Global Information Assurance Certification Paper

GIAC Security Essentials Certification (GSEC)
Version 1.4

Internet Email: Defense in Depth
Howard Edin

Introduction

On March 26, 1999, the Melissa virus rapidly spread through the Internet, overloading many corporate email systems (Foley and Bowman). Amazingly a similar virus, spread through email, did even more damage over a year later. In May 2000 the "Love Letter" virus clogged email servers and destroyed data on many computers (Hooper).

Companies often take email security for granted. Then a significant disruption like the "Love Letter" virus strikes and a business discovers how important and vulnerable email is. Most businesses have implemented virus protection to try and protect their systems. However, this is not a comprehensive defense because it only recognizes one avenue of vulnerability. In addition to viruses, the use of Internet email presents other threats such as Spam, unauthorized reception of programs and the exchange of sensitive information.

This paper will outline possible threats to an organization and how Internet email processing can be viewed in layers. By increasing the number of layers handling Internet email greater protection is given to the main mail server and clients. Each layer can address security threats, creating an in-depth email defense. Using a layered methodology with an SMTP mail proxy and SMTP mail gateway will greatly enhance an organizations email protection and reliability.

Overview of Email Threats

A good defense strategy depends on understanding the possible threats to a business. Several key areas in email vulnerabilities can be identified:

- Viruses or malicious code arriving as email attachments threaten data and impair the operation and reliability of the mail servers. The last two major virus outbreaks have cost businesses billions of dollars to clean up (Festa and Wilcox). Both of these viruses arrived as email attachments and exploited poor security on the receiving mail program to execute and propagate themselves.

- Spam, a popular term for Unsolicited Bulk Email (UBE), is a significant problem for two reasons: the cost to the organization to process this mail (Hoffman) and the potential liability issues with offensive material.

- Information disclosure or unsecured exchanges of private information is a security issue for the business. In contrast to most threats this is of internal origin and can be very difficult to control. Many organizations are affected by the Heath Insurance Portability

and Accounting Act (HIPAA) which sets strict guidelines on what information must be exchanged securely (Department of Health and Human Services).

- Unauthorized reception or transmission of information or computer programs can be a productivity problem and a reliability issue for the company. This includes games, executable files, operating system (OS) components, etc.

- Denial of service (DoS) attacks targeted at the company's Internet email system can disrupt email service. If the primary internal email server also acts as the email gateway to the Internet, any attack may impair or disable internal email delivery.

There is no perfect defense against all these threats. Each organization must determine what level of importance should be assigned to each threat. By examining how Internet email operates a defensive structure can be formulated.

The Basic Email Process

Internet email is based on the Simple Mail Transfer Protocol (SMTP) originally defined in Request For Comment (RFC) 821. Its original purpose, still true today, is the efficient and reliable exchange of messages regardless of host system types (Postel).

The transfer of email using SMTP is straightforward. A sending system communicates to the receiving server using simple commands in clear text. These commands, like MAIL FROM, RCPT TO, and DATA are used to transfer information in plain text format. The adoption of open standards to move mail from system to system provides an easy way to communicate.

These standards have provided the framework for today's rich Internet email. The exchange of email with imbedded attachments, formally called Multipurpose Internet Mail Extensions or MIME, was defined in RFC 1521 (Borenstein and Freed). Extensions to the original SMTP specification continue to be added.

The SMTP email exchange process can be viewed from a layered perspective. The typical Internet email exchange looks like this:

Sending Mail Program		Receiving Mail Program
Mail Server	← →	Mail Server

The sending computer's mail application sends a message to the local mail server. This mail server then uses SMTP to deliver the message to the receiver's local mail server, which in turn, delivers it to the receiver's mail application.

This arrangement does not provide enough protection against potential email threats. Even with anti-virus software running on the server and the local computer there are several issues with this configuration:

- Internet email arrives directly to the email server accessed by internal users. The system may be vulnerable to exploits or attacks at either the mail application level or the OS level. DOS attacks launched against the mail server can potentially interrupt all email processing both internal and external (Microsoft Security Bulletin MS02-25).

- New virus or Trojan programs will arrive directly at the email server. Even with regularly updated virus signature files the system is at risk from both new viruses and executable files.

- Many email systems do not have an adequate level of control for restricting attachments, implementing user policies, and deploying centralized messaging encryption. Information may be flowing in and out of the organization without proper control.

