realtimepublishers.com™

Tips and Tricks Guide™ To

Secure Messaging

Jim McBee

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 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 might sound somewhat impossible 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 and the release of a new edition possible in a far shorter timeframe than is the case with conventional printed books. Because Realtimepublishers publishes 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, 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 six email security topics. 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: Strategies for Defending Email Infrastructure
- Topic 2: Policies and Procedures
- Topic 3: Architecture and Deployment Considerations
- Topic 4: Antivirus and Anti-Spam Strategies and Best Practices
- Topic 5: Firewall Strategies and Best Practices
- Topic 6: Protecting and Controlling Sensitive Information in Email

Copyright Statement

Topic 1: Strategies for Defending Email Infrastructure

Q 1.1: Why is email security important?

A: Email use has replaced the telephone as the predominant business tool for communicating with employees, customers, suppliers, and management. In 2004, IDC estimated that each day there are more than 30 billion email messages crossing the Internet; this estimate does not include email communication within an organization, where the messages never cross the Internet in the first place. The Radicati Group estimated that in 2004 there were nearly 900 million active mailboxes on the Internet with approximately half of these being used for business purposes.

Clearly, email systems and email communication has become extremely pervasive in business environments as a result of the ease with which information can quickly and efficiently be disseminated to one or many people anywhere in the world. The Meta Group estimates that approximately 75 percent of all corporate knowledge is now communicated via email.

✎ 74 percent of business people surveyed recently believed that losing email service presents more of a hardship than losing telephone service (Source: META Group survey).

As email systems have evolved and become more pervasive, the types of information that are being moved between users have changed. In early LAN-based mail systems, the typical email contained simple information and fairly innocuous data such as “What are you doing for lunch?” or “Let’s meet to discuss the Acme proposal.” Today, emails frequently and routinely contain business-critical, personal, and confidential information. As employees and managers rely on this information being available, email administrators have increasingly focused on maintaining higher levels of availability and performance for their users.

In the past, email administrators gave little thought to the possibility of dangerous email content arriving on their users’ desktops via email. That all changed on March 26, 1999 when a macro virus called Melissa spread almost overnight to tens of thousands of computers by exploiting the Exchange Global Address List (GAL). Some mail systems were crippled because the message transfer agent (MTA) queues had hundreds of thousands of messages to be delivered to both local and remote mailboxes. In 2000, the second act came in the form of the Love Bug virus (know as such because the subject line contained the words I Love You). Computer Economics estimated that the cost to fight and eradicate the Love Bug virus came close to $9 billion dollars.

Unfortunately, the standard method of email communication over the Internet uses the Simple Mail Transfer Protocol (SMTP); SMTP was designed to transmit 7-bit ASCII character data between two IP hosts using the simplest and most efficient method possible. Security was an afterthought with SMTP, and security systems are almost never implemented by default in an SMTP-based mail system. Further, the threats against an email system and its users have emerged just as quickly as the growth of the number of mailboxes.

Consequently, typical email administrators are spending more and more of their time not only focusing on the availability of their mail systems but also ensuring that both the system and the information are secure. The now invasive nature of spam and phishing attacks requires new mechanisms for ensuring that only valid messages reach the intended recipients.

Evolving Threat Landscape

Email administrators and IT managers now look back to the days of email-based threats such as the I Love You virus or Melissa and feel a little bit nostalgic. Although these threats are just as potentially damaging as today's threats, viruses such as the Love Bug virus were much simpler to defend against than the new threats that email administrators and users face today. The line between a virus, worm, and Trojan horse has been so blurred that they are now lumped into a single category called malware. Spam and phishing scams not only threaten productivity but also have the potential to perform identity theft. Further, as a result of email systems being so easy to use, accidental or malicious disclosure of sensitive or damaging information has become even easier to accomplish.

As if to add insult to injury, spammers are now teaming with virus writers to use viruses and worms to spread spam by building distributed networks of spam bots. Some security experts believe that floods of spam or viruses may be used to do political or economic harm in the future. Although this scenario might seem unlikely, imagine the costs involved in an organization of 1000 employees, each of whom have to spend 30 minutes of their day sorting through the junk in their mailbox in order to find the legitimate messages. For argument's sake, let's say that each of those employees costs the company $20.00 per hour. So, for each employee, it costs the organization $10.00 per day for that person to sort through their spam, or a total of $10,000 per day. And that is 30 minutes of that person's day that they should have spent doing their job.

Malware Comes to Town

The collective term for hostile or malicious email content is now simply malware; *malware* is any program that is potentially harmful to the user, the user's computer, the email system, or the network. The malware category includes viruses—a *virus* is a small program that attaches itself to another program running on a computer. Anytime that "infected" program is executed, it attempts to infect other programs.

Worms

This malware category contrasts with its cousin the worm; the worm goes out looking for computers (usually with a specific vulnerability) on which to replicate itself. Once the worm has been executed, the original infected program does not have to continue to run because the program is loaded in memory and begins searching the network for other computers to infect. A worm may also open backdoors on the compromised computer and lower security settings. The infamous SQL Slammer is an example of a worm that exploits a known weakness; the SQL Slammer worm managed to infect computers worldwide within 10 minutes of its initial release. A Yankee Group study found that in 2003, more than 80 percent of businesses that had Internet connections experienced disruption in business due to worms and viruses.

> ✎ The terms virus and worm are often (and sometimes incorrectly) used interchangeably and often worms are generically categorized as viruses.

Trojan Horse

An off-shoot of viruses and worms is the Trojan horse. A *Trojan horse* is a program that appears to be harmless or to do one thing, thereby tricking the user into opening the program. Once the program is running, it executes hostile code or attempts to collect information.

Blended Threat

Blended threats are threats that have characteristics of a virus, worm, and/or Trojan horse. These threats use a combination of malicious code and exploitation of vulnerabilities to propagate. Most of the email-based threats and other viruses that are common today are blended threats. Malware such as the Sober virus and its variants can access a user's personal address book, but uses its own SMTP engine to propagate. A client infected with Sober can connect internal email servers (bypassing firewalls and SMTP inspection systems) and deliver infected messages directly to internal mail servers as well as external mail servers. This virus can clog SMTP queues, delay delivery of legitimate email, consume Internet bandwidth, and generate complaints from your ISP.

Phishing

Phishing schemes have emerged as one of the biggest threats to users personally. Phishing comes from the idea that you toss out a line and see who will grab it. The messages usually look reasonably legitimate and may actually appear to come from an organization from which you do business. Figure 1.1 shows an example of a phishing scheme.

