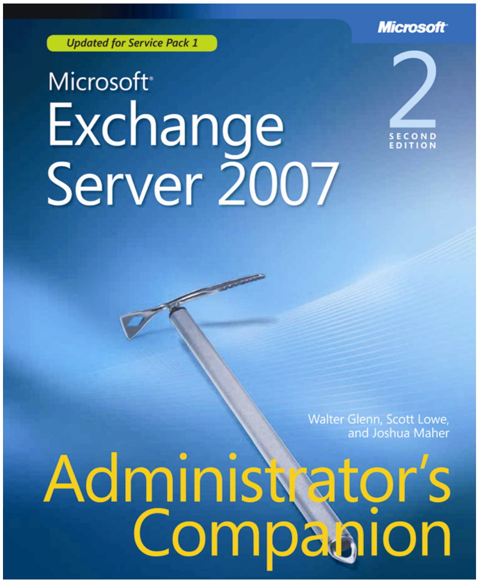

Microsoft® Exchange Server 2007 Administrator's Companion, Second Edition

Walter Glenn, Scott Lowe, and Joshua Maher

To learn more about this book, visit Microsoft Learning at
http://www.microsoft.com/MSPress/books/12754.aspx

9780735625907

Table of Contents

Part I
Introduction

Part II

Planning Your Deployment

Part III
Installation and Deployment

Part IV

Management

Part V

Maintenance

Part VI
Security

Part VII

Clients

Part VIII

Appendices

Chapter 19

Exchange Server Security Basics

Security incidents, including hacking, virus attacks, spyware outbreaks, and identity theft, have rocked the computing world. Due to the e-mail server's reliance on access to the outside world, e-mail has become a target for miscreants everywhere, who try to use this medium to gain access to an organization. As such, security has become so central to the administrator's role that a large portion of this book is devoted to a discussion of it.

This chapter offers ideas about how to add complexity and create hindrances to those who wish to attack your network over port 25. It is never foolproof, but the more you invest in security, the more secure your e-mail server will be. However, if you have good strategies in place and adequate tools to assist you, you can anticipate and thwart most attacks.

Real World Think Globally When Diagnosing a Security Problem

Recently, a U.S. firm with national visibility in its industry was attacked by a group based outside of the U.S. The attacking group used its Exchange server to send out spam messages (in its own language) to addresses all over the world. At first, this problem looked like a virus, but then the company realized the attackers had planted a program on the Exchange server that was launching the outgoing e-mails.

By the time the firm figured out the problem, outbound SMTP queues had nearly 100,000 messages sitting in them, ready to be sent. Besides the obvious concern that the people receiving the spam would be unhappy, there were also a multitude of other negative possible consequences that could have occurred as a result of this problem:

- **A tarnished reputation** By "allowing" this activity to take place, the company proved to those that received the spam that inadequate security measures were being taken. Whether this statement actually reflected reality would be a moot point to those whose perceptions of this company changed.
- **Lawsuits** By sending out spam, the company opened itself up to lawsuits that could prove to be costly and further harm the company's reputation.

The Scope of Security

Everyone has heard the old phrase "a chain is only as strong as its weakest link." You can easily apply that thinking to security: a network is only as secure as its least secured component. Always consider e-mail to be one of those weak links on your network because it is an obvious entry point. Attackers use e-mail to wreak havoc because it's easy: no matter how well you secure your network, chances are good that you have port 25 open on your firewall and that a Simple Mail Transport Protocol (SMTP) server is ready to work with e-mail when it comes in.

When you begin thinking about security strategies, always answer the following question: What am I securing Exchange Server 2007 against? The answers to this question are varied and can be grouped into four categories:

- Protection against social engineering attempts
- Physical security
- Administrative security
- SMTP security

You learned about social engineering in depth in Chapter 18, "Security Policies and Exchange Server 2007." In this chapter, the other three security categories are covered.

Motivations of a Criminal Hacker

Although a lot of literature has been written about the technical aspects of securing a network, not much is available about who your enemies are and what motivates them to attack. Before you can determine how to protect your organization, you must learn to think like a hacker, figure out where you're vulnerable, and then develop a game plan to reduce your exposure. If you can understand who would want to do you harm and what can be gained from such harm, you can better protect your company and your information. Make the following assumptions:

- You do have professional adversaries.
- You are on their target list.
- You will be attacked some day.
- You cannot afford to be complacent.

One of the most difficult realities for an organization to accept is the presence of adversaries who might attempt to harm it by using technology. It's also possible that you really do not have adversaries in this traditional sense. Today, attackers look for any system that has an exploitable weakness that they can turn to their advantage. Often, attackers look at weakly secured systems as bases from which to launch more sophisticated attacks.

The motivations of attackers can be varied and complex. Hackers are often motivated, in part, by their invisibleness. Today's more sophisticated hackers are often also motivated by prospect of a big payday. On the Internet, a hacker can "peek" into a company's private world—its network—and learn a lot while remaining anonymous.

Some individuals are just curious to see what they can learn about your company or individuals within your company. These hackers, sometimes referred to as "script kiddies," often don't have any malicious intent and are unaware that their actions violate security policy or criminal codes. That does not mean that these "casual hackers" are any less dangerous, however.

Others hackers are simply trying to help. You've probably been in this category once or twice yourself. In your zeal to be helpful, you bypass security policies to fix problems or accomplish emergency assignments. You might even believe that your efforts are more efficient than following established guidelines and policies. Nevertheless, the bypassing of known security policies is one element of hacking a network.

Some individuals act with malicious intent, engaging in acts of sabotage, espionage, or other criminal activities. They can become moles, stealing information to sell to competitors or foreign groups. Some simply enjoy destroying the work of others as well as their own work. Others act out of revenge for a real or perceived wrong committed against them, or believe they are acting in line with a strongly held belief system. Still others are more methodical and hardened and turn hacking into a career: they might even take employment just to do your company harm.

How Hackers Work

Hackers start by learning that an e-mail server exists, which generic scanning tools can tell them. Coupled with the public information of your Domain Name System (DNS) records, hackers can quickly know a lot about your network.

Finding company information is easy for anyone. You can do it. Simply open a command prompt and type **nslook**. Set the type of the record you're looking for to a mail exchanger (MX) record by typing **set type=mx**. Type a domain name. This example uses Microsoft .com. Figure 19-1 shows the results.