Relying on a single email server to protect against all the potential threats to an organization is a poor defense. No single system or software package can address all security needs. The principle tenant of Defense In Depth is to have many layers of protection. In the interest of security and reliability, Internet email delivery should be separated from internal email functions. This is accomplished through the use of layered SMTP processing systems.

Layered Email Systems

A layered email system is comprised of two or more layers of SMTP email processing systems. Using three layers provides considerably more protection from potential threats. From an application perspective a three-layer SMTP processing system would process email like this:

		Receiving Mail Program
		[Layer 3] Mail Server
Sending Mail Program		[Layer 2] SMTP Gateway
Mail Server	← →	[Layer 1] SMTP Proxy

Each layer in the chain examines email with a specific function. Ideally the host and software at each level is different, so that no single vulnerability can affect the entire system. However, that is not always practical. It is essential that each layer implement defense in depth: hardening of the host OS and up-to-date patching of the application software.

Adding processing layer(s) to existing infrastructure is simple. An existing mail server is configured to forward all mail on to another SMTP server, perhaps a Layer 2 SMTP gateway. The function and possible implementation solutions for each layer are discussed in the following sections.

Layer 1: SMTP Mail Proxy

The SMTP mail proxy is the contact point between external and internal email systems. The proxy's primary function is to transfer email in a controlled and secure manner between the external system and the internal SMTP gateway. This proxy should be capable of performing several checks on the email exchange. These include:

- Determining if the remote host is allowed to connect. This is usually done using a deny list of addresses or domain names. This is a rudimentary form of Spam control. Performing this blocking at the proxy immediately reduces the load on subsequent email systems. This function can also be done at the SMTP gateway, which may be easier to administer (since email and firewall administration is often done by different groups).

- Determining if the sending domain exists, and/or verifying the sending host is who they say they are (reverse dns lookup). This is a form of Spam control that attempts to check if the message is spoofing it's source. This may not be particularly effective and can block legitimate messages from misconfigured external mail and DNS servers.

- Ensuring that the email message and messaging connection conforms to RFC standards. Monitoring and ensuring that the email connection is following established SMTP standards safeguards the internal systems from hacking attempts or incorrectly formatted email.

The proxy should be configured to forward all incoming email to an internal SMTP gateway. The proxy should also be configured to only accept outgoing SMTP connections from the SMTP gateway (level 2). This prevents rogue internal mail servers from operating within the organization.

There are two ways to implement the proxy: using an application level proxy firewall or a dedicated mail proxy system. In the first case, the firewall will function as a proxy server by it's very design. In the second case a standalone mail proxy system is used in conjunction with a circuit or packet level firewall.

Application level Firewall: Figure 2 illustrates the method of implementing a multi-level Internet email system using a proxying firewall:

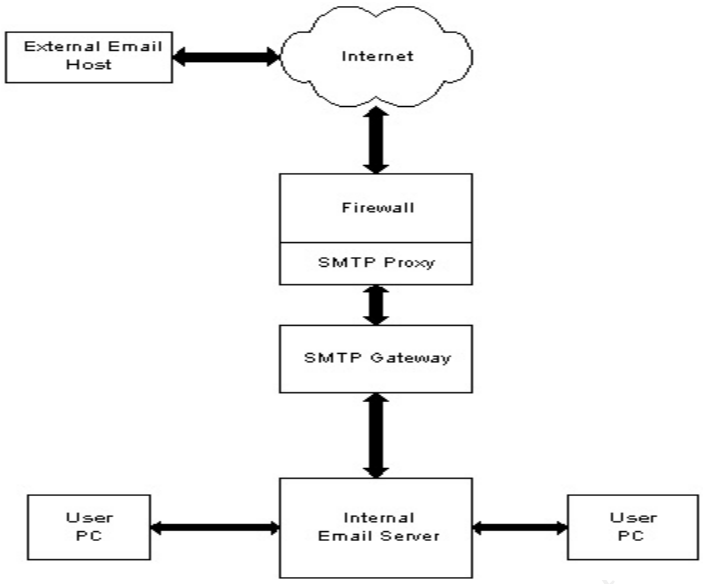

Figure 2: Application Level Proxy Firewall in Multi-Layer Email System

In this scenario the firewall is the first point of contact for incoming Internet email. When the firewall software detects the connection is for SMTP (port 25) it passes the data stream through an internal SMTP proxy. This is the most important feature of an application level proxy firewall; it examines the mail exchange and ensures that proper commands and data formats are being used as SMTP email is exchanged. A packet level firewall that simply controls the ports available to systems will not provide this type of protection (Wack).

