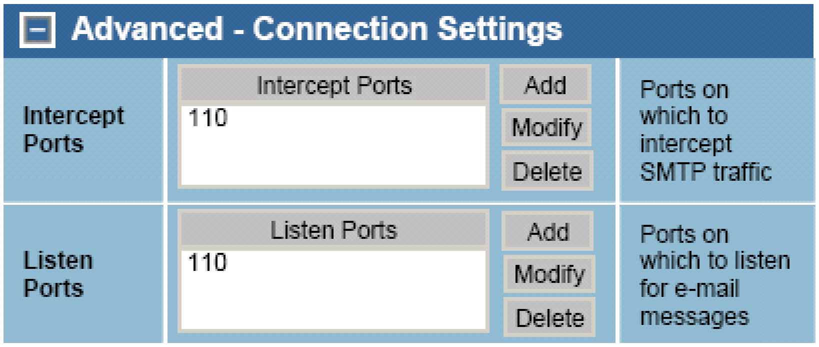

Advanced - Connection Settings

Intercept Ports	Intercept Ports 110	Add Modify Delete	Ports on which to intercept SMTP traffic
Listen Ports	Listen Ports 110	Add Modify Delete	Ports on which to listen for e-mail messages

Figure 8.6: Configuration settings, such as the ports with intercepted traffic, should be logically organized for easy access by systems administrators.

Effective reporting on content analysis, the results of that analysis, and the configuration of the secure content system is a fundamental requirement for systems administrators. Reporting on application status, system performance, significant events and configuration status covers a basic set of reporting requirements for system managers.

Features such as reporting, adaptability, and comprehensive coverage of threats are obvious aspects along which one might compare different secure content systems. Another characteristic that is not as often discussed is the way the systems may be bypassed and thwarted.

Compensating for Simple Bypass Techniques

> This section discusses techniques that can bypass some secure content applications. The objective is not to explain how to bypass but to point out to network and systems administrators the limits of secure content systems. In this section, the content tries to strike a balance between informing administrators of potential weaknesses and disclosing information that could be readily used to bypass such systems. For this purpose, high-level descriptions of techniques are provided but detailed procedures are not.

The history of information security is riddled with examples of countermeasures put in place to neutralize threats only to have those countermeasures rendered ineffective. For example, when signature-based antivirus programs were made available, virus writers responded by encrypting viruses. Antivirus developers then deployed tools to detect encrypted viruses, and virus writers again responded with even more sophisticated techniques, including the development of polymorphic viruses. The lesson is not that deploying countermeasures is futile but that even with countermeasures in place, you must be aware of ways determined individuals can circumvent those countermeasures.

Methods for bypassing secure content mechanisms include:

- Using IP addresses instead of domain names
- Using anonymous proxy sites
- Passing banned pages through a Web service provider
- Remotely controlling a device off the network with a remote control package

These are simple techniques that have equally simple countermeasures. For example, lists of banned sites should include both IP addresses as well as their domain names. Similarly, popular anonymous proxy sites can be banned along with other categories of unwanted content.

In the case of Web service providers, such as aggregators or search engines, it may be possible to access banned content. For example, to access the home page for Realtimepublishers.com, one could browse to:

http://www.realtimepublishers.com

but if that were blocked, one might be able to retrieve the cached version from Google using the following URL:

http://72.14.203.104/search?q=cache:UjoNwhAmuLcJ:www.realtimepublishers.com/+realtimepublishers&hl=en&gl=us&ct=clnk&cd=1&client=firefox-a

Secure content devices could undermine this technique by scanning the full length of a URL looking for well-formed embedded URLs passed as parameters to a site. This countermeasure would work with other Web service providers, such as translation services and redirectors.

With enough time, money, and expertise, most security systems can be circumvented. The goal is not to find a secure content tool that can never be bypassed but to find one that adequately addresses the threats present in an organization without costing more than the value of the resources protected. When choosing a secure content application, network and systems administrators should ensure that simple countermeasures cannot undermine the effectiveness of the system.