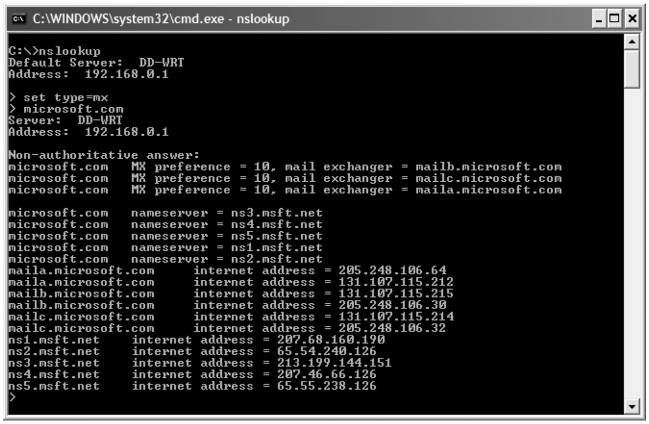

```
C:\>nslookup
Default Server:  DD-WRT
Address:  192.168.0.1

> set type=mx
> microsoft.com
Server:  DD-WRT
Address:  192.168.0.1

Non-authoritative answer:
microsoft.com   MX preference = 10, mail exchanger = mailb.microsoft.com
microsoft.com   MX preference = 10, mail exchanger = mailc.microsoft.com
microsoft.com   MX preference = 10, mail exchanger = maila.microsoft.com

microsoft.com   nameserver = ns3.msft.net
microsoft.com   nameserver = ns4.msft.net
microsoft.com   nameserver = ns5.msft.net
microsoft.com   nameserver = ns1.msft.net
microsoft.com   nameserver = ns2.msft.net
maila.microsoft.com     internet address = 205.248.106.64
maila.microsoft.com     internet address = 131.107.115.212
mailb.microsoft.com     internet address = 131.107.115.215
mailb.microsoft.com     internet address = 205.248.106.30
mailc.microsoft.com     internet address = 131.107.115.214
mailc.microsoft.com     internet address = 205.248.106.32
ns1.msft.net    internet address = 207.68.160.190
ns2.msft.net    internet address = 65.54.240.126
ns3.msft.net    internet address = 213.199.144.151
ns4.msft.net    internet address = 207.46.66.126
ns5.msft.net    internet address = 65.55.238.126
>
```

Figure 19-1 Using the NSLookup tool to find the public MX records for Microsoft.com

Next, the hacker determines the platform of your SMTP server in one of two ways. In the first approach, the hacker can use Telnet to open a session to your server over port 25 and then read the banner. Under Exchange Server 2007, the banner no longer identifies the version of Exchange Server being run, but does still indicate that the server is running the Microsoft ESMTP service. By removing the version number, Microsoft makes it harder for hackers to determine the exact version of Exchange that you are using. Note, Exchange Server 2007 is the only version that, by default, lacks this identifying information. However, a hacker can still figure out what he wants to know. It will take a couple of service packs and another major version of Exchange before this default omission really begins to bear fruit. Figure 19-2 gives you a look at an ESMTP conversation that takes place with an Exchange Server 2007 server.

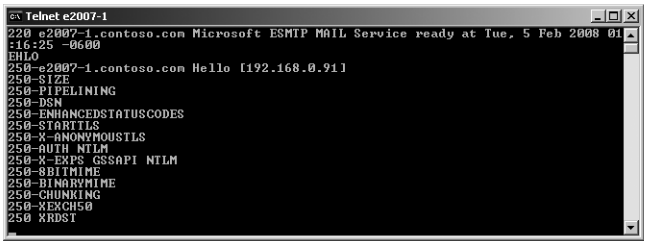

```
220 e2007-1.contoso.com Microsoft ESMTP MAIL Service ready at Tue, 5 Feb 2008 01
:16:25 -0600
EHLO
250-e2007-1.contoso.com Hello [192.168.0.91]
250-SIZE
250-PIPELINING
250-DSN
250-ENHANCEDSTATUSCODES
250-STARTTLS
250-X-ANONYMOUSTLS
250-AUTH NTLM
250-X-EXPS GSSAPI NTLM
250-8BITMIME
250-BINARYMIME
250-CHUNKING
250-XEXCH50
250 XRDST
```

Figure 19-2 Opening a Telnet session to a server running Exchange Server 2007

Under older versions of Exchange Server, the exact version of the Exchange server being run is displayed (see Figure 19-3). The main version number, 6.0, means Exchange Server 2003. An Exchange 2000 Server registers with a main version number of 5.0. A SendMail server has its name and the version of SendMail software used by the company displayed in the header as well as the operating system (OS). Using this kind of information, a hacker can target his efforts by looking for exploits that will work for your specific system.

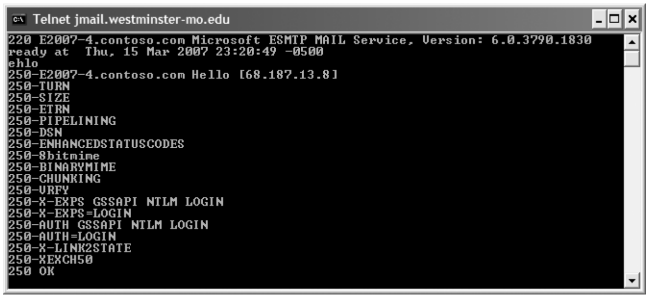

```
220 E2007-4.contoso.com Microsoft ESMTP MAIL Service, Version: 6.0.3790.1830
ready at  Thu, 15 Mar 2007 23:20:49 -0500
ehlo
250-E2007-4.contoso.com Hello [68.187.13.8]
250-TURN
250-SIZE
250-ETRN
250-PIPELINING
250-DSN
250-ENHANCEDSTATUSCODES
250-8bitmime
250-BINARYMIME
250-CHUNKING
250-VRFY
250-X-EXPS GSSAPI NTLM LOGIN
250-X-EXPS=LOGIN
250-AUTH GSSAPI NTLM LOGIN
250-AUTH=LOGIN
250-X-LINK2STATE
250-XEXCH50
250 OK
```

Figure 19-3 Opening a Telnet session to a server running Exchange Server 2003

More Info Although Exchange Server 2007 is the first version of Exchange Server that, by default, does not display versioning information in a telnet window, you can manually configure older versions of Exchange Server to act the same way. Refer to *http://technet.microsoft.com/en-us/library/bb124740.aspx* for more information.

The second way to determine your e-mail server platform is to send a bogus e-mail to your server. This is accomplished by sending a message to an unlikely e-mail address such as pancake@contoso.com. The nondelivery report (NDR) that is returned has the e-mail server information located somewhere in the NDR. The following sample is a message header sent to the lab Exchange server at contoso.com. Notice that the Exchange server version is included right in the NDR's Sent by line:

```
Delivery has failed to these recipients or distribution lists:
 pancake@contoso.com
The recipient's e-mail address was not found in the recipient's e-mail system.
Microsoft Exchange will not try to redeliver this message for you.
Please check the e-mail address and try resending this message, or provide the
following diagnostic text to your system administrator.
--------------------------------------------------------------------------------
Sent by Microsoft Exchange Server 2007
Diagnostic information for administrators:
Generating server: e2007-1.contoso.com
pancake@contoso.com
#550 5.1.1 RESOLVER.ADR.RecipNotFound; not found ##
Original message headers:
Received: from e2007-1.contoso.com ([192.168.0.91]) by e2007-1.contoso.com
 ([192.168.0.91]) with mapi; Tue, 5 Feb 2008 01:25:12 -0600
Content-Type: application/ms-tnef; name="winmail.dat"
Content-Transfer-Encoding: binary
From: Cat Francis <cat.francis@contoso.com>
To: "pancake@contoso.com" <pancake@contoso.com>
Date: Tue, 5 Feb 2008 01:25:06 -0600
Subject: Test
Thread-Topic: Test
Thread-Index: AQHIZ8g79IUM/OhzKk2PKwL9+dATWg==
Message-ID: <1772808B96DEC14094F0236A00882DD7A43089@e2007-1.contoso.com>
Accept-Language: en-US
Content-Language: en-US
X-MS-Has-Attach:
X-MS-TNEF-Correlator: <1772808B96DEC14094F0236A00882DD7A43089@e2007-
1.contoso.com>
MIME-Version: 1.0
```

Note that, even with Service Pack 1 installed, NDR messages still report just Exchange Server 2007 as the server version. By looking at just the NDR, there is no indication that SP1 is deployed on the server.

