



Configuring SMTP in Microsoft[®] Exchange 2000 Server



Patricia Anderson
Simon Attwell

Microsoft®

Configuring SMTP in Microsoft® Exchange 2000 Server



Patricia Anderson
Simon Attwell

Copyright

Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2002 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Hotmail, MS-DOS, MSDN, Outlook, Visual Basic, Visual C++, Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Published: December 2002

Applies To: Exchange 2000 Server SP3

Editor: Brendon Bennett

Technical Reviewers: Pretish Abraham, Simon Attwell, Max Ciccotosto, Wayne Cranston, Ade Famoti, Scott Landry, Andrew Moss, Bill Witten

Artist: Kristie Smith

Production: Stephanie Schroeder

Table of Contents

Introduction	1
What Will You Learn from This Book?.....	1
Who Should Read This Book?	2
How Is This Book Structured?	2
Terminology.....	3
Chapter 1	
Overview of Exchange and SMTP.....	7
Understanding SMTP and Exchange.....	7
Receiving Internet Mail.....	9
Sending Internet Mail	10
Chapter 2	
SMTP Elements.....	13
SMTP Virtual Server	13
Inbound Mail Settings on the SMTP Virtual Server.....	14
Relay Restrictions	15
Default Relay Restrictions.....	16
Outbound Mail Settings on the SMTP Virtual Server	17
SMTP Connectors.....	17
The Function of an SMTP Connector	19
Uses for a Connector	20
Simplified Administration	21
Limited Internet Exposure	21
Isolated Route for Communicating with Other Domains	21
Load Balancing with Multiple Bridgehead Servers	22
Use Specific SMTP or ESMTP Commands.....	22
Schedule and Customize Outbound Connections	22
Chapter 3	
SMTP Dependencies	25
Internet Information Services.....	25
Active Directory	26
DNS.....	26

How DNS Queries Work.....	27
Querying a DNS Server.....	27
The Role of DNS in Receiving Internet Mail.....	28
The Role of DNS in Sending Internet Mail.....	29
Using DNS to Send Internet Mail.....	29
Forwarding Internet Mail to a Smart Host.....	31
Recipient Policies.....	31

Chapter 4

SMTP Deployment Scenarios..... 33

Common Deployment Scenarios.....	34
Using a Single Exchange Server in Its Default Configuration.....	35
Basic Configuration.....	35
Inbound Internet Mail.....	36
Outbound Internet Mail.....	36
Using a Dual-Homed Exchange Server as an Internet Gateway.....	36
Basic Configuration.....	37
Inbound Internet Mail.....	38
Outbound Internet Mail.....	38
Security Considerations.....	39
Using a Bridgehead Server Behind a Firewall.....	40
Basic Configuration.....	40
Inbound Mail Flow.....	42
Outbound Mail Flow.....	42
Using a Windows 2000 SMTP Relay Server in a Perimeter Network.....	42
Basic Configuration.....	43
Inbound Mail Flow.....	46
Outbound Mail Flow.....	46
Custom Deployment Scenarios.....	46
Using a Network Service Provider to Send and Receive Mail.....	46
Supporting Two SMTP Mail Domains and Sharing an SMTP Mail Domain with Another System.....	47
Supporting Two SMTP Mail Domains.....	48
Sharing an SMTP Mail Domain with Another System.....	49
Supporting Additional Mail Systems.....	58

Chapter 5

Configuring Exchange to Send and Receive E-Mail..... 59

Verifying SMTP Port Settings.....	59
Setting Up Your Exchange Server to Receive Internet Mail.....	63
Configuring Recipient Policies.....	63
Verifying That Recipient Policies Do Not Contain an SMTP Address Matching the FQDN of an Exchange Server.....	63

Verifying That Recipients Can Receive Mail from Other SMTP Domains	64
Configuring the SMTP E-Mail Addresses for Your Users.....	66
Configuring Inbound Settings on SMTP Virtual Servers	68
Configuring the Inbound Port and IP Address.....	69
Verifying Default Relay Restrictions on Your Inbound SMTP Virtual Server.....	70
Verifying DNS Set Up for Inbound Mail.....	73
Using Nslookup to Verify DNS Configuration.....	73
Using Telnet to Ensure Internet Accessibility.....	75
Setting Up Your Exchange Server to Send Internet Mail	75
Configuring Outbound Settings on SMTP Virtual Servers.....	75
Verifying That the Outbound TCP Port Is Set to 25.....	76
Allowing Anonymous Access on the Outbound Virtual Server.....	78
Configuring a Smart Host on an SMTP Virtual Server	79
Configuring an SMTP Connector	80
Creating an SMTP Connector.....	80
Configuring an Address Space.....	82
Configuring DNS for Outbound Mail	84
Method 1: Using Internal DNS Servers for External Name Resolution	84
Method 2: Configuring External DNS Servers on an SMTP Virtual Server.....	86
Using Nslookup to Verify DNS Configuration.....	88
Configuring Advanced Settings	89
Configuring Advanced Inbound Settings	89
Configuring Access Controls and Security Settings	89
Setting Global Message Filters	92
Specifying Message Limits.....	93
Configuring Advanced Outbound Settings.....	95
Configuring Internet Mail Message Formats.....	95
Configuring Outbound Message Limits	96
Configuring Advanced Settings on the SMTP Connector	98
Configuring Notification of Delivery Reports	103
Using Distribution Lists in Multi-Domain Environments	105

Chapter 6

Security Considerations	107
Securing Your Infrastructure	108
IIS Lockdown Wizard.....	108
Firewalls	108
Virtual Private Networks.....	109
Securing Your Exchange Server	110
Disabling Open Relaying on All SMTP Virtual Servers	110

Implementing Sender Filtering for Your SMTP Mail Domain on Inbound Gateway Servers	111
Preventing Anonymous Access on Internal SMTP Virtual Servers and Dedicated SMTP Virtual Servers for IMAP and POP Clients	111
Controlling Unsolicited Commercial E-Mail	112
Using Message Filters	112
Identifying Spoofed Mail.....	113

Chapter 7

Troubleshooting Mail Flow 115

Using Telnet.....	116
Understanding Non-Delivery Reports.....	118
Using the SMTP and X.400 Queues.....	130
Understanding the SMTP Queues.....	131
Understanding the X.400 Queues	134
Viewing the Properties of a Queue	135
Viewing the Messages in a Queue.....	136
Using Message Tracking Center.....	136
Using Event Viewer	137
Viewing the Application Log	137
Viewing the System Log.....	138
Configuring Diagnostic Logging for the SMTP Protocol	139
Modifying Logging Settings	139
Enabling Debugging Level Logging.....	140
Configuring Diagnostic Logging for the X.400 Service (MSEExchangeMTA)	141

Chapter 8

Reference..... 143

SMTP Commands and Definitions	143
Understanding the Internal SMTP Transport Mechanisms	146
Receiving Internet Mail.....	147
Sending Internet Mail	148
Event Sinks.....	149
Common Ports Used by Exchange	150

Appendix A

Additional Resources 155

Technical Papers.....	155
Microsoft Knowledge Base Articles	155
Other Useful Resources.....	157



Introduction

This book explains how to configure Microsoft® Exchange 2000 Server for sending and receiving Internet mail. Although Exchange supports many Internet protocols and features, this book mainly focuses on Internet mail and SMTP (Simple Mail Transfer Protocol).

What Will You Learn from This Book?

Essentially, this book provides detailed answers to the following questions:

- What is SMTP, and what is its purpose? (Chapter 1)
- How does SMTP work in Exchange? (Chapter 1)
- What are the basic building blocks of SMTP, and how do I manage the protocol in Exchange? (Chapter 2)
- What are the components upon which SMTP relies? How do these components affect the operation of SMTP? (Chapter 3)
- How is SMTP commonly deployed to support various organizations or meet special requirements (such as sharing an SMTP mail domain or supporting two SMTP mail domains)? (Chapter 4)
- How do I configure SMTP, Exchange, and Domain Name System (DNS) to support Internet mail delivery? (Chapter 5)
- What measures can I take to secure my infrastructure and my Exchange servers? (Chapter 6)
- What tools and processes can I use to troubleshoot and diagnose mail flow problems? (Chapter 7)

Who Should Read This Book?

This book assumes that you have a basic understanding of Microsoft Windows® 2000 architecture and a working knowledge of Domain Name System (DNS). While practically anyone with a technical background can benefit from reading this book, it's designed to produce maximum benefits for the following professionals:

- IT Professionals—those people who are responsible for installation, maintenance, and administration of software in the enterprise. This includes managers, system administrators, system engineers, system operators, and database administrators (DBAs).
- System Architects—those people who are responsible for planning and crafting overall business strategies and solutions.

How Is This Book Structured?

This book has six chapters and one appendix. For best results, review these chapters in order, as each chapter builds upon the concepts revealed in preceding chapters.

Chapter 1, “Overview of Exchange and SMTP”

This chapter presents an overview of SMTP and explains how SMTP enables message flow in an Exchange organization.

Chapter 2, “SMTP Elements”

This chapter explains the two building blocks of SMTP: the SMTP virtual servers and the SMTP connectors.

Chapter 3, “SMTP Dependencies”

This chapter explains the components that are necessary for SMTP to function properly: IIS, DNS, the Active Directory® directory service, and recipient policies.

Chapter 4, “SMTP Deployment Scenarios”

This chapter presents common and customized SMTP deployment scenarios.

Chapter 5, “Configuring Exchange to Send and Receive E-Mail”

This chapter contains procedures for configuring Exchange and SMTP to send and receive mail.

Chapter 6, “Security Considerations”

This chapter focuses on security considerations for SMTP.

Chapter 7, “Troubleshooting Mail Flow”

This chapter demonstrates how to test SMTP communication using telnet and how to interpret non-delivery reports. This chapter also shows you how to use the SMTP and X.400 queues, the Message Tracking Center, Event Viewer, and the SMTP log to diagnose mail flow problems.

Chapter 8, “Reference”

This chapter contains reference material about SMTP commands, the functions of the internal SMTP transport components, and event sinks.

Appendix, “Additional Resources”

This section contains links to additional resources that will help you maximize your understanding of Exchange and SMTP.

Terminology

Before reading this book, familiarize yourself with the following terms:

A (address) record

An address resource record in DNS; specifically, a DNS record that associates a host name with an IP address.

authoritative domains

Domains for which Exchange is exclusively responsible for message delivery; Authoritative domains are replicated to the IIS metabase that SMTP uses for mail delivery.

bridgehead server

A computer that connects servers that use a single communications protocol so that information can be passed from one server to another. In Exchange 2000, a bridgehead server is a connection point from a routing group to another routing group, remote system, or other external system.

connector

A component that enables information to flow between two systems. For example, connectors support message transfer, directory synchronization, and calendar querying between Exchange and other messaging systems. When connectors are in place, the basic user experience is maintained on both messaging systems. The exchange of mail and other information between Exchange and other messaging systems is transparent to the user, even if the two systems function differently.

Domain Name System (DNS)

A TCP/IP standard name service that allows clients and servers to resolve names into IP addresses and vice versa. Domain Name Service in Windows 2000 allows clients and servers to automatically register themselves without the need for administrators to manually define records.

Exchange System Attendant

A component of Exchange responsible for many scheduled Exchange tasks. Exchange System Attendant includes the metabase update service that replicates SMTP configuration information from Active Directory to the IIS metabase.

fully qualified domain name (FQDN)

A DNS domain name that completely and unequivocally identifies the position of the domain in relation to the domain namespace. A FQDN such as host.example.com, typically ends with a trailing period (.), indicating the root of the DNS tree. For some Exchange functions, you will use the Exchange server's host name, which is the FQDN minus the domain portion.

global catalog server

A domain controller that contains a partial replica of every domain in Active Directory. A global catalog holds a copy of every object in Active Directory, but with a limited number of each object's attributes.

Internet Information Services (IIS)

In Windows 2000, the software services that support Internet-related functions ranging from Web site creation to development of server-based Web applications. IIS includes support for the widely used Internet protocols Simple Mail Transfer Protocol (SMTP), Network News Transfer Protocol (NNTP), Hypertext Transfer Protocol (HTTP), and File Transfer Protocol (FTP).

IIS metabase

A repository that contains configuration metadata used by IIS.

Lightweight Directory Access Protocol (LDAP)

The primary access protocol for Active Directory. Lightweight Directory Access Protocol (LDAP) version 3 is defined by a set of Proposed Standard documents in Internet Engineering Task Force (IETF) Request For Comments (RFC) 2251.

message transfer agent (MTA)

The Exchange component responsible for routing messages. Depending on the destination, the MTA routes messages to other MTAs, to the Exchange store, to Exchange connectors, and to third-party gateways. Also referred to as the X.400 protocol in Exchange 2000 System Manager (used by Exchange when transferring mail to X.400 systems, such as Exchange 5.0 or Exchange 5.5 servers.)

mail exchanger (MX) resource record

In the transmission of e-mail over the Internet, the MX resource record is a DNS record that associates an e-mail domain with the FQDN of one or more SMTP servers.

non-authoritative domains

Domains for which the Exchange organization is not exclusively responsible. If an Exchange server receives mail destined for a non-authoritative domain, it searches in Active Directory to find a recipient. However, if the Exchange server does not find a recipient in Active Directory, it looks for another route (for example, through a connector or DNS) to attempt delivery. If both methods fail to locate a route for the message, a non-delivery report (NDR) is generated with a 5.4.0 error code, indicating a name resolution issue.

relay

The ability to send mail to an external domain. By default, Exchange permits only authenticated users to relay mail. If you change your Exchange server configuration to allow anonymous users to relay mail, unauthorized users will be able to send unsolicited commercial e-mail.

Simple Mail Transfer Protocol (SMTP)

An Internet standard for transporting and delivering electronic messages. Based on specifications in RFC 2821 and RFC 2822, Microsoft SMTP service is included in the Windows 2000 operating system. SMTP is the default transport for Exchange 2000.

1

Overview of Exchange and SMTP

Before configuring your Exchange organization to send and receive mail, you should have a good understanding of how SMTP enables message flow in Exchange 2000.

Exchange 2000 uses SMTP to deliver internal mail between Exchange servers and routing groups. Similarly, Exchange uses SMTP to deliver Internet mail outside the Exchange organization. This chapter provides a detailed overview of SMTP, including how SMTP works in Exchange 2000, and explains the process of sending and receiving Internet mail.

Understanding SMTP and Exchange

SMTP is the Internet standard for transporting and delivering electronic messages. Based on specifications in RFC 2821 and RFC 2822, the Microsoft SMTP service is included in Windows 2000.

The Windows SMTP service is a component of Internet Information Services (IIS) and runs as part of Inetinfo.exe. Exchange 2000 relies on the Windows 2000 SMTP service as its native transport protocol; therefore Exchange uses SMTP to route all internal and external messages.

When Exchange is installed, it modifies the SMTP service by extending the underlying SMTP functionality. Exchange extends SMTP functionality by:

- Moving management of the SMTP service (by means of SMTP virtual servers) from the IIS administrative console to Exchange System Manager.
- Implementing support for link state information. Exchange uses link state routing to determine the best method for sending messages between servers, based on the current status of messaging connectivity and cost.

- Extending SMTP to support the command verbs used to support link state routing and other Exchange functionality. The following commands are added when Exchange is installed:
 - X-EXPS GSSAPI
 - X-EXPS=LOGIN
 - X-EXCH50
 - X-LINK2STATE

For a list of all the SMTP commands and their definitions, see “SMTP Commands and Definitions” in Chapter 8.

- Setting up an Exchange Installable File System (IFS) store driver to allow message retrieval from and delivery to the Exchange store.
- Setting the disk location where messages are queued to `\exchsrvr\mailroot\vs 1\pickup`. This is the location of the first SMTP virtual server on the Exchange server. If you add a second SMTP virtual server, a new location (`\exchsrvr\mailroot\vs 2\pickup`) is created.
- Implementing support for advanced queuing. Exchange enhances the queuing capabilities of Windows 2000. The advanced queuing engine handles underlying transport functions in Exchange.
- Enhancing message categorization. Message categorization is a process performed by the message categorizer, a component of the advanced queuing engine. The message categorizer sends lightweight directory access protocol (LDAP) queries to the global catalog server to retrieve configuration information stored in Active Directory. The message categorizer retrieves recipient policy information and Exchange virtual server information to enable message delivery. It uses this information to validate the recipient address, to verify that message limits are not exceeded, and to ultimately determine how the message is delivered using Exchange routing and SMTP.

An important concept to understand about Exchange 2000 and SMTP is the interaction among Exchange, Active Directory, and the IIS metabase. With Exchange System Manager, any the configuration changes you make (such as to your recipient policies and SMTP virtual servers) are written to Active Directory, allowing for easy and remote administration. However, because the SMTP service reads its settings from the IIS metabase, Exchange System Attendant replicates this information from Active Directory into the local server’s IIS metabase.

Receiving Internet Mail

If the following conditions exist, an Exchange 2000 server is able to receive Internet mail in its default configuration:

- There is a constant connection to the Internet.

Note Dial-up connections to the Internet require special configuration. For more information see “Using a Network Service Provider to Send and Receive Mail” in Chapter 4.

- The external DNS servers for your domain must have mail exchanger (MX) resource records pointing to your mail servers. Your ISP or the administrative contact for your domain may need to set this up for you. For information about how to verify your MX records, see “Using Nslookup to Verify DNS Configuration” in Chapter 5.
- Your mail server must be accessible to other servers on the Internet. For information about how to verify that your mail server is accessible on the Internet, see “Using Telnet to Ensure Internet Accessibility” in Chapter 5.
- Your recipient policies must be set up correctly. To receive Internet mail, you must have a recipient policy configured that contains an address space matching the SMTP domain. Also, your Exchange organization must be responsible for delivering mail to this address (this is the default setting). For example, to accept Internet mail for `sfine@example.com`, you must have a recipient policy that contains `@example.com`. However, there are some exceptions to this rule; for example, you can create a connector that allows relaying to a specified domain. For information about how to configure your recipient policies, see “Configuring Recipient Policies” in Chapter 5.

