



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Security Overhaul at a Mid-sized Company

Bruce Brooks

March 29, 2003

GIAC Security Essentials Certification (GSEC) Practical Assignment 1.4b

Introduction

After sending two engineers to GIAC training and seminars my company decided to put our skills to the test. Our goal was to identify possible security vulnerabilities and offer solutions to better secure the corporate network. Viruses and Trojan applications, Nimda, Code Red and others which had disabled our email services multiple times over the past years, were of particular concern.

The purpose of this paper is to demonstrate how we enhanced my organizations network security. The primary focus is firewall architecture and how we secured connections and data passed to Microsoft Exchange and IIS servers, these were my primary responsibilities. Workstation and server hardening were the other engineer's primary focus. Although we had our separate goals, both of us contributed to the end result, a more secure network.

Assessment

The company had an agreement with an internet service provider (ISP) for a T1 line with class C internet address block. These addresses were used for internal systems and an external system in the DMZ of a firewall. There was a router, supplied by the ISP, and a Gauntlet Firewall in place. The DMZ consisted of one Windows NT4 server with DNS installed and configured. The servers which hosted the company's website, web access for email, Exchange messaging and a custom billing system accessed through IIS were behind the firewall. Authentication for email, IIS applications and user logon was a Windows 2000 based domain. Figure 1 depicts this network.

This is a mid-sized company where 20% of employees reside in corporate office, with the remaining 80% residing at client's sites. This division of the workforce had caused us some problems in the past. Since the majority of users are outside the LAN, workstation control was out of the LAN administrators hands. The off-site staff used either company supplied laptops, computers supplied by the client or their home computers to check company email and access IIS services for time and billing management. Virus control was not in our control, it was left to the diligence of the individual employee or the administrators of the client's network.




Figure 1

During our analysis it was found that we were able to fully scan the system within the DMZ, as expected. The results for systems within the LAN were not expected. We did not actively try to compromise any systems; the goal was to gather data to better understand our current security posture. We used tools and online probes obtained from the following sites:

- <http://www.hackerwatch.org/probe/>
- <http://www.dslreports.com/scan/>
- <http://www.nessus.org/>
- <http://www.eeye.com/html/Research/Tools/nmapnt.html>

Scans were run from multiple systems on the internal network, as well as from the internet, and the reports revealed most systems were vulnerable to known attacks. The online scanners, which use the same reconnaissance tools as potential hackers such as Nmap or Nessus, allow for an easy overview. These are particularly effective in user awareness. We passed the online scanner URL's to all company employees, in an effort to raise their awareness, but to also illustrate the status of their systems.

Using the Nessus scanner from outside the network we found that most servers in the LAN, which were a mix of Windows NT4 and 2000, had multiple

vulnerabilities. Below is a sample of the highest risks reported from Nessus for an IIS server hosting Outlook Web Access (OWA):

Vulnerability found on port netbios-ssn (139/tcp)

. It was possible to log into the remote host using a NULL session. The concept of a NULL session is to provide a null username and a null password, which grants the user the 'guest' access

. All the smb tests will be done as "/"

This vulnerability allows a potential hacker to connect to tcp port 139, using session layer protocols to connect to administrative shares, \$IPC for example. Proper filtering at the firewall would have blocked this port. On a properly configured WWW server, there is not a need to have port 139 open, only port 80 should have been advertised.

Vulnerability found on port www (80/tcp)

It is possible to get the source code of the remote ASP scripts by doing the request :

```
GET
/null.htw?CiWebHitsFile=/default.asp%20&CiRestriction=none&CiH
iliteType=Full
```

ASP source codes usually contain sensitive informations such as logins and passwords.

Solution : if you need the fonctionnality provided by WebHits, then install the patch available at :
<http://www.microsoft.com/technet/security/bulletin/ms00-006.asp>

If you do not need this functionality, then unmap the .htw extensions
from webhits.dll using Internet Service Manager MMC snap-in.

Risk factor : Serious
[CVE : CVE-2000-0097](#)

"Malformed Hit-Highlighting Argument" Vulnerability, as defined in Microsoft Q251170 article. This vulnerability allows a potential hacker to view the web server directory and perhaps alter the content on the server.

This could have been avoided by applying proper hot-fixes and service packs.