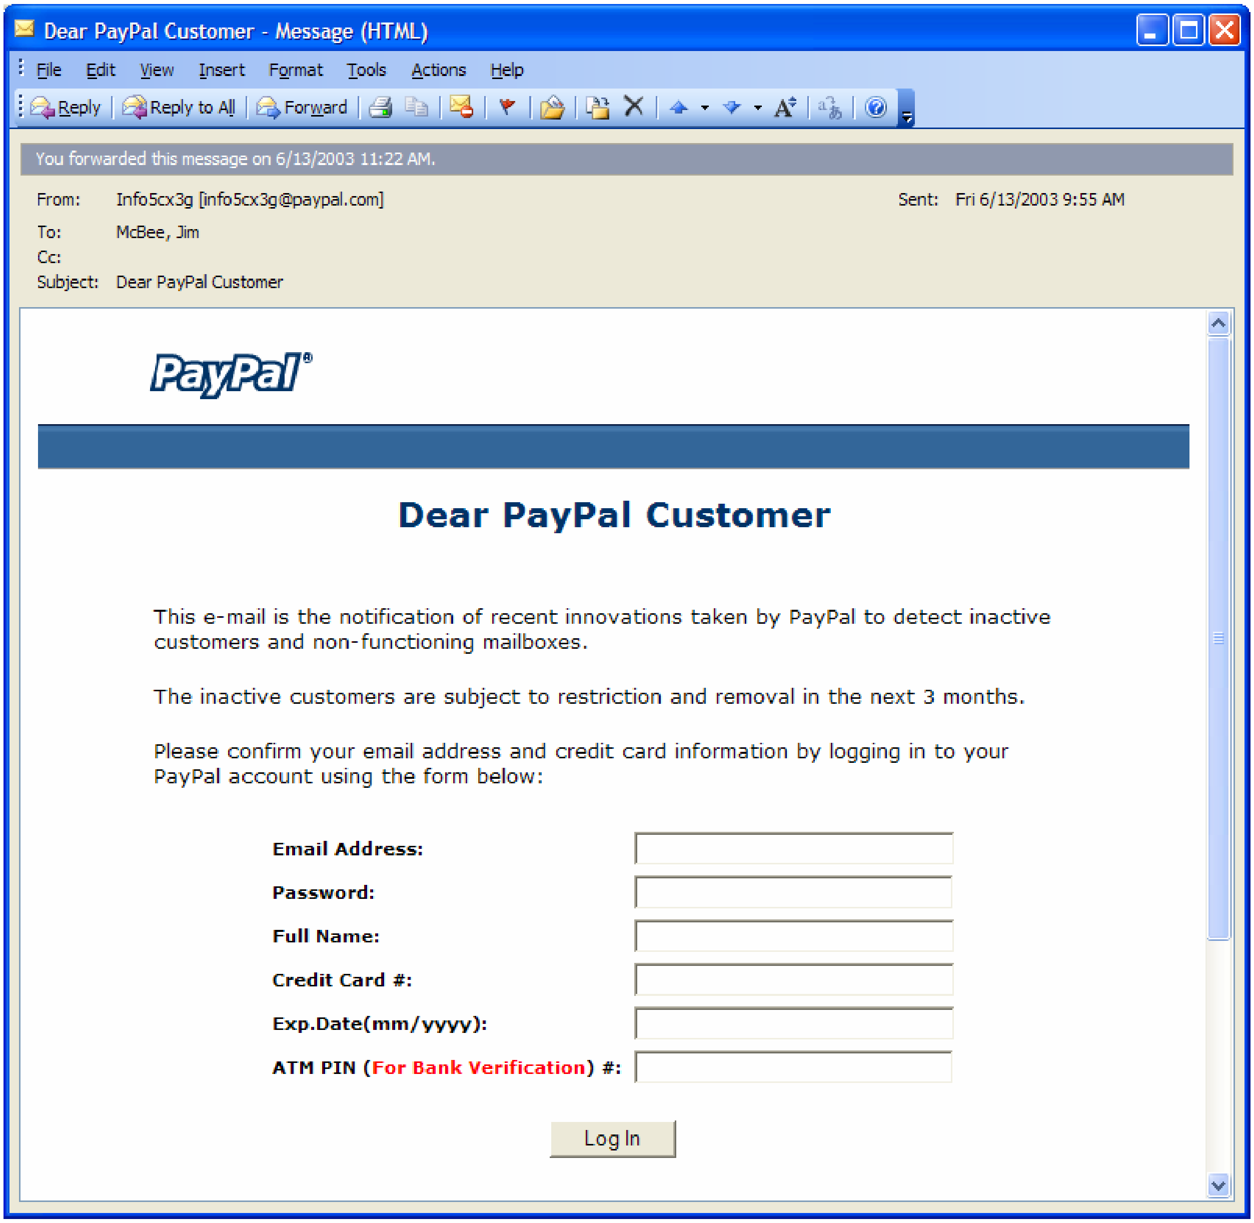
Dear PayPal Customer - Message (HTML)
File Edit View Insert Format Tools Actions Help
Reply
Reply to All
Forward
PayPal®

You forwarded this message on 6/13/2003 11:22 AM.

From: Info5cx3g [info5cx3g@paypal.com]
Sent: Fri 6/13/2003 9:55 AM
To: McBee, Jim
Cc:
Subject: Dear PayPal Customer

Dear PayPal Customer

This e-mail is the notification of recent innovations taken by PayPal to detect inactive customers and non-functioning mailboxes.

The inactive customers are subject to restriction and removal in the next 3 months.

Please confirm your email address and credit card information by logging in to your PayPal account using the form below:

Email Address:	
Password:	
Full Name:	
Credit Card #:	
Exp.Date(mm/yyyy):	
ATM PIN (For Bank Verification) #:	

Log In

Figure 1.1: An example of a phishing scheme.

Often, phishing schemes can be exposed because you are receiving a message from a bank with which you do not do business or there are a lot of typographical errors in the message. The message in Figure 1.1 was received right at a time in which the credit card I used on PayPal was about to expire, so it did make me think twice. However, not even your bank will ever ask you for your ATM PIN, so that request was the final clue that this message was fraudulent.

For the technically adventurous, you can view the HTML source of a message. When I did so for the message that Figure 1.1 shows, the HTML revealed that when the logon button was clicked, it would take you to the link shown in Figure 1.2. In the figure, I have highlighted the actual domain name to which the information is sent. Notice that the author of this message attempted to mask the intent of this message by putting www.paypal.com in the front of the URL, but followed that with many 1s and 0s.

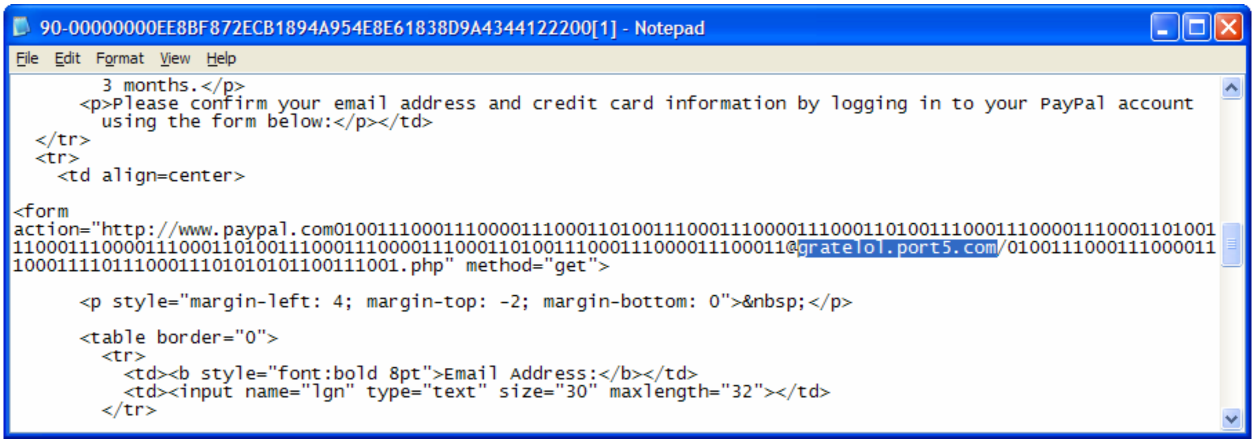
90-00000000EE8BF872ECB1894A954E8E61838D9A4344122200[1] - Notepad
File Edit Format View Help

```
        3 months.</p>
      <p>Please confirm your email address and credit card information by logging in to your PayPal account
        using the form below:</p></td>
  </tr>
  <tr>
    <td align=center>

<form
action="http://www.paypal.com01001110001110000111000110100111000111000011100011010011100011100001110001101001
1100011100001110001101001110001110000111000110100111000111000011100011@gratelol.port5.com/0100111000111000011
10001111011100011101010101100111001.php" method="get">

      <p style="margin-left: 4; margin-top: -2; margin-bottom: 0">&nbsp;</p>

      <table border="0">
        <tr>
          <td><b style="font:bold 8pt">Email Address:</b></td>
          <td><input name="lgn" type="text" size="30" maxlength="32"></td>
        </tr>
```

Figure 1.2: The "real" URL exposed.