Examples of application level proxy firewalls are:

- CheckPoint Firewall-1
- Cisco Secure PIX Firewall
- Symantec Enterprise Firewall (SEF), formerly Raptor Eagle Firewall

These products can perform application level proxying. Typically the proxy's settings can be tuned for best performance. For example, the Symantec Enterprise Firewall (version 7.x) SMTP proxy has the following configurable options through the administration interface:

- **Soft Recipient Limit**: Limits the number of recipients per message. If the sending mail system attempts to exceed the set value of recipients for a single message the sender will be told to retry the message send.
- **Hard Recipient Limit**: As above except the entire message will be rejected if the recipient count exceeds the setting.
- **Hide Internal Domain**: Any internal domains listed will be have the mail header rewritten to hide the domain name.

- **Recipient Domains:** What domains the proxy will accept mail for.
- **Check Sender Domain**: A DNS record for the connecting host must exist. Disabled by default.
- **Reject Source Routes**: If enabled the system will reject mail using source routing syntax. Disabled by default; this works in conjunction with the **Recipient Domains** field to restrict mail relaying.
- **ESMTP**: Per command control of ESMTP, AUTH, ATRN, ETRN, EXPN, and VRFY. By default most of the enhanced SMTP (ESMTP) commands are disabled – only ESMTP and AUTH are enabled.

Certain commands like EXPN and VRFY represent security holes because they can reveal mail aliases (Postel). In addition to the settings listed above, several other tunable parameters are in the firewall's main text configuration file:

- **smtpd.max_body_line_length**: Disabled by default; when enabled it is set at 1024 characters but can be adjusted as desired. Enabling this option may prevent buffer-overflow attacks. The majority of email will not be affected by this option so I always enable it.
- **smtpd.length.Content-Type** and **smtpd.length.Content-Disposition :** Disabled by default; when enabled these provide size restrictions on two portions of an SMTP message.
- **smtpd.check_sender_regexp:** Configurable regular expression checking of sender address. Disabled by default; when enabled the **smtpd.bad_sender_regexp** setting will be tested against all messages.

These are not as obvious to change and are only documented in the **config.sg** file. Other firewall products offer similar control features for handling SMTP email. These features provide some Spam control and ensure the connection is following accepted standards.

The inherent design of application level proxy firewalls leads to a very secure system. When implementing SMTP restrictions, the firewall logs should be watched closely for abnormal amounts of rejected messages from known sources; this may indicate a miss-configuration at either end. For instance, I worked on an issue where email was not making it from an external vendor to the internal customer. The problem was line length restrictions at the firewall. Although the firewall configuration could have been changed, the correct solution was for the vendor to properly format the message they were generating.

Dedicated Mail Proxy: If it is not possible to use a proxying firewall, an alternative is to deploy a dedicated email system on the service network of the firewall as shown in figure 3 below.

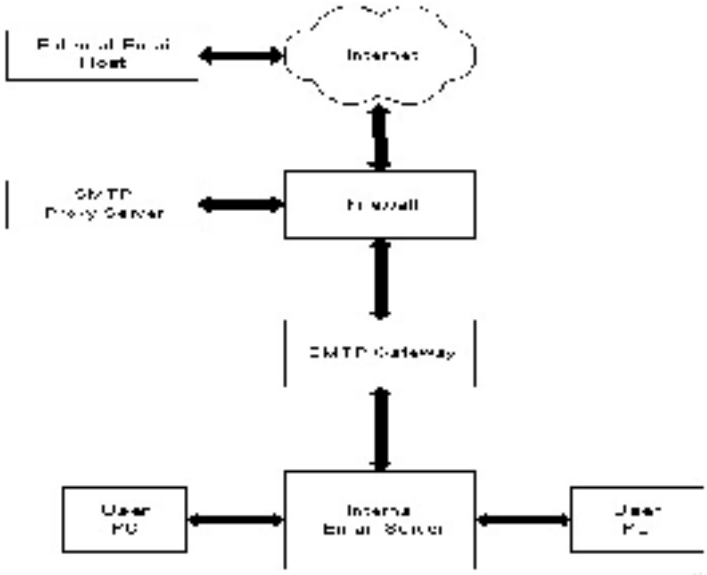

Figure 3: Firewall with external SMTP Proxy server

In this scenario the firewall redirects SMTP port connections to a SMTP proxy server on a service leg of the firewall. This provides the desired isolation of internal email systems from external host connections.