Using NmapNT from the internet, we scanned all systems on the internal network. The sample below is from our Exchange messaging server:

Interesting ports on (207.x.x.5):

(The 1513 ports scanned but not shown below are in state: closed)

Port	State	Service
21/tcp	open	ftp
25/tcp	open	smtp
80/tcp	open	http
110/tcp	open	pop-3
113/tcp	open	auth
135/tcp	open	loc-srv
137/tcp	filtered	netbios-ns
138/tcp	filtered	netbios-dgm
139/tcp	filtered	netbios-ssn
143/tcp	open	imap2
443/tcp	open	https
515/tcp	open	printer

This illustrates one common problem found on almost all systems. Services that are not needed are installed, enabled and allowed through the firewall. Proper filtering at the firewall would not have advertised all of these ports as open.

There were three mission critical services which all users needed access to, either from the internal LAN or from the internet:

1. Exchange Messaging, MS Exchange MAPI client (internal) and POP3 (internal/external) were the methods for checking mail.
2. Internet Information Server (IIS) for access to the intranet.
3. Outlook Web Access (OWA) to check email, calendar, public folders or contacts from outside the LAN.

A major problem was IIS, it was configured with Basic Authentication; passwords were sent in the clear. Any person, with the right tools, sitting between the users' session and our internal network would easily be able to grab the username and password, allowing access to email, internal file servers and beyond. IIS was used not only for OWA access, but also for our billing and time management system.

The only defensive measure taken on the Exchange server was the installation of Norton Group shield. Although this allowed for active virus scanning of both incoming and outgoing email, virus scanning can only accomplish so much. The

problems we encountered with Code Red were not stopped by the virus scanner, the virus definitions were only updated by Symantec after we were infected.

Analysis of the firewall revealed that it had not been updated or closely monitored since it was installed 3 years prior. The rule sets configured for the firewall allowed the majority of ports to be open for servers within the internal network and some filtering for the rest of the LAN.

These findings were collected in to a security assessment document which was presented to management. Included in this assessment were recommendations which detailed how to make the network more secure and easier to manage.

1. Implement a secure firewall architecture.
2. Harden all servers.
3. Implement secure access to the three services defined as mission critical.

The management agreed that the problems needed to be rectified. They prioritized the engineers' duties to deal with the problem(s) immediately.

Implementation

Firewall

We decided on a two tier firewall architecture. Using an external firewall to filter ports and inspect packets was not enough. In order to offer secure services to users outside the LAN an application layer firewall was also installed. By using a two tier design were able to provide a "defense in depth" approach to securing our internal LAN while also providing the same to our external users' sessions.

External Firewall

Using a firewall which was easy to understand and troubleshoot was a top priority. The active firewall was a NAI Gauntlet firewall running on UNIX. Since this server's hardware was not able to be upgraded and the firewall software was being retired by the OEM, we chose to purchase a firewall appliance. After researching the current offerings on the market we choose a Sonic Pro 200 firewall appliance. This firewall was network address translation (NAT) capable, had built-in VPN capabilities, and was simple to setup and administer.

Current network addressing was using a registered Class-C address block for all internal and external systems (see Figure 1). We decided to change the internal address block to a RFC 1597 compliant block, private or non-routable addresses. Since we were planning on using a private address block, not assigned by our ISP for internal addressing, we needed to implement NAT on the firewall. During the transition from one ISP to another we decided to install the new firewall and assign a 10.x.x.x addressing scheme to the internal network.

Once the firewall was in place, using new internal/external IP addresses, we configured rules to allow appropriate port traffic to and from servers on the LAN and DMZ to the WAN. Instead of following the rules used on the old firewall, which opened all ports and only closed those which are known to cause problems, we created a rule-set which denied all ports unless explicitly opened. The main advantage to this approach, the firewall needs to be touched less. If a new vulnerability is exposed which is not within the scope of services for that server, we do not need to create new rules to block the vulnerability. The first rule is to deny all traffic from WAN and DMZ to the LAN.

With the proper one-to-one NAT rules and port mappings, all critical services were working from within the internal network and from the internet. A sample one-to-one NAT rule would be to allow port 25 (SMTP) from the public address x.x.x.7 to the private address of 10.1.2.7, the Exchange server. Now only port 25 traffic is allowed from the WAN to the specified server. This does not differ greatly from Figure 1, other than the TCP/IP address scheme had changed.