Email administrators must be concerned that users will be defrauded or have their identities stolen by such a scam, but organizations need to be concerned as well. In the not-too-distant future, a user that has been scammed by a phishing scheme is going to have a creative lawyer that is going to hold their email provider accountable for receiving the email in the first place.

Spyware

Spyware is another form of malware that can be spread via email. *Spyware* is any program or technology that will gather information about a user's activity when they use their computers. In the early days of spyware, Internet marketing companies would use this software to determine a user's Web surfing habits and generate pop-ups that were targeted to that user's interests; this is also called *adware*. If those annoying pop-ups were not bad enough, users can just as easily be tricked into loading spyware that performs keystroke logging or that scans the user's hard disk for sensitive information and transmits that information somewhere else. Email masquerading as generic spam or a phishing scheme can trick a user into compromising his or her work computer just as easily as a home computer.

Spam

It used to be the case that receiving one or two spam messages per day was annoying. These days, a primary SMTP address can receive anywhere from 100 to 500 unwanted messages per day. *Spam*, or *unsolicited commercial email* (UCE) as it is officially known, is any message that ends up in your mailbox that you neither wanted nor requested. Laws have been passed and lawsuits brought against spammers have been won, but still the amount of spam received by the average email user continues to increase.

Estimates of just how bad the problem is vary widely between analysts. The IDC estimates that as much as 70 percent of all SMTP mail coming into an organization is unwanted or invalid messages. Without some type of filter or blocking mechanism, these messages end up in the user's Inbox. By some estimates, users spend as much as 30 minutes per day weeding through their Inbox to isolate the good mail and delete the spam. This percentage of most people's productive work day is essentially wasted. In addition, there is a cost factor associated with accepting and storing (even temporarily) this onslaught of spam.

Although spam is not necessarily going to crash your mail servers or put your organization out of business, it can still be detrimental. Some spam messages include graphic pictures and offensive words; a user of delicate constitution may well be offended to the point of hiring a lawyer or filing a harassment claim if an organization is not making a concerted effort to detect and eliminate spam.

Denial of Service

Day zero of a virus/worm outbreak is the day the virus is initially identified and protection against that virus is developed. In any given day, antivirus vendors may identify six or more new viruses. Although most of these don't quickly go "global," a select few end up spreading rapidly. Within a few days of being discovered, the Sober.Z variant of the Sober worm was clogging email queues, filling up hard disks, and slowing mail delivery for both commercial organizations and large email providers.

This type of Denial of Service (DoS) is just as effective as a direct attack by a malicious hacker who is intent on delaying mail delivery for an organization. Attackers may target a specific organization by sending significant numbers of large email messages in an attempt to delay messaging services for users.

A fairly new type of "attack" that is now being launched by spammers is called a directory harvesting attack (DHA). A DHA can be run against any SMTP mail server that validates inbound or outbound mail against a valid list of recipients. The Windows 2003 SMTP service, for example, does not perform this type of a validation against a directory service, but once Exchange 2003 Server is installed on Windows 2003, the SMTP service will perform validations of mailboxes against Active Directory (AD).

There are two possible approaches to a DHA. The first involves the sender going through all possible alphabetical characters looking for email that is delivered successfully. For example, a sending SMTP service might start with a@somorita.com, b@somorita.com, and so until it has gone through zzzzzzzz@somorita.com or some other maximum number characters. Each time a successful message is delivered, that address is recorded. A second and possibly more efficient mechanism is when the DHA attacking SMTP server goes through a dictionary of common surnames and given names looking for successes.

DHAs end up using not only a lot of bandwidth but also SMTP server and directory resources. And, in the end, some percentage of your user community's email addresses is compromised.

Information Leaks

Your email system contains your organization's sensitive and private information; unfortunately, this reality makes accidental or intentional disclosure of information easy. Passing along information via email to unauthorized parties (either internally or externally) is as simple as a few mouse clicks. Even accidental disclosure becomes as simple as a few incorrect keystrokes.

Given the volume of email that some users send on a daily basis, they frequently overlook to whom the message is addressed. This may be due to the user incorrectly choosing the Reply All function of the email client or accidentally including an incorrect recipient in the To, Cc, or Bcc lines.

Administrative Concerns

Email administrators often take a very technical approach to improving email security by implementing SMTP gateways, firewalls, appliances, content scanners, and third-party service providers. Throwing all your energies towards a technical solution is all well and good but does not help address some other issues that may plague your organization.

The first of these is administrative-level access to mailboxes. Most users don’t realize the level of trust that must be placed in their IT department. As many administrators in the IT department have full administrative access to the network, servers, and data, they have the ability to view email and often an end user’s mail. With great power comes great responsibility.

> It is not uncommon to read in an IT industry newsletter or magazine about a case in which a network administrator has been fired for accessing information that they should not be reading. This includes recreational “mailbox surfing;” mailbox surfing should always be treated as a serious offense.

The need to have mailbox-level access for users’ mailboxes is often necessary in most organizations. This may be because a hostile or sensitive message was sent to everyone and needs to be extracted, information may need to be archived from users’ mailboxes, or due to the investigation of a user that requires access to his or her email. Access of a user’s mailbox should never be treated as a trivial matter. Procedures should be documented and published for when a mailbox will be accessed by anyone other than the owner of the mailbox. Management oversight and authorization should always be obtained prior to accessing a user’s mailbox for any purpose. In some larger organizations, a representative from Human Resources or the Information Security department must be present.

Administrators in small and midsized organizations and organizations that have branch offices frequently overlook physical security of servers. It is not uncommon to find an email server sitting on someone’s desk or in a spare cubicle. If an intruder gets physical access to your servers, the information on those servers is easily compromised.

Another area that is often overlooked with respect to administrative procedures is storage of backup media. Although we all know that the backup media has a complete copy of data (at least, you hope it does), often the media is not treated as securely as the servers themselves. In some situations, backup tapes for large companies are stored on a shelf in the hallway outside the data center door. Once an unencrypted tape is in the wrong hands, it is trivial matter to restore the data to an alternative computer and access it.

> Backup media should always be stored in locked containers such as a tape backup storage container or a safe. If an intruder acquires your backup tapes, he or she can access the data at their leisure. Backup media should have at least the level of protection that your servers have been afforded, if not more.

Liabilities and Punitive Damages

For most technical staff thinking about vulnerabilities with their systems, lawsuits or government regulations are not the first topic to come to mind. You normally view systems in terms of capacity, availability, performance, and features—not in terms of lawsuits and government oversight.

