realtimepublishers.com™

Tips and Tricks Guide™ To

Secure Content Appliances

Dan Sullivan

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

Copyright Statement

[**Editor's Note:** This eBook was downloaded from Content Central. To download other eBooks on this topic, please visit http://www.realtimepublishers.com/contentcentral/.]

Topic 1: Business Justification for Secure Content Appliances

Q1.4: 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).

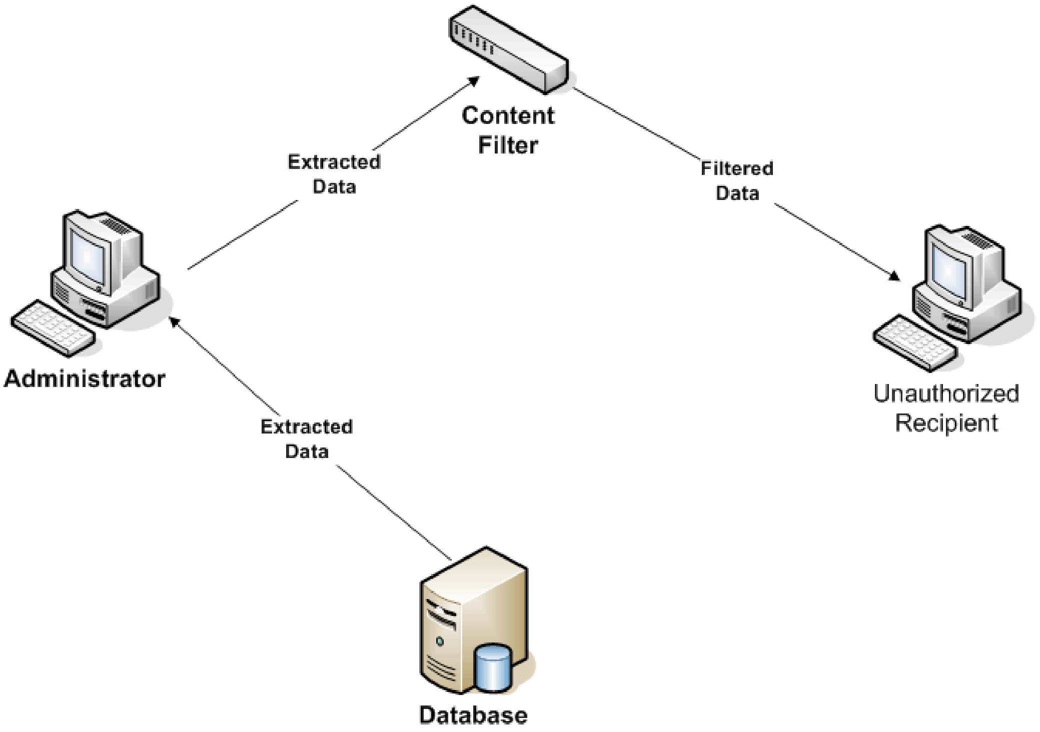

Figure 1.4: Extracted data analyzed by a content filtering mechanism can prevent the transmission of protected data.

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).