One-to-one NAT by itself is no more secure than we were before, ports needed to be opened to particular servers which exposed them to known and unknown vulnerabilities. The primary advantages of a firewall with NAT is that it will effectively mask your internal address block and requires less live IP addresses from the ISP. By running nMapNT again, it was demonstrated that even with a new firewall, more strict rules and NAT; ports are still visible to those with the tools:

Interesting ports on (x.x.x.7):

(The 1520 ports scanned but not shown below are in state: filtered)

Port	State	Service
25/tcp	open	smtp
80/tcp	open	http
110/tcp	open	pop-3

TCP Sequence Prediction: Class=truly random
Difficulty=9999999 (Good luck!)

There is no way to hide or stealth an open port where an external client is allowed to connect to a service. Understanding your systems and only advertising those ports which you need to advertise is a must. Compared to the earlier scans of our email server with 13 advertised ports, this scan shows that we reduced the number of vulnerable ports. To demonstrate how we are still vulnerable, using telnet we can open a connection to port 25, this will show what email platform we are running:

```
C:\telnet
Microsoft Telnet> open mail.mycompany.com 25
```

220 mail.mycompany.com Microsoft ESMTP MAIL Service, Version:
5.0.2195.5329

With this information in hand, a potential hacker is now aware that we are running MS Exchange 2000 and can now concentrate his efforts on known exploits. A firewall appliance provided the first step in securing our internal network. By opening only the ports needed to this server, it was much cleaner to a port scan, but vulnerabilities could still be leveraged.

Internal Firewall

The second step was installing a firewall with proxy capabilities. Since two of the resources all employees need to access, the Exchange server and IIS server(s) were Microsoft solutions, we chose Microsoft Internet Security and Acceleration server (ISA). One primary advantage with this proxy capable firewall is the ability to analyze HTTP, SMTP and RPC packets, insuring they contain legal requests. Since the ISA server acts as a proxy for the Exchange and IIS servers, while providing stateful inspection of the traffic, packets are dropped which do not resemble known traffic before they reach the internal server. The exterior firewall examines packets at the network layer, while ISA examines packets at the application layer as well.

This allowed us to setup a one-to-one NAT rule on the exterior firewall allowing only ports, 25, 110, 80, and 443 to the ISA server. Simplifying management of this firewall, we no longer needed to maintain one-to-one NAT rules for each server on the internal network, since this is accomplished with the interior firewall (ISA).

The ISA server was configured to also protect our internal network by providing content and virus scanning, further protecting our internal network from malicious code contained within packets. We setup filters to inspect packets and block .vbs, .wsh, .doc and .exe, which are commonly used to spread malicious code. This, in conjunction with GFI Download Security, which includes virus scanning, allowed a two tiered process for protecting both incoming and outgoing packets from viruses or known threats.