Now that the hacker knows which e-mail server software you run, he or she checks known databases to find vulnerabilities to exploit. The known vulnerabilities for Exchange Server 2007 are listed in Microsoft's Security Bulletins and can be found at *www.microsoft.com/security/default.mspx*. On older versions of Exchange, some of the vulnerabilities could involve Microsoft Internet Information Services (IIS) because IIS managed the SMTP service for Exchange. In Exchange Server 2007, SMTP is a core part of Exchange itself, which helps to reduce the attack potential on your server. Other vulnerabilities may involve Microsoft Outlook Web Access (OWA), again because of the involvement of IIS managing the HTTP connectivity to the Exchange server. At a minimum, be aware of any vulnerabilities that exist for Exchange Server 2007 and test and install the updates when they are released.

Generally speaking, the e-mail administrator can expect the following kinds of attacks:

- **Buffer overflows** Buffer overflows send a larger quantity of data to the server than is anticipated. Depending on how the overflow is executed, it could cause the server to stop working or it might run malicious code from the attacker.
- **Data processing errors** These are not common currently, but the concept is that a small program is sent directly to the server and the server runs it. More common today is sending these programs to a network though e-mail as attachments. Depending on their function and purpose, these programs can be viruses, Trojans, or worms (discussed at length later in this chapter).
- **HTML viruses** These do not require user intervention to run unattended scripts.
- **Custom programs written to run against port 25 (SMTP)** The more common types of programs that attack port 25 include e-mail flooding programs or programs that contain their own SMTP engine that use the port for their own malicious purposes.
- **Denial of Service (DoS)** A Denial of Service attack is an attack on a network that is undertaken in an effort to disrupt the services provided by a network or server.
- **Cross-site scripting** Cross-site scripting is a vulnerability whereby an attacker places malicious code into a link that appears to be from a trusted source.
- **Spam and phishing expeditions** Spam, or junk mail, is a well-known e-mail malady and affects just about everyone that uses the communication medium. A particular type of spam, called a phishing e-mail, attempts to lure unsuspecting users into clicking on unsafe web links. These links point to web forms that ask the user to provide sensitive personal information.

Here are some broad actions you can take to guard against the attacks just described, plus others:

- **Physical access to the server** Lock the doors, and use some type of biotech authentication.
- **Viruses, Trojans, and worms** Use antivirus software and regularly scan your servers and workstations. Use the Exchange Server 2007 Edge Transport server role on at least one Exchange server.
- **Loss of data** Perform regular backups.
- **Unauthorized use of user accounts** Conduct user training on information security policies and require complex passwords.
- **Denial of Service attack** Harden the TCP/IP stack and the router.
- **Platform vulnerabilities** Install all software patches and engage in service that offers minimization. Microsoft has released excellent free software for updating its patches on your servers. This software is called Windows Server Update Services (WSUS).

More Info A discussion of WSUS is outside the scope of this chapter, but you can learn more about WSUS on Microsoft's Web site at *http://www.microsoft.com/windowsserversystem/updateservices/default.mspx*

As is the case with Windows Server 2003, WSUS is not installed by default on Windows Server 2008, but is still available as a free download.

The rest of this chapter is intended to help you secure Exchange Server 2007 against these types of attacks. However, a brief discussion of physical security of your Exchange server is in order.

Physical Security

Physical security is a topic not often mentioned in many security books, particularly in books only about Exchange, but it is a topic worth mentioning. Servers can be left on desks running in a corner cubicle or in an unlocked server room. However, it is always best practice to store your servers in a secure location using door locks and, in some instances, motion detectors and/or other physical security measures.

When you limit physical access to a server, you limit who can log on locally to the server, who can use portable storage to introduce a new virus or malicious program on your network, and who can retrieve information directly from the server. Limiting physical access is one of the easiest and most elementary methods of securing your server against internal attacks that exist.

Most administrators reading this book already have these physical security measures in place. Those who haven't physically secured your servers should do so at their earliest opportunity. Limiting physical access to a server can go a long way toward protecting your information from would-be attackers.

Administrative Security

In previous versions of this book, this section talked extensively about the use of administrative groups as a way to achieve some semblance of administrative security for your Exchange organization. In Exchange Server 2007, however, Microsoft has mostly done away with administrative groups, leaving only a single administrative group named Exchange Administrative Group (FYDIBOHF23SPDLT) in which only Exchange Server 2007 servers reside. This administrative group is present only to support coexistence with legacy Exchange servers.

Note The name of the Exchange administrative group, Exchange Administrative Group (FYDIBOHF23SPDLT), is pretty convoluted. Likewise, Exchange Server

2007's legacy routing group, named Exchange Routing Group (DWBGZMFD01QNBJR), is also fairly convoluted. Have you wondered at all why Microsoft chose these particular names? First, Microsoft had to be careful that it didn't choose a name that already exists in a customer's legacy Exchange organization. Second, the Exchange team decided that a little creativity was in order. Look carefully at the two names. Both have the same number of characters with each letter and number occupying the same positions. To make a long story short, if you look at the administrative group's name, you find you can go to the previous letter (or number) in the alphabet for each character in the name and spell "EXCHANGE12ROCKS." Likewise, for the routing group, go to the next letter of the alphabet for each letter in the routing group name and you also get "EXCHANGE12ROCKS." It's really nice to see the product team having so much fun with a product that is generally considered all business!

Why did the Exchange team eliminate administrative groups from the Exchange equation? With the complete overhaul of the management interface and its new "area of responsibility" focus, administrative groups simply aren't necessary and can add to the overall complexity of managing Exchange. Figure 19-4 gives you a side-by-side look at the legacy Exchange System Manager and the Exchange Server 2007 Exchange Management console. With their absence in Exchange Server 2007, you need to use a way other than administrative groups to achieve administrative security. In this section, you learn two methods by which you can add users to act in various Exchange administrative capacities.

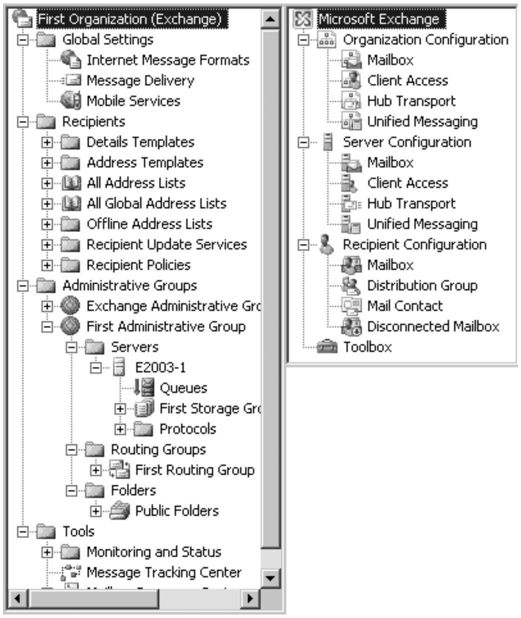

Figure 19-4 The Exchange Server 2003 Exchange System Manager is on the left and the Exchange Server 2007 Exchange Management Console is on the right.