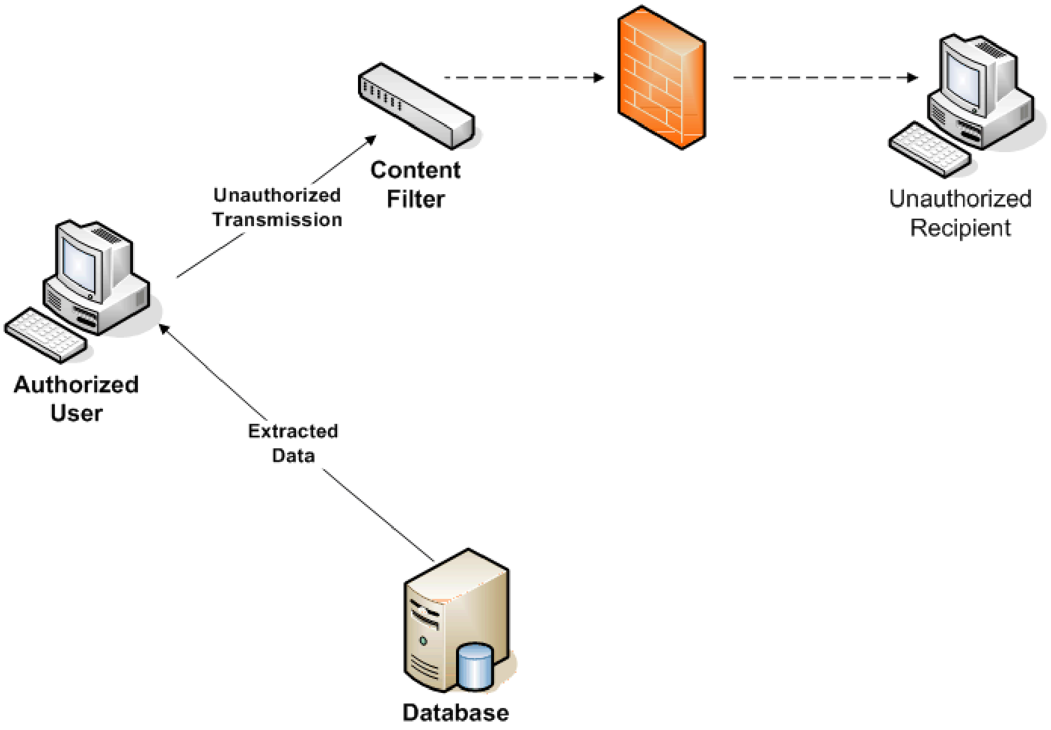

Figure 1.5: Depending on the content and the location of the recipient, a content filter can prevent protected data from being transmitted outside an intranet.

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.

Topic 2: Policies and Procedures for Secure Content Management

Q2.3: 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 a 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).

Topic 3: System Architecture and Secure Content Management

Q3.3: 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.

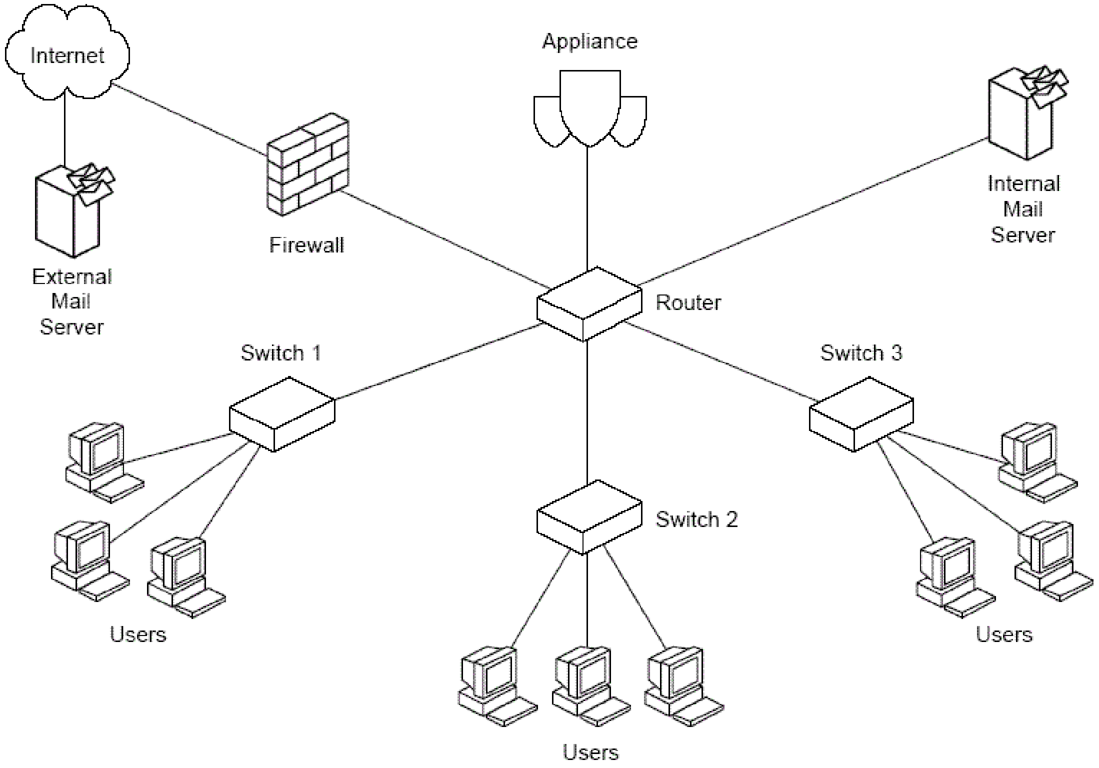

Figure 3.6: In explicit proxy mode, the appliance can be placed anywhere in the network because traffic is routed as needed to the appliance.

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).