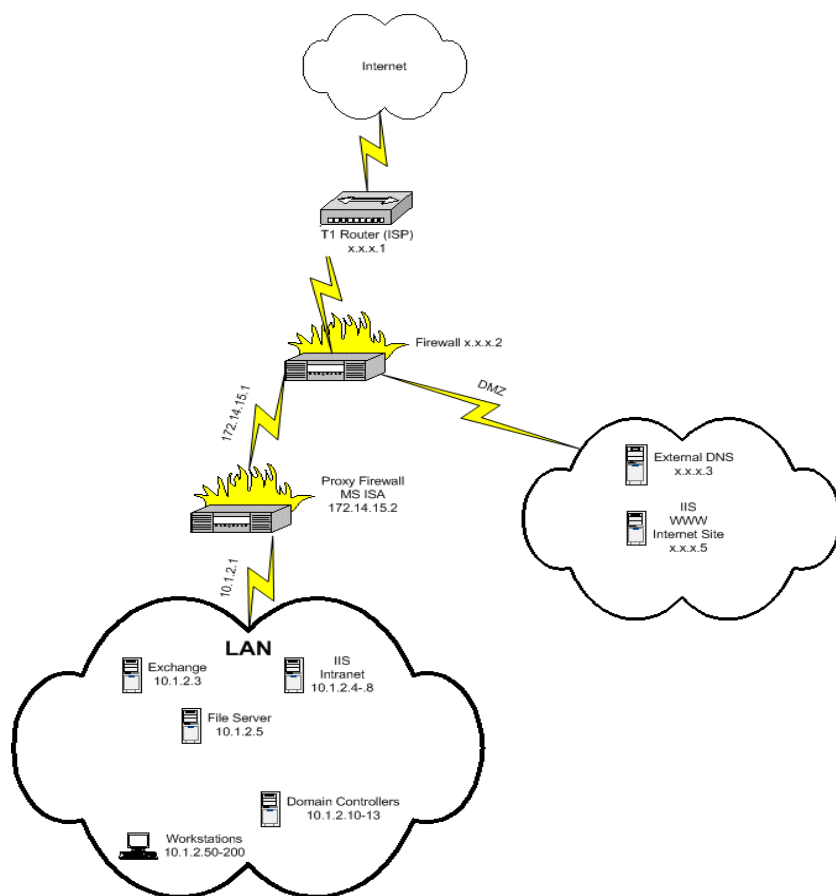


Figure 2

Figure 2 represents the physical layout of the network once the firewalls had been installed and the TCP/IP address scheme had been changed. By moving the public website, www.mycompany.com, to the DMZ we isolated this server from the internal network. Since this server does not host sites used by employees nor does it contain company proprietary information, there was not a need to secure it in the same fashion as other servers, an up-to-date backup was sufficient.

Beyond acting as a content and virus scanner, the MS ISA server provided another critical function, protecting our internal network by acting as a proxy for services. An obvious example is a HTTP proxy, for clients on the LAN, allowing inspection of HTTP traffic from the internet. Malicious code embedded within websites is becoming more prevalent. ISA offered protection from this growing threat. Further discussions of ISA proxies will be expanded in Securing Email.

Hardening

One of the goals was to implement proper hardening techniques. We used publicly available hardening guides from the NSA, SANS and Microsoft. The

recommendations detailed in these guides were only used on newly reformatted servers, which were upgraded to Windows 2000.

Steps taken to harden systems included:

1. Choose a highly complex password with at least 21 characters for the local and domain Administrators account.
2. Disable services not needed but which are turned on by default: Alerter, ClipBook Server, Computer Browser, DHCP Client, Directory Replicator, Messenger, Netlogon, Network DDE, Plug and Play, RPC locator, Server/Workstation, SNMP Trap, TCP NetBIOS Helper and Telephony Service.
3. Disabling default accounts, Guest Account, IUSR or IWAM for all but IIS servers.
4. Use Organizational Units (OU) to segregate users, workstations, servers and Domain Controllers allowing for more granular GPOs.
5. Assigning Windows 2000 GPOs to enforce other security measures such as password complexity, auditing rules and default local security policies.
6. Running tools to further secure the IIS infrastructure, IISlockdown and URLscanner.

Note: The above steps are not meant to provide an end-to-end process in the hardening techniques used, instead an acknowledgement that they are needed and applied. For more detail in the hardening techniques used, please see the references.

The ISA server is the one server which had more specialized hardening applied since it was directly exposed to the internet. Also, since we are relying on it to supply another layer of security to the internal network, this server needed to be addressed separately. The same basic hardening techniques mentioned above were used, but more needed to be done:

- Disable Microsoft Client, File and Printer sharing and Netbios over TCP, on the external network interface card (NIC) and the internal NIC.
- Do not install applications unless they are needed. We found that we did not need any applications from the core OS except Notepad.
- Do not install Terminal Services, IIS, SMTP or other Windows Components which are not absolutely needed.
- Disable services which are not needed Print Spooler, RunAs and Telnet, should not be needed. We disable services one at a time to assess whether or not they are needed.

Many of these steps are detailed in, "ISA Server Security Checklist", by Thomas Shinder. He also makes the point to modify DNS properties for the external NIC or to make registry changes on the ISA server to circumvent this vulnerability:

“The problem with this is that some of the NIMDA and Code Red like Worms are beginning to use **www** instead of an IP address in the HTTP request header. That means that even if you don’t have a Destination Set for **www**, the unqualified **www** request turns into a FQDN by virtue of the DNS interface settings.”




Figure 3

In Figure 3, the correct DNS suffix would be mycompany.com. By substituting a bogus suffix, a call to www does not get passed to www.mycompany.com, instead www.test.local cannot be resolved by DNS and is therefore dropped. Registry changes can be made on the ISA server to not resolve these requests:

HKLM\SYSTEM\CurrentControlSet\Services\W3Proxy\Parameters\

- Value name: SkipNameResolutionForPublishingRules
- Value name: SkipNameResolutionForAccessAndRoutingRules
- Value type: DWORD
- Change value data: 1

Securing Email

The first step in securing our email services was complete, limiting the ports mapped to the email servers and hardening the email server and IIS. Content and virus filtering was added to the interior firewall allowing another line of defense.