Inbound Internet mail flows through an Exchange 2000 server in the following manner. (For detailed information about internal transport mechanisms, see “Understanding the Internal SMTP Transport Mechanisms” in Chapter 8.)

1. The sending SMTP server queries DNS to locate the IP address of the recipient’s SMTP mail server.
2. The sending SMTP server then initiates a conversation on the recipient’s SMTP server (on port 25). On an Exchange gateway, the recipient’s SMTP server is the SMTP virtual server that is configured to accept inbound Internet mail.
3. Ideally, the inbound SMTP server only accepts the incoming message if it is destined for a recipient of its SMTP mail domain. These recipients are defined in the recipient policies (unless the server is open to relay, which is strongly discouraged).

Note If you leave your system open for relay, unauthorized users can use your servers to send mail to external addresses. As a result, your system may be blacklisted—a process that blocks mail from servers suspected of sending unsolicited commercial e-mail. For more information about relaying, see “Relay Restrictions” in Chapter 2. For instructions about how to verify your relay restrictions, see “Verifying Default Relay Restrictions on your Inbound SMTP Virtual Server” in Chapter 5.

4. When the message is accepted, the SMTP virtual server uses the transport mechanisms within Exchange to determine the method for delivering the message. Exchange locates the recipient in Active Directory and determines which server in the Exchange organization will deliver the message. For detailed information about the internal components of SMTP, see “Understanding the Internal SMTP Transport Mechanisms” in Chapter 8.
5. Finally, the SMTP virtual server uses its internal transport mechanisms to deliver the message to the appropriate Exchange server.

Sending Internet Mail

Assuming there is a constant Internet connection, there are two basic methods Exchange uses to send Internet mail:

- Use DNS directly to contact the remote mail server
- Route mail through a smart host that assumes responsibility for DNS name resolution and mail delivery

Before each of these methods is described in detail, you should have a general understanding of how outbound mail flows in an Exchange organization.

Outbound Internet mail flows through an Exchange 2000 server in the following manner. (For detailed information about internal transport mechanisms, see “Understanding the Internal SMTP Transport Mechanisms” in Chapter 8.)

1. An internal user sends a message to a recipient in a remote domain.
2. To determine if the recipient is local or remote, the SMTP virtual server on the sender’s Exchange server uses internal transport functions to query the global catalog server for the recipient address. If the recipient’s address on the message is not in a recipient policy, it will not be stored in Active Directory; therefore, Exchange would determine that the message is destined for a remote domain.
3. If necessary, the Exchange server delivers the message to the appropriate SMTP virtual server.

4. The SMTP virtual server uses its IIS metabase information to determine the method for delivering a message to a remote domain.
5. The SMTP virtual server on the Exchange server then does one of two things:
 - Uses DNS to look up the IP address for the target domain, and then attempts to deliver the message.
 - Forwards the message to a smart host that assumes responsibility for the DNS resolution and delivery.

2

SMTP Elements

This section describes the basic building blocks of SMTP. Specifically, these include the following:

SMTP virtual servers

SMTP virtual servers provide the Exchange mechanisms for managing SMTP. Each SMTP virtual server represents an instance of the SMTP service running on the Exchange server. You use Exchange System Manager to configure SMTP virtual servers that control the behavior of SMTP.

SMTP connectors

An SMTP connector is used to designate an isolated route for mail. You can use SMTP connectors to establish a gateway for Internet mail or a smart host, or to connect routing groups internally. Connectors allow you to define specific options for the designated mail route.

SMTP Virtual Server

Essentially, an SMTP virtual server is an SMTP protocol stack (a process or server that both receives e-mail and acts as a client for sending e-mail). Each SMTP virtual server represents an instance of the SMTP service on a server. An SMTP virtual server is defined by a unique combination of an IP address and port number. The default SMTP virtual server uses all available IP addresses on the server and uses port 25 for inbound connections. A single physical server can host many virtual servers.

You use Exchange System Manager to control most of the SMTP settings. The property settings of the SMTP virtual server control inbound mail and, to a lesser degree, outbound mail settings.

Important Because an SMTP virtual server plays a critical role in mail delivery, use caution when modifying its property settings. For example, the default SMTP virtual server sends messages within a routing group. Additionally, if the server is a domain controller, Active Directory uses this virtual server for SMTP directory replication. Therefore, instead of modifying the default SMTP virtual server, it is recommended that you either create an additional SMTP virtual server or create an SMTP connector to override the default virtual server settings.

Inbound Mail Settings on the SMTP Virtual Server

You can use the virtual server's property settings to configure the following inbound settings:

Inbound ports and IP addresses

The SMTP virtual server listens on its assigned IP address for incoming communications and accepts inbound connections on its assigned port. To configure these settings, use the **General** tab of the SMTP virtual server's properties.

Important The SMTP service defines port 25 as its standard port. Do not change this setting.

Note Upon installation in its initial configuration, the default virtual server connects to the remote SMTP server on port 25 to send outbound mail. This is a separate setting. To configure this setting, use the **Outbound Connections** button on the **Delivery** tab.

Relay restrictions

To prevent unauthorized users from using your server to send messages to external addresses, use the **Relay** button on the **Access** tab. By default, the default SMTP virtual server only relays messages for authenticated users. For more information about relay restrictions, see "Relay Restrictions" later in this chapter.

Security

You can require Transport Layer Security (TLS), an implementation of Secure Sockets Layer (SSL), on incoming connections.

You can also configure other settings such as inbound connection restrictions, performance tuning, and handling of delivery reports notifications.

Relay Restrictions

Relaying is the ability to forward mail to domains other than your own. More specifically, relaying occurs when an inbound connection to your SMTP server is used to send e-mail to external domains. By default, your Exchange server accepts mail from users and sends it to an external domain. If your server is open for relaying, or if relaying is unsecured on your server, unauthorized users can use your server to send unsolicited commercial e-mail.

Therefore, to secure your SMTP virtual server, it is crucial that you set relay restrictions.

It is important to understand the difference between authenticated relaying and anonymous or open relaying.

Authenticated relaying

Authenticated relaying allows your internal users to send mail to domains outside of your Exchange organization, but requires authentication before the mail is sent. By default, Exchange only allows authenticated relaying.

Anonymous relaying

Anonymous relaying allows any user to connect to your Exchange server and use it send mail outside your Exchange organization.

The following examples demonstrate how Exchange 2000 accepts and relays mail using authenticated relaying:

Example 1

An anonymous user connects to the SMTP virtual server and attempts to deliver mail to an internal user in the Exchange organization.

In this situation, the SMTP virtual server accepts the message because it is destined for an internal domain and because the user exists in Active Directory.

Example 2

An anonymous user connects to the SMTP virtual server and attempts to deliver mail to an external user in an external domain.

In this situation, the SMTP virtual server rejects the mail because it is destined for an external domain for which the Exchange server is not responsible. Because the user is not authenticated, the SMTP virtual server does not relay this mail outside of the Exchange organization.

Example 3

A user connects to the SMTP virtual server using a POP or IMAP client (for example Microsoft Outlook[®] Express), authenticates, and then attempts to send a message to a user in an external domain.

In this situation, Outlook Express connects directly to the SMTP virtual server and authenticates the user. Although the message is destined for a remote domain, the SMTP virtual server accepts and relays this mail because the user is authenticated.

By using the relay control features of Exchange 2000, you can prevent third parties from relaying mail through your server. Relay control allows you to specify a list of incoming remote IP address and subnet mask pairs that have permission to relay mail through your server. Exchange checks an incoming SMTP client's IP address against the list of IP networks allowed to relay mail. If the client is not allowed to relay mail, only mail addressed to local recipients is allowed. Relay control can also be implemented by domain—however, this requires implementation of reverse DNS resolution, which is controlled at the SMTP virtual server level.

Default Relay Restrictions

By default, the SMTP virtual server allows relaying only from authenticated users. This configuration is designed to prevent unauthorized users from using your Exchange server to relay mail. As illustrated in Figure 1, the virtual server's default configuration allows only authenticated computers to relay mail.

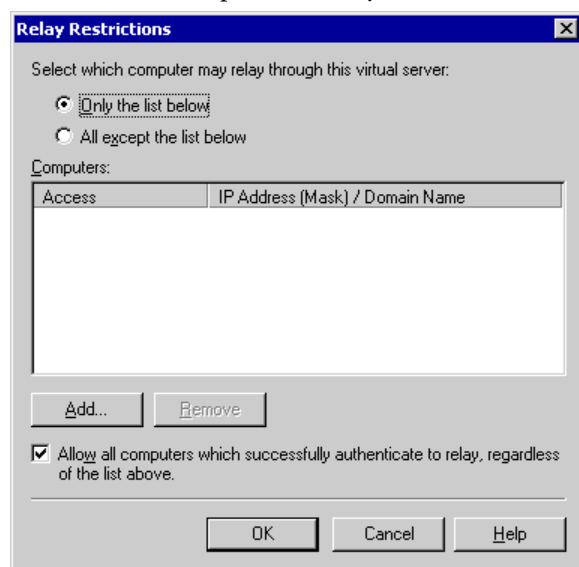


Figure 1 Default relay restrictions

Unsolicited commercial e-mail generally comes from a spoofed or forged address and is often relayed using a server that is not secured for relay. For this reason, Exchange 2000 allows only authenticated users. Be very cautious when changing this setting—many Internet providers will block servers that allow open relaying.

Outbound Mail Settings on the SMTP Virtual Server

If you want your SMTP virtual server to send mail directly to the Internet, you can configure outbound mail settings. Specifically, you can configure your virtual server to use an external DNS server to resolve external addresses and send mail directly to mail servers outside of your organization.

Important Because an SMTP virtual server plays a critical role in mail delivery, use caution when modifying its property settings. For example, the default SMTP virtual server sends messages within a routing group. Additionally, if the server is a domain controller, Active Directory uses this virtual server for SMTP directory replication. Therefore, instead of modifying the default SMTP virtual server, it is recommended that you either create an additional SMTP virtual server or create an SMTP connector to override the default virtual server settings.

In many instances, it is preferable (but not required) to set up an SMTP connector to handle outbound mail. For more information about SMTP connectors, see “SMTP Connectors” later in this chapter.

Note If you use an SMTP connector, it will override some of the outbound mail settings and control outbound mail delivery.

To control outbound delivery on your virtual server, you can configure the following settings:

- Outbound port
- Outbound restrictions
- Outbound delivery options
- Outbound security
- Performance tuning
- Notification of delivery reports

For information about how to configure these settings, see “Configuring Outbound Settings on SMTP Virtual Servers” in Chapter 5.

SMTP Connectors

SMTP connectors are used primarily to connect to other mail systems or to define additional options for an SMTP Internet gateway. SMTP connectors can also be used to connect a routing group to another routing group internally. Essentially, SMTP connectors allow you to designate an isolated route for messages to flow either to a specific domain or over the Internet.

As mentioned earlier, one advantage to using a connector is that you can specify additional configuration settings to affect mail delivery. These settings include:

Outbound mail delivery

When you configure a connector, you can route mail in one of two ways:

- Use DNS to route all outgoing mail through the connector. If you use DNS to route outgoing mail, the SMTP connector uses DNS to resolve the IP address of the remote SMTP server, and then it delivers the mail.
- Specify a smart host (another server to which the connector routes all mail). The smart host takes responsibility for DNS resolution and delivers the mail.

Local bridgeheads

An SMTP virtual server hosts a connector. When you create a connector, you designate at least one Exchange server and SMTP virtual server as bridgeheads. The connector inherits size restrictions and other settings from the SMTP virtual server; however, you can override these settings on the connector. You can also designate multiple bridgeheads for load balancing, performance, and redundancy.

Address space

The address space defines the mail addresses or domains for the e-mail you want routed through a connector. For example, an address space of * (asterisk) encompasses all external domains—this connector is used to route all external e-mail. If you created a second connector with an address space of *.net, Exchange would route all mail sent to a domain with a .net extension through the second connector. This is because Exchange would select the connector that has the most similar address space. This setting is configured on the **Address** tab of the SMTP connector's properties.

Scope

You can select either an entire organization or a routing group for the connector's scope. The scope is also defined on the **Address** tab of the SMTP connector's properties.

Delivery restrictions

You can restrict who can send mail through a connector. By default, mail is accepted from everyone. These settings are configured on the **Delivery** tab of the SMTP connector's properties.

Note By default, you cannot restrict mail unless you change the registry key settings. If you chose to enable delivery restriction, be aware that restricting delivery is extremely process-intensive and can impact server performance. For information about how to enable delivery restrictions, see "Specifying Delivery Restrictions" in Chapter 5.

Content restrictions

You can specify what types of messages are delivered through a connector. These settings are configured on the **Content Restrictions** tab of the SMTP connector's properties.

Delivery options

If you connect to a network service provider to retrieve your mail, you can configure a connector to run on a specified schedule and implement advanced queuing and dequeuing features. These settings are configured on the **Delivery Options** tab of the SMTP connector's properties.

SMTP communication

You can control how the connector uses SMTP to communicate with other SMTP servers. Specifically, you can specify whether the connector uses SMTP or Extended Simple Mail Transfer Protocol (ESMTP) commands to initiate a conversation with another server and control the use of the ERTN and TURN commands (these commands are used to request that another SMTP server send any e-mail messages it has). These settings are configured on the **Advanced** tab of the SMTP connector's properties.

Outbound security

You can also ensure that any mail flowing through the connector is authenticated. This is useful if you want to establish a secure route for communicating with a partner company. With this setting, you can establish an authentication method and require TLS encryption.

The Function of an SMTP Connector

SMTP relies on DNS to determine the IP address of its next destination server. To send mail directly to an external mail server, an SMTP connector must use DNS to resolve external domain names. Alternatively, the connector can simply forward mail to a smart host that assumes responsibility for DNS name resolution and delivery. For more information about the dependency of SMTP on DNS, see "DNS" in Chapter 3.

After you set up an SMTP connector, as long as the address space matches, the servers no longer route the mail directly; instead, the servers route the mail through the SMTP connector. (These servers are called gateway or bridgehead servers.)

To illustrate this point, let's assume that you want all external mail routed through a connector to a bridgehead server (which is the only server that communicates with the Internet). To configure this, you would create a connector on the bridgehead server with an address space of * (asterisk), which specifies all external domains. When e-mail is sent to an external domain, Exchange automatically routes it to this connector, rather than an SMTP virtual server sending the external mail directly. If you have more than one connector, Exchange first attempts to route mail through the connector that has the most similar address space (which is the most restrictive address space).

Note In a mixed mode environment, if you have an Exchange 5.5 Internet Mail Connector, Exchange 2000 treats this connector as a valid route. If you experience problems sending or receiving Internet e-mail, check the MTA queues on the Exchange 5.5 server and the X.400 queues on the Exchange 2000 server. Exchange 2000 uses the MTA to communicate with legacy versions of Exchange. For more information, see "Using the SMTP and X.400 Queues" in Chapter 7.

Uses for a Connector

Because of Exchange 2000 virtual server functionality, it is not necessary to create an SMTP connector to allow for mail flow, to connect it to other servers in an Exchange organization, or to connect it to the Internet. Furthermore, you don't need a connector if all of your Exchange 2000 servers connect to the Internet and successfully perform Domain Name System (DNS) lookups for Internet addresses.

However, although it is not essential for Internet mail delivery, an SMTP connector provides the following benefits:

- Provides simplified administration
- Provides limited your exposure to the Internet
- Establishes an isolated route for communicating with another domain or another mail system
- Routes mail to another mail system or relays mail to another domain
- Allows multiple bridgehead servers for load balancing
- Allows you to control how SMTP is used to communicate with other servers
- Permits scheduled connection times with customized settings

The following sections provide detailed information about each of these benefits.

Simplified Administration

A connector provides more administrative control over how Internet mail flows out of your organization. You can use a connector, or set of connectors, to limit the available routes for outgoing Internet mail. Also, because you need only check the SMTP queues and other configurations on a single server, using a single server as a gateway simplifies troubleshooting processes.

Limited Internet Exposure

One of the primary benefits of creating an SMTP connector is that you can route all inbound or outbound external SMTP mail through a particular server or set of gateway servers. By designating an isolated route for Internet mail using a connector, you limit your exposure to the Internet.

To use an SMTP connector to route Internet mail, specify one server or a set of servers as your gateway to the Internet, create an SMTP connector, and then set those servers as the source bridgehead servers of the connector.

Isolated Route for Communicating with Other Domains

You can also use a connector to establish an isolated route for communicating with other domains. This can be useful when you want to use secure communications with a particular company.

In previous versions of Exchange, there were options to configure settings per e-mail domain. Although these options are not available in Exchange 2000, you can create multiple SMTP connectors, set address spaces for these connectors, and then specify the settings that you want for those domains.

For example, suppose you want to use SSL to secure all e-mail sent to the military, but you do not want to use SSL for other e-mail communications. To achieve this, you would need two SMTP connectors:

- One with an address space of SMTP:*.mil
- One with an address space of SMTP:*

Because Exchange routes all mail through the connector that most closely matches the address space, all mail destined for the .mil domain would initially try to pass through the *.mil connector, if it was available. You could then specify that the *.mil connector only send mail to one server (a smart host) and that it use SSL and require authentication.

Load Balancing with Multiple Bridgehead Servers

When you have a single connector hosted by multiple bridgeheads, the servers using the connector will randomly select the bridgehead server they use, thereby load balancing requests across the bridgehead servers. The situation is different if you have multiple connectors with the same address space, each with a single bridgehead. The servers that use these connectors use a method based on the server GUID to determine which of the available connectors they will use. The algorithm may not evenly distribute the server selections across the available connectors. So, to achieve load balancing, you are better off using a single connector sourced to multiple bridgehead servers.