Key Features of a Secure Content System

Up to this point, this chapter has summarized the key features one should look for in a secure content system, including comprehensive coverage, a secure and reliable platform, support for open standards, easy customization, the ability to align technical deployment along organizational structures, adequate reporting, and countermeasures to bypass techniques. These are essentially the "whats" of secure content systems; the next section will examine the "how" or implementation options of such applications.

Implementation Options of Secure Content Systems

The market is currently providing three types of secure content systems:

- Appliances
- Software applications
- Services

Each has advantages and drawbacks.

Implementation Option 1: Secure Content Server Appliance

A server appliance is a bundled solution that includes hardware and software designed for ease of installation and maintenance. A key advantage of a server appliance is that standardized hardware and software is configured in an optimal way prior to shipping. System and network administrators do not have to learn and implement lengthy installation and configuration procedures.

> This discussion focuses on server appliances, which should not be confused with network appliances, also known as thin clients. Network appliances are low-cost personal computers (PCs) with minimal hardware (for example, no disk drive or CD drive). Software is downloaded from the network as needed.

Standardized Hardware

In the case of a secure content server appliance, the system would be configured with enough processing power, memory, and secondary storage to meet specific performance ranges. For example, an entry-level secure content appliance might use:

- A mid-range Celeron processor
- 512MB RAM
- A 60GB to 100GB disk drive
- Dual high-speed network interfaces

This configuration could not process high volumes of traffic and would not provide reliability features needed in a high volume, mission-critical environment. Higher-performance appliances would use server-scale processors and improve reliability using components such as:

- One or more Intel Xeon-class processor
- 1GB to 4GB RAM
- Hot-swappable SCSI drives in a RAID configuration
- Dual power supplies
- Dual high-speed network interfaces

Appliance servers typically also have standard video interfaces, USB ports, CD or DVD drives, and other equipment that does not directly impact the content processing performance of the device but are used for monitoring, management, and maintenance. Standard rack-mountable chassis make appliances easy to add to existing hardware environments. In addition to standardizing hardware, appliance servers also standardize software.

Standardized Software

Standardizing software can significantly reduce the time and resource required to manage systems. Of course, there are different levels of standardization. At one level, an IT department might select a family of software as a standard. One company might standardize on Microsoft software but run Windows 2000 (Win2K), Windows XP, and Windows Server 2003 (WS2K3). Another organization might standardize on Linux for servers but run Red Hat, SUSE and Debian. Although similar, the OSs with the Windows and Linux categories can be different enough to require slightly different configurations, different patches, and may suffer from different security vulnerabilities. For a systems manager, these differences add up to managing multiple types of implementations in spite of similarities.

At another level of software standardization, an organization could select one OS and one set of software applications. For example, a Microsoft shop might standardize on WS2K3, SQL Server 2005, and Internet Information Services (IIS) 6.0. Even in this scenario, different hardware and varying application requirements can lead to different configurations. One server may require an FTP server to run while another does not; one server might run applications needing one version of .NET components while another server needs a different set. Again, even though the organization has standardized on software, no two servers are guaranteed to be configured identically.