The ISA server allowed us to configure Server Secure publishing rules which act as proxies for certain services. A majority of our organization was on the road, outside the LAN, for most of their week; access to email, contacts and calendars was paramount. POP and OWA were used by the users outside the LAN and MAPI was used by internal users. Use of a VPN was considered, but this was not

practical for many users due to restrictions on our clients' networks or workstations.

ISA server allowed us to proxy the email services, providing a layer of insulation from the internet. By enabling POP, SMTP and OWA publishing on the ISA server, we did not need to place an Exchange Front-end server in the DMZ or open direct port access to the internal Exchange server. ISA server publishing also provided some other added benefits:

- An attack aimed at our advertised SMTP server may disable the ISA server, but the internal Exchange server will continue to service internal users.
- If a problem occurs on the internal Exchange server, the SMTP proxy will accept incoming mail until the Exchange server is online.
- Server banners are masked. Connecting to port 25 using telnet no longer advertises our mail server as an Exchange 2000 server.

Although Exchange 2000's version of OWA is greatly improved over Exchange 5.5, many users desired the ability to use a full MAPI client outside the network. This can be accomplished using ISA Server Secure Mail Wizard with the RPC application filter.

We decided to use this publishing service within ISA, but not to open the needed ports on the exterior firewall. Since we have a VPN capable firewall, with the requisite client licenses, we decided on the following course of action. Use a VPN, ending on the LAN interface of the exterior firewall, but not tunneling directly to the LAN. This allows external users to access to the ISA server, where RPC publishing is enabled, but does not give full access to our internal LAN. Now if a client establishes a VPN connection to the external firewall, they are able to use a full MAPI client for messaging.

Using a full MAPI client required one change on the Exchange 2000 server, defining how the client authenticates. Since the ISA server does not publish the Domain Controllers on the internal network and Exchange server requires the client to authenticate with a DC directly, authentication of clients needs to be handled by the Exchange server. By making the registry change below, the Exchange server will proxy domain authentication for the Outlook client:

- HKLM\System\CurrentControlSet\Services\MSExchangeSA\Parameters
 - Value name: No RFR Service
 - Value type: DWORD
 - Change value data: 0x1

Even though we are using an ISA server to publish OWA, we needed to configure this to be more secure. In our original network, OWA connections were authenticated using Basic Authentication. This was an obvious weakness since

username and passwords are sent in the clear. Secure Sockets Layer (SSL) provided the means to overcome this vulnerability.

Using SSL to Secure IIS Services

Although Microsoft Certificate Server would allow us to install the needed certificate services, we chose to use Verisign as a Root Certificate provider for our installation. When using a MS Certificate server, the certificate also needs to be installed on each client computer which connects. Some of our users connect on computers on which they do not have administrative control. Using a Verisign certificate did not require any additional client configuration.

Microsoft article Q320291 explains the steps in detail for installing a certificate for OWA on Exchange 2000 server. This allowed us to create a SSL connection between client computers and the Exchange IIS server, but we are isolated behind an ISA server. There are two methods available for publishing OWA through the ISA server; Server Publishing or Web Publishing. For OWA we chose to use server publishing. Now the ISA server passes SSL traffic directly to the Exchange 2000 server. The steps to setup this publishing rule on ISA:

1. Open the ISA MMC, expand the computer node then the publishing node.
2. Activate the Server Publishing Wizard, in the Action Menu select NEW and RULE. Create a name for the rule, we used OWA Rule.
3. Enter the IP address of the internal Exchange Server, for IP of Internal Server.
4. Enter the IP address of the external NIC on ISA, for IP of External Server.
5. Select the Protocol Setting tab, use the drop-down menu to select HTTPS Server.

Configuring DNS records for webmail.mycompany.com, where the internal DNS record is the IP address of the Exchange server and the external DNS record is the IP of the external NIC of ISA, is all that was left. Now, when a user goes to the URL <https://webmail.mycompany.com> it will result in a secure session for OWA. Now either from the LAN or internet authentication is encrypted, from client to server.

With OWA communicating and authenticating over SSL it was time to turn to our billing and time management servers. All users in the company needed secure access to these servers as well. While Server Publishing was a simple task with ISA, Web Publishing can be a bit more difficult. Using separate IP addresses on the external NIC of the ISA server and configuring DNS records accordingly for each site, greatly simplifies Web Publishing.

Microsoft article Q324167 explains the steps needed to install a certificate on the ISA server. The difference between Server Publishing (used for OWA) and Web Publishing (used for internal websites) is the location of the certificate for SSL.