Use Specific SMTP or ESMTP Commands

You can use a connector to control how SMTP is used to communicate with other servers. To initiate SMTP sessions, you can choose whether your server uses the Extended Simple Mail Transfer Protocol (ESMTP) commands or SMTP commands, and you can control what type of commands your server issues.

When you configure an SMTP connection, the following communication options are available:

- Send or do not send server-side or client-side ETRN/TURN commands.

TURN is an SMTP command that allows the client and server to switch roles and send mail in the reverse direction without having to establish a new connection. ETRN is an ESMTP command sent by an SMTP server to request that another server send any e-mail messages it has. You can use these commands if you depend on a network service provider to hold your mail for you and deliver it upon request.

- Request ETRN/TURN from specific servers.
- Send HELO (an SMTP command) instead of EHLO (an ESMTP command).

HELO is an SMTP command sent by a client to identify itself, usually with a domain name; EHLO is an ESMTP command with which a server identifies its support for ESMTP commands.

Schedule and Customize Outbound Connections

You can use a connector to open an outbound connection at specified times. This is useful if you use a network service provider to deliver your outbound mail, or if you have limited bandwidth and want to control when external mail is sent.

You can also configure a connector to:

- Allow high, normal, or low message priorities for a domain.
- Allow system or non-system messages.

- Use different delivery times for oversized messages.
- Queue mail for remote triggered delivery.
- Set specific delivery restrictions.

3

SMTP Dependencies

To function properly, SMTP depends on the following components:

- Internet Information Services (IIS)
- Active Directory
- DNS
- Recipient policies

This chapter provides detailed information about each of these components and their interaction with SMTP.

Internet Information Services

IIS provides a framework process for Internet services such as HTTP, SMTP and NNTP. Do not confuse IIS with HTTP because several other services, such as SMTP, depend on IIS to function.

The installation of IIS provides:

- The framework process known as the IIS Admin Service, which allows administration of services through the IIS snap-in.
- Administrative consoles or snap-ins for the MMC.
- The IIS metabase, which is the configuration repository for IIS.
- Common files, which are shared libraries that provide socket connection pooling, registration, and management of these Internet services

Exchange setup requires that the World Wide Web Service, SMTP Service, and NNTP Service be installed. This prerequisite ensures that all the necessary components are installed prior to the installation of Exchange. Exchange leverages the core SMTP service through an event infrastructure. (For more information about event infrastructures, see the MSDN® Web site (<http://msdn.microsoft.com/>). After Exchange is installed, the SMTP service is only dependant on the IIS Admin Service. You can disable the World Wide Service without affecting the SMTP service; however, you cannot use the **Add/Remove Windows Component** option in **Add or Remove Programs** to disable the IIS Admin Service or to remove the IIS component entirely.

Installing IIS creates a number of virtual directories under the World Wide Web Service that are not required for any Exchange component, including Microsoft Outlook Web Access. To remedy this, Microsoft provides the IIS Lockdown Wizard, a security tool that removes unnecessary virtual directories, enhances file security, and processes real-time URL requests against user-defined configurations. To increase protection in the unlikely event that the World Wide Web Service is started in error, you should deploy the IIS Lockdown Wizard on every Exchange server and domain controller. You can download the IIS lockdown wizard from the Microsoft Download Center (<http://go.microsoft.com/fwlink/?LinkId=12281>).

Active Directory

Exchange 2000 is tightly integrated with Windows 2000 and Active Directory. Exchange stores all of its configuration information in Active Directory, including information about recipient policies, SMTP virtual server configuration, user mailboxes, and much more. However, SMTP reads its settings from the IIS metabase. Therefore, to supply IIS with the information it needs for SMTP functionality, Exchange System Attendant (an Exchange component) replicates the configuration information from Active Directory to the IIS metabase.

DNS

Although a complete analysis and discussion of DNS is beyond the scope of this book, this section provides information about the relationship between DNS and SMTP in Exchange. Because Exchange 2000 relies on DNS for name resolution, DNS plays a crucial role in Internet mail flow.

SMTP depends on DNS to determine the IP address of its next internal or external destination server. Generally, internal DNS names are not published on the Internet. Therefore, SMTP must be able to contact a DNS server that can resolve external DNS names to send Internet mail, as well as a DNS server that can resolve internal DNS names for delivery within the organization. For information about how to configure DNS for sending and receive mail, see Chapter 5, “Configuring Exchange to Send and Receive Mail.”

The following sections provide a general overview of DNS queries and an explanation of the role DNS plays in sending and receiving mail.

How DNS Queries Work

When a DNS client needs to resolve the name of a server, it queries the DNS servers. Each query that the client sends essentially asks the DNS server to provide the following information:

- A fully qualified domain name (FQDN)
- A specified query type, which can either specify a resource record by type or a specialized type of query operation. For mail servers, the query type specified is MX (mail exchanger resource record).
- A specified class for the DNS domain name.

Note For Windows DNS servers, this should always be specified as the Internet class (IN).

For example, the name specified could be a computer’s FQDN, such as “host-a.example.microsoft.com.”, and the query type specified to look for an MX resource record by that name. Think of a DNS query as a client asking a server a two-part question: First, “Do you have any MX resource records for a computer named ‘hostname.example.microsoft.com.’?” followed by “If so, can you resolve this MX record to an A (host) record and resolve its IP address?” When the client receives an answer from the server, it reads and interprets the answered MX resource record and gets the A record, thereby resolving the computer’s IP address.

Querying a DNS Server

When the DNS server receives a query, the server first checks to see if it can answer the query authoritatively, based on resource record information contained in a locally configured zone on the server. If the queried name matches a corresponding resource record in the local zone, the server answers authoritatively and uses this information to resolve the queried name.

If no zone information exists for the queried name, the server then checks to see if it can resolve the name using locally cached information from previous queries. If a match is found, the server answers with this information. Again, if the preferred server can provide the requesting client with a positive matched response from its cache, the query is completed.

If no zone or cached information exists for the queried name, the query process uses recursion to fully resolve the name. Recursion is the process in which a DNS server queries other DNS servers on behalf of the requesting client to fully resolve the name, then sends an answer back to the client. By default, the DNS Client service requires that the server use recursion to fully resolve names on behalf of the client before returning an answer. In most cases, the DNS server is configured (by default) to support the recursion process, as illustrated in the Figure 2.

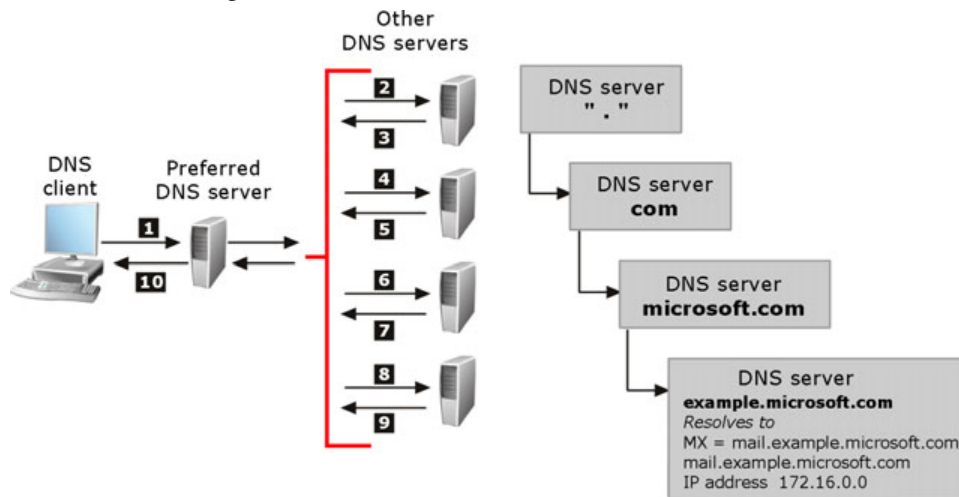


Figure 2 How DNS resolves a query for an MX record and finds the IP address

For more information, see the Windows 2000 DNS online help.

The Role of DNS in Receiving Internet Mail

In order to receive Internet mail, your external DNS servers must have an MX record pointing an A record with the IP address of your mail servers. To ensure that your MX records are configured correctly, you can use the Nslookup utility. To verify that your server is accessible on port 25 to other servers on the Internet, you can use telnet.

For specific information about Nslookup, see “Using Nslookup to Verify DNS Configuration” in Chapter 5.

For specific information about Telnet, see “Using Telnet to Ensure Internet Accessibility” in Chapter 5.

The Role of DNS in Sending Internet Mail

As mentioned earlier, Exchange uses one of two methods for sending Internet mail:

- Uses DNS for external name resolution
- Forwards mail to a smart host that assumes responsibility for mail delivery and name resolution

To send Internet mail using DNS rather than forwarding mail to a smart host, the Exchange server resolves the receiving domain and IP address of the recipient’s SMTP server. The server then uses SMTP over TCP port 25 to establish a conversation with the recipient’s SMTP server and deliver the mail.

Using DNS to Send Internet Mail

When using DNS, the most important thing to remember is that all servers in the DNS search order must be able to resolve external domains (also referred to as Internet domains). Because it is likely you will use internal servers for internal name resolution, you have three possible setup options:

- Set up your internal DNS servers as caching servers that use root hints for Internet domains. Root hints point to DNS servers that are authoritative for the zone containing the domain root and top-level domains. Root hints help DNS servers locate the correct server to resolve a domain name.
- Set up the internal DNS servers with forwarders to external DNS servers. A forwarder is a DNS server designated by an internal server to be used for resolving external DNS names. (To set up a forwarder, in the DNS console, select the DNS server. On the **Action** menu, click **Properties**, click the **Forwarders** tab, and then select the **Enable forwarders** check box. Add IP addresses for other DNS servers that act as forwarders for this server.)
- Configure the SMTP service to use external DNS servers. (To configure an external DNS server, right-click your SMTP virtual server, click **Properties**, and then click the **Delivery** tab. Click **Advanced**, and then click **Configure** to set up an external DNS server.) If you decide to use this method, ensure that you have the latest service pack for Exchange (SP3) installed.

The following steps show how Exchange uses DNS to resolve an external IP address. In this example, an internal client in the domain example.com sends a message to a recipient in the remote domain contoso.com. To route the message, Exchange uses DNS to resolve the IP address of the SMTP server in the contoso domain and deliver the message to the recipient at contoso.com. This process is also illustrated in Figure 3.

1. After the SMTP server in the domain example.com receives the message destined for the recipient at contoso.com, the SMTP virtual server contacts the appropriate DNS server and sends an MX query for the external domain of contoso.com.
2. The DNS server locates an A record that is associated with the MX record for contoso.com, and then uses that A record to determine the IP address. For more information about how the DNS server locates the A record, see “Querying a DNS Server” earlier in this chapter.
3. The DNS server returns the IP address of 172.23.234.23 for the mail server in contoso.com to the SMTP virtual server.
4. The SMTP virtual server opens a connection on port 25 of the remote SMTP server at the IP address of 172.23.234.23 and delivers the mail.

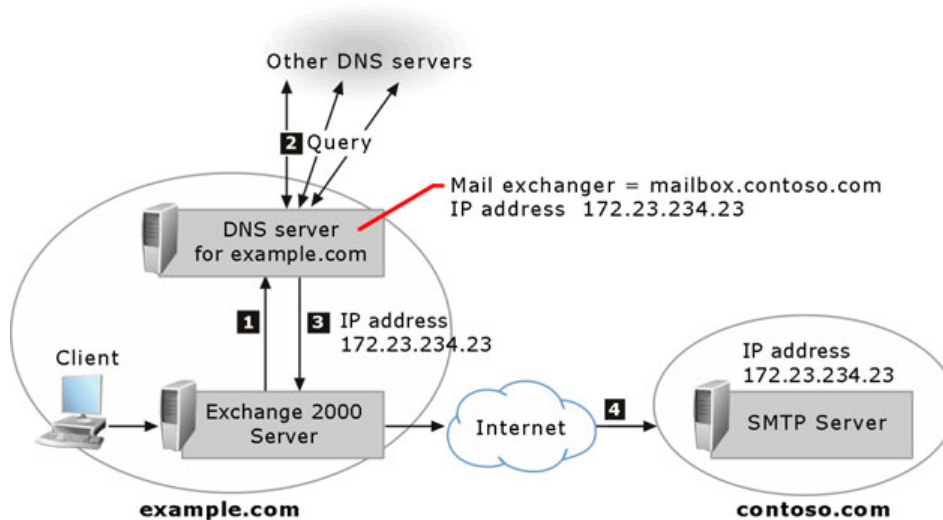


Figure 3 How Exchange uses DNS to resolve external IP addresses

Forwarding Internet Mail to a Smart Host

A smart host is a server or mail process that handles delivery of Internet mail. A smart host does not have to be an Exchange server—it can be any SMTP process or server that takes responsibility of delivering mail, either by sending it to another SMTP server or by using DNS to deliver the mail directly. In scenarios where there is a persistent connection to the Internet, a smart host is not required. However, often the smart host is an anti-virus scanner or a Windows 2000 SMTP service that is in a perimeter network.

Using a smart host for DNS resolution is similar to using a DNS server except that the smart host assumes responsibility of resolving the IP address and sending the mail (as detailed in steps 2-4 in the previous section).

For information about how to configure the Windows 2000 SMTP service in a perimeter network, see Microsoft Knowledge Base article Q293800, “XCON: How to Set Up Windows 2000 as a SMTP Relay Server or Smart Host” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=293800>).

For information about how to set up Exchange 2000 behind Microsoft Internet Security & Acceleration (ISA) Server, see the technical paper, *Configuring and Securing Exchange 2000 Server and Clients* (<http://go.microsoft.com/fwlink/?LinkId=10733>).

Recipient Policies

A recipient policy establishes the default e-mail addresses that use a specific protocol (such as SMTP) for a set of users. E-mail addresses are used to define the valid formats for addressing inbound e-mail to the Exchange system. The default recipient policy sets the mail domain for which the virtual server accepts incoming e-mail. It specifies the default SMTP and X.400 addresses for all Exchange 2000-based mailbox-enabled objects.

Any SMTP domains specified in the recipient policies are replicated into the IIS metabase and set as authoritative local domains. This means that SMTP accepts inbound mail for these domains. The only time an SMTP address is not considered local is, when adding the address to the recipient policy, you clear the **This Exchange Organization is responsible for all mail delivery to this address** check box in **SMTP Address Properties**.

A recipient policy can contain more than one e-mail address for a specified protocol (such as SMTP or X.400). For example, if all users in your Exchange organization have an external e-mail address of @example.com, but you want all your Seattle users to have two external mail addresses—one with @example.com and an additional e-mail address of @seattle.example.com—you can set up a recipient policy for all users in your Seattle office and add an additional address of @seattle.example.com. To achieve this you would perform the following procedure.

1. In Exchange System Manager, create a new recipient policy: In the console tree, expand **Recipients**, right-click **Recipient Policies**, point to **New**, and then click **Recipient Policy**.
2. In **New Policy**, click **E-mail Addresses** and then click **OK**.
3. In recipient policy properties, on the **General** tab, under **Filter Rules**, click **Modify** to create a filter that specifies all the users in the sales department.
4. On the **E-mail Addresses (Policy)** tab, click **New**.
5. In **New E-mail Address**, click **SMTP Address**.
6. In **SMTP Address Properties**, in **Address**, type **@seattle.example.com**. When you do this, you add an additional SMTP address of @seattle.example.com in addition to the @example.com SMTP address.
7. Specify a primary address.

When you have more than one address type, you must specify one address as primary. The primary address is the one that appears in the **From** line in outgoing e-mails. If you wanted the return e-mail address of the Seattle users to always appear as @seattle.example.com, you would set this address as the primary address.

For instructions about how to create recipient policies, see “Configuring Recipient Policies” in Chapter 5.

4

SMTP Deployment Scenarios

Now that you know all about how SMTP relates to Exchange, you're probably interested in all the ways you can deploy SMTP in your Exchange organization. To help you with this, Chapter 4 presents both some common and customized SMTP deployment scenarios.

Regardless of what scenario applies most to you, consider the following tips as you contemplate your own implementation of SMTP:

- If your organization contains multiple servers, you should include bridgehead servers when planning your deployment.
- Firewalls offer the most security for Internet connectivity.
- SMTP connectors offer a simple way to route outgoing Internet mail.
- The default SMTP virtual server in its default configuration is sufficient for most scenarios.
- If you use multiple SMTP virtual servers, be careful when configuring them. By default, multiple virtual servers cannot communicate with one another. For proper mail flow, you need to configure them appropriately so that mail can be routed between them. Additionally, each SMTP virtual server must be configured with a unique IP address and port combination. Generally, all SMTP virtual servers require port 25, so you must assign unique IP addresses to them.

Note Some companies configure multiple virtual servers on a bridgehead server, with one network interface card accepting inbound Internet mail and another routing outbound Internet mail. For more information about this configuration, see "Using a Dual-Homed Exchange Server as an Internet Gateway" later in this chapter.

Common Deployment Scenarios

This section presents some common SMTP deployment scenarios. The scenarios are presented in order of complexity, starting with the simplest configuration (a single Exchange server in its default configuration). Table 1 summarizes each of these common scenarios.

Table 1 Summaries of common SMTP deployment scenarios

Topology	Best for	Advantages	Considerations
Single Exchange server in its default configuration	Small business with a small user base	Using the default configuration requires no set-up on the Exchange server	This topology does not offer the more robust protection of a firewall. Your Exchange server is exposed on the Internet.
Dual-homed Exchange server	Small business with a small user base	Offers a secure configuration when behind a firewall	This topology should be used in conjunction with a firewall. Otherwise, your Exchange server is still exposed on the Internet. Consider using Internet Protocol security (IPSec) policies to filter ports on the Internet NIC.
Using an Exchange bridgehead server behind a firewall	Any size company	Using a designated bridgehead server for Internet mail isolates Internet traffic. A firewall protects your intranet.	N/A

Table 1 Summaries of common SMTP deployment scenarios (continued)

Topology	Best for	Advantages	Considerations
Using an Exchange bridgehead server to send mail to a relay server on a perimeter network	Medium to large companies with multi-server environments	Offers the same advantages as an Exchange bridgehead server behind a firewall, but adds an extra layer of security by isolating your SMTP server from the intranet. A simple SMTP relay server, rather than an Exchange server handling Internet mail, is in an isolated network. Your user information is secured on your Exchange server behind a firewall.	This topology involves more configuration and set up.