Configuring this design involves more work in setting up the firewall rules and hardening the mail proxy host. It also may not provide the level of security an application level proxy firewall can. For instance, a proxy firewall might block the sending mail system from issuing a particular SMTP command that would otherwise pass through a packet firewall. In contrast, a dedicated mail proxy has several advantages: it would be inexpensive to deploy, will perform better under heavy load, and will have more flexibility in handling current and future SMTP extensions.

The email proxy can be implemented using a message transport agent (MTA) such as Qmail or sendmail. Both products have extensive features for processing email, but sendmail is often regarded as insecure (Scambray et al. 326). Qmail contains features well suited for use as a proxy. For instance, it can be invoked from a connection control program like tcpwrappers to provide connection restrictions. In Qmail the EXPN command is unimplemented and the VRFY command does not reveal any information. Qmail also hides its identity when connected to by a mail server (Anonymous).

Either method of implementing a SMTP proxy can also hide the identity of internal systems and software. For example, consider this sanitized connection sequence:

```
> telnet mail.somecompany.com 25
Trying 10.10.10.176...
Connected to mail.somecompany.com.
Escape character is '^]'.
220 mail.somecompany.com ESMTP Sendmail 8.9.3.Anti_Relay.Anti_Spam ;
Fri, 10 May 2002 09:45:55 -0700 (PDT)
helo test.com
```

```
250 mailer.somecompany.com Hello bogus.test.com [1.1.1.1], pleased to
meet you
quit
221 mail.somecompany.com closing connection
Connection closed by foreign host.
```

In this connection example the remote host reveals it is using Sendmail 8.9.3 software. It may be possible to exploit that software. On the other hand, consider the following connection sequence:

```
> telnet mail.foobar.com 25
Trying 10.10.10.162...
Connected to mail.foobar.com.
Escape character is '^]'.
220 mail.foobar.com Generic SMTP handler
helo test.com
250 mail.foobar.com talking to bogus.test.com ([1.1.1.1])
quit
221 mail.foobar.com Service closing connection
Connection closed by foreign host.
```

On the surface it is difficult to determine what type of mail software answered the connection request; it is an SMTP proxy at the firewall.

It may be good design practice to alter the SMTP welcome message regardless of the system used. This aids in hiding the software's identity from SMTP scanners. As previously noted, proxy firewalls and Qmail do not directly reveal their identity. However, the widely used sendmail program does reveal information with a default configuration. This can be changed; the welcome string is modified using the 'smtpgreetingmessage' option in the **`sendmail.cf`** configuration file (Costales 752).

First layer defense separates the foreign mail hosts from the internal mail hosts. While the SMTP proxy provides a high level of isolation it lacks the anti-virus software and SMTP control features described in the next section.

Layer 2: SMTP Mail Gateway

The SMTP mail gateway is the second level of defense and should be the central control point for all email entering and leaving the business. The SMTP gateway's role is to stop any known viruses from passing through, control email attachments allowed in or out, and regulate who may send and receive email.

Virus scanning should be done on all messages passing through the system. This includes attachments that contain file archives, which may contain additional files themselves. Virus scanning will prevent known malicious code from entering or leaving the system. Attachment blocking can stop any type of email payload and is an effective defense against new or unknown types of malicious code. Blocked attachments should include scripting code (.VBS, .VBP, .SCR**)** and common executables **(**.COM, .EXE, .BAT). This ensures that new virus code that slips past the virus scanner cannot make it into the main email system.

Effective attachment blocking requires new or revised email policies in the company. It can be argued that no executable should be allowed into the organization without review. The policy can be restrictive to users expecting code updates or other legitimate programs. Although this may place an additional administrative burden on the mail administrator(s), it effectively safeguards the company from virus or Trojan infection via SMTP email. This is particularly important as the sophistication of viruses increases and the cost of cleaning them up rises (Fisher). Blocking email attachments can also reduce some support costs. By preventing the introduction of games, screen savers and incorrect system files the internal PCs may require less desktop support.

The SMTP gateway is located on the internal protected network and is the relay point for all internal hosts. Several email domains may exist inside the corporate network and the SMTP gateway should be configured as the central routing point (or 'smart host'). Internal hosts allowed to send Internet email are controlled through relay restrictions on the gateway. Blocking of external email senders should be done at the gateway and would represent a single system to administer for blocking Spam.