Regulatory compliance is not a term that most in the technology business were used to hearing until recently. There are now a number of laws that directly or indirectly affect the operation of an email system. These laws were passed with the intent of requiring organizations to provide better protection and accountability of the information they process. The Health Insurance Portability and Accountability Act (HIPAA) is intended to protect protected health information (PHI), including the security of information that is transmitted via email. For example, a doctor can write an email to another doctor containing the word diabetes; diabetes can be mentioned many times and that email message does not have to conform to HIPAA security requirements. However, if that message contains any of, but not limited to, the following, HIPAA guidelines must be followed:

- Uniquely identifying information such as a patient's name, medical record number, or Social Security number
- Phone/fax numbers
- Email address or mailing address
- Photograph or fingerprints

Violation of the HIPAA requirements can bring about lawsuits from the federal government as well as from the patient. However, failure to meet regulatory compliance with your email system is not the only place that may cause your organization to be sued. Misuse of an email system by users, such as sending information that is inappropriate in most work environments may bring about lawsuits. Inappropriate content may include jokes, pictures, or copyright-protected content such as MP3 or MPEG files.

Aspects and Trends in Email Security

The email security process, software, and choices available are evolving rather quickly. The terms that have been known for years for antivirus software, anti-spam software, and content inspection software are now morphing in to a single category of software called message hygiene software. The number of vendors and service providers in this field are rapidly growing and a surprising number of partnerships between companies that formerly competed are now blossoming to form new products with new capabilities. A common theme among these vendors is that you can use a single software package for all your message hygiene needs and with that unified management of antivirus and anti-spam technologies.

Network and email administrators must consider that keeping their messaging systems secure and protecting their information is not a destination to which they arrive after the purchase of a few products and some configuration changes. Messaging security is an ongoing journey that requires adaptation of methods, software, policies, and procedures.

Topic 2: Policies and Procedures

Q 2.1: What should email training include?

A: One of the most valuable exercises in security that you can put your organization through is to develop a well-crafted, concise training program for use of your company's messaging system. Putting together any training program for users on an email system is difficult. Most of your user community will already know how to use their email client; some can use their email client as well or better than people in the IT department.

For this reason, you will have a difficult road to face when bringing your users back in for training that includes use of the email system. Even for a class that only lasts an hour or two, the users' time is valuable and getting buyoff from company management to conduct this training and have people out of their jobs is a difficult concept to sell. Thus, the content of the training class is very important.

> ✎ Some organizations have set up mandatory refresher classes that users must take once each year and they must sign an updated acceptable use policy statement. If they fail to take this class, their user accounts are disabled.

Refresh the Basics

In any email course, it never hurts to start out with the basics of using the email system, covering topics that all users are familiar with. Reiterating these topics periodically can help reduce security problems and Help desk calls. The basics include:

- Using distribution lists, especially distribution lists with large memberships
- Recognizing the characteristics of a phishing scheme, messages that may contain a virus, or messages that may carry a Trojan horse
- Using rules such as automatic forwarding of email and out of office replies
- Exploring the difference between Reply and Reply All
- Checking the To, Cc, and Bcc lines prior to sending a message and discussing how some email clients (such as Outlook 2003) will automatically fill in email addresses when you partially type an address that has been recently used
- Reminding users that the email system is the property of their employer and as such the IT department may have access to the contents of anyone's mailbox if a valid need arises

Imposed Limits

Any limits that are imposed by the servers, firewalls, or message hygiene systems should be covered so that the users are reminded of these limitations. Limitations may include:

- File attachment types that are not permitted
- Message system mailbox limits such as the maximum size to which a mailbox may grow
- Maximum inbound and outbound message sizes
- Maximum number of recipients that a single message can contain
- Archival or auto-deletion processes that run and what they do—such as purging messages from the Deleted Items folder or archiving to an archival system any message in the Inbox older than 90 days

Acceptable Use

The acceptable use policy is growing in popularity in many organizations because it helps to define what a user is allowed to do and more importantly what they are not allowed to do. Ideally, users sign the acceptable use policy to indicate that they have read it and understand the restrictions that are being imposed. During training, users should be reminded of items including:

- Whether the work email account can be used for personal use
- Where users should not use their work email address, such as Internet sites where they register for contests or questionable e-commerce sites—these are an invitation to be spammed
- Use of the email system for joke email or inappropriate content such as off-color pictures
- User of the email system for sending copyrighted material such as music files or video files (this not only causes the mail system to use a lot of unnecessary disk space, but it can expose an organization to lawsuits)

Information Security

Finally, users should be reminded of the nature of their job and the nature of their business. Any user that processes sensitive information—such as accounting information, sales information, and health care information—must understand the care they must take when sending this type of information via email. Accidental or intentional disclosure is not acceptable.

Organizations that must abide by a regulatory requirement such as the Health Insurance Portability and Accountability Act (HIPAA) must have a subsection of the training that outlines their responsibilities when processing this type of information and the ramifications not only to the organization but also to their job if such information is disclosed. Some users may only be concerned about their own jobs, but violation of some laws could expose a user to prosecution, jail time, or lawsuits. This is especially true when the organization has done due diligence to prevent such disclosures and the responsibility falls on the user.

Topic 3: Architecture and Deployment Considerations

Q 3.1: What is the best placement of Exchange servers, front-end servers, and domain controllers on a network?

A: Since the release of Exchange 2000 and the concept of front-end and back-end servers was introduced, there has been much discussion about the placement of these servers such that remote users can access their email services. After all, the front-end server is essentially nothing more than a Simple Mail Transfer Protocol (SMTP) and Hyper Text Transfer Protocol (HTTP) "pass through" server, and the obvious place to put this server is a screened subnet (aka the perimeter network or a demilitarized zone—DMZ). Even some organizations' written security policies state that all Web services accessed by external resources must be in the DMZ.

The intent of the DMZ network is to provide a place where resources (such as a Web server or a mail relay server) can be accessed by users outside of your network. Servers sitting in a DMZ should have no critical or sensitive data on them; thus, if they are compromised, the loss of data or service is minimized. However, this can give network security people a false sense of security.

The Old Way

When Exchange 2000 was first released, many Exchange gurus as well as Microsoft recommended placing the Exchange front-end servers in the DMZ and then place the Exchange back-end servers and the domain controllers on the internal network. This architecture is illustrated in Figure 3.1. In Figure 3.1, the front-end server sits in the DMZ and specifically allows Outlook Web Access (OWA) clients using Secure Socket Layer (SSL) to connect to the front-end server; the front-end server then passes the OWA clients' requests to the back-end servers on the internal network. If only it were that simple, though; a number of additional ports must be opened.

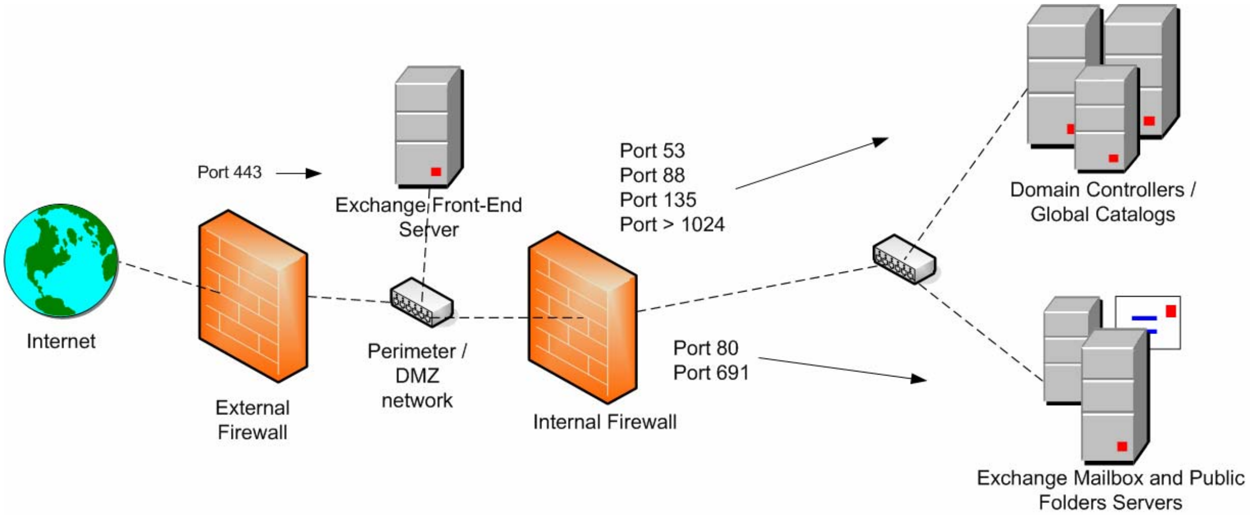

Figure 3.1: Traditional way of thinking about front-end and back-end server placement on the network.