Note For small companies that want a full-featured network solution that provides a unified setup for e-mail, group scheduling, fax, database, as well as a shared Internet connectivity for an environment of up to fifty computers, Microsoft Small Business Server may be an appropriate solution. For more information about Small Business Server, see the Small Business Server Web site (<http://www.microsoft.com/sbserver/default.asp>).

Using a Single Exchange Server in Its Default Configuration

This scenario describes how Exchange delivers Internet mail in its default configuration.

Basic Configuration

In this scenario, you need the following:

- A persistent connection to the Internet
- A DNS server that can resolve external domain names, and a DNS server on the Internet with an MX record that points to your Exchange server
- A recipient policy configured with the SMTP mail domain for which you want the Exchange server to receive mail

Inbound Internet Mail

When using a single Exchange server in its default configuration, incoming Internet mail flows into the Exchange server in the following manner.

1. The remote SMTP server queries its own DNS server to resolve the MX record for your mail domain and to obtain the IP address of your Exchange server.
2. The remote SMTP server then connects to your default SMTP virtual server on port 25.
3. Your default SMTP server verifies that the domain on the incoming message matches an SMTP domain in its recipient policies.
4. Your default SMTP server then accepts the message and delivers it to the recipient.

Outbound Internet Mail

When using a single Exchange server in its default configuration, outgoing Internet mail flows out of the Exchange server in the following manner.

1. A user sends a message to an external user.
2. From its recipient policy information, the default SMTP virtual server determines that the message is destined for a remote domain.
3. Because the user is authenticated, the default SMTP virtual server accepts the message for outbound delivery. Remember, the default SMTP virtual server allows relaying for authenticated users only.
4. The default SMTP virtual server queries its DNS server to resolve the MX record of the remote mail server.
5. DNS returns an IP address for the remote mail server.
6. The default SMTP virtual server connects to the remote SMTP server on port 25 and initiates delivery.

Using a Dual-Homed Exchange Server as an Internet Gateway

This scenario describes a supported configuration of a dual-homed Exchange server that acts as a gateway server for the Exchange organization. This server can handle mail individually, or it can act as a bridgehead for other servers in the organization.

For security purposes, you should use this configuration behind a firewall.

Basic Configuration

The basic configuration consists of a mail gateway configured with two network interfaces; this gateway acts as the single connection point between your intranet and the Internet.

The following lists provide general configuration requirements for the two virtual servers and the SMTP connector:

Note If you configure two virtual servers on a single Exchange server, be sure to use a unique combination of IP addresses and ports. Do not configure either virtual server to use the default value of all available IP addresses.

Virtual server 1

- Configure virtual server 1 as the bridgehead server for the SMTP connector.
- Configure virtual server 1 to use external Domain Name System (DNS) servers, through the external DNS server list.
- Bind virtual server 1 to an intranet IP address on port 25.
- Enter the local company domain (for example, winery_co.co).

Virtual server 2

- Configure virtual server 2 so it does not relay mail (this is the default configuration). For specific instructions, see “Verifying Default Relay Restrictions on Your Inbound Virtual Server” in Chapter 5.
- Configure virtual server 2 to allow anonymous access (this is the default configuration). For specific instructions, see “Allowing Anonymous Access on the Outbound Virtual Server” in Chapter 5.
- Bind virtual server 2 to an Internet IP address on port 25.
- Select the local company domain (for example, winery_co.co).

SMTP connector

- Configure the SMTP connectors to use DNS to route to each address space on the connector.
- Home the SMTP connector to virtual server 1 by specifying it as the bridgehead server.
- Create an address space of * (asterisk) or an equivalent.
- Use two network interface cards (NICs)—an internal NIC and an external NIC.
- Verify that there is no IP routing configuration between the two networks on your server. (This is the default configuration.)

For more information about how to configure an SMTP connector, see “Configuring an SMTP Connector” in Chapter 5.

Inbound Internet Mail

Messages flow into an Exchange organization in the following manner.

1. Messages that originate from the Internet use the Internet IP address to send mail to recipients in the local domain.
2. Virtual server 2 monitors this Internet IP address for mail and receives all incoming Internet messages. Because virtual server 2 is not configured to relay mail, it rejects mail that is not directed to the company's domain (for example, winery-co.co).
3. When virtual server 2 receives a message from the Internet that is intended for a host inside the local domain, it contacts Active Directory through the internal NIC to determine where to send the message. Therefore, messages received by virtual server 2 are sent directly to the internal host.

Note Although virtual server 2 monitors an external IP address for incoming mail, it uses whatever IP address is appropriate for routing messages, based on the entries in the routing table. Virtual server 2 uses only internal DNS services for name resolution. Virtual server 2 is not configured with an external list of DNS servers, so it does not resolve external addresses. It rejects all messages with addresses to a domain other than the company's domain (in this case, winery-co.co).

Outbound Internet Mail

Mail flows out of an Exchange organization in the following manner.

1. A user sends a message to an external recipient.
2. Because this message is outbound, it uses the SMTP connector homed on virtual server 1.
3. When virtual server 1 receives a message for a remote domain, it uses the list of external DNS servers to find the IP address of the message recipient, and then uses the external NIC to deliver the external mail. (Generally, external Internet IP addresses are not available on an internal DNS server.)

Important Although virtual server 1 is configured to monitor the intranet IP address, it uses the Internet NIC for external mail.

Figure 4 illustrates the flow of mail through a dual-homed server.

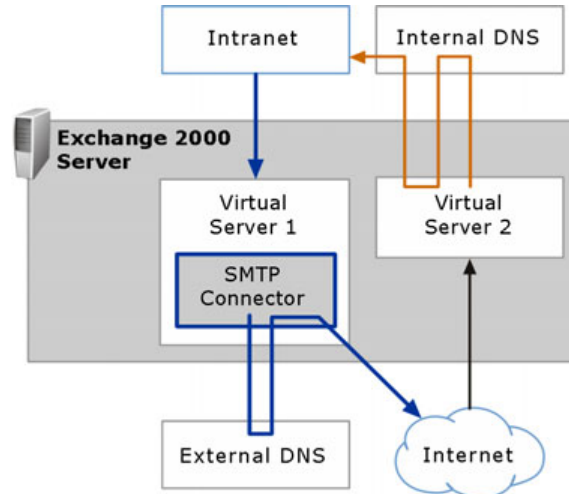


Figure 4 Internet mail flow through a dual-homed Exchange gateway server

Security Considerations

To increase the security of a dual-homed gateway server configuration, consider the following recommendations:

- Use Internet Protocol security (IPSec) policies to filter ports on the Internet NIC. For more information about IPSec policies, see the *Microsoft Exchange 2000 Server Resource Kit* (<http://go.microsoft.com/fwlink/?linkid=6544&clcid=0x409>) or the Microsoft Windows 2000 online documentation.
- Strictly limit the users you allow to log on to the server.

Using a dual-homed Exchange server as a gateway server in this configuration allows a company to limit its vulnerability by minimizing the entry points from the Internet to its intranet. By preventing the virtual server on the Internet from relaying messages to other Internet hosts, you ensure that the virtual server routes only mail that is addressed to valid internal recipients. Because virtual server 1 uses an external list of DNS servers to route only outbound Internet mail (not for internal mail), external DNS server issues won't affect internal mail traffic. By separating your incoming Internet mail, internal mail, and outgoing Internet mail processes, the points of failure for any of the three processes remain distinct and more manageable.

Using a Bridgehead Server Behind a Firewall

Generally, if your organization contains multiple Exchange servers, you should use a bridgehead server to provide Internet connectivity to a routing group or an Exchange organization.

Figure 5 illustrates this topology.




Figure 5 Providing Internet connectivity to a routing group

If you use a bridgehead server, it is not necessary for every Exchange server to have Internet connectivity. This increases security because only the bridgehead server is exposed to the Internet.

Important Because gateway servers usually have different security requirements than internal computers, you must examine them carefully for security risks.

Basic Configuration

The basic configuration consists of an Exchange bridgehead server that is connected to the Internet and has the appropriate DNS configuration. An SMTP connector is installed on the bridgehead server and provides outgoing message delivery over the Internet. Furthermore, to protect the internal network, a firewall filters incoming Internet traffic and routes mail from the internal and external IP addresses.

The following lists provide general configuration requirements for the DNS servers, the Exchange bridgehead server, the Exchange member servers, and the firewall:

DNS

Exchange relies on the existing DNS servers in its organization. Specifically, Exchange uses internal DNS to route internal messages and relies on the internal DNS server to forward and resolve external addresses through an external DNS server. To configure DNS in this way, ensure the following conditions are met:

- In order for the bridgehead server to be identified as the domain’s mail server, the organization’s external DNS server must contain an MX record for that bridgehead server. This allows inbound mail to be directed to the bridgehead server.
- The organization’s internal DNS server must have a forwarder to its external DNS server.
- The Exchange server should point to the internal DNS server.

For more information about how to configure DNS in this way, see “Verifying DNS Set Up for Inbound Mail” and “Configuring DNS for Outbound Mail” in Chapter 5.

Exchange bridgehead server

- The Exchange bridgehead server has an Internet connection through the firewall on port 25.
- The default SMTP virtual server is configured to send and receive Internet mail with the following default settings:
 - An IP address of port 25, the standard SMTP port.
 - Configured to allow anonymous access. This is because Internet SMTP servers that send mail to this domain will not expect to authenticate.
 - Configured to not relay mail.
- The SMTP connector that is hosted by the SMTP virtual server is configured with an address space of * (asterisk) to force all outgoing mail to use the bridgehead server.

Exchange member servers

- These servers do not have a direct connection to the Internet.
- These servers use the default settings on the SMTP virtual server.

Firewall

The firewall is configured in accordance with your organizational guidelines and vendor specifications.

Note A complete discussion about firewall configuration is outside the scope of this book. There are many ways you can configure a firewall to work with an SMTP relay server. You can allow either the firewall or the SMTP relay server to perform network address translation (between internal and external addresses). For the purposes of this book, mail flow through the firewall is treated as if it were transparent.

Inbound Mail Flow

Mail flows into an Exchange organization in the following manner.

1. Incoming mail flows through the firewall on port 25.
2. The SMTP virtual server allows the connection from the remote server over port 25, accepts the incoming message, and then routes the mail to the Exchange server that hosts the user's mailbox.

Outbound Mail Flow

Mail flows out of an Exchange organization in the following manner.

1. An internal user sends a message to a recipient in an external domain.
2. The internal user's Exchange server sends mail to the SMTP connector on the bridgehead.

Because the connector is configured with an address space of * (which denotes all external domains), each Exchange server in the routing group sends external e-mail through the SMTP connector on the bridgehead server.

3. The SMTP connector uses DNS to resolve the IP address of the recipient's e-mail server and route the mail directly to the recipient's SMTP server.

Using a Windows 2000 SMTP Relay Server in a Perimeter Network

Many organizations use a stand-alone Windows 2000 SMTP server in a perimeter network as a mail relay server for incoming and outgoing Internet mail. In this configuration, your Exchange organization is in an internal domain behind the firewall, while the SMTP server is in a separate domain in a perimeter network. Internal Exchange bridgehead servers route outgoing mail through a connector to the SMTP relay server, which assumes responsibility for DNS resolution and mail delivery. Similarly, you can configure the SMTP relay server to accept incoming Internet mail and route it internally.

Figure 6 illustrates this topology.

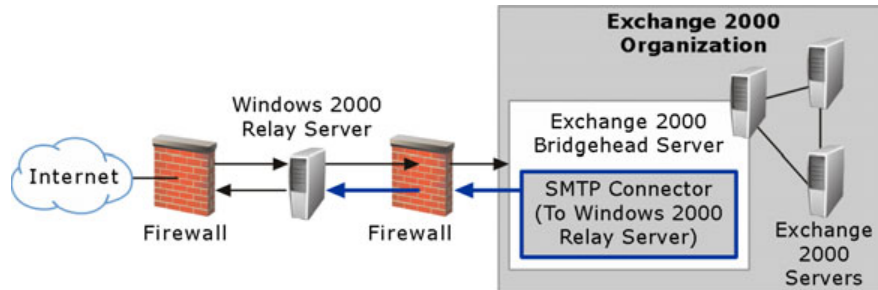


Figure 6 Windows 2000 relay server in a perimeter network

Advantages to using an SMTP relay server in a perimeter network include:

- Limited risk of Internet exposure. The internal network protects your Exchange servers that contain your user information and other configuration data.
- Additional security. You can install virus-scanning software to scan incoming mail before it reaches your internal network.

Basic Configuration

The basic configuration consists of the following:

Windows 2000 SMTP relay server

The SMTP relay server is configured with a default public domain. It is also configured to relay messages for only SMTP mail domains within the Exchange organization—it does not relay messages to other domains. The following procedure provides detailed steps about how to configure the SMTP relay server.

► To configure a Windows 2000 server as a relay server or smart host

1. Verify that SMTP is installed on the Windows 2000 server. To verify that SMTP is installed:
 - a. In Control Panel, double-click **Add/Remove Programs**, and then click **Add/Remove Windows Components**.
 - b. Under **Components**, highlight to select **Internet Information Services (IIS)**, and then click **Details**.
 - c. Under **Subcomponents of Internet Information Services (IIS)**, verify that the **SMTP Service** check box is selected. If the check box is not selected, select it, click **OK**, and then complete the installation instructions.

2. In Internet Services Manager, add the SMTP mail domain for which you want the Windows server to relay. To add the SMTP domain:
 - a. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Internet Services Manager**.
 - b. Expand the server you want, and then expand the default SMTP virtual server. By default, the default SMTP virtual server has a local domain with the fully qualified domain name for the server.
 - c. To create the inbound SMTP mail domain, right-click **Domains**, point to **New**, and then click **Domain**.
 - d. In **New SMTP Domain Wizard**, click **Remote** as the domain type, and then click **Next**.
 - e. In **Name**, type the domain name of your SMTP mail domain for your Exchange organization.
 - f. Click **Finish**.
3. Configure the SMTP mail domain you just created for relay:
 - a. In Internet Services Manager, right-click the SMTP mail domain, and then click **Properties**.
 - b. Click **Allow the Incoming mail to be Relayed to this Domain**.
 - c. Click **Forward all e-mail to smart host**, and then type the IP address in square brackets ([]) or the FQDN of the Exchange server responsible for receiving e-mail for the domain. For example, to enter an IP address, type [123.123.123.123].
 - d. Click **OK**.
4. Specify the hosts that you want to openly relay to all domains:
 - a. In Internet Services Manager, right-click **Default Virtual Server** and click **Properties**.
 - b. On the **Access** tab, click **Relay**.
 - c. Click **Only the list below**, click **Add**, and then add the hosts that you want to use the SMTP server to send mail.
 - d. Under **Single computer**, specify the IP address of the Exchange bridgehead server that you want to relay using this SMTP server. Click **DNS Lookup** to find the IP address of the specific server.

For more information about how to configure a Windows 2000 server as a relay server or smart host, see the Microsoft Knowledge Base article Q293800, "XCON: How to Set Up Windows 2000 as a SMTP Relay Server or Smart Host"

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=293800>).

DNS

- Your external DNS server is configured with an MX record that points to the IP address of your SMTP relay server's domain.
- All Exchange servers point to your internal DNS server

Exchange bridgehead server

The Exchange bridgehead server is connected to the Internet through the firewall on port 25.

SMTP virtual server

The SMTP virtual server is configured to send and receive Internet mail with the following default settings:

- IP address of port 25, the standard SMTP port.
- Configured to allow anonymous access. This is because Internet SMTP servers that send mail to this domain will not expect to authenticate.
- Configured to not relay mail.

SMTP connector

- The SMTP virtual server hosts the connector.
- The connector is configured with an address space of * (asterisk) to force all outgoing mail to use the Exchange bridgehead server.
- The connector is configured to use the SMTP relay server as a smart host to relay mail.
- All other settings remain at their default values.

Other Exchange member servers

- Member servers do not have a direct connection to the Internet.
- All member servers use the default SMTP virtual server with its default settings.

Firewall

The firewall is configured according to your organizational guidelines and vendor specifications.

Note A complete discussion about firewall configuration is outside the scope of this book. There are many ways you can configure a firewall to work with an SMTP relay server. You can allow either the firewall or the SMTP relay server to perform network address translation (between internal and external addresses). For the purposes of this book, mail flow through the firewall is treated as if it were transparent.

Inbound Mail Flow

When using a relay server in a perimeter network, inbound Internet mail flows into the Exchange organization in the following manner.

1. Incoming Internet mail flows through port 25 on the firewall.
2. Mail is then sent to port 25 of the SMTP relay server in the perimeter network.
3. The SMTP relay server routes the mail back through the firewall to the Exchange bridgehead server.
4. The Exchange bridgehead server uses SMTP and internal routing to deliver mail to the Exchange server that hosts the user's mailbox.

Outbound Mail Flow

When using a relay server in a perimeter network, outbound Internet mail flows out of the Exchange organization in the following manner.

1. An internal user submits a message to a remote user.
2. The Exchange server on which the user's mailbox resides forwards mail to the SMTP connector on the Exchange bridgehead server.
3. The SMTP connector relays the mail through the firewall to the SMTP relay server in the perimeter network.
4. The SMTP relay server uses DNS to find the MX record and IP address of the remote user's SMTP server.
5. The SMTP relay server sends mail back through the firewall to port 25 of the remote user's SMTP server.