The Built-in Exchange Administrative Groups

When you run the initial installation of Exchange Server 2007, six Active Directory universal security groups are created, each with specific rights to various parts of the Exchange organization. Five of the six groups, shown in Figure 19-5 inside Active Directory Users And Computers, pertain directly to management of the Exchange organization and are as follows:

- **Exchange View-Only Administrators** This role allows you to view configurations on all Exchange objects, but not to make any changes to those configurations.
- **Exchange Servers** This role provides the following rights:
    - ❑ Members of this group have all of the rights of Exchange View-Only Administrators.
    - ❑ Members of this group have access to server-based Exchange configuration information and to the Active Directory objects that are server-related.
    - ❑ Members of this group may perform server-based administration, but cannot perform operations at the global Exchange organization level.
    - ❑ Members of this group are also members of the local Administrators group on each server on which Exchange Server 2007 is installed.
- **Exchange Recipient Administrators** This role provides the following rights:
    - ❑ Members of this group have all of the rights of Exchange View-Only Administrators.
    - ❑ Members of the group are also allowed to configure any object related to recipients and public folders, including contacts, groups, public folder objects, Unified Messaging mailbox settings, Client Access mailbox settings, and any other recipient Exchange property found in Active Directory.
- **Exchange Public Folder Administrators** This role provides the following rights:
    - ❑ Members of this group have all of the rights of Exchange View-Only Administrators.
    - ❑ Members of this group are also allowed to manage public folders.
- **Exchange Organization Administrators** This role provides the following rights:
    - ❑ Members of this group have all of the rights of Exchange Recipient Administrators, plus more.
    - ❑ Members of this group also have all of the rights of Exchange Public Folder Administrators.

- Users assigned to this group are allowed to view and administer all aspects of the Exchange organization, including servers, recipients, public folders, and organizational configuration.
- Members of the role are considered the owners of all Exchange-related Active Directory objects.
- During Exchange Server 2007 installation, this group is added to the membership of the server's local Administrators group. If you install Exchange Server 2007 on a domain controller, which is not recommended, Exchange Organization Administrators have additional rights by virtue of the local Administrators group having more rights on a domain controller.

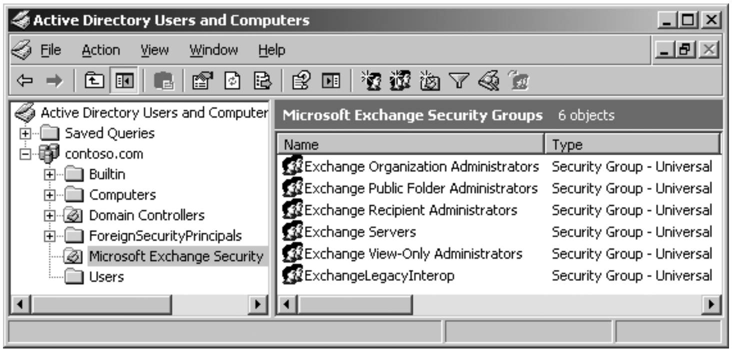
Active Directory Users and Computers
File
Action
View
Window
Help
Active Directory Users and Computer
Saved Queries
contoso.com
Builtin
Computers
Domain Controllers
ForeignSecurityPrincipals
Microsoft Exchange Security
Users
Microsoft Exchange Security Groups 6 objects

Name	Type
Exchange Organization Administrators	Security Group - Universal
Exchange Public Folder Administrators	Security Group - Universal
Exchange Recipient Administrators	Security Group - Universal
Exchange Servers	Security Group - Universal
Exchange View-Only Administrators	Security Group - Universal
ExchangeLegacyInterop	Security Group - Universal

Figure 19-5 The Exchange Server 2007 built-in security groups

If you want to add a full Exchange administrator to your organization, all you have to do is add the appropriate user account to the Exchange Organization Administrators group. The same holds true for the other security groups.

The Add Exchange Administrator Wizard

Exchange Server 2007 also provides an easy way to add additional Exchange administrators with each administrator role having responsibility for only a specific part of the Exchange organization, such as a single server, a group of servers, or only able to manage recipients. You will find that this administrative delegation method is far more flexible and effective than administrative groups were in the past.

The best way to demonstrate how the Add Exchange Administrator operation works is to see it in action. To start the process, open the Exchange Management Console and select the Organization Configuration option, as shown in Figure 19-6.

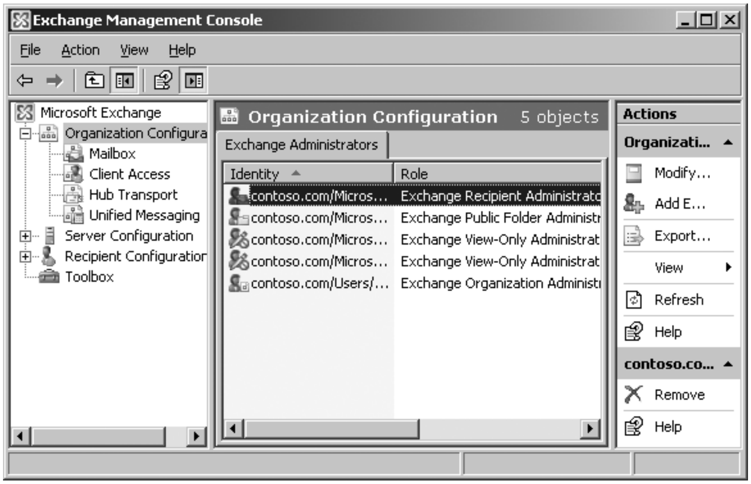

Identity	Role
contoso.com/Micros...	Exchange Recipient Administratc
contoso.com/Micros...	Exchange Public Folder Administr
contoso.com/Micros...	Exchange View-Only Administrat
contoso.com/Micros...	Exchange View-Only Administrat
contoso.com/Users/...	Exchange Organization Administr

Figure 19-6 The Organization Configuration window

Note that the work pane shown in Figure 19-6 shows you the groups that already have some level of permission to the Exchange organization. To add additional Exchange administrators, from the Action pane, choose Add Exchange Administrator. This selection displays a one-page wizard, shown in Figure 19-7.

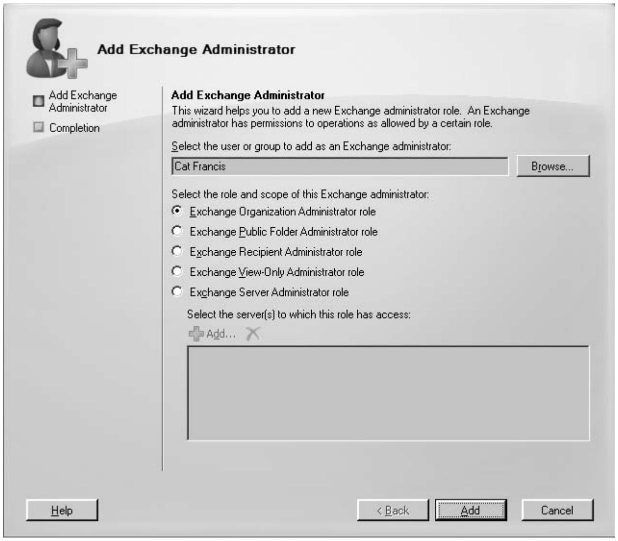

Add Exchange Administrator

This wizard helps you to add a new Exchange administrator role. An Exchange administrator has permissions to operations as allowed by a certain role.