The SMTP gateway can be implemented using either commercial software or open source solutions. Examples of full-featured commercial SMTP gateways include:

- Tumbleweed Secure Mail (MMS)
- Network Associates WebShield SMTP
- Trend Micro InterScan

Many of these products include additional features that further enhance the functionality of the gateway. A full description of these features is beyond the scope of this paper, but a brief summary of each follows.

Content Filtering: Nearly all SMTP gateway software provides some measure of content filtering. This can be very useful in reducing or eliminating many types of offensive or Spam messages. However some legitimate messages may be blocked. Content filtering is accomplished by examining the SMTP message's text for specific words or phrases. More sophisticated filtering programs can assign a risk factor to the message based on the type of words or phrases within the email. This allows for greater customization and flexibility.

Typically messages that violate the content policy are quarantined. This allows for release and transmission of the message. The quarantine queue should be monitored to verify the software is working correctly and is not blocking words with a common usage within the company. For example, I worked with a version of WebShield SMTP that scanned for words in the mime-encoded portion of the SMTP message. This is not correct behavior (the MIME portion should be decoded prior to word scanning) and led to a surprising number of quarantined messages.

A good introduction to content filtering is available on the SANs website (Cohen). Implementation of a content filtering policy should be coordinated with a company's Human Resources and/or legal departments.

Per user control of email: Common email systems do not have enough control over individual users receiving and sending Internet email. Some gateways, such as Tumbleweed MMS, provide a highly configurable level of control. Rules are implemented that allow only designated sending and receiving addresses access to Internet email.

Message Encryption: A messaging encryption feature is useful for protecting sensitive information. Secure MIME, or S/MIME as it is commonly referred to, is another extension to SMTP. Its purpose is to apply security measures, such as encryption, to SMTP messages (Galvin). As defined in RFC 1847, no single method of encryption is specified. Many S/MIME implementations use a certificate much like HTTPS. Instead of deploying certificates to individual users the gateway can act as a centralized certificate server and S/MIME decryption point. This can be particularly useful in a large organization.

As an alternative to commercial software it is possible to construct an SMTP gateway from open source components. One solution would be to use sendmail as the message transport agent (MTA). In sendmail version 8.11 a new content filter application programming interface (API) was introduced called milter. This API was created for incorporating filtering software as mail messages are processed (Filtering Mail With Sendmail). Several open source projects that use this interface are available by searching for 'milter' on the sourceforge.net website. As of this writing no single open source program

The SMTP gateway represents an important element in the email defense strategy. Combining virus scanning, attachment restrictions and relay restrictions, it regulates the flow of email in and out of the organization. In addition, it can perform message filtering and encryption to further secure the business's email.

Layer 3: Internal Email Server

The third level of defense is the main internal email server. These are the systems currently serving the company's business. These servers should be configured to relay Internet email only through the SMTP gateway. In addition, these systems should incorporate two basic defensive measures: virus scanning and message relay restrictions.

Virus scanning should be done on the internal email server; this prevents viruses of internal origin from infecting other internal hosts. The SMTP gateway will not be able to provide anti-virus protection to internal users homed on the same email server. It is essential to use virus-scanning software at both level 2 and level 3. A good defense against viruses is created if different anti-virus products are used at each level.

Typically the internal email server(s) are Microsoft Exchange Server or Lotus Notes, both have configurable options to limit message relaying. Out of office messages or other automated replies are typically not selective and will respond to any email sent to the recipient's mailbox. Both cases represent a security concern because these replies may be used in social engineering. For instance, John Doe announces he is out of the office; the message may disclose contact information including help desk contacts and presents the best time to socially engineer his password.

Leakage of information through automatic email replies is an often-overlooked security concern. Prior defensive layers cannot adequately prevent this (although content filtering would help). Out of office replies or automatic rules that relay internal messages to an external address can be blocked from leaving the company at the main email server. By restricting these messages to internal use only, email information disclosure will be reduced.

Final Defense: Client Email Software

The end user is the reason the email system exists. They are also the least controllable element in the process. Despite the extensive mail system protection prior to the client, other vectors of attack exist. There are numerous and ever changing hazards including malicious websites, instant messaging (IM) applications, and peer-to-peer software sharing programs. Three specific areas should be addressed on the client side:

- Antivirus software should be installed and kept up to date.

- Browsers and email clients should be patched. It is essential that organizations maintain the client software with vendor updates. Both the Melissa and Love Letter viruses exploited security holes in client software.