Custom Deployment Scenarios

This section presents two custom deployment scenarios, including overviews of the general configuration requirements for each one.

Using a Network Service Provider to Send and Receive Mail

If your Exchange server uses a dial-up connection to send and retrieve Internet mail, you must have a dial-up account to your network service provider. Furthermore, you must configure the Windows 2000 Routing and Remote Access Service (RRAS) to dial and authenticate with the network service provider on demand. For more information about configuring RRAS, see the Microsoft Windows 2000 online help.

If you want to use a network service provider's SMTP server as a smart host (also known as a relay server) to deliver outbound e-mail, then you can verify addresses on outgoing mail when you send it. Mail can be sent on demand, or you can set up a specific delivery schedule. To configure these settings, use the **Delivery options** tab in the SMTP connector's properties.

To retrieve e-mail from the smart host, on the **Advanced** tab of the SMTP connector's properties, click **Request ETRN/TURN when sending messages**. As mentioned earlier, ETRN is an ESMTP command sent by an SMTP server to request that another server send any e-mail messages it has. TURN is an SMTP command that allows the client and server to switch roles and send mail in the reverse direction without having to establish a new connection. This ability to switch during a SMTP session is useful because you can send mail and then issue the TURN command to receive mail without having to re-establish a new connection. Additional times can be specified for retrieval purposes only.

If you want to send e-mail directly to remote domains without using the network service provider's e-mail server as a smart host, you can configure the SMTP connector to use DNS to send mail. However, you can still retrieve mail from your network service provider. To retrieve mail from your network service provider, select the **Request ETRN/TURN from different server** option on the **Advanced** tab of the SMTP connector's properties. If you configure the SMTP connector in this way, you are required to set up a schedule for retrieval.

Supporting Two SMTP Mail Domains and Sharing an SMTP Mail Domain with Another System

There are special situations (mergers and acquisitions, in particular) that necessitate the support of two namespaces and the sharing of a namespace with another system.

To help explain this situation, consider the merger of two fictitious companies: Northwind Traders and Fourth Coffee. Northwind Traders (northwindtraders.com) acquires Fourth Coffee (fourthcoffee.com). The process of consolidating domain namespaces is as follows.

1. Northwind Traders configures its Exchange organization to accept mail for the non-local domain of fourthcoffee.com. For information about accepting mail for multiple domains, see "Supporting Two SMTP Mail Domains" later in this chapter.
2. Both systems eventually share the SMTP mail domain northwindtraders.com.
3. Finally, the users are migrated to a single Exchange organization, and the old organization or system is removed.

Supporting Two SMTP Mail Domains

Supporting two SMTP mail domains is common during the initial phase of a merger or acquisition.

To elaborate on how one Exchange organization can support two SMTP mail domains, consider the same merger scenario involving Northwind Traders and Fourth Coffee. In the initial phases of the acquisition, Northwind Traders continues to use its local SMTP mail domain of northwindtraders.com. However, to allow Fourth Coffee employees to receive e-mail with their original address, Northwind Traders must also accept mail for the non-local mail domain of fourthcoffee.com. Figure 7 illustrates how both the domains of fourthcoffee.com and northwindtraders.com are supported.

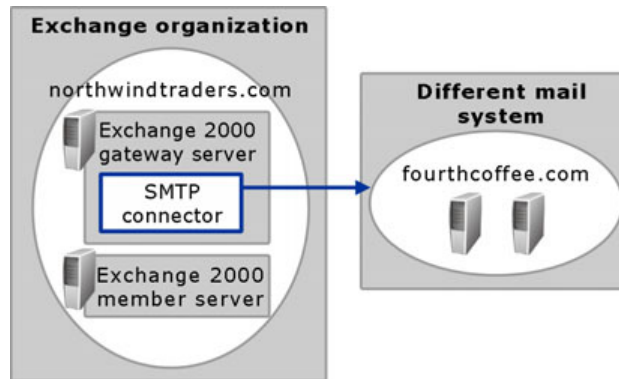


Figure 7 Supporting two SMTP mail domains

To accept mail for the non-local domain of the newly acquired company, Fourth Coffee, an administrator at Northwind Traders creates an SMTP connector to fourthcoffee.com. This connector is configured with an address space of the SMTP domain used by Fourth Coffee (fourthcoffee.com) and configured to relay messages to this domain. To do this, the administrator opens the SMTP connector's properties, clicks the **Address space** tab, and then selects the **Allow messages to be relayed to this domain** check box.

Important You must configure this connector on each bridgehead server that accepts incoming Internet e-mail for the fourthcoffee.com domain.

Additionally, for the mail domain (fourthcoffee.com) that the administrator wants to accept mail, he ensures that an MX record exists on the Internet DNS server. This MX record should point to the IP address of the gateway server that accepts inbound mail. For more information about DNS, see "DNS" in Chapter 3.

Sharing an SMTP Mail Domain with Another System

Sharing an SMTP mail domain between an Exchange 2000 organization and another e-mail system or another Exchange 2000 organization is common during the final stages of a merger or acquisition. To continue with the previous scenario, assume that Northwind Traders is in the final stages of consolidating its systems with those of the newly acquired company, Fourth Coffee. Mailboxes in both the Exchange 2000 organization (which contains all of the employees of Northwind Traders) and the other system (which contains all of the employees of Fourth Coffee) now use the same SMTP domain of northwindtraders.com in their addresses. Ideally, the best way to share an SMTP mail domain is to allow Exchange to accept incoming mail from the Internet, locate a matching recipient in the Exchange organization, and then forward the mail to the users on the other mail system. Figure 8 illustrates a shared domain with another system.

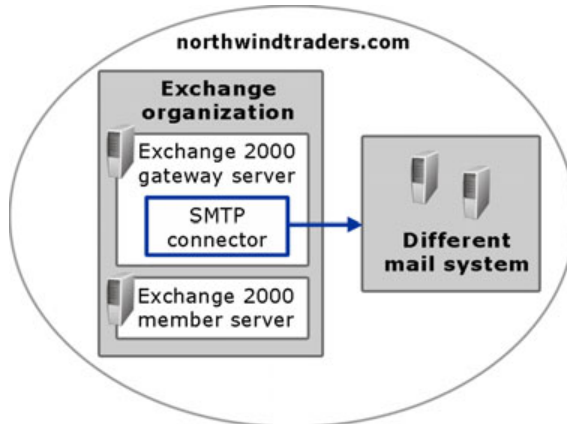


Figure 8 Sharing an SMTP domain

If Exchange functions as the first mail server, there are two methods you can use to configure Exchange to share an SMTP address space:

Note Before using either of these methods, ensure that you are running the latest service pack of Exchange (SP3).

Method 1: Selective Name Sharing

In method one, the mail systems share only selected SMTP address spaces—Exchange remains authoritative over the others. This is the preferred method because it is the most flexible. Also, you must use this method if any of the following conditions exist in your environment:

- You create contacts in Active Directory for sending mail to external recipients.
- The target SMTP addresses of those external recipients matches any of the SMTP domains that are configured in Exchange 2000 recipient policies. For example, if the address @northwindtraders.com is configured on one of your recipient policies, and you want to create contacts with a target address of @northwindtraders.com, you must use this method to share the @northwindtraders.com SMTP mail domain.

Method 2: Share All Address Spaces

Although method two is less flexible, it is easier to configure in small environments. However, you cannot use this method if contacts exist in Active Directory for the external recipients on the other mail system. For information about using contacts in a shared SMTP domain, see Microsoft Knowledge Base article Q319759, “XADM: How to Configure Exchange 2000 Server to Forward Messages to a Foreign Messaging System That Shares the Same SMTP Domain Name Space” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=319759>).

Method 1: Sharing Selected Namespaces

Method one offers excellent flexibility because you can create contacts in Active Directory and more easily migrate users to a single system. This method uses two basic principles:

- **An SMTP connector is created with an address space of the remote domain, fourthcoffee.com** The connector allows messages to be relayed to this domain. This permits Exchange to accept inbound messages for this domain.

Important You must configure this connector on each bridgehead server that accepts incoming Internet e-mail for the fourthcoffee.com domain.

- **Exchange is non-authoritative over the domain** If Exchange is authoritative over a domain, it assumes that all the addresses in the domain exist in its organization. Therefore, if messages cannot be resolved locally, Exchange never attempts to send the messages through an external connector. By configuring Exchange to be non-authoritative for the domain, if the user cannot be found locally, Exchange routes the message through the connector to the remote system.

Note In this case, because this SMTP mail domain is non-authoritative, it is irrelevant that Exchange accepts messages that are inbound for domains it is authoritative over. The connector configuration ensures that the Exchange organization accepts mail for this domain—this is because the connector is configured with an SMTP address space of the remote domain and allows relaying to this domain. Exchange only accepts inbound e-mail for the shared SMTP domain because the connector to the remote e-mail system allows messages to be relayed to this address space. Because Exchange is non-authoritative for the shared mail domain, if you remove the connector, Exchange stops accepting inbound mail for this SMTP domain. Therefore, if you remove the connector, remember to change the recipient policy and make Exchange authoritative for this SMTP mail domain.

There are three main steps to using method one (each step is detailed further in the sections following).

Step 1

Determine if Exchange is authoritative over the SMTP mail domain you want to share.

Step 2

Configure the recipient policy for the SMTP mail domain you want to share. How you do this depends on whether the SMTP mail domain exists on the default recipient policy, on another recipient policy, or if it does not yet exist on a recipient policy.

Step 3

Create an SMTP connector to route mail to the other mail system or host.

Step 1: Determine if Exchange is Authoritative Over the SMTP Mail Domain You Want to Share

Before you configure your recipient policy for the SMTP mail domain you want to share, you must determine if Exchange is authoritative over the domain.

Remember, depending on whether Exchange 2000 is authoritative or non-authoritative, Exchange treats e-mail messages differently for particular SMTP addresses. Because Exchange does not forward messages that it cannot resolve locally for an authoritative domain, you must ensure that Exchange is not authoritative over the SMTP mail domain you want to share.

► **To view the setting that determines whether Exchange is authoritative**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Recipients**, and then click **Recipient Policies**.
3. In the details pane, right-click a recipient policy, and then click **Properties**.

4. Click the **E-Mail Addresses (Policy)** tab, select an SMTP address, and then click **Edit**. A dialog box similar to Figure 9 displays.

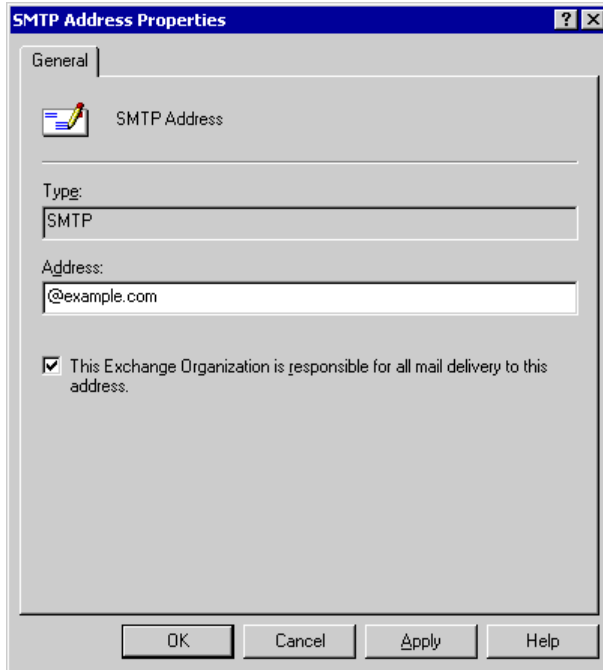


Figure 9 The **SMTP Address Properties** dialog box for an authoritative domain

5. If the **This Exchange Organization is responsible for all mail delivery to this address** check box is selected, then Exchange is authoritative for the address. If the check box is cleared, then Exchange is not authoritative for the address.

For more information about authoritative and non-authoritative SMTP domains in Exchange 2000, see Microsoft Knowledge Base article Q315591, "XCON: Authoritative and Non-Authoritative Domains in Exchange 2000" (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=315591>).

Step 2: Configure the Recipient Policy for the SMTP Mail Domain You Want to Share

When configuring the recipient policy for the SMTP mail domain you want to share, there are three possible scenarios you may encounter:

Scenario 1

The SMTP mail domain you want to share exists on the default recipient policy.

Scenario 2

The SMTP mail domain exists you want to share exists on another recipient policy.

Scenario 3

The SMTP mail domain you want to share does not exist on a recipient policy.

Scenario 1: Configuring the Shared SMTP Domain If It Exists on the Default Recipient Policy

You cannot set Exchange to be non-authoritative over the default recipient policy's primary SMTP address space. In order to prevent Exchange from being authoritative over this domain, you need to change the default recipient policy by adding a new primary address space strictly for internal use. This address could be similar to @localhost, signifying that it is used solely for internal mail flow within your Exchange organization. After you add the new address space, you must make it non-authoritative by editing it for the SMTP mail domain that you want to share.

To configure Exchange to share a mail domain that exists as the primary address space on the default recipient policy, you must perform the following tasks.

1. On the default recipient policy, add a new primary address space over which Exchange is authoritative, and then make the shared address space non-authoritative.
2. Create a second recipient policy that has the same search filter as the default recipient policy. Then, assign the second recipient policy a higher priority than the default recipient policy so the reply-to or return address is displayed as the shared address space.

This step is necessary because Exchange uses the primary address space as the reply-to address that is displayed in outgoing mail. Because you want outgoing messages to display the shared namespace on the reply-to line, you must create another recipient policy that is also non-authoritative but has a higher priority; therefore, Exchange uses this address space on the return address of outgoing mail. Because the new recipient policy is not the default recipient policy, you can make this address space non-authoritative.

Perform the following procedure to create a new primary address space on the default recipient policy and make the shared address space non-authoritative.

► **To modify the default recipient policy**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Recipients**, and then click **Recipient Policies**.
3. In the details pane, right-click your default recipient policy, and then click **Properties**.
4. Click the **E-Mail Addresses (Policy)** tab, and then click **New**.
5. In **New E-mail Address**, click **SMTP Address**, and then click **OK**.
6. In **SMTP Address Properties**, in the **Address** box, type **@localhost** or some other address space for which the Exchange organization can be authoritative. You can use **@localhost** or your Active Directory domain if it is different from your Internet domain. This address space is strictly for internal use.
7. Verify that the **This Exchange Organization is responsible for all mail delivery to this address** check box is selected.
8. Click **OK**.
9. On the **E-mail Addresses (Policy)** tab, click the new SMTP address you just created, and then click **Set as Primary**.
10. Click the SMTP address space that you want to share (for example, northwindtraders.com), and then click **Edit**.
11. To make Exchange non-authoritative for this SMTP address, clear the **This Exchange Organization is responsible for all mail delivery to this address** check box.
12. Click **Apply**.
13. A message displays asking if you want to update all corresponding recipient e-mail addresses. Click **Yes**.
14. On the **E-mail Addresses (Policy)** tab, click **OK**.

Changing the default recipient policy in this way causes Exchange to use the new primary address as the return or reply-to address in outgoing e-mails. In the example above, all users in this policy now have a return e-mail address that matches the new primary address space of @localhost. Because you want all your users to have the return address of the shared mail domain (in this case, northwindtraders.com), you must create a new recipient policy with a higher priority recipient policy that contains the northwind.com address space. Exchange uses the higher priority recipient policy on the return address. Furthermore, because this recipient policy is not the default recipient policy, you can make it non-authoritative. (Remember, this address space must be non-authoritative in order for Exchange to route it through the connector to the external system.)

Perform the following procedure to create a higher priority recipient policy so that outgoing e-mails display the correct return (reply-to) address.

► **To create a higher priority recipient policy with the shared mail domain**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Recipients**, right-click **Recipient Policies**, point to **New**, and then click **Recipient Policy**.
3. In **New Policy**, select the **E-Mail Addresses** check box, and then click **OK**.
4. On the **General** tab, in the **Name** box, type an appropriate name, such as “User Addresses.”
5. Under **Filter rules**, click **Modify**.
6. In **Find Exchange Recipients**, select or clear the appropriate check boxes to specify all applicable users. If you want to apply the policy to all users, click **OK**.
7. On the **E-mail Addresses (Policy)** tab, click the SMTP mail domain that you want to share, and then click **Set as Primary** (leaving the @local domain as a secondary proxy).
8. Click **Apply**.
9. A message displays asking if you want to update all corresponding recipient e-mail addresses. Click **Yes**.
10. On **E-mail Addresses (Policy)** tab, click **OK**.

Scenario 2: The SMTP Domain You Want to Share Exists on Another Recipient Policy

If the SMTP domain that you want to share is not on the default recipient policy, you can simply make the address space non-authoritative.

► **To modify an existing recipient policy for the SMTP domain you want to share**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Recipients**, and then click **Recipient Policies**.
3. In the details pane, right-click the recipient policy that has the SMTP address space you want to share, and then click **Properties**.
4. On the **E-mail Addresses (Policy)** tab, click the SMTP address space, and then click **Set as Primary**.
5. Click the SMTP address space that you want to share, and then click **Edit**.
6. To make Exchange non-authoritative for this SMTP address, clear the **This Exchange Organization is responsible for all mail delivery to this address** check box.
7. Click **Apply**.
8. A message displays asking if you want to update all corresponding recipient e-mail addresses. Click **Yes**.
9. On **E-mail Addresses (Policy)** tab, click **OK**.

Scenario 3: The SMTP Domain You Want to Share Does Not Exist on a Recipient Policy

If the SMTP domain that you want to share does not exist on a recipient policy, you can create a new recipient policy with the address space and make it non-authoritative.