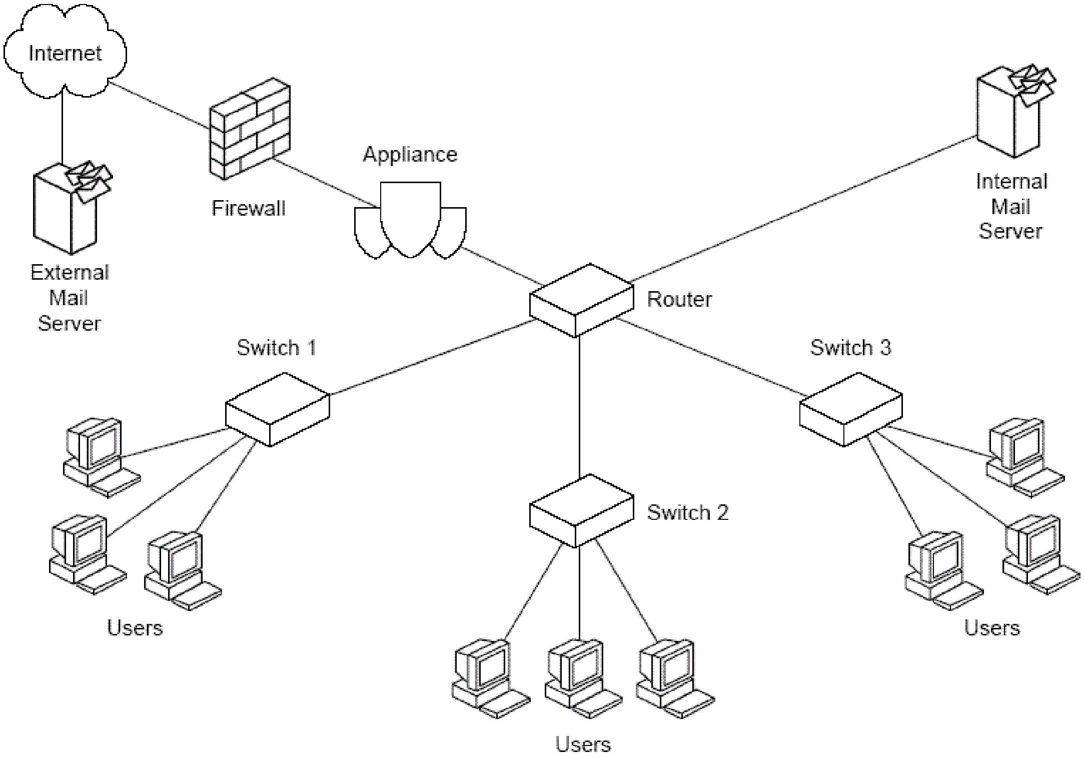

Figure 3.7: In both transparent bridge mode and transparent router mode, the secure content appliance should be placed just inside the firewall.

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.

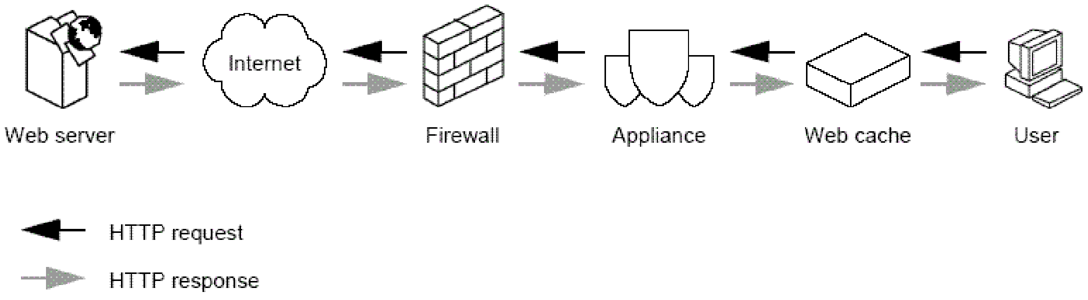

Figure 3.8: The secure content appliance is placed between the intranet firewall and the Web cache to ensure content that reaches the cache has been appropriately filtered.

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.

Q3.4: 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.

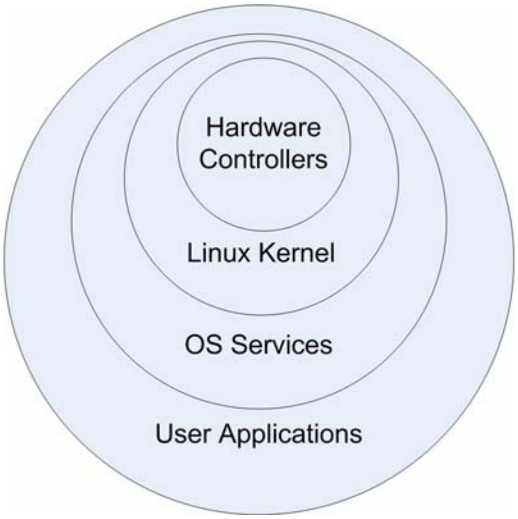

Figure 3.9: The Linux OS uses four major subsystems that provide a ringed architecture.

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/).

Topic 4: Secure Content Appliance Performance

Q4.3 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.

Content Central

Content Central is your complete source for IT learning. Whether you need the most current information for managing your Windows enterprise, implementing security measures on your network, learning about new development tools for Windows and Linux, or deploying new enterprise software solutions, Content Central offers the latest instruction on the topics that are most important to the IT professional. Browse our extensive collection of eBooks and video guides and start building your own personal IT library today!

Download Additional eBooks!

If you found this eBook to be informative, then please visit Content Central and download other eBooks on this topic. If you are not already a registered user of Content Central, please take a moment to register in order to gain free access to other great IT eBooks and video guides. Please visit: http://www.realtimepublishers.com/contentcentral/.