- The end users need to be educated on company policy for email usage. Clear email policies and end user education will help ensure the integrity and security of the business.

End user education is very important. The list of topics that could be covered should include how to handle Spam, what information not to disclose, what the company policies are toward personal mail and attachments, and what the company is doing to ensure secure messaging.

Summary

Viruses are a significant concern to corporate email and have caused considerable disruption in the past. However, they represent only one of many threats to a business's Internet email security.

Internet email is based on SMTP and has been in use for two decades. The simple transfer of messages from sender to receiver can be viewed in layers. A stronger defense can be built by isolating the main mail server from Internet connections. Using a layered methodology to handle SMTP email provides a Defense in Depth.

Many different methods of implementing this strategy are possible. Commercial and open source software solutions exist at each level. Each layer can address specific security concerns. While the installation and maintenance of additional software may initially increase the cost of email, it is offset with improved security and reliability.

References

All illustrations are by the author.

Anonymous. "Maximum Linux Security", 1999.
URL: http://mandrake.petra.ac.id:8888/info/max/BkPg155x164.htm, (August 25, 2002).

Borenstein, N. and Freed N., "MIME (Multipurpose Internet Mail Extensions) Part One: Mechanism for Specifying and Describing the Format of Internet Message Bodies", September 1993.
URL: http://www.ietf.org/rfc/rfc1521.txt, (22 July 2002).

Cohen, Bradley. "Introduction to Enterprise Content Filtering", 25 April 2001.
URL: http://rr.sans.org/email/content_filter.php , (3 May 2002).

Costales, Bryan. sendmail, Second Edition, Sebastopol: O'Reilly & Associates, Inc, January 1997.

Festa, Paul and Wilcox, Joe, "Experts estimate damages in the billons for bug", May 5, 2000.
URL: http://news.com.com/2100-1001-240112.html?txt, (30 June 2002).

"Filtering Mail With Sendmail". 2000.
URL: http://www.sendmail.com/partner/resources/development/milter_api/ , (18 May 2002).

Fisher, Dennis. "Network Defenses Can't Foil Viruses", 4 Mar 2002.
URL: http://www.eweek.com/article/0,3658,s=712&a=23564,00.asp , (10 May 2002)

Foley, Mary and Bowman, Lisa. " 'Melissa' virus swamps corporate e-mail", March 26, 1999.
URL: http://zdnet.com.com/2100-11-514149.html, (27 July 2002).

Galvin, J. "Security Multiparts for MIME: Mulipart/Signed and Multipart Encrypted", October 1195.
URL: http://www.ietf.org/rfc/rfc1847.txt, (20 July 2002).

Gaspar, Suzanne. "Fighting Back Against Spam", NetworkWorld 13 May 2002. 48-50.

Hardin, John. "Enhancing E-Mail Security With Procmail". 1.155. 13 May 2002.
URL: http://www.impsec.org/email-tools/procmail-security.html (17 May 2002).

Hopper, D. Ian, "ILOVEYOU computer bug bites hard spreads fast", May 4 2002.
URL: http://www.cnn.com/2000/TECH/computing/05/04/iloveyou.01/, (20 July 2002)

Hoffman, Paul. "Unsolicited Bulk Email: Definitions and Problems", 5 Oct.1997.
URL: http://www.imc.org/ube-def.html, (3 Aug. 2002).

Microsoft Security Bulleting MS02-025. "Malformed Mail Attribute can Cause Exchange 2000 to Exhaust CPU Resources (Q320436)".

URL: http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS02-025.asp, (6 Aug. 2002).

Postel, Jonathan. "Simple Mail Transfer Protocol", August 1982.
URL: http://www.ietf.org/rfc/rfc0821.txt?number=821, (11 May 2002).

Scambray, Joel, Mcclure, Stuart, and Kurtz, George, Hacking Exposed: Network Security Secrets & Solutions Second Edition, Osborn/McGraw-Hill, 2001.

Symantec Enterprise Firewall and Symantec Enterprise VPN Configuration Guide, Symantec Corporation, 2001.

United States. Department of Health and Human Services. Standards for Privacy of Individually Identifiable Heath Information. 28 Dec. 2000.
URL: http://www.hhs.gov/ocr/part1.html, (30 Aug. 2002).

Wack, John, "Guidelines on Firewalls and Firewall Policy", January 2002.
URL: http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf , (4 May 2002).