Over the years, a number of weaknesses have been found in the architecture that Figure 3.1 shows. Front-end servers must be domain members and have the ability to communicate with all back-end Exchange servers as domain controllers and DNS servers. The following ports will need to be opened between the Exchange front-end servers and the domain controllers that the Exchange server will be configured to use:

- User Datagram Protocol (UDP) and Transmission Control Protocol (TCP) port 53 (DNS, if DNS is on the domain controller)
- UDP and TCP port 88 (Kerberos)
- TCP port 135 (Remote Procedure Calls—RPCs)
- TCP port 389 (Lightweight Directory Access Protocol—LDAP)
- TCP port 3268 (LDAP for Global Catalog—GC)
- TCP port above 1024 for RPC communication

The following ports may need to be opened between the Exchange front-end server and all back-end Exchange servers:

- TCP port 25 if front-end server is used for SMTP
- TCP port 80 if front-end server is used for OWA/ActiveSync
- TCP port 110 if front-end server is used for Post Office Protocol 3 (POP3)
- TCP port 143 if front-end server is used for Internet Messaging Access Protocol 4 (IMAP4)
- TCP port 691 for link state

Exchange can be configured to use specific domain controllers on the internal network by configuring the domain controllers to be used on the properties of each Exchange server's directory access property page (see Figure 3.2).

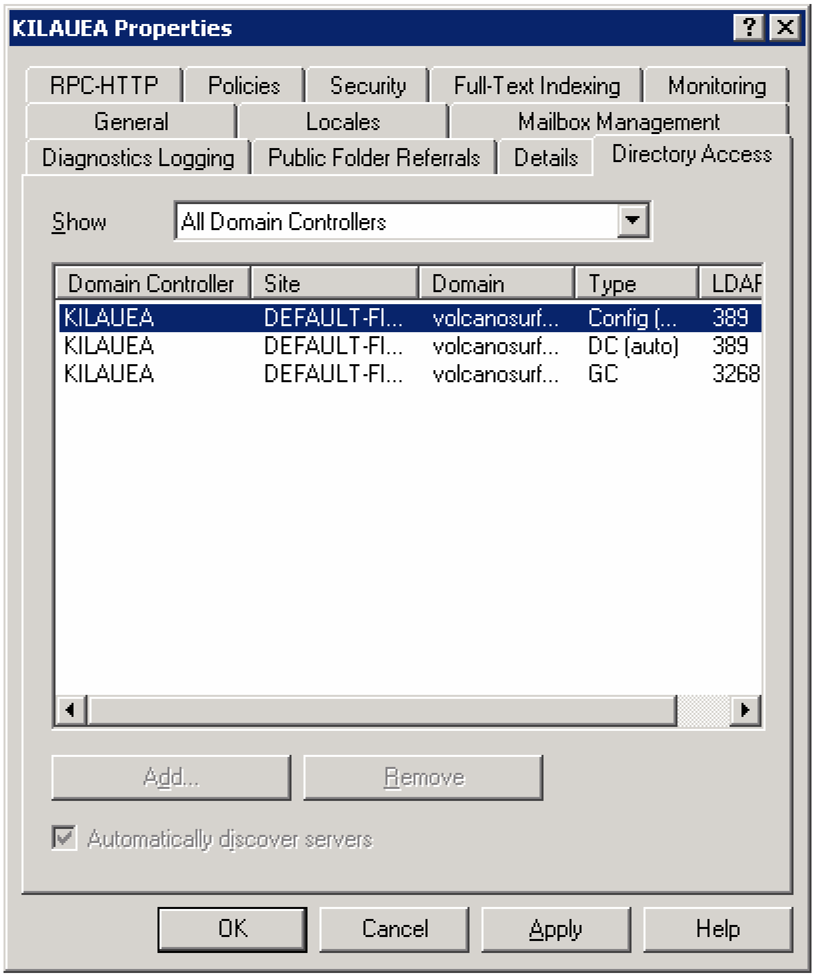

Show: All Domain Controllers

Domain Controller	Site	Domain	Type	LDAF
KILAUEA	DEFAULT-FI...	volcanosurf...	Config (...	389
KILAUEA	DEFAULT-FI...	volcanosurf...	DC (auto)	389
KILAUEA	DEFAULT-FI...	volcanosurf...	GC	3268

☑ Automatically discover servers

Figure 3.2: Configuring an Exchange server to use specific domain controllers.

Thus, TCP and UDP ports must be opened between the front-end servers and domain controllers for Kerberos, LDAP, and RPCs. The port that must be opened above 1024 to the domain controllers/GC servers is a dynamic port that may be different each time a domain controller is rebooted. You can avoid opening the ports above 1024 by statically mapping the RPC port for domain controller services to a specific port on the domain controller/GC servers. To do so, create a registry value called TCP/IP Port in the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\NTDS\Parameters registry key.

Set this registry value to some value up to 5000; I usually pick a standard value for all my domain controllers that is just below 5000. Microsoft recommends the port be below 5000 in Microsoft article 298369 “How to Configure a Global Catalog Server to Use a Specific Port When Servicing MAPI Clients.” You can find out more about static port mappings that are used with Exchange and domain controllers in Microsoft article 270836 “Exchange Server Static Port Mappings.”

As many network administrators have found, this approach is often difficult to configure. In addition, this approach does not necessarily provide any better security than if the front-end server resided directly on the internal network. The reason is that a large number of ports to critical services must be opened between the front-end servers in the DMZ and domain controllers and Exchange servers on the internal network. If a front-end server is compromised, the compromised front-end server has almost unfettered access to the domain controllers and Exchange servers.

Using Internet Protocol Security (IPSec) can reduce the number of ports that have to be opened between the front-end server in the DMZ and the servers in the internal network. The following UDP and IP protocols must be opened in the firewall to allow IPSec to function through the firewall:

- UDP port 500 for Internet Key Exchange (IKE)
- IP protocol ID 50 for Encapsulating Security Protocol (ESP)
- IP protocol ID 51 for Authentication Header (AH)

However, the IPSec policy must be tightly and precisely configured so that only necessary hosts and services are accessible by the hosts in the DMZ; otherwise, you will have a bigger security problem than you did with the additional ports opened. Thus, Microsoft and most Exchange system designers recommend against putting Exchange servers into the DMZ network.

Improving Security for Front-End Servers

The approach that is recommended by Microsoft for placement of front-end servers and improving Exchange security for servers that must be accessed via Internet clients calls for all Exchange servers to be placed on the internal network. Figure 3.3 shows an example of a network that has implemented this approach.