Select the user or group to add as an Exchange administrator: Cat Francis

Select the role and scope of this Exchange administrator:

- ● Exchange Organization Administrator role
- ○ Exchange Public Folder Administrator role
- ○ Exchange Recipient Administrator role
- ○ Exchange View-Only Administrator role
- ○ Exchange Server Administrator role

Select the server(s) to which this role has access:

Figure 19-7 The Add Exchange Administrator Wizard

There are three selections that you must make in order to complete this wizard.

First, select the user or group to which you want to grant Exchange administrative rights. Next, select the role and scope that should apply to the new Exchange administrator. Finally, if you've selected the Exchange Server Administrator role, select at least one server to which this new user or group has access. Click Add, and from the Select Exchange Server window, choose the desired servers. Figure 19-8 shows what the screen looks like after you select the Exchange Server Administrator role and add a managed server.

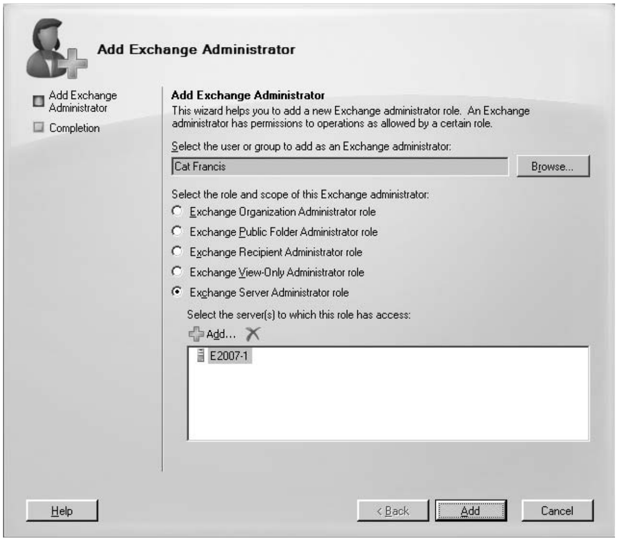

Add Exchange Administrator

This wizard helps you to add a new Exchange administrator role. An Exchange administrator has permissions to operations as allowed by a certain role.

Select the user or group to add as an Exchange administrator: Cat Francis

Select the role and scope of this Exchange administrator:

- ○ Exchange Organization Administrator role
- ○ Exchange Public Folder Administrator role
- ○ Exchange Recipient Administrator role
- ○ Exchange View-Only Administrator role
- ● Exchange Server Administrator role

Select the server(s) to which this role has access:

E2007-1

Figure 19-8 Selecting the Exchange Server Administrator role

Note When you add someone to the Exchange Server Administrator role, you must manually add that user or group to each managed server's local Administrators group.

In reality, when you run the Add Exchange Administrator operation, the resulting command simply adds the selected users to one of the groups that you learned about in the section "The Built-in Exchange Administrative Groups." The only role for which this does not hold true is for the Exchange Server Administrator role. When users or groups are assigned to this role, the user or group is assigned Full Control permission on the specified server object and all child objects.

Management Shell

Adding additional users or groups to manage your Exchange organization in Exchange Server 2007 is a whole lot easier than it ever was in previous versions of Exchange Server. During the initial Exchange Server 2007 installation, a number of universal security groups are created. Each of these groups corresponds to security roles that can be granted in Exchange Server 2007 and are listed here:

- **Exchange View-Only Administrators** This role allows you to view configurations on all Exchange objects, but not to make any changes to those configurations.
- **Exchange Servers** This role provides the following rights:
    - ❑ Members of this group have all of the rights of Exchange View-Only Administrators.
    - ❑ Members of this group have access to server-based Exchange configuration information and to the Active Directory objects that are server-related.
    - ❑ Members of this group may perform server-based administration, but cannot perform operations at the global Exchange organization level.
    - ❑ Members of this group are also members of the local Administrators group on each server on which Exchange Server 2007 is installed.
- **Exchange Recipient Administrators** This role provides the following rights:
    - ❑ Members of this group have all of the rights of Exchange View-Only Administrators.
    - ❑ Members of the group are also allowed to configure any object related to recipients and public folders, including contacts, groups, public folder objects, Unified Messaging mailbox settings, Client Access mailbox settings, and any other recipient Exchange property found in Active Directory.
- **Exchange Public Folder Administrators** This role provides the following rights:
    - ❑ Members of this group have all of the rights of Exchange View-Only Administrators.
    - ❑ Members of this group are also allowed to manage public folders.
- **Exchange Organization Administrators** This role provides the following rights:

- Members of this group have all of the rights of Exchange Recipient Administrators, plus more.
- Members of this group also have all of the rights of Exchange Public Folder Administrators.
- Users assigned to this group are allowed to view and administer all aspects of the Exchange organization, including servers, recipients, public folders, and organizational configuration.
- Members of the role are considered the owners of all Exchange-related Active Directory objects.
- During Exchange Server 2007 installation, this group is added to the membership of the server's local Administrators group. If you install Exchange Server 2007 on a domain controller, which is not recommended, Exchange Organization Administrators will have additional rights by virtue of the local Administrators group having more rights on a domain controller.

The following command adds a user account that can manage the Exchange Server 2007 server named E2007-4:

```
Add-ExchangeAdministrator -Identity 'contoso.com/Users/David So'
   -Role 'ServerAdmin' -Scope 'E2007-4'
```

If you add someone using Exchange Server Administrator role, you need to manually add the selected user or group to the built-in local administrators group on the target server.

This command adds a user to the Exchange Recipient Administrators role:

```
Add-ExchangeAdministrator -Identity 'contoso.com/Users/David So'
   -Role 'RecipientAdmin'
```

This command adds a user to the Exchange View-Only Administrators role:

```
Add-ExchangeAdministrator -Identity 'contoso.com/Users/David So'
   -Role 'ViewOnlyAdmin'
```

This command adds a user to the Exchange Organization Administrators role:

```
Add-ExchangeAdministrator -Identity 'contoso.com/Users/David So'
   -Role 'OrgAdmin'
```

This command adds a user to the Exchange Public Folder Administrators role:

```
Add-ExchangeAdministrator -Identity 'contoso.com/Users/David So'
   -Role 'PublicFolderAdmin'
```

Table 19-1 comes from Microsoft's documentation on the role of roles in Exchange Server 2007 and provides a concise look at exactly what each administrative role accomplishes.