For access to our internal IIS servers, we chose to only install the certificate on the ISA server. This only allows for SSL communication to the ISA server, HTTP requests are then passed to the internal network. By using SSL bridging we could have an additional layer of encryption, protecting against anyone on the internal network from intercepting usernames or passwords which are being sent in the clear.

Next Steps

The steps taken have offered a secure method for users to access email and internal websites. The main purpose of our project had been completed:

1. Multiple firewalls were used to isolate the LAN and provide application layer filtering to help mitigate viruses and Trojan applications.
2. Server hardening was instituted.
3. Access to mission critical servers was secured using SSL and ISA server.

Vigilance will still be required, there are always going to be next steps. We are investigating the next steps to continue adding layers of security.

Since the original installation of our ISA server, feature release one (FR1) was distributed. Although we were able to use some of the added functionality, late in our deployment, some advantages were not leveraged. We plan to implement a more secure network by making some changes in the near future to include:

- Switch from Server Publishing to Web Publishing for all sites.
 - This will enable us to use URLScanner, included with FR1, to scan all HTTPS for malicious code.
 - Installation of additional certificates will be required on the ISA server and the IIS servers.
 - SSL Bridging will have to be enabled to allow SSL traffic to pass between the ISA server and the IIS servers. This will prevent even internal users from intercepting any sessions and easily reading the data.
- Use authentication forwarding for the secure websites.
 - Currently all authentication is handled at the web servers themselves.
 - FR1 allows the ISA server to authenticate the source and then pass it to the web server for authentication, where authentication can be refused on either server.
 - This allows for a two tier authentication process and further isolates internal servers from malicious attacks.

Investigate a solution to offer users VPN access to the LAN. Currently the VPN ends at the ISA server, to allow full Exchange client functionality. There is a need for some users to securely access the LAN from the internet. The main problem

has been with the Sonic client which was supplied with the firewall. This client will not simply pass through the ISA server. Using the Microsoft (Windows 2000 and XP) VPN client with a properly configured IP Security Policy is our current direction.

We plan to purchase RSA SecureIDs to add another layer of security. These can be used in conjunction with the ISA server for user authentication or with a VPN client. While we test different VPN solutions mentioned above, the goal is to incorporate SecureID. This will enhance our overall security posture by adding a layer to authentication from the internet while providing remote users with “something they have”.

© SANS Institute 2003, Author retains full rights

References

“Address Allocation for Private Internets”, by IETF, March 1994

<http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc1597.html>

“HOW TO: Securely Publish Multiple Web Sites by Using ISA Server in Windows 2000”, by Microsoft Corporation

<http://support.microsoft.com/default.aspx?scid=kb;en-us;Q300435&sd=tech>

“ISA Server Security Checklist”, Dr. Thomas W Shinder

<http://www.tacteam.net/isaserverorg/isachecklist.htm>

“Filtering SMTP Traffic with the SMTP Filter”, by Microsoft Corporation

<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/prodt ech/Windows/SecWin2k/06basewn.asp>

“XCCC: Turning On SSL for Exchange 2000 Server Outlook Web Access”
Q320291, by Microsoft Corporation

<http://support.microsoft.com/default.aspx?scid=kb%3Ben-us%3B320291>

“HOW TO: Export, Install, and Configure Certificates to Internet Security and Acceleration Server”, by Microsoft Corporation

<http://support.microsoft.com/default.aspx?scid=kb;en-us;Q324167>

“Configuring and Securing Microsoft Exchange 2000 Server and Clients”, by Microsoft Corporation

<http://www.microsoft.com/isaserver/techinfo/deployment/ISAandExchange.asp>

“Step-by-Step Guide to Using the Security Configuration Tool Set”, 16 February 2000 , by Microsoft Corporation

<http://www.microsoft.com/windows2000/techinfo/planning/security/secconfsteps.asp>

Timothy M. Mullen, “Hardening Windows 2000 in the Enterprise Part One: Seeing the Forest In Spite of the Trees”, May 21, 2001

<http://online.securityfocus.com/infocus/1296>

Gavin Reid, “Hardening an IIS 4.0 Web Server”, October 2001

<http://www.secadministrator.com/Articles/Index.cfm?ArticleID=22282>

“Securing Microsoft Windows 2000 Server”, by Microsoft Corporation

<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/prodt ech/Windows/SecWin2k/06basewn.asp>

National Security Agency, Security Recommendation Guides, Windows 2000 Guides

<http://nsa1.www.conxion.com/win2k/index.html>