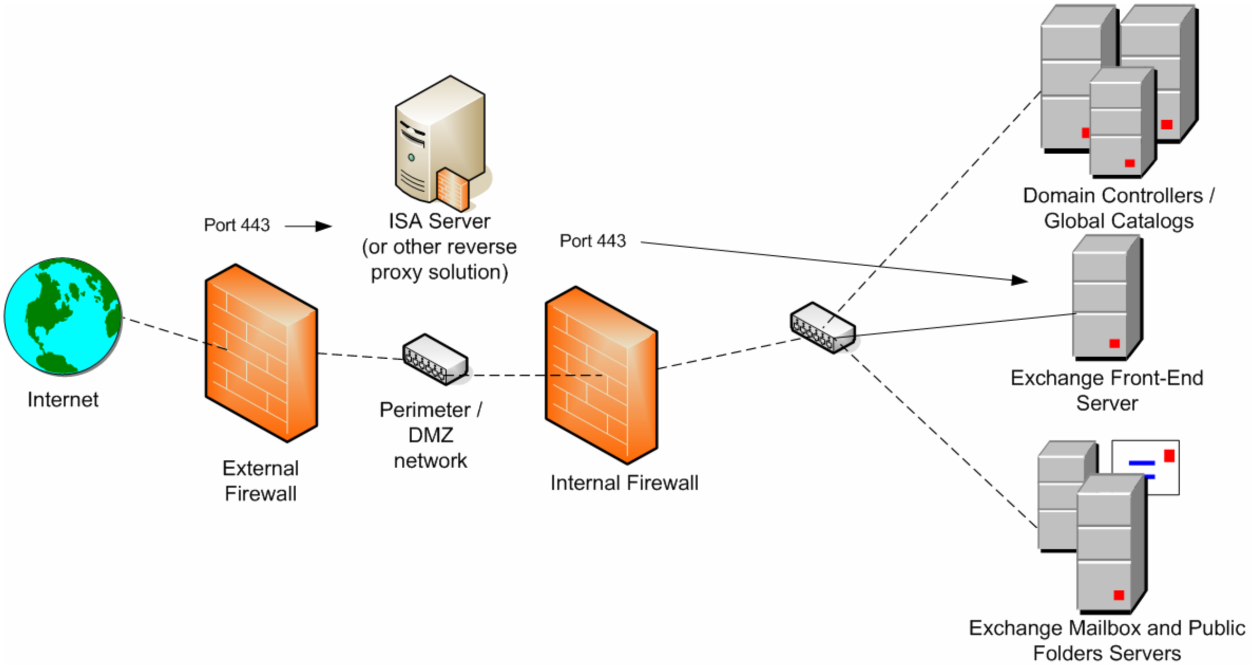

Figure 3.3: Moving the front-end servers to the internal network.

The approach illustrated in Figure 3.3 works well for any Web publishing solution, not just for Exchange servers. By using a solution that is capable of Web server publishing (or reverse proxy), the actual servers can be safely located on the internal network. The only host to which the client must have access is the server that handles the reverse proxy functions.

> 🕮 For more information on the basics and concepts of reverse proxies, see the SANS Institute's paper "A Reverse Proxy is a Proxy By Any Other Name" at http://www.sans.org/rr/whitepapers/webservers/302.php.

In the solution that Figure 3.3 shows, a Microsoft ISA server is placed in the DMZ. The only port necessary to be opened from the outside is port 443 (HTTPS). The ISA server accepts the inbound HTTPS requests, decrypts the request, examines the URL, determines the correct server to pass the HTTP data on to and the validity of the data, optionally re-encrypts the data, and passes it on to the internal front-end server. This solution can be put into place even if the solution you use as a reverse proxy server is not your primary firewall solution.

Topic 4: Antivirus and Anti-Spam Strategies and Best Practices

Q 4.1: How do I get started with a system for protecting against spam and viruses?

A: Most organizations already have some type of email security system in place. You may be relying on client-based antivirus software, server-based antivirus software, a Simple Mail Transfer Protocol (SMTP) gateway that handles virus inspection, or some combination of these.

Before moving on, understand that any single solution is probably not sufficient for protecting your servers and users from the crafty devils that are now writing viruses, attacking systems, blasting out spam, or developing phishing schemes. I recently worked with an organization in that needed help with a virus cleanup because the only antivirus solution was on the Exchange server. Although the company had a client-side antivirus solution, it was not distributed to all clients. A few of the clients on the network were infected with a variant of the Sober worm. The company was broadcasting viruses to the Internet almost continually.

At a minimum, a complete solution should address protection for the mail stores on your mail server and protect the client from viruses that they may be exposed to via email messages or a document or a Web page.

At this point, you may be at the “chicken or egg” intersection. Should you develop policies first or put the enforcement mechanism in place first. A definite way to make enemies of your user community is to place restrictions on what they are currently able to do without first informing them. If a user was able to send messages with a 10MB attachment last week, but this week, 10MB attachments are blocked, the user will undoubtedly complain loudly. Users may not be happy with restrictions that are put in place, but if they are informed about them and they are made to understand that these restrictions are necessary to maintain the health and security of the system, they will be more accepting.

The best starting point is to develop a usage policy that will define the restrictions and security constraints for your organization. Defined restrictions should include:

- File and content types that are blocked
- File size and message storage restrictions
- Message content that will be inspected and blocked (such as inappropriate language or inappropriate, non-business content)
- Internal network access types that are allowed (client types and versions)
- Internet email access (client types and versions)

Once the policies are in place and the date that the policy becomes effective is published, you can start enforcement. Choosing mechanisms for email security (and policy enforcement) is also tricky. The process of choosing, building, and tuning the best email security platform for your company is a delicate balancing act between the best possible security, best possible usability, and the most reasonable cost. Given your budget for email security, decide your tolerance for locking down your messaging system and the hardware/software/services products that will help you to do so.

Topic 5: Firewall Strategies and Best Practices

Q 5.1: What features should I look for in a firewall?

A: Choosing a firewall is one of the most important decisions an organization makes with respect to keeping their network safe from outside intrusion. Organizations that have been connected to the Internet for a number of years may be on their third- or fourth-generation firewall. Today's firewalls have far more features than their predecessors did a few years ago. This is due, in part, to customer demand for more functionality and partially due to the evolving threats that are an inherent part of being connected to the Internet.

For a person that does not work with firewalls on a regular basis, figuring out which firewall features are necessary can be a mind-numbing task. For this reason, discuss your security and protection needs with a person knowledgeable in firewalls. Otherwise, you may select a solution that does not meet your existing needs or won't accommodate changes you may require in the near future.

Basic Firewall Features

Firewalls started out as devices that were not much more complex than a simple router. A term that would describe an early firewall is a filtering router. This device could view inbound and outbound IP addresses and decide whether the source and/or destination IP address should be forwarded or discarded. Many routers now actually include this feature.

Application protocol filtering firewalls can examine not only the source and destination IP address of inbound/outbound IP data but also the source and destination port numbers that the application data is using. So, for example, a firewall that is capable of this type of filtering might have a rule that allows inbound port 443 (Hyper Text Transfer Protocol—HTTP—using Secure Sockets Layer—SSL) but only to a specific IP address on the internal network.

Another evolution of firewalls is the stateful inspection feature. *Stateful inspection* (also known as dynamic packet filtering) is a feature of more advanced firewalls through which the firewall keeps track of inbound and outbound packets over time and analyzes whether the requests and responses are valid for the type of traffic that is being generated. The firewall keeps track of the "state" of each connection and determines whether the traffic is valid for a connection in a particular state.