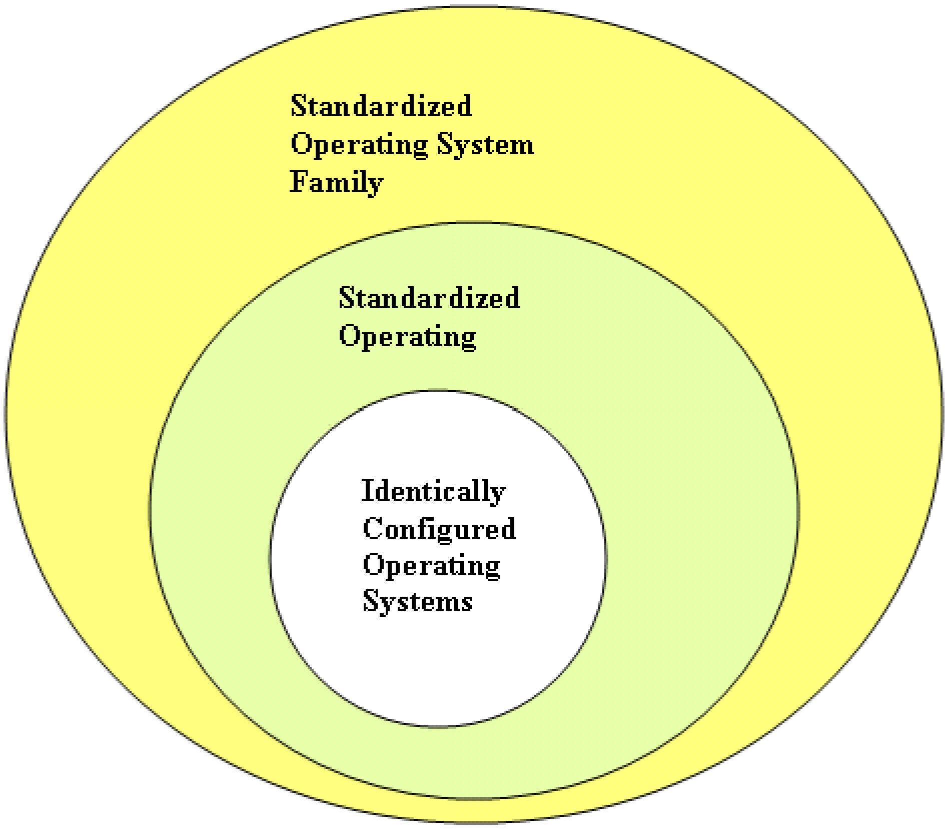

Figure 8.8: Standardizing software can mean several things; the more similar the instances of software installations are, the greater the benefits of standardization.

Another level of standardization is when the software installed on different machines is identical. This setup is virtually impossible unless the hardware and OS are identical. For example, a secure content appliance is a mass produced product; the vendor will use the same processor, video controller, disk drive, and other hardware components. These appliances can then use the same drivers for these components, run the same OS and applications—such as anti-spam filtering software—and be maintained at the same patch level. Standardizing on hardware and software results in significant benefits to appliance users.

Advantages of Appliances

The benefits of an appliance help both vendors and appliance users. Vendors do not have to troubleshoot problems in unusual or rarely used configurations, there is no need to support multiple OSs, and the core application, in this case secure content applications, can be thoroughly tested on a single target platform.

Appliance users benefit from that fact that the appliance is preconfigured at the factory, so minimal configuration is required. The hardware installation is minimal; typically, it is a matter of installing a rack- mounted device. Software maintenance and patching can be automated by the vendor, easing the workload of administrators. Finally, as network traffic increases or if network architecture changes and additional appliances are needed, identically configured devices can be added to the network without adding significantly to maintenance.

Implementation Option 2: Software Applications for Secure Content Management

Securing content depends on the use of multiple applications, including antivirus, anti-spam, anti-phishing, and URL blocking software and other programs. With software at the heart of securing content, one option is to deploy secure content applications on existing hardware.

One advantage of this approach is that it combines multiple functions on a single hardware platform. Systems administrators have fewer machines to manage and there is potentially less capital cost.

There are two primary disadvantages to a software-on-shared-server approach. The first is that because the application is not running on a dedicated server, performance may vary. For example, if a content filtering application were run on the same hardware as a Web server, when the demand for Web access is high, both the Web server and the content filter would have large workloads and would compete for the same computing resources. If the content filtering application shared a server with an application with peak demand periods—for example, a data exchange server that replicated data to remote offices several times a day—performance of the content filter would vary according to the other application's schedule.