► To create a new recipient policy for an SMTP mail domain that does not exist on a recipient policy

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Recipients**, right-click **Recipient Policies**, point to **New**, and then click **Recipient Policy**.
3. In **New Policy**, select the **E-Mail Addresses** check box, and then click **OK**.
4. On the **General** tab, in the **Name** box, type a name for your new policy.
5. On the **E-Mail Addresses (Policy)** tab, click the SMTP address space, and then click **New**.
6. In **New E-mail Address**, click **SMTP Address**, and then click **OK**.
7. In **SMTP Address Properties**, in the **Address** box, type the SMTP address space that you want to share.
8. To make Exchange non-authoritative for this SMTP address, clear the **This Exchange Organization is responsible for all mail delivery to this address** check box.
9. In **SMTP Address Properties**, click **OK**.
10. On **E-mail Addresses (Policy)** tab, click **OK**.

Step 3: Create an SMTP Connector to Route Mail to the Other Mail System

Now that Exchange 2000 is non-authoritative for the shared SMTP domain, when Exchange 2000 cannot find a matching address in Active Directory, it attempts to locate an external path to this domain. To find this path, Exchange first searches for a connector and then checks Domain Name System (DNS). Unless the MX record for that domain already points to the server on which the other mail system resides (in many cases the MX record points to the Exchange 2000 server itself), you must create an SMTP connector to route the mail to a specific host.

Important You must configure this connector on each bridgehead server that accepts incoming Internet e-mail for the fourthcoffee.com domain.

► To create an SMTP connector to route mail to a specific host

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, right-click **Connectors**, point to **New**, and then click **SMTP Connector**.

3. On the **General** tab, type an appropriate name, and then click the **Forward all mail through this connector to the following smart hosts** option. In square brackets ([]), type the fully qualified domain name (FQDN) or the IP address of the server to which e-mail for the shared SMTP address space is to be routed.
4. Click **Add** to configure your bridgehead servers, and then select your Exchange gateway servers that accept Internet mail for this domain.
5. Click the **Address Space** tab, click **Add**, click **SMTP**, and then click **OK**.
6. In **E-mail domain**, type the SMTP address space without the “at” symbol (@), for example, **fourthcoffee.com**, and then click **OK**.

Warning It is important to enter the specific SMTP mail domain. Do not type * (asterisk) on the SMTP connector. Setting * causes Exchange to accept mail for all external domains and then relay it externally. This configuration allows open relaying for anyone on the Internet and is extremely insecure.

7. Because Exchange 2000 must also receive messages for this domain, on the **General** tab, click **Allow messages to be relayed to these domains**. This setting makes it possible for all SMTP virtual servers that are listed under **Local Bridgeheads** to accept messages for domain.
8. Click **OK**.

After you configure these settings, when Exchange 2000 cannot locate a local address match in that SMTP domain, Exchange forwards the mail to the host that has the matching address space, as specified on the SMTP connector.

Method 2: Sharing All Address Spaces

Method two involves sharing all address spaces or SMTP mail domains. Although this configuration is easier to perform, it is much less flexible. In this configuration, Exchange 2000 is authoritative for all address spaces. You cannot have any contacts in your directory that have a target address matching a domain over which Exchange 2000 is authoritative.

► To share all address spaces in your Exchange organization

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand <Server Name>, expand **Protocols**, and then expand **SMTP**.
3. Right-click your SMTP virtual server, and then click **Properties**.
4. In the SMTP virtual server's **Properties**, click the **Messages** tab.

5. In the **Forward all messages with unresolved recipients to host** box, type the IP address, in square brackets ([]), or the FQDN of the server that will receive unresolved mail.
6. Click **OK**.
7. Repeat this procedure for the default SMTP virtual server on all Exchange 2000 servers, except for any virtual server that is acting as an inbound gateway for the other system. It is recommended that no mailboxes reside on this server.

Remember, this setting only affects authoritative domains. Therefore, in an authoritative domain, any message sent to an unresolved address is forwarded to the server that is specified on the SMTP virtual server. Any non-authoritative domain in Exchange 2000 is not affected by this setting. Any message sent to an unresolved address in a non-authoritative domain is routed to a matching SMTP connector, if present. If no matching SMTP connector is located, the message is sent to the server that is specified in the MX record found in DNS.

Supporting Additional Mail Systems

As described in the preceding scenarios, the other mail system that receives mail forwarded by Exchange may perform the same tasks as Exchange and forward mail to a third e-mail system. To avoid mail looping, it is essential that the last e-mail system (to which mail is forwarded) is authoritative for the domain. In other words, the final receiving mail system must search for a matching recipient; if the system does not find a matching recipient, it generates a non-delivery report (NDR) for the message. Mail looping occurs when the receiving system searches for a match in its recipients and then forwards the mail back to the original system when a match is not found.

If Exchange is the last system in this configuration, by default, it will return an NDR for any unresolved messages. However, it is preferable to create custom recipients in Active Directory for all recipients that reside on a different mail system. These recipients should have target addresses similar to *@subdomain.example.com*, where *subdomain* provides additional address information to distinguish the address space from the typical *@example.com* namespace; for example, *@microsoft.example.com*

5

Configuring Exchange to Send and Receive E-Mail

You understand how SMTP relates to Exchange. You've been introduced to all the ways you can deploy SMTP in your Exchange organization. Now it's time to get down to business. Chapter 5 contains procedural information about how to configure your Exchange 2000 organization to send and receive Internet mail. Specifically, you will learn how to:

- Verify that all the SMTP commands have been properly installed on your Exchange server.
- Configure Exchange to send Internet mail.
- Configure Exchange to receive Internet mail.

Verifying SMTP Port Settings

For mail to flow properly, SMTP must be installed correctly on the Exchange server with all of the necessary commands. If you experience mail problems, you should first verify the basic functionality of your SMTP installation.

When an Exchange server uses SMTP to communicate, it must have access to port 25. When SMTP is configured correctly, Exchange provides extended SMTP verbs to allow for proper communication. These verbs are controlled in the IIS metabase and in Exchange event sinks.

To determine whether or not the proper extended Exchange verbs are loaded, you can perform a telnet test. To perform this test, telnet to port 25 of your Exchange server's IP address. For example, type the following text at a command prompt:

telnet <server IP address> 25

where *server IP address* is the IP address of your Exchange server, and **25** indicates a connection to TCP port 25. The following example shows a telnet command to connect to port 25 on a server with an IP address of 172.16.0.0:

```
telnet 172.16.0.0 25
```

Next, type **ehlo <server name>**, where *server name* is the fully qualified domain name of your Exchange server. Your Exchange server then responds by listing the SMTP and ESMTP verbs that it supports.

Example 1 lists the verbs you will receive if SMTP is loaded properly. If SMTP is not configured properly, you will see only the verbs listed in Example 2.

Example 1 SMTP extended verbs (if Exchange event sinks are loaded properly)

```
ehlo example.com
250-mail1.example.com Hello [172.16.0.0]
250-TURN
250-ATRN
250-SIZE 5242880
250-ETRN
250-PIPELINING
250-DSN
250-ENHANCEDSTATUSCODES
250-8bitmime
250-BINARYMIME
250-CHUNKING
250-VERFY
250-X-EXPS GSSAPI NTLM *
```

```
250-AUTH GSSAPI NTLM
240-X-EXPS=LOGIN *
250-X-LINK2STATE *
250-XEXCH50 *
250 OK
```

* These extended verbs should be displayed.

When Exchange SMTP is not loaded properly, or the IIS metabase is corrupt, the extended Exchange verbs do not appear in the server's response. Example 2 lists the verbs you will receive if Exchange SMTP is not loaded properly.

Note The verbs listed in Example 2 are the same as the verbs you would see if you had never installed Exchange.

Example 2 SMTP extended verbs (if Exchange 2000 sinks are not loaded)

```
ehlo example.com
250-mail1.example.com Hello [172.16.0.0]
250-TURN
250-ATRN
250-SIZE 5242880
250-ETRN
250-PIPELINING
250-DSN
250-ENHANCEDSTATUSCODES
250-8bitmime
250-BINARYMIME
250-CHUNKING
250-VRFY
250-AUTH GSSAPI NTLM
250 OK
```

If you receive only the SMTP verbs listed in Example 2, the SMTP service for Windows 2000 is installed, but SMTP in Exchange is not loaded properly. Note that all verbs starting with "X" ("X" = eXtended) are missing.

Other incomplete lists can also indicate that Exchange is not properly loaded or that there is a possible corruption of the IIS metabase. Corruption of the IIS metabase can occur for any of the following reasons:

- Re-installing Exchange 2000
- Reinstalling Windows 2000 Server
- Removing or disabling IIS If there is corruption to the IIS metabase, you must perform one of the following tasks: Properly load Exchange SMTP
- Run the SMTP Reinstall Tool

For Exchange 2000 SP2 or later, you can use the SMTP Reinstall Tool (Smtpreinstall.exe) to restore the missing Exchange 2000 ESMTP verbs and the Exchange 2000 versions of the files. You can find Smtpreinstall.exe in the \Server\Support\Utils\i386 folder on the Exchange 2000 SP2 or later CD.

Note You cannot use the SMTP Reinstall Tool on an Exchange cluster. In this case, you must reinstall Exchange.

► **To properly load Exchange SMTP**

1. Reapply the latest Windows 2000 service pack.
2. Reinstall Exchange. Reinstalling Exchange replaces any missing files without affecting the settings on the Exchange server.
3. Reapply any Exchange service packs and any other Exchange-related program updates (for example, any Exchange updates available from the Microsoft Web site at <http://www.microsoft.com/exchange>).

Note Subscribe to the Hotfix and Security Bulletin Service to automatically receive notifications about any security-related Exchange updates. You can register for the Hotfix and Security Bulletin Service at (<http://go.microsoft.com/fwlink/?linkid=12322&clcid=0x409>).

► **To run smtpreinstall.exe**

1. Copy Smtpreinstall.exe to the \Exchsrvr\Bin folder.
2. Run Smtpreinstall.exe from the \Exchsrvr\Bin folder.
3. Restart the computer when prompted to do so.

Setting Up Your Exchange Server to Receive Internet Mail

This section explains how to set up your Exchange server to receive Internet mail. Specifically, you will learn how to:

- Configure recipient policies.
- Configure inbound SMTP virtual server settings.
- Verify DNS setup for inbound mail.

Configuring Recipient Policies

Exchange uses recipient policies to determine which messages should be accepted and internally routed to mailboxes in your organization. Recipient policies that are configured improperly can disrupt message flow for some or all recipients in your messaging system. To ensure that your recipient policies are configured properly, verify the following:

- Verify that recipient policies do not contain an SMTP address that matches the fully qualified domain name (FQDN) of any Exchange servers in your organization. For example, if you have @exchangeserver.example.com listed as an SMTP address and as a domain name on any recipient policy, it prevents mail from routing to other servers in the routing group.
- Verify that the domain for which you want to receive SMTP mail is listed on a recipient policy—either on the default policy or another recipient policy. By verifying this, you ensure that your users can receive mail from other SMTP domains.
- Verify that you configured the necessary SMTP e-mail addresses to receive e-mail for additional domains. If you are not receiving e-mail for all of your SMTP domains, you may need to configure additional SMTP addresses for your recipients. For example, some of your users may currently receive e-mail addressed to contoso.com, but you also want to them to receive e-mail addressed to adatum.com.

Verifying That Recipient Policies Do Not Contain an SMTP Address Matching the FQDN of an Exchange Server

Perform the following procedure to verify that your recipient policies are configured correctly and match your mail domain (for example, @example.com) rather than the FQDN of your Exchange server (for example, @exchange.example.com).

- ▶ **To verify that your recipient policies do not contain addresses that match the FQDN**
 1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
 2. In the console tree, expand **Recipients**, and then click **Recipient Policies**.

3. In the details pane, right-click a recipient policy that is configured on the server, and then click **Properties**.
4. On the **E-Mail Addresses (Policy)** tab of that policy, view the SMTP addresses configured by that policy and ensure that none of the SMTP addresses match the FQDN of any Exchange servers in your organization (Figure 10).

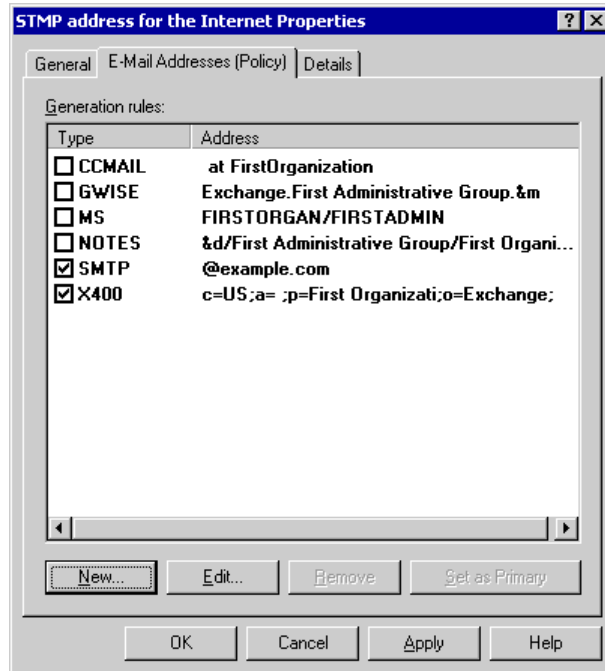


Figure 10 SMTP addresses on a recipient policy

5. Repeat steps 3 and 4 of this procedure for each recipient policy configured on this server.

For more information about why recipient policies cannot match the FQDN of Exchange servers, see Microsoft Knowledge Base article Q288175, “XCON: Recipient Policy Cannot Match the FQDN of Any Server in the Organization, 5.4.8 NDRs”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=288175>).

Verifying That Recipients Can Receive Mail from Other SMTP Domains

In order to receive e-mail from other SMTP domains, your recipient policy must correctly specify the domain for which you want to receive mail.

Important By default, the SMTP domain name on the default recipient policy is the name of the domain in which Active Directory resides. This default SMTP domain name is not always the same name you want to use for SMTP mail.

For example, if your organization is a large distributed corporation, you can use a unique SMTP address to create distinct e-mail addresses for the recipients in each division. For example, users in different divisions at the company Blue Yonder Airlines could have addresses such as `someone@administration.blueyonderairlines.com` and `someone@marketing.blueyonderairlines.com`.

Perform the following procedure to confirm that recipients in your organization are able to receive mail from other SMTP domains.

► **To verify that your users can receive e-mail from other SMTP domains**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Recipients**, and then click **Recipient Policies**.
3. In the details pane, right-click a recipient policy that is configured on this server, and then click **Properties**.
4. On the **E-Mail Addresses (Policy)** tab of that policy, view the SMTP addresses configured by that policy, and then ensure that the domain you want to receive SMTP mail is listed as an address. Verify that the check box next to the address is selected.
5. Double-click the SMTP address you want, and then, in **SMTP Address Properties**, verify that the **This Exchange Organization is responsible for all mail delivery to this address** check box is selected (Figure 11).

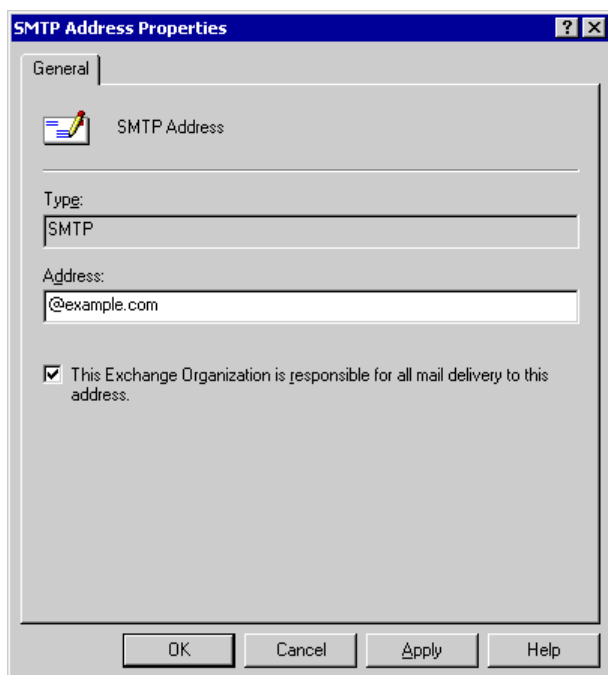


Figure 11 The **SMTP Address Properties** dialog box

Note If you have more than one recipient policy configured on a server, the SMTP e-mail address you are attempting to verify may be located on another recipient policy.

6. If you have more than one recipient policy configured on a server, repeat steps 3 through 5 of this procedure for each recipient policy.

Configuring the SMTP E-Mail Addresses for Your Users

Use the following procedure to ensure that each user's e-mail address is correctly configured on a recipient policy. Remember that Exchange only accepts e-mail for addresses that are configured correctly in a recipient policy. These addresses are stored in Active Directory and the IIS metabase where the message categorizer checks for address and configuration information.

► **To configure the necessary SMTP e-mail addresses for your users**

1. On the **Start** menu, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Recipients**, and then click **Recipient Policies**.
3. In the details pane, right-click the recipient policy you want to modify, and then click **Properties**.
4. On the **E-Mail Addresses (Policy)** tab, click **New**.
5. In **New E-mail Address**, click **SMTP Address**, and then click **OK**.
6. In **SMTP Address Properties**, in the **Address** box, type the information required by the address type you selected. For example, to route mail to Example Corporation, type **@example.com** (Figure 12).