Table 19-1 Exchange Server Administrative Roles

Role	Members	Member of	Exchange permissions
Exchange Organization Administrators	Administrator, or the account that was used to install the first Exchange 2007 server	Exchange Recipient Administrator, Administrators local group of <*Server Name*>	Full control of the Microsoft Exchange container in Active Directory
Exchange Recipient Administrators	Exchange Organization Administrators	Exchange View-Only Administrators	Full control of Exchange properties on Active Directory user object
Exchange Server Administrators		Exchange View-Only Administrators, Administrators local group of <*Server Name*>	Full control of Exchange <*Server Name*>
Exchange View-Only Administrators	Exchange Recipient Administrators, Exchange Public Folder Administrators	Exchange Recipient Administrators, Exchange Server Administrators	Read access to the Microsoft Exchange container in Active Directory. Read access to all the Windows domains that have Exchange recipients.
Exchange Public Folder Administrators	Exchange Organization Administrators	Exchange View-Only Administrators	Ability to administratively manage public folders.

SMTP Security

By default, an SMTP server attempts to make a TCP port 25 connection to your Exchange server via an anonymous connection. Anonymous does not mean that a user account set up in your Active Directory proxies the connection request, as is the case with the IIS Anonymous user account, IUSR_<*machinename*>. In the SMTP world, anonymous means that no user name or password is required for the remote SMTP service to make a port 25 connection. Hence, any SMTP server on the Internet can make, by default, a port 25 connection to your Exchange server.

To make SMTP more secure, you could require either Basic or Integrated Windows Authentication (IWA) before the SMTP Virtual Server (VS) could accept an inbound connection. But this configuration isn't practical on the Internet because you can't predict who will be connecting to your Exchange server in the future and thus can't assume that the user has an appropriate user name and password to make a connection. Moreover, not many messaging administrators are interested in implementing such a security measure at their end. So even though an anonymous connection to port 25 on your Exchange server represents a vulnerability, it is one that must be managed using a different approach than removing anonymous connections.

How do you protect against these kinds of attacks? With Exchange Server 2007, you can use an Edge Transport server that offloads the security burden from your primary Exchange servers. You learn about implementing the Edge Transport server in Chapter 20, "Antivirus and Anti-spam." This chapter also discusses how the Edge Transport server can help improve the overall security of your Exchange infrastructure. However, more traditional ways of protecting Exchange also apply even when Edge Transport servers are used.

Perhaps the most common way to protect an Exchange infrastructure is through the use of two firewalls. A dual firewall topology allows you to protect your internal Exchange servers while also filtering incoming e-mail against potential attacks. The area between the two firewalls is called the *perimeter network*. The philosophy is to put up a line of defense against potential attacks. Hence, you're willing to sacrifice your Exchange servers in the perimeter network, but not willing to sacrifice your Exchange servers on the internal network. Because the Exchange servers in the perimeter network do not host any important information—no mailboxes or public folders—they can be both sacrificed during an attack and easily rebuilt. And because they act only as relay servers, they can be used to sanitize incoming e-mail over port 25.

Take a look at Figure 19-9. Note that there are three network levels. Starting from the top, each network becomes more trusted, with the External, or Internet, zone being completely untrusted. The Perimeter network is more trusted as it resides behind at least one organizational firewall and generally houses servers that can be considered "expendable." In this diagram, the external firewall has port 25 open in order to facilitate incoming SMTP traffic. Mail is routed to the Exchange Server 2007 Edge Transport server where it is processed for viruses, checked using various spam filters, and run through various incoming transport rules. Your external MX records must point to this Edge Transport server. There is another important note in this diagram. Note that the external firewall also provides the ability to scan incoming content for viruses and spyware. When possible, always run your e-mail through a similarly configured firewall even before that mail hits the Edge Transport server's content-scanning engines. Many of today's security appliances, such as the Cisco ASA and Sonicwall's family of firewalls, provide this additional protection.

From a software perspective, also consider running Microsoft Forefront Security for Exchange Server. Forefront has the ability to scan every incoming message with up to five completely separate virus scanners. By instituting this multilayer security infrastructure, all incoming mail is scanned by many different virus scanning engines, some hardwarebased and some software-based, which results in a much higher likelihood you will be protected against even the newest viruses.

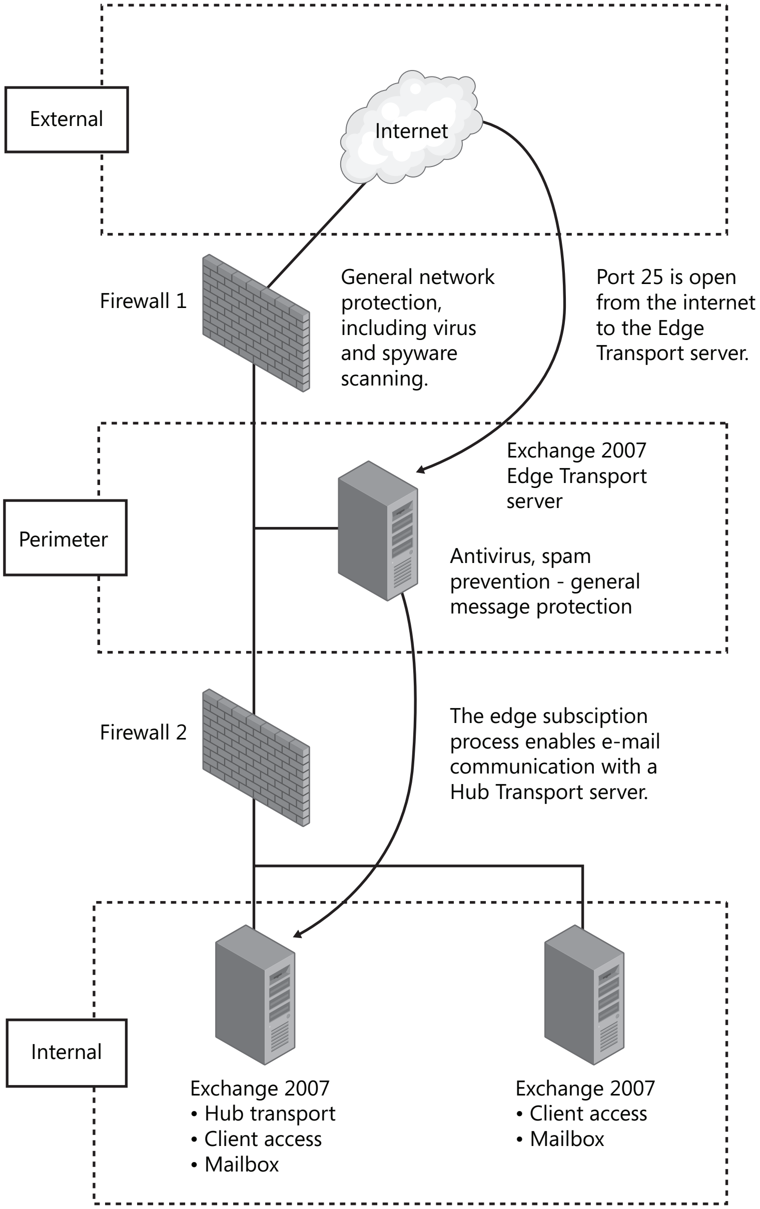
External
Internet
Firewall 1
General network protection, including virus and spyware scanning.
Port 25 is open from the internet to the Edge Transport server.
Perimeter
Exchange 2007 Edge Transport server
Antivirus, spam prevention - general message protection
Firewall 2
The edge subsciption process enables e-mail communication with a Hub Transport server.
Internal
Exchange 2007
• Hub transport
• Client access
• Mailbox
Exchange 2007
• Client access
• Mailbox

Figure 19-9 One way to secure your Exchange infrastructure

However, even the best virus-scanning infrastructure on the planet does not always protect you. Think back to some of the major viruses in the last few years, which were able to spread worldwide very quickly, usually in a matter of hours. It is almost impossible for any antivirus company to get the virus, study it, write a definition for it, and then push out the new definition for that virus before it spreads worldwide. You can tell an Edge Transport server, however, to quarantine or delete any message that contains certain types of attachments and, in effect, block most viruses based on their type of content rather than on a comparison to a virus definition file.essent.