Any firewall that you choose should at least offer these features. Unless you are looking at very obscure and immature products, all firewalls on the market today support these features.

Common Firewall Configurations

There are a couple of common configurations for firewall deployments. The first of these (and usually the simplest) is a single firewall that provides an internal network connection as well as an external network connection, and a demilitarized zone (DMZ) network connection (the DMZ is often called a screened subnet or a perimeter network).

Figure 5.1 shows an example of a single firewall with a DMZ port. The DMZ contains servers that should be accessible by users from the Internet. Internet users would be able to access the servers in the DMZ but would not necessarily be able to access servers on the internal network.

However, the DMZ servers may be capable of connecting through the firewall to internal servers for data that a user from the Internet requests. Servers with data (such as Exchange servers or domain controllers) should not be placed in the DMZ because the possibility of these servers being compromised is slightly higher than for those on the internal network.

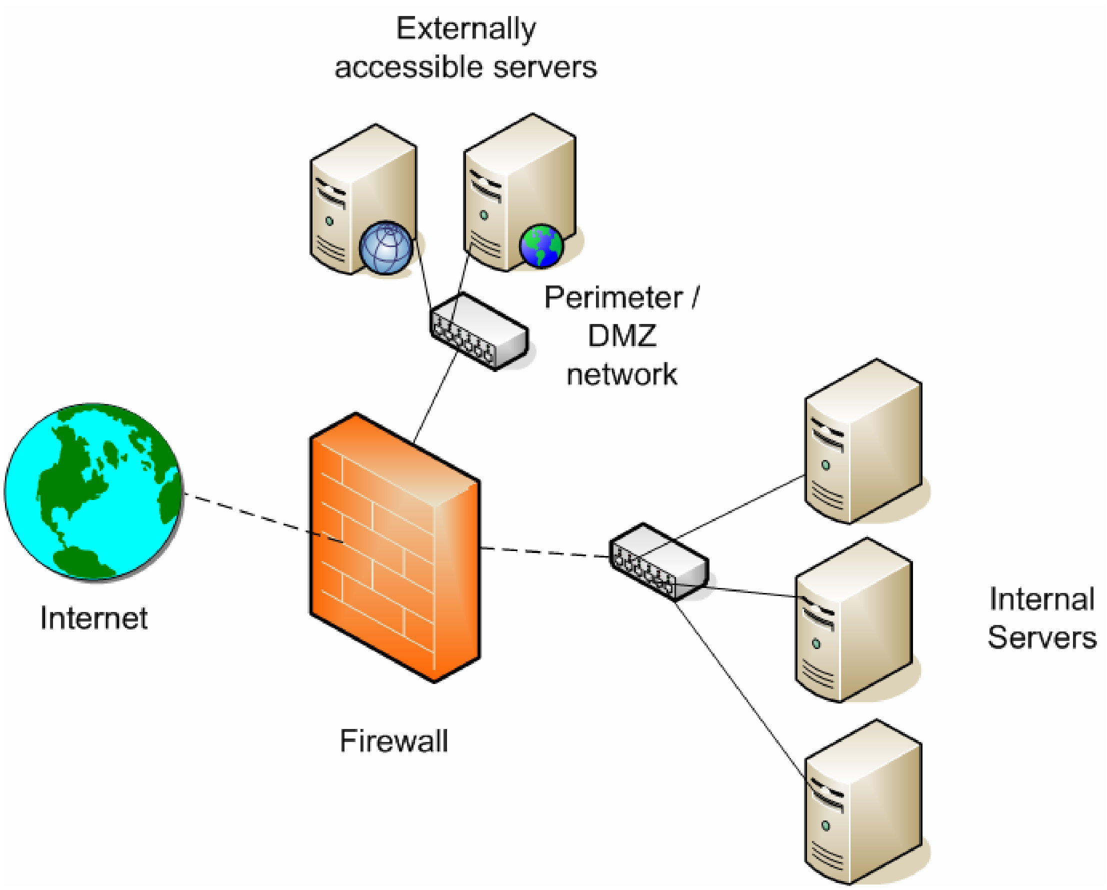

Figure 5.1: Single firewall with a DMZ port.

A second possible configuration for firewalls is to use more than one firewall. Figure 5.2 shows a multi-firewall configuration. Often, when using multiple firewalls, the internal firewall will be from a different vendor than the external firewall. This setup helps mitigate the possibility of an attacker penetrating the network as the result of a weakness in one firewall because the weakness would have to exist on both vendors' firewalls. Sometimes the external firewall is a simple firewall appliance and the internal firewall is a software-based and/or application-layer firewall and thus allows more detailed and complex configurations.

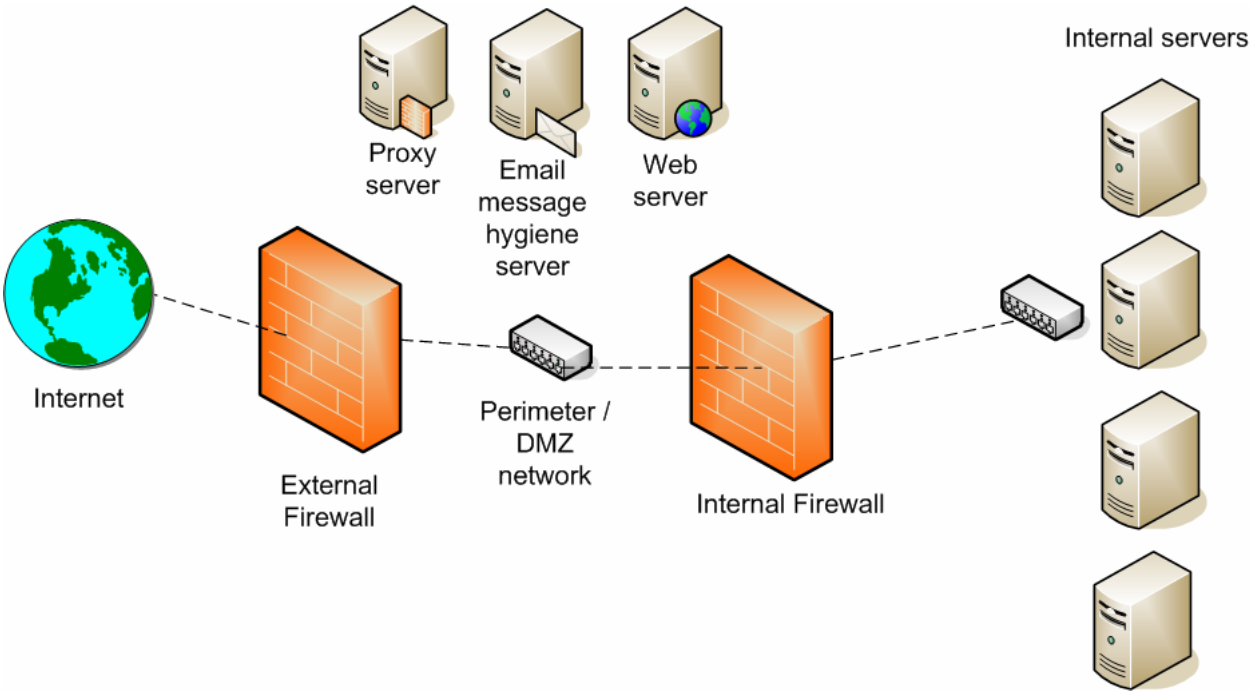

Figure 5.2: Multiple firewall setup.