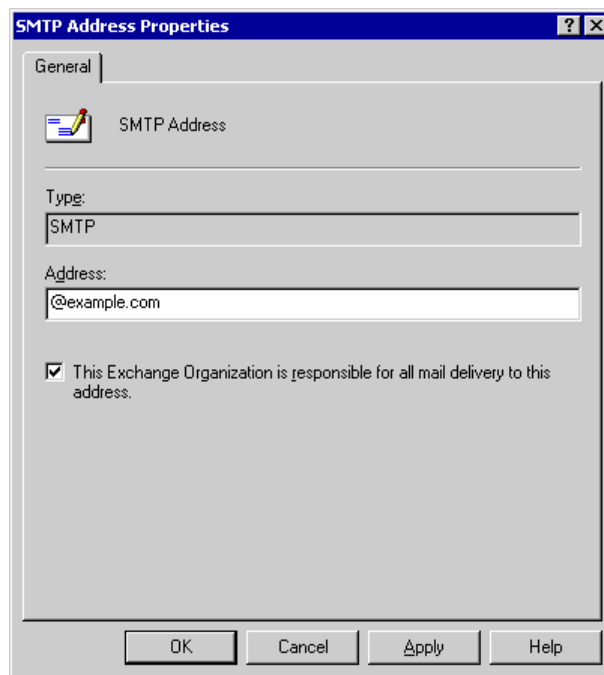


Figure 12 The **SMTP Address Properties** dialog box

7. Ensure that the **This Exchange Organization is responsible for all mail delivery to this address** check box is selected, and then click **OK**.

Note The **This Exchange Organization is responsible for all mail delivery to this address** check box determines whether or not Exchange is authoritative over the selected domain. If Exchange is authoritative for a domain, it accepts all mail for the domain; if it does not locate a valid recipient in Active Directory, Exchange returns an NDR for the message.

8. To keep track of information about the recipient policy you modified, in the recipient policy properties, click the **Details** tab. Under **Administrative note**, type information about the address you added to the recipient policy.
9. On the **E-Mail Addresses (Policy)** tab, under **Generation rules**, select the address you added, and then click **Apply**.

Important When you click **Apply**, Exchange may prompt you to update all corresponding recipient e-mail addresses to match the changes you made. If you click **Yes**, the changes made to the recipient policy are applied to the recipients defined for the policy on the next cycle of Recipient Update Service. E-mail addresses that were previously configured for these recipients are demoted to secondary addresses.

If you want this e-mail address to apply only to a subset of users, create a new recipient policy with a filter that selects the subset of recipients you specify. If the filter is too complex or only a small number of users require the additional address, you can create a filter that creates e-mail addresses that apply only to individual recipients.

Caution All SMTP mail domains for which Exchange accepts mail should have a recipient policy configured for them; however, that recipient policy need not apply to every user. You can add new SMTP e-mail addresses, but it is imperative that the SMTP e-mail addresses do not match the FQDN of any Exchange server in your organization. If an SMTP e-mail address matches a server's FQDN, remote and local e-mail will stop flowing.

For information about how to configure recipient policies, see Microsoft Knowledge Base article Q260973, "XCON: Setting Up SMTP Domains for Inbound and Relay E-Mail in Exchange 2000 Server" (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=260973>).

For more information about how to correct problems with SMTP proxy addresses, see Microsoft Knowledge Base article Q140933, "XFOR: SMTP Proxy Address Generated Incorrectly" (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=140933>).

Configuring Inbound Settings on SMTP Virtual Servers

This section discusses how to configure your SMTP virtual server to receive Internet mail. To configure your SMTP virtual server to receive Internet mail, you must perform the following tasks:

- Configure the inbound port as 25 and specify the IP address.
- For security reasons, verify the relay restrictions on your inbound virtual server. By default, relay settings allow only authorized users to relay mail.

Important You should verify that your SMTP virtual server settings are correct. You should also be familiar with the consequences of specific configuration choices when troubleshooting SMTP-related message flow issues.

Configuring the Inbound Port and IP Address

The inbound port is the port where the SMTP virtual server listens for incoming communications; the IP address is the address to which incoming requests are sent. By default, the default SMTP virtual server uses port 25 and all available IP addresses to listen for incoming requests.

- ▶ **To configure the inbound port and IP addresses on the SMTP virtual server**
 1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
 2. Expand **Servers**, expand <*Server Name*>, expand **Protocols**, and then expand **SMTP**.
 3. Right-click **Default SMTP Virtual Server**, and then click **Properties**.
 4. In **Default SMTP Virtual Server Properties**, click the **General** tab (Figure 13).

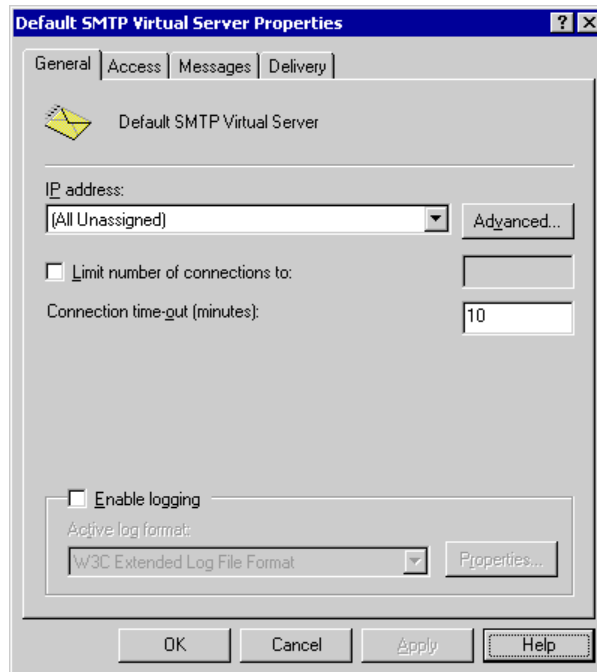


Figure 13 The **General** tab in the **Default SMTP Virtual Server Properties** dialog box

5. Under **Default SMTP Virtual Server**, verify the following settings:

- **IP address** The default setting is (**All Unassigned**). You should not change this setting unless you want to configure multiple SMTP virtual servers. (This is the IP address used for incoming connections.)
- If you have either multiple network interface cards (NICs) or multiple IP addresses assigned to a single NIC for this SMTP virtual server to listen on, and you want to select individual IP addresses, click **Advanced**, and then specify ports other than the default.

Note Use the **Advanced** option carefully. Other servers (on the Internet, for example) expect to communicate with your server on the default TCP port 25.

Verifying Default Relay Restrictions on Your Inbound SMTP Virtual Server

By default, the default SMTP virtual server allows only authenticated users to relay e-mail. This is the preferred setting because it prevents unauthorized users from using your Exchange server to send e-mail to external domains. The most secure relay configuration requires authentication for anyone connecting from the Internet and attempting to relay.

As mentioned earlier, bridgehead servers that are connected to the Internet and that accept Internet mail must generally accept anonymous connections; however, by default, these bridgehead servers do not allow anonymous relaying. Enabling anonymous relaying is strongly discouraged. If you allow anonymous relaying, other users can use your server to send unsolicited commercial e-mail. Subsequently, this would cause other Internet servers to blacklist your server.

► **To verify relay restrictions on an SMTP virtual server**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. Expand **Servers**, expand <*Server Name*>, expand **Protocols**, and then expand **SMTP**.
3. Right-click **Default SMTP Virtual Server**, and then click **Properties**.

4. In **Default SMTP Virtual Server Properties**, click the **Access** tab (Figure 14).



Figure 14 The **Access** tab in the **Default SMTP Virtual Server Properties** dialog box

5. Under **Relay restrictions**, click **Relay** to verify relay restrictions. The **Relay Restrictions** dialog box displays (Figure 15).

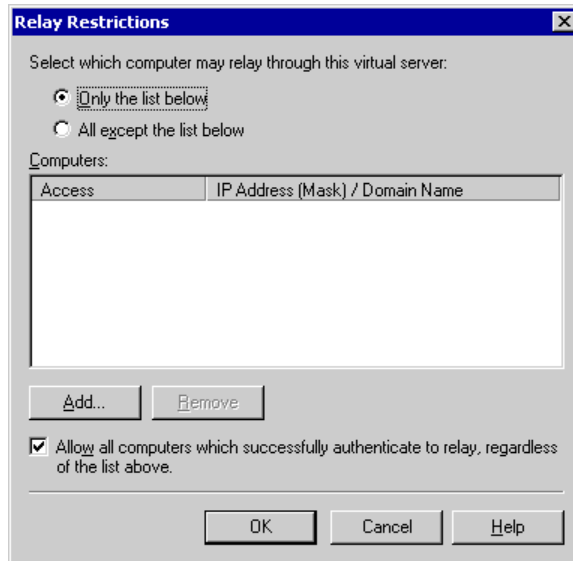


Figure 15 Default relay restrictions

6. In **Relay Restrictions**, verify the following settings:
 - Verify that the **Only the list below** button is selected. To list only those hosts you want to allow to relay mail, click **Add**, and then follow the instructions. If you click **All except the list below**, your server may appear to be a server that is a source of unsolicited e-mail on the Internet.
 - Verify that the **Allow all computers which successfully authenticate to relay, regardless of list above** check box is selected. This setting allows you to deny access to all users who do not authenticate. Any remote POP and IMAP users accessing this server will authenticate to send mail. If you do not have users who access this server through POP or IMAP, you can clear this check box to prevent relaying entirely, thereby increasing security.

Verifying DNS Set Up for Inbound Mail

DNS plays a vital role in Internet mail delivery. In order to receive Internet mail, the following settings are necessary:

- A mail exchanger (MX) record for your mail server must exist on your DNS server. You can use the Nslookup utility to determine if your MX records are configured correctly.
- In order for external DNS servers to resolve your mail server's MX record and contact your mail server, your mail server must be accessible from the Internet. You can use the telnet program to determine if other servers can access your mail server.
- Your Exchange server must be configured to contact a DNS server or to resolve DNS names.
- Your DNS server must be configured correctly.

The following sections explain how to verify each of these settings.

Note It is recommended, although not required, that you use the DNS Server service in Windows 2000. There are other DNS server software suites, but the DNS Server service has been thoroughly tested, and is therefore the most reliable for Windows 2000.

The guidelines in the following sections apply to the DNS Server service in Windows 2000.

Using Nslookup to Verify DNS Configuration

For Exchange to receive Internet mail, the external DNS servers for your domain must contain an MX record pointing to your mail servers that accept Internet mail. Ensure that the mail servers you use as bridgehead servers or Internet mail servers have an MX record on your external DNS servers.

To verify that your MX records are configured correctly, you can use the Nslookup utility on the mail server that accepts Internet mail.

► **To verify that your MX records are configured correctly**

1. At a command prompt, type **Nslookup**, and then press ENTER.
2. Type **server <IP address>**, where *IP address* is the IP address of your external DNS server.
3. Type **set q=MX**, and then press ENTER.
4. Type **<domain name>**, where *domain name* is the name of your domain, and then press ENTER.

The MX record for the domain you entered should be displayed. If the MX record is not displayed, DNS is not configured properly.

Example 3 shows how MX records appear for the fictitious domain, example.com.

Example 3 MX records for example.com

```
C:\> nslookup
Default Server: pdc.corp.example.com
Address: 192.168.6.13
> server 172.31.01.01
Default Server: dns1.example.com
Address: 172.31.01.01
> set q=mx
> example.com.
Server: dns1.example.com
Address: 10.107.1.7
example.com MX preference = 10, mail exchanger = mail1.example.com
example.com MX preference = 10, mail exchanger = mail2.example.com
example.com MX preference = 10, mail exchanger = mail3.example.com
example.com MX preference = 10, mail exchanger = mail4.example.com
example.com MX preference = 10, mail exchanger = mail5.example.com
mail1.example.com internet address = 172.31.31.01
mail2.example.com internet address = 172.31.31.02
mail3.example.com internet address = 172.31.31.03
mail4.example.com internet address = 172.31.31.04
mail5.example.com internet address = 172.31.31.05
```

In Example 3, the pre-configured DNS server was behind a proxy server. Therefore, an external or Internet DNS server with a known IP address of 172.31.01.01 was used to perform the query. Next, the query type was set to MX to locate the mail exchangers for example.com. In this example, five SMTP servers are equally balanced, each with its own IP address. However, your domain may only have a single entry, as seen in Example 4.

Example 4 Single DNS mail exchanger record

```
nwtraders.com MX preference = 10, mail exchanger =
mailbox.nwtraders.com
mailbox.nwtraders.com internet address = 10.57.22.3
```

Using Telnet to Ensure Internet Accessibility

If servers on the Internet cannot reach your mail server, you cannot receive Internet mail. You can use telnet to verify that your mail server is accessible by other servers on the Internet.

After you verify that your MX records are set up correctly, you can then ensure that other servers on the Internet can access your Exchange server. To do this, from a location outside of your intranet, use telnet to connect to your mail server on port 25. You need use a computer that has a direct access to the Internet, so that when you connect, you can validate connectivity. If the server has multiple NICs or IP addresses, you must use telnet to connect to the Internet-facing IP address.

► **To verify that your server is accessible on the Internet**

1. At a command prompt, type **telnet <your mail server> 25**, and then press ENTER.
2. Verify that you receive a response similar to the one shown in Example 5.

Example 5 shows the results of a telnet session to the mail server for Northwind Traders, mailbox.northwindtraders.com.

Example 5 Resulting telnet to mailbox.northwindtraders.com

```
C:\> telnet mailbox.northwindtraders.com 25
220 corp.northwindtraders.com Microsoft ESMTMP MAIL Service, Version:
5.0.
2195.1600 ready at Tue, 5 Sep 2002 11:52:36 -0400
```

Setting Up Your Exchange Server to Send Internet Mail

This section explains how to configure your Exchange server to send Internet mail. Specifically, you will learn how to:

- Configure outbound settings on SMTP virtual servers.
- Configure an SMTP connector.
- Verify DNS set up for outbound e-mail.

Configuring Outbound Settings on SMTP Virtual Servers

The outbound settings control the ports and IP addresses through which outbound mail is sent. Connectors configured on bridgehead servers that route mail to the Internet use these settings. Most of these settings are configured on the **Delivery** tab in the virtual server properties.

To configure your SMTP virtual server to deliver outbound mail you must:

- Ensure that the outbound port is set to port 25 (this is the default setting).
- Allow anonymous access for your outbound connection (this is the default setting).
- Set external DNS servers for SMTP to use, if desired. You can configure the SMTP virtual server to use an external DNS server; however, it is easier and more common to rely on your internal DNS servers to forward mail to trusted external DNS servers.

Note If you want to configure external DNS servers on your SMTP virtual server, ensure that you have the latest Exchange 2000 service pack installed (SP3).

Verifying That the Outbound TCP Port Is Set to 25

To configure the outbound port your server uses to deliver Internet mail, use the **Delivery** tab in the SMTP virtual server properties. If you use the same gateway servers to send and receive Internet mail, the inbound and outbound ports should be set to port 25.

► **To verify your outbound port is set to use port 25**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. Expand **Servers**, expand <*Server Name*>, expand **Protocols**, and then expand **SMTP**.
3. Right-click **Default SMTP Virtual Server**, and then click **Properties**.

4. In **Default SMTP Virtual Server Properties**, click the **Delivery** tab. On this tab, you can specify outbound settings such as retry timers, outbound security and connection limits, and other advanced settings (Figure 16).



Figure 16 The **Delivery** tab in **Default SMTP Virtual Server Properties**

5. On the **Delivery** tab, click **Outbound connections** to set the TCP port that the server will use to connect to remote servers. The **Outbound Connections** dialog box displays (Figure 17).




Figure 17 The **Outbound Connections** dialog box

6. In **Outbound Connections**, verify that the **TCP port** is set to **25**. Remote servers on the Internet expect your server to use TCP port 25. Changing the **TCP port** is not recommended.

Allowing Anonymous Access on the Outbound Virtual Server

For your outbound SMTP virtual server, you should enable anonymous access (unless you connect directly to a smart host). Remote servers on the Internet do not expect your server to authenticate.

Note Generally, configuring a smart host works better on a connector. Configuring a smart host on an SMTP virtual server is not the preferred method.

► **To allow anonymous access on your outbound SMTP virtual server**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. Expand **Servers**, expand *<Server Name>*, expand **Protocols**, and then expand **SMTP**.
3. Right-click *<Your Outbound SMTP Virtual Server>*, and then click **Properties**.
4. Click the **Delivery** tab.
5. Click **Outbound Security** to select what type of authentication the server will use with remote servers.

6. In **Outbound Security**, click **Anonymous Access** (Figure 18).



Figure 18 The **Outbound Security** dialog box

Note If you connect to a smart host (configured by clicking **Advanced** on the **Delivery** tab), the smart host may require you to authenticate. To see if authentication is required, contact the owner of the smart host or your ISP.

Configuring a Smart Host on an SMTP Virtual Server

Problems may occur if you set the smart host at the virtual server level, rather than at the SMTP connector level. When you configure the smart host at the virtual server level, consider the following restrictions:

Note The following smart host settings are located in the **Advanced Delivery** dialog box. To access this dialog box, in *<Your Outbound SMTP Virtual Server> Properties*, on the **Delivery** tab, click **Advanced**.

- If your Exchange organization contains more than one computer running Exchange, you should not type any data in the **Smart host** box. Mail flow between servers may not work.
- If an IP address is listed in the **Smart host** box, it should be enclosed in square brackets (for example, [10.0.0.1]).
- If an IP address is listed in the **Smart host** box, verify that it does not match the IP address of this Exchange server.

- If a name is listed in the **Smart host** box, it should be a FQDN. For example, “Server Name” is not a FQDN; however, servername.contoso.com is a FQDN.
- If a name is listed in the **Smart host** box, it should not be the FQDN of this server.
- If you do not have a smart host within your network, contact your ISP to find out what IP address or FQDN you should enter here.
- If you do enter a smart host, select the **Attempt direct delivery before sending to smart host** check box. Selecting this check box may help reduce queuing on this server.
- Using multiple smart hosts and load balancing requests across them requires a specific configuration.

Configuring an SMTP Connector

SMTP connectors are an efficient way to route Internet mail. This section describes how to create and configure a connector to send Internet mail.

Creating an SMTP Connector

Use the following procedure to create an SMTP connector.