Note Be aware of two issues regarding traditional antivirus servers. First, many products offered by the major antivirus vendors perform content scanning at the same time as the virus scanning. While there may be no problem with this method of scanning e-mail, be aware of a distinction between content scanning and antivirus scanning, which highlights the need to perform both types of scanning in the perimeter network, a capability enabled through the use of Exchange Server 2007's Edge Transport server. Second, everyone may not be able to afford to purchase everything required in order to achieve the configuration outlined in this chapter—namely, a separate Exchange server running Edge Transport as well as firewalls/security appliances that perform virus scanning functions. These ideas are presented to highlight the concepts being discussed. Other, less expensive (and potentially less secure) options include:

- Using a single firewall with multiple interfaces and creating a perimeter network using firewall rules
- Using a single firewall and running the Edge Transport server on the internal interface alongside your other Exchange servers
- Skipping the installation of the Edge Transport server and delivering mail directly to an internal Hub Transport server

Once scanned and approved, the e-mail is sent to an internal Hub Transport server. The internal Exchange Server 2007 Hub Transport server should be configured to accept inbound e-mail only from the perimeter network's Edge Transport server. Inbound mail that has been approved by the Edge Transport server also rides on the standard SMTP TCP port 25, so you need to open this port on your internal firewall as well. To do this in the most secure way possible, create a firewall rule that only allows port 25 traffic specifically between the Edge Transport server and one of your internal Hub Transport servers. Then, secure the communication tunnel using IPsec, which is discussed further in Chapter 21. The internal Exchange server should also be running its own antivirus software, preferably from a vendor that is different from the one the servers are using in the perimeter network. The whole point of implementing this model is to ensure that port 25 traffic is as well protected as possible.

In order to use an Edge Transport server, subscribe the Edge Transport server to the Active Directory domain. The subscription process establishes one-way replication of recipient and configuration information from your Active Directory into an Active Directory Application Mode (ADAM) instance running on your Edge Transport server. Further, the Edge Subscription process creates the SMTP Send connectors required to enable mail flow from your Exchange servers to the Internet through an Edge Transport server. If you are using the recipient lookup or safe list aggregation features of the Edge Transport server, subscribe the Edge Transport server to the organization.

More Info The complete process for installing, configuring, and subscribing the Edge Transport services is covered in Chapter 20, "Antivirus and Anti-spam."

No system is foolproof, but this dual firewall topology has multiple advantages:

- By passing incoming e-mail through the Edge Transport servers content filtering services, you filter for code types that virus scanners don't.
- By passing your e-mail through a virus scanner, you do your best to ensure that all known viruses are cleaned out. Not passing your e-mail through an updated antivirus scanner after running it through a content scanner is unwise because older viruses might not be caught by the content scanner.
- By passing all of your outgoing e-mail through the Exchange Server 2007 Edge Transport server, the IP address (private or public) of the internal Exchange Server 2007 server does not need to be published in the public DNS records. This means that an attacker attempting to Telnet into your server is never able to reach it directly. Also, if you configure the internal Exchange Server 2007 server to accept e-mail only from perimeter network–based Exchange servers, any attempts to make port 25 connections to the internal Exchange server from any other IP address will fail.

If a hacker decides to bring down your perimeter Exchange servers, you've really lost nothing of value other than your time in getting the servers functioning again. Your company might lose some money due to the inability to communicate via e-mail, but it hasn't lost any current data. This is an important point. The server that hosts your data is the one most protected. And the ones most exposed do not host important data. If those servers are lost, at least all the business-critical data is saved on the internal Exchange Server 2007 Server. For many companies, this is an acceptable level of risk to assume. This is the beginning stage of a defense that provides multiple layers of protection, starting with expendable services with the really important data protected in a variety of different ways.

As explained throughout this chapter, no answer is perfect, and this security scenario does have a few major holes, such as doing nothing to protect against messages sent to the Exchange server via Outlook Web Access. Port 25 is well protected but port 80 access to your Exchange server is wide open. If you want to learn more about OWA, refer to Chapter 24, "Supporting Outlook Web Access."

The second major hole in this model is one that cannot be plugged: messages are continuing to flow to your internal Exchange server. As long as a packet can reach your internal Exchange server, there is always the potential for harm. So remember the 80 percent rule: you can make your data only about 80 percent secure. But don't let that discourage you from implementing appropriate security strategies.

Computer Viruses

This section expands on computer viruses in general and discusses some implications for viruses on Exchange Server 2007.

What Is a Virus?

A *virus* is a piece of code that attaches itself to other programs or files. When these files run, the code is invoked and begins replicating itself. The replication occurs over the network. Viruses can now exploit the vulnerabilities of nearly every platform.

Some viruses reside in memory after the original program is shut down. When other programs are executed, the virus attaches itself to these new programs until the computer is shut down or turned off. Some viruses have a "dormant" phase and appear only at certain times or when certain actions are performed.

There are many types of viruses. Some overwrite existing code or data. Others include the ability to recognize whether an executable file is already infected. *Self-recognition* is required if the virus is to avoid multiple infections of a single executable, which can cause excessive growth in size of infected executables and corresponding excessive storage space, contributing to the detection of the virus.

Resident viruses install themselves as part of the operating system upon execution of an infected host program. The virus remains resident until the system is shut down. Once installed in memory, a resident virus is available to infect all suitable hosts that are accessed.

A *stealth virus* is a resident virus that attempts to evade detection by concealing its presence in infected files. For example, a stealth virus might remove the virus code from an executable when it is read (rather than executed) so that an antivirus software package sees only the noncompromised form of the executable.

Computer viruses can spread by the use of e-mail and usually appear in e-mail attachments. If the virus can find its way into the messaging stream, it uses the client capability to send and receive e-mail to replicate itself quickly and do its damage as fast as possible.

An essential aspect of protecting your messaging system against viruses is user education. Users should learn to be guarded about which attachments they are allowed to open. Your information security policies should also outline the types of e-mails and attachments that users are allowed to open. For example, users should be forbidden to open attachments in two instances: when they were not expecting the attachments, and when the attachments arrive from unrecognizable aliases.

Finally, whenever possible, consider a centralized antivirus service that updates the distributed clients from a centrally managed server. Most such solutions provide you with ways to more granularly manage each client and proactively fix problems that may take place.

Trojans

A *Trojan* (also known as a Trojan horse) is a malicious program embedded inside a normal, safe-looking program. The difference between a virus and a Trojan is that the Trojan is embedded and the virus is attached to the file or executable.