The second disadvantage is that there may be conflicts in software dependencies between applications. For example, one application may require a particular patch to a shared library to correct a bug, but that same patch is incompatible with the secure content application. In another case, the optimal system parameter for one application is not optimal for other applications. As a result, systems administrators have to try to balance the requirements of each application with the relative value of each application to the organization. Some questions systems managers will confront include:

- How do I prioritize processes from different applications? Should they all be equal or should one application have a higher priority over another?
- If multiple software library versions are required, how should they be organized?
- If two applications have different administrators and both require root privilege, is it acceptable for each administrator to have potential control over the other application?
- How can upgrades be scheduled for one application to minimize the down time of the other applications?
- If additional capacity is required, should one application move to another server?

Servers that support multiple applications certainly have their role in IT. They are ideal in development environments, in remote offices that do not warrant multiple servers, and—when virtualization is available—in some cases in which dedicated servers (or at least virtual servers) are needed.

Does Hardware Virtualization Provide the Best of Both Worlds?

Virtualization is a method for making a single hardware platform appear as if it were several different platforms. This has traditionally been accomplished using software such as VMware (http://www.vmware.com/), but hardware vendors are adding support for virtualization to processors. (See for example, an introduction to Intel's virtualization technology at http://www.intel.com/cd/ids/developer/asmo-na/eng/218153.htm?page=4.)

Virtualization technology solves problems related to software dependencies because applications with conflicting configuration requirements can be run in different virtual machines. Unfortunately, virtualization does not solve the problems related to sharing computing resources among multiple applications.

In addition to appliances and software-only solutions, some enterprises may consider service providers for content filtering.

Implementation Option 3: Outsourced Service

Rather than install an appliance or software on a non-dedicated server, organizations may consider using a service that provides secure content filtering. This option is especially appropriate when email services are hosted by a third party or if an organization generates small to moderate amounts of Internet content.

A basic email filtering service works by routing all inbound and outbound email through a third-party's content filter mechanism. This approach shares several characteristics with an appliance-based approach, especially minimal maintenance. With well-defined and enforced SLAs, outsourced providers can meet the performance demands of an organization. As one of the service provider's core competencies is content filtering, one would expect that they would keep URL filtering black lists up to date, provide for white lists, maintain secure servers, and deploy the latest antivirus systems.

A disadvantage of services is that content has to be sent to a third party site before it is sent on to its ultimate destination. If the content filtering appliance or server is down, content will not go through. Because the server is not managed internally, the IT department cannot respond directly to the problem. Business continuity planning will have to take into account this potential disruption and formulate options, such as allowing email through without filtering or switching to a backup third-party service.

Specialized services, such as email-only content filtering provided with email hosting services, will not protect other content sent via ftp, HTTP and other protocols. Finally, services may use a single secure content appliance or other server to filter traffic from multiple customer sites. Performance may vary unless strict SLAs are in place. Organizations with modest amounts of content and a tolerance for some degree of disruption and possible swings in performance may be best suited for a service-oriented approach to secure content filtering.

Choosing Among the Options

It should be no surprise to anyone in IT that there is no single secure content option that is best for every situation. The following lists summarize the advantages and disadvantages of each of the three approaches.

Secure Content Appliance Server

Advantages

- Known performance
- Stability
- Easily scaled
- Minimal software (fewer potential vulnerabilities)
- Consistent performance
- Automated maintenance

Disadvantages

- No control over software installed on appliance

Software Applications

Advantages

- Able to use existing, underutilized servers

Disadvantages

- Performance can vary
- Must balance conflicting dependencies with other software
- Must lock down and maintain security of server OS
- Internal staff responsible for upgrades, patches, and so on

Secure Content Service

Advantages

- No additional servers on site
- Management is outsourced

Disadvantages

- Performance may vary
- Service may not filter all protocols used to transfer content
- Additional complexity of business continuity planning

In general, secure content appliances provide a balance between the minimal maintenance provided by service providers and the flexibility of installing software on existing hardware. Secure content services are appropriate for small organizations; software on shared servers is best for organizations with the staff and resources to maintain both the content filtering system and the security of the OS. Secure content appliances will well serve other organizations. Regardless of which secure content method is used, there are a number of management issues that will have to be addressed.

What About Filtering at the Client?