Appliances vs. Software Firewalls

A single firewall in the configuration in Figure 5.1 may be a firewall "appliance." An appliance firewall is term given to a "blackbox" type solution where the hardware and software are bundled and sold as a single unit. These appliances are often designed as "hardware only" solutions and built to take up as little as 1U of rack space. Although they are not really "hardware only," the illusion of it being hardware only is applicable because you purchase the appliance with the software preloaded on an internal hard drive or firmware.

Hardware-based firewall appliances are desirable for some organizations because they are plug-and-play solutions; you simply plug in the device and turn it on. Of course, things are never that simple, but that is the way they are often marketed. One of the most common firewall appliances is the Cisco PIX firewall that Figure 5.3 shows. Other common hardware firewalls include SonicWall, Nokia, NetScreen, and Watchguard.

Figure 5.3: Cisco PIX firewall.

Appliance firewalls are useful because they are shipped as a self-contained unit and usually have out-of-band management features that allow them to be configured via a modem or terminal connection. Firewall appliances are usually not as versatile as a software-based firewall.

Although all firewalls run some type of software and they run on some type of hardware, a “software only” firewall is a product on which you to choose your own hardware on which the software will be installed. Software firewalls usually run on top of a Windows or UNIX operating system (OS). Often, with a software-based firewall, you will have more choices as to what types of add-on products you can purchase and are allowed to work with the firewall.

Application-Layer Firewalls

Simply put, an *application-layer* firewall is a firewall that can inspect inbound and outbound content at the application layer of the Open Systems Interconnect (OSI) or Department of Defense (DoD) protocol models. An application-layer firewall can examine data such as HTTP headers or SMTP commands in order to determine whether that type of traffic is allowed. For example, you can configure an application-layer firewall so that it will allow outbound HTTP from your users, but will block HTTP data that contains instant messaging content.

Firewall products such as Microsoft ISA Server and Check Point Firewall-1 can even inspect Remote Procedure Call (RPC) traffic to determine whether it is valid traffic for an Exchange Server/Outlook client session and either permit the communication or not.

Features, Features, Features

When you start shopping for a firewall, you want to make sure that it will meet your immediate and near-term needs. It is useful to be able to also project your long-term networking needs, but the firewall should at least accommodate you for the next 1 to 2 years.

Most any modern firewall product you are going to consider is going to include basic features such as IP filtering and port filtering, but what other features may be important? The following list highlights features that you might require or just find useful:

- Make sure the firewall will support the maximum number of concurrent connections you plan to have.
- Will the firewall provide virtual private network (VPN) support and if so, which VPN protocols do you need (Point to Point Tunneling Protocol—PTP, Layer-2 Tunneling Protocol—L2TP, Internet Protocol Security—IPSec)?
- Make sure the management interface is easy to use and is consistent with other interfaces that your IT department currently manages. For example, if you have UNIX and Cisco router administrators on staff, they will easily adapt to a command-line interface of a product such as the Cisco PIX. However, if you are entirely a Windows shop, GUI-based interfaces such as the Microsoft ISA server may be more applicable.
- Is high availability, scalability, and load-balancing a factor for your network? Make sure the firewall you choose can be load balanced.
- Does the firewall offer application-layer inspection? It should be able to inspect the various protocols you will be supporting such as inspecting HTTP headers or SMTP traffic.
- Does the firewall act as a Web caching proxy server for outbound HTTP requests?
- Can the firewall act as a reverse proxy server for inbound HTTP/HTTPS requests?
- Does the firewall provide logging, intrusion reports, and usage reports?
- What additional features can be added to the firewall that may meet your security requirements? These might include virus, hostile content, and spam filtering as well as Web site blocking.

Topic 6: Protecting and Controlling Sensitive Information in Email

Q 6.1: What are common methods for data to leak out of a company?

A: Email systems have provided both an incredible platform for moving important data between an organization’s knowledge workers and a method of quickly disseminating confidential information to unauthorized parties. Email has become one of the most common vectors for sensitive information being leaked to external organizations whether the intent is malicious or accidental. Although this is information that many IT managers and even company executives would like to ignore, a large portion of security breaches, systems being compromised, or data being leaked actually comes from within the organization. The 2002 Computer Crime and Security Survey conducted by CSI and the San Francisco office of the Federal Bureau of Investigation (FBI) Computer Intrusion Squad estimate that approximately 60 percent of security breaches come from within an organization’s network.

> Once information leaves your organization, you no longer have any control of the dissemination of that information.

Content filtering vendor SurfControl conducted a poll recently that found that nearly 40 percent of email users have received confidential information that was not meant for them. Another 15 percent of respondents admit accidentally sending confidential information.

A 2004 report by Jupiter Research found that large organizations were reporting email forwarding as being among their top three security breaches. Some IT administrators, managers, and executives will immediately think of confidential information that could leak out of the organization and harm the organization. Others may not immediately feel like they have any information that could be compromised. However, if you think about it, all organizations have accounting information, internal plans for the future, competitors, and human resources information. What are some examples of information leaks that can easily occur via email? The following list highlights answers to that question:

- An employee may intentionally seek to steal trade secrets of an organization by using the company email system to send client lists, sales information, pricing information, insider information, marketing plans, and so on to a competitor.
- Instead of using the internal email system, an employee may use their own personal Web-mail account to upload information to the Internet and send that information to others, thus bypassing internal content inspection controls that you have put in place.
- Information that is passed between any two users within your company can accidentally be misaddressed and include someone on the outside.
- Future plans or internal problems may be leaked to the media by a disgruntled employee.
- Human resources information may be sent intentionally or unintentionally to other employees or outsiders.

- Employees send confidential information to someone that they believe to be trustworthy, but the recipient misuses that information.

Ramifications of Data Leaks

No one wants their internal company problems or sensitive information to appear on the front page of the *Wall Street Journal*, but leaks happen and they happen frequently. What are some of the ramifications of important data being leaked to outside sources?

- Leaked information can cause enormous financial losses. If the company is a publicly held company, a leak can (and usually does) extend to the price of the company's stock.
- Public embarrassment and a loss of customer confidence can occur when information is disclosed.
- Reputations and credibility can be damaged.
- Legal action can be taken if the disclosure causes a law to be violated such as a privacy law. This legal action can come in the form of civil or criminal action based on the situation and the laws that were violated.
- Employee confidence can be eroded if preliminary company plans are leaked to the media. In a few situations, this has caused key employees to resign in order to avoid future problems.
- In the case of government information, such as military plans or intelligence gathering, lives can be placed at risk.

Depending on the type of information that is leaked and to whom the information is disclosed, a company may not survive. For example, information that is leaked to a competitor who can make better use of something such as marketing information and thus beat your company to market with your new "key" product.

Is S/MIME the Answer?

The S/MIME standards provide an email user with properly equipped client software to digitally sign a message (to authenticate the sender) and to digitally encrypt the message (to protect the content during data transmission and while the message is stored). Upon initial examination, you may be tempted to think that implementing S/MIME is the total answer to your data protection needs, and S/MIME messages that are encrypted will protect the message from unauthorized access. However, once a message and its attachments are sent to someone using S/MIME encryption or signing, the recipient still has complete and total control of that messaging content.

The Bottom Line

What is the bottom line? Employees must be trustworthy. Although mechanisms—such as authentication, SSL, S/MIME encryption, content inspection, and enterprise rights management—can be put in place to support your security policies, a determined user that is given access to this information can still disclose it. Anyone given access to sensitive information must be made to understand the ramifications of accidental or intentional information disclosure. Furthermore, there must be consequences that employees face if they are either careless or malicious with sensitive information.