When the normal program runs, the malicious code runs as well and can cause damage or steal critical information. An example of a Trojan is a word-processing program that, when executed, allows the user to compose a document while, in the background, malicious code is running that deletes files or destroys other programs.

Trojans generally are spread through e-mail or *worms*, which are programs that run by themselves. The damage that Trojans can cause is similar to that of a virus: from nominal to critical. Trojans are particularly frightening because in most cases, users are unaware of the damage the Trojan is causing. The malicious work is being masked by the Trojan effect of the program.

Worms

As just mentioned, worms are programs that run by themselves. They do not embed or attach themselves to other programs nor do they need to do this to replicate. They can travel from computer to computer across network connections and are self-replicating. Worms might have portions of themselves running on many different computers, or the entire program might run on a single computer. Typically, worms do not change other programs, although they might carry other code that does.

The first network worms were intended to perform useful network management functions by taking advantage of operating system properties. Malicious worms exploit system vulnerabilities for their own purposes. Release of a worm usually results in brief outbreaks, shutting down entire networks.

The damage that worms can cause, like Trojans and viruses, ranges from the nominal to the critical. The type and extent of damage must be assessed individually for each worm. However, worms can install viruses and Trojans that then run their own code.

An attack that combines a worm, Trojan, and/or virus can be a very difficult attack to survive without significant damage. The impact of viruses, Trojans, and worms on your messaging system and network should not be underestimated. Because they use e-mail to exploit system vulnerabilities, installing antivirus software is simply not enough. You must also ensure that known vulnerabilities in all your operating systems are updated. Don't focus only on your servers. Every device should be updated with the most recent updates from each vendor as soon as possible. Most environments will want to test these updates before installing them. But after they have been tested, install them.

Junk E-Mail

Junk e-mail is a huge issue. One client with whom this author recently worked installed its first e-mail filtering software and found that it had 46 percent fewer inbound e-mails.

Exchange Server 2007's new Edge Transport role has new capabilities that can help to significantly reduce the amount of junk e-mail that enters your environment. The Edge Transport role server has the following agents that help to protect your e-mail infrastructure. The information in Table 19-2 is right from Microsoft's Edge Transport server documentation.

Many of these features are discussed in the next chapter, Chapter 20, "Antivirus and Antispam," and Chapter 21, "Securing Exchange Server 2007 Messages."

Table 19-2 Edge Transport Agents

Agent name	Description
Connection Filtering Agent	Performs host IP address filtering based on IP Allow Lists, IP Allow List providers, IP Block Lists, and IP Block List providers.
Address Rewriting Inbound Agent	Modifies recipient SMTP addresses in inbound messages based on predefined address alias information. Address rewriting can be useful in scenarios where an organization wants to hide internal domains.
Edge Rule Agent	Processes all messages received over SMTP to enforce transport rules defined on the Edge Transport server.
Sender ID Agent	Determines whether the sending SMTP host is authorized to send messages for the SMTP domain of the message originator.

Table 19-2 Edge Transport Agents

Agent name	Description
Recipient Filter Agent	Verifies that the recipients specified during the SMTP session through the RCPT TO: command are valid and not on the list of blocked SMTP addresses and domains.
Sender Filter Agent	Verifies that the sender specified in the MAIL FROM: command and in the message header is valid and not on the list of blocked SMTP addresses and domains.
Content Filter Agent	Uses Microsoft SmartScreen technology to assess the contents of inbound messages in order to assign an SCL rating for junk e-mail processing based on transport and store thresholds.
Protocol Analysis Agent	Interacts with Connection Filtering, Sender Filtering, Recipient Filtering, and Sender ID agents to determine Sender Reputation Level (SRL) rating and to take action based on rating thresholds.
Attachment Filtering Agent	Filters messages based on attachment file name, file name extension, or MIME content type to block potentially harmful messages or remove critical attachments.
Address Rewriting Outbound Agent	Modifies sender SMTP addresses in outbound messages based on predefined address alias information. Address rewriting can be useful in scenarios where an organization wants to hide internal domains.
Forefront Security for Exchange Routing Agent	Responsible for connecting into the Transport stack to ensure that the scanning process scans messages prior to delivery to Hub Transport servers.

Security Tools Provided by Microsoft

In order to help you deploy and maintain the most secure Exchange infrastructure possible, Microsoft provides a number of tools designed to remove malware, make sure that your environment is properly configured, and help you configure a multitude of security settings.

- **Malicious Software Removal Tool** The Microsoft Windows Malicious Software Removal Tool checks computers running Windows XP, Windows 2000, and Windows Server 2003 for infections by specific, prevalent malicious software—including Blaster, Sasser, and Mydoom—and helps remove any infection found. When the detection and removal process is complete, the tool displays a report describing the outcome, including which, if any, malicious software was detected and removed. Microsoft releases an updated version of this tool on the second Tuesday of each month, and as needed to respond to security incidents. On a regular basis, run the

Malicious Software Removal Tool on your Exchange server to make sure your system is free of threats.

More Info To download the Microsoft Software Removal Tool, visit *http://www.microsoft.com/security/malwareremove/default.mspx*.

- **Microsoft Baseline Security Analyzer** The Microsoft Baseline Security Analyzer (MBSA) is a tool that analyzes your existing environment and, in particular, analyzes how you have configured a number of Microsoft products, including Windows 2000 SP3; Windows XP and Windows Server 2003; Office XP, 2003 and 2007; Exchange 2000, 2003 and 2007; SQL Server 2000 SP4; and SQL Server 2005. With this information, Microsoft compares your configuration against a list of best practices and provides you with a report of action items that you can take to improve the security of your environment.

More Info To download the Microsoft Baseline Security Analyzer, visit *http://www.microsoft.com/technet/security/tools/mbsa2/default.mspx*.

- **Security Configuration Wizard** Windows Server 2003 Service Pack 1 includes the Security Configuration Wizard (SCW), a tool designed to reduce the attack surface of your Windows servers. SCW helps administrators to create security policies that are consistent with the practice of least privilege. In this case, that means running the fewest possible services on a server in order to reduce the number of services that can be used to attack the computer.
- **Microsoft Exchange 2007 Anti-Spam Migration Tool** The Exchange 2007 Anti-Spam Migration Tool is designed to ease the administrative burden involved in transitioning from Exchange Server 2003 to Exchange Server 2007, particularly for those administrators who have deployed Exchange Server 2003 anti-spam services and want to maintain the service configuration under Exchange Server 2007. This tool converts the Exchange Server 2003 anti-spam service settings into PowerShell commands that can be used to appropriately configure anti-spam service settings in Exchange Server 2007

More Info To download the Microsoft Exchange 2007 Anti-Spam Migration Tool, visit *http://www.microsoft.com/downloads/details.aspx?FamilyId=805EAF35-EBB3-43D4-83E4-A4CCC7D88C10&displaylang=en*.

This tool is not available for use on Windows Server 2008.

Summary

This chapter discussed how hackers think, how to secure incoming SMTP e-mail, and how to secure Administrator access to your Exchange server. Also discussed were the differences between a virus, a Trojan, and a worm, and a method was outlined for securing inbound SMTP traffic. Two other areas in this book were also referenced that discuss sender filtering and securing OWA. The next chapter discusses how to secure e-mail messages using encryption and certificates.