We have not discussed client-based filtering because, as a primary solution, it is not appropriate for enterprise-scale networks. Client-based filtering is popular with home PC users who are concerned about malicious software, PUPs, and inappropriate content.

Desktop antivirus software, personal firewalls, and similar products are certainly appropriate for enterprise use; they just do not constitute the full solution to the secure content problem. These applications play an important role in the "defense-in-depth" strategy used in security conscious organizations. For example, laptops may be protected by network-based services when connected to the corporate network, but when accessing the Internet from employee's homes or at coffee shop, they must be protected with client-based applications.

For enterprise network managers, desktop solutions present two problems. First, they can be bypassed or turned off by determined users. Second, deploying and maintaining a large number of desktops is more challenging than updating a single appliance. For example, even if a systems administrator could push a patch or antivirus signature file to every client on the network, the administrator still needs to account for devices not connected at that time and deliver the new software when they connect to the network. Of course, by that time, the device could already be infected with the malware the patch was designed to counter.

Management Issues in Secure Content Management

When securing content, IT managers should consider:

- Appropriate use
- Auditing and reporting
- Quarantining and deleting content
- Black list and white list maintenance
- Capacity planning

Management decisions about these areas should be formalized in policies and procedures.

Appropriate Use

Appropriate use policies define what employees, contractors, business partners, and other users of IT resources are allowed to do in general terms. (Appropriate use does not delve into detailed access control issues, such as which users have write access to the Human Resources database). These policies should describe what types of Web browsing are allowed, what are legitimate uses of email, the boundaries of personal use of IT resources, and the consequences of not adhering to the policies.

Auditing and Reporting

Systems administrators should be aware of significant events, such as a spike in virus-ridden emails or a large number of attempts to access a blocked site. They should also establish regular procedures for reviewing secure content filtering logs to stay abreast of patterns in usage and detect emerging problems. In addition, industry-specific regulations may require further auditing and reporting procedures.

Quarantining and Deleting Content

When questionable content is detected in the email stream, Web traffic, or other protocol, the content should be isolated by either sending the content to a quarantine area, deleting the content, or blocking further access to a problematic Web site. System managers should establish policies for how long to keep quarantined content and how to notify users if their messages or other content has been blocked, quarantined, or deleted.

Black List and White List Maintenance

Black lists are often updated automatically by secure content vendors, but organization may add their own set of sites to those lists that are publicly available. They may also want to allow access to sites that are typically blocked on black lists. In the case of white lists, it is important for organizations to track which sources are essentially allowed a free pass to send content to their networks and update that list regularly.

Capacity Planning

As the number of users grows and new Internet-based services are provided and adopted, the volume of content passing through an organization's network will grow. While planning for additional servers, routers, firewalls, and other network equipment, network administrators should plan for additional appliances, servers, or services to meet the demand.

> 🕮 A comprehensive security plan is based on well-defined policies. For help developing these policies, see the SANS Institute's Security Policy Project at http://www.sans.org/resources/policies/.

Best Practices in Secure Content Management

The purpose of this guide has been to describe the problems facing organizations when dealing with Internet content, and outline solutions to those problems. This chapter concludes with a brief overview of some best practices to incorporate with your secure content strategy.

First, deploy comprehensive content filtering. Even when client-side applications are used to prevent infections by viruses, worms, and other malware, these tools complement, they do not replace, network-based content filtering. Comprehensive filtering includes blocking malicious software, blocking spam and phishing attempts, and preventing access to inappropriate Web content.

Second, organizations should develop and enforce secure content policies. Users should be aware of what is considered appropriate use and the consequences of not following those policies. Policies should be reviewed periodically and adjusted to changing requirements.

Next, secure content administrators should align filtering practices with specific organizational goals, including:

- Protection of information assets
- Efficient operations
- Maintaining a non-threatening work environment
- Protecting against loss of intellectual property, proprietary data, and private information
- Maintaining regulatory compliance

The combination of comprehensive coverage and well-defined and enforced policies aligned with broader organizational goals and responsibilities underlie successful secure content practices.