- **To create an SMTP connector**
1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
 2. In the console tree, go to **Connectors** by performing one of the following steps:
 - Under the **Exchange organization**, click **Connectors**.
 - If you do not have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, and then click **Connectors**.
 - If you have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, expand **Routing Groups**, expand *<Routing Group Name>*, and then click **Connectors**.
 3. Right-click **Connectors**, point to **New**, and then click **SMTP Connector**.

4. In **Properties**, on the **General** tab, in the **Name** box, type a name for the connector (Figure 19).

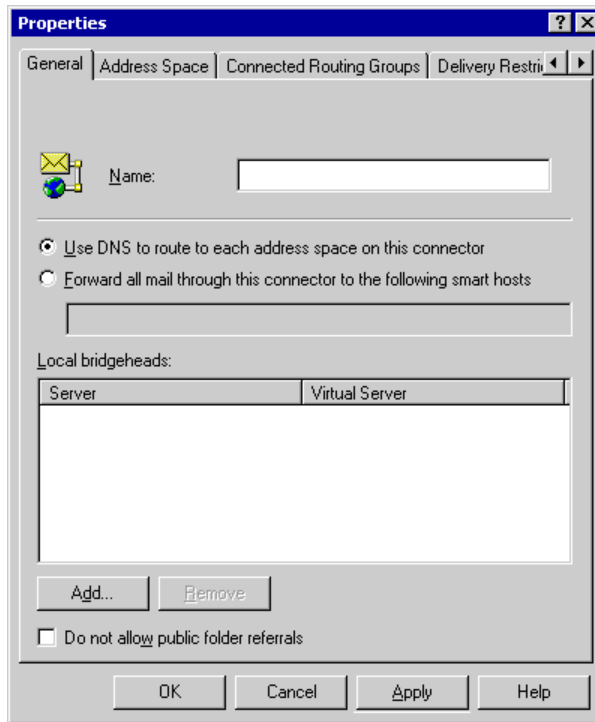


Figure 19 SMTP connector properties

5. Select one of the following check boxes:
- If you want this connector to use DNS names to route mail directly to the remote server, select **Use DNS to route to each address space on this connector**. By selecting this option, the connector uses the DNS settings configured on the SMTP virtual server hosting the connector. If you select this check box, verify the following information:
 - Verify that you can use Nslookup to successfully resolve names on the Internet. For information about how to use Nslookup to verify DNS configuration, see “Using Nslookup to Verify DNS Configuration” later in this chapter.
 - If you use an external DNS server to resolve names, and this server is configured at the SMTP virtual server level (that is, using a different DNS server than the one specified on your network connection), ensure that you have the latest Exchange 2000 service pack installed (SP3), and then use these servers to test name resolution on the Internet.

- If you want to route mail to a smart host that assumes responsibility for DNS name resolution and mail delivery, select the **Forward all mail through this connector to the following smart hosts** check box. This option is often used if you route mail to a Windows SMTP server or another server in your perimeter network. If you select this check box, verify the following information:
 - If you list an IP address for the smart host, enclose the IP address in square brackets (for example, [10.0.0.1]).
 - If you specify an IP address for the smart host, it should not match the IP address of this server.
 - If you specify a name for the smart host, the name should be a FQDN. For example, “Server Name” is not a FQDN; however, servername.contoso.com is a FQDN.
 - If a name is specified, it should not be the FQDN of this server.
 - If you do not have a smart host within your network, contact your ISP to find out what IP address or FQDN you should enter here.
- 6. Under **Local bridgeheads**, click **Add** to define at least one bridgehead server and SMTP virtual server. To send outbound mail, the connector uses the outbound port configured on the SMTP virtual server.

Configuring an Address Space

A connector’s address space defines the domain or range of domains to which a connector sends mail. You can specify which address groups that a specific connector will handle. If you use multiple SMTP connectors to route Internet mail, at least one connector should have an address space of * (asterisk). The asterisk represents all external domains.

► To specify an address space for the connector

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, go to **Connectors** by performing one of the following steps:
 - Under the **Exchange organization**, expand **Connectors**.
 - If you do not have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, and then expand **Connectors**.
 - If you have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, expand **Routing Groups**, expand *<Routing Group Name>*, and then expand **Connectors**.
3. Right-click the SMTP connector, and then click **Properties**.
4. In the SMTP connector **Properties**, click the **Address Space** tab.

5. Click **Add**. The **Add Address Space** dialog box displays (Figure 20).

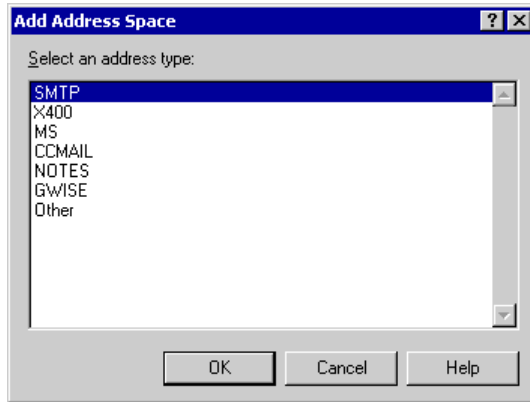


Figure 20 The **Add Address Space** dialog box

6. Under **Select an address type**, click **SMTP**, and then click **OK**. The **Internet Address Space Properties** dialog box displays (Figure 21).




Figure 21 The **Internet Address Space Properties** dialog box

Important In **Internet Address Space Properties**, in the **E-mail domain** box, there is a default value of *. The * represents all addresses. At least one connector in your organization should have this address space to ensure that all external domains are routed to the Internet.

7. In **Internet Address Space Properties**, in the **E-mail domain** box, type an e-mail domain for the connector. In the **Cost** box, leave the default cost of 1.

Important Do not list your “inbound” domains on an SMTP address space for a connector. Your inbound domains are listed in your recipient policies (for more information, see “Configuring Recipient Policies” earlier in this chapter). If some or all of your inbound domains are listed, you may receive NDRs that indicate a mail loop (these NDRs may have the diagnostic code 5.3.5). By specifying domains on the **Address Space** tab, you can configure these domains as routable domains.

8. Click **OK** to return to the **Address** tab.
9. Under **Connector scope**, verify that **Entire Organization** is selected. With this option selected, all Exchange servers in the organization can route mail through this connector to the Internet.
10. If you want mail to be relayed through your system to the domains you specified, select the **Allow messages to be relayed to these domains** check box.
11. Click **OK**.

Configuring DNS for Outbound Mail

There are two methods you can use to configure DNS for outbound mail.

Method 1

You can configure Exchange to rely on your internal DNS servers. These servers resolve external names on their own or use a forwarder to an external DNS server.

Method 2

You can configure Exchange to use a dedicated external DNS server.

Method 1: Using Internal DNS Servers for External Name Resolution

In method one, Exchange relies on your DNS servers to resolve domain names. Generally, you configure your Exchange servers as DNS clients of your internal DNS server. On your internal DNS server, configure an external forwarder to point to trusted external DNS servers.

The following sections explain how to configure:

- DNS settings on the Exchange server.
- Settings on the DNS server.

Configuring DNS Settings on the Exchange Server

The Exchange server should typically specify a local DNS server—in other words, the Exchange server should “point” to a DNS server in its own domain.

To specify which DNS server that the Exchange servers will point to, you must access the **Internet Protocol (TCP/IP) Properties** dialog box.

► **To access Internet Protocol (TCP/IP) Properties for a server**

1. Click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**.
2. Double-click **Local Area Connection**, and then, in **Local Area Connection Status**, click **Properties**.
3. In **Local Area Connection Properties**, under **Components checked are used by this connection**, double-click **Internet Protocol (TCP/IP)**.
4. In **Internet Protocol (TCP/IP) Properties**, verify that DNS is configured correctly.

The Exchange server should point to the primary DNS server for your domain. If you have multiple local DNS servers, you can configure Exchange to point to any of them. However, it is recommended that Exchange point to the primary DNS server for that domain.

Configuring Settings on the DNS Server

Use the following guidelines to configure your DNS server (to access the DNS console, click **Start**, point to **Administrative Tools**, and then click **DNS**):

- Ensure that the DNS server points to its IP address. To confirm this setting, access the **Internet Protocol (TCP/IP) Properties** dialog box for the DNS server. For instructions about how to access this dialog box, see the procedure “To access Internet Protocol (TCP/IP) Properties for a server” in the previous section.
- The DNS server should contain forward lookup zones for each of the domains being hosted. To configure forward lookup zones, in the **DNS** console, expand the DNS server, expand **Forward Lookup Zones**, right click forward lookup zone you want, click **Properties**, and then use the settings on the **General** tab. For each forward lookup zone:
 - **Allow dynamic updates** should be set to **Yes**.
 - **Type** should be set to **Active Directory Integrated**.

- The DNS server should contain reverse lookup zones for each IP subnet range being hosted. To configure reverse lookup zones, in the **DNS** console, expand the DNS server, expand **Reverse Lookup Zones**, right-click reverse lookup zone you want, click **Properties**, and then use the settings on the **General** tab. For each reverse lookup zone:
 - **Allow dynamic updates** should be set to **Yes**.
 - **Type** should be set to **Active Directory Integrated**.
- Configure your DNS server to include forwarders to external (Internet) DNS servers. This setting allows your DNS server to receive queries for external names, forward the query to the remote server, and deliver the response to the requestor. To configure this setting, open the **DNS** console, right-click your DNS server, click **Properties**, click the **Forwarders** tab, and then configure forwarders to external DNS servers.

Note If the **Enable Forwarders** check box on the **Forwarders** tab is unavailable, the DNS server was configured as a root DNS server. If this is the case, to configure forwarders, you must remove the “.” (period) zone, restart the DNS console, and then configure the forwarders.

For more information about DNS in relation to Windows 2000 and Active Directory, see Microsoft Knowledge Base article Q298448, “Windows 2000 DNS and Active Directory Information and Technical Resources”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=298448>).

Method 2: Configuring External DNS Servers on an SMTP Virtual Server

This section explains how to configure external DNS servers on an SMTP virtual server. By configuring external DNS servers, you can use different servers to deliver mail and to resolve names for Windows 2000.

Note If you use external DNS servers to resolve names, and mail delivery is slow or nonexistent, you should use the default DNS servers that are listed in your network connection in Windows 2000.

► To configure external DNS servers on an outbound SMTP virtual server

Note Before you configure an external DNS server, ensure that you are running the latest Exchange 2000 service pack (SP3).

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand <*Server Name*>, expand **Protocols**, and then expand **SMTP**.
3. Right-click <*Your Outgoing SMTP Virtual Server*>, and then click **Properties**.

4. Click the **Delivery** tab, and then click **Advanced**. The **Advanced Delivery** dialog box displays (Figure 22).

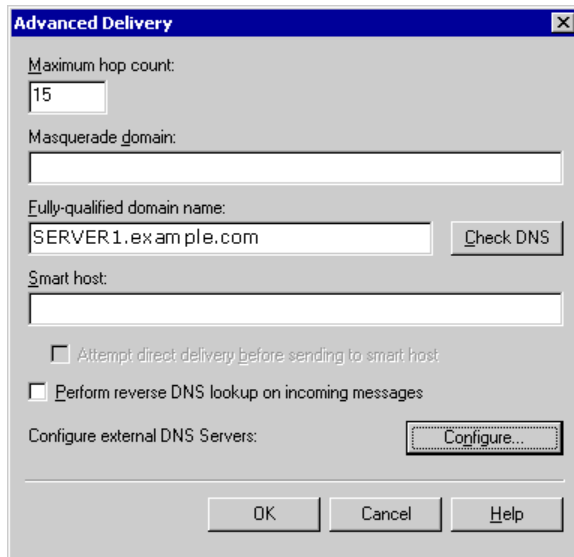


Figure 22 The **Advanced Delivery** dialog box

5. In **Advanced Delivery**, click **Configure**. The **Configure** dialog box displays (Figure 23).

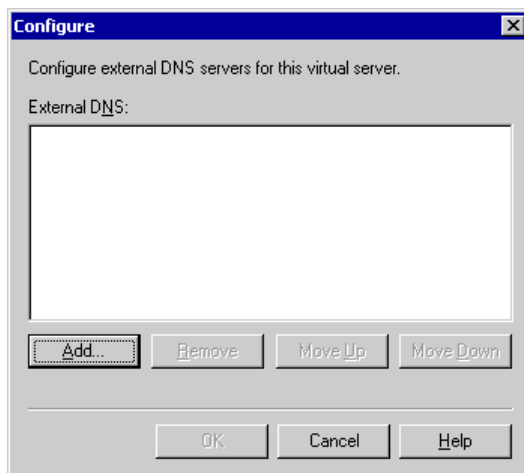


Figure 23 The **Configure** dialog box

6. In **Configure**, click **Add**, type the IP address of the external DNS server you want to use, and then click **OK**.
7. In **Configure**, under **External DNS**, verify the IP address is correct, and then click **OK** twice to apply the settings.

Using Nslookup to Verify DNS Configuration

In order for Exchange to send Internet mail, the DNS servers that Exchange uses for your domain must be able to resolve external domain names

To verify that your DNS servers can resolve external domain names, you can use the Nslookup utility.

► **To verify that your DNS server can resolve external DNS names**

1. At a command prompt, type **Nslookup**, and then press ENTER.
2. Type **server <IP address>**, where *IP address* is the IP address of your external DNS server.
3. Type **set q=MX**, and then press ENTER.
4. Type **<domain name>**, where *domain name* is the name of an external mail domain, and then press ENTER.

The mail exchanger (MX) resource record for the domain you entered should be displayed. If it the MX record is not displayed, DNS is not configured to resolve external domain names.

Example 6 shows how the DNS server for example.com resolves the IP address of the external domain contoso.com.

Example 6 Using Nslookup to verify DNS configuration

```
C:\> nslookup
Default Server:  pdc.corp.example.com
Address:  192.168.6.13
> server 10.255.255.255
Default Server:  dns1.example.com
Address:  10.255.255.255
> set q=mx
> contoso.com.
Server:  dns1.example.com
Address:  192.168.10.10
contoso.com  MX preference = 10, mail exchanger = mail1.contoso.com
contoso.com  MX preference = 10, mail exchanger = mail2.contoso.com
```

```
contoso.com    MX preference = 10, mail exchanger = mail3.contoso.com
  mail1.contoso.com    internet address = 192.168.255.011
  mail2.contoso.com    internet address = 192.168.255.012
  mail3.contoso.com    internet address = 192.168.255.013
```

In Example 6, the pre-configured DNS server was behind a proxy server. Therefore, an external or Internet DNS server with a known IP address of 10.255.255.255 was used to perform the query. Next, the query type was set to MX to locate the mail exchangers for contoso.com. In this example, three SMTP servers are equally balanced, each with its own IP address.

Configuring Advanced Settings

This section explains how to configure some of the advanced settings that control Internet mail delivery. Although these settings are not essential for mail flow, they can assist you in performance tuning, controlling access to your SMTP virtual servers, and many other areas.

Specifically, you will learn how to:

- Configure advanced inbound settings.
- Configure advanced outbound settings.
- Handle undeliverable mail.
- Use distribution lists in multi-domain environments.

Configuring Advanced Inbound Settings

This section shows you how to configure advanced settings for inbound mail. Specifically, you will learn how to configure access controls and other security settings, how to configure message filters, and how to set limits for incoming messages.

Configuring Access Controls and Security Settings

For SMTP virtual servers, you can specify what types of connections are accepted or denied, and you can require user authentication before mail delivery. If you support IMAP or POP clients that connect from the Internet, authentication methods are useful.

However, on an SMTP virtual server that acts as an Internet gateway, you cannot require authentication if you want to receive mail from users on the Internet.

► **To configure access controls and authentication methods**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand <*Server Name*>, expand **Protocols**, and then expand **SMTP**.
3. Right-click **Default SMTP Virtual Server**, and then click **Properties**.
4. Click the **Access** tab, and then, under **Access control**, click **Authentication** to specify the ways in which users must be authenticated prior to sending mail to this server. The **Authentication** dialog box displays (Figure 24)



Figure 24 The **Authentication** dialog box

5. In **Authentication**, the following check boxes are available:
 - **Anonymous access** Usually, you would use this check box for servers that are directly connected to the Internet. If you select this check box, other servers on the Internet will not authenticate to this server prior to sending mail. For increased security, disable anonymous access on your internal SMTP virtual servers that do not accept incoming Internet mail. For similar security purposes, you can also disable anonymous access on dedicated SMTP virtual servers used for remote IMAP and POP users. However, you must allow anonymous access on your Internet gateway servers.

Note If the **Anonymous access** check box is not selected on your Internet gateway servers, you may not receive incoming mail from the Internet. However, for internal SMTP virtual servers or SMTP virtual servers used exclusively by IMAP and POP users, you can clear this check box because they must authenticate.

- **Basic authentication** Use this check box for mail clients (such as Microsoft Outlook) that use Post Office Protocol version 3 (POP3) or Internet Message Access Protocol version 4rev1 (IMAP4) to connect to the server. To send e-mail, these clients authenticate to the server.
-

Important If you select the **Basic authentication** check box, user names and passwords are sent across the network in clear text. This information can be easily intercepted on the Internet. If you use basic authentication, consider implementing Transport Layer Security (TLS) for more security.

- **Requires TLS encryption** Use this check box if you have a digital certificate, typically in a high-security environment. If you select this check box, in the corresponding **Default domain** box, you must type the Windows 2000 domain name that the user should authenticate against if he or she does not specify a domain. For more information about TLS encryption, see the Exchange online documentation.
 - **Integrated Windows Authentication** This check box is used only by Windows user accounts. Using the NTLM protocol, user names and passwords are encrypted and are then passed to the SMTP virtual server for authentication purposes.
-

Note By default, the **Anonymous access**, **Basic authentication**, and **Integrated Windows Authentication** check boxes are selected. If you are using a single default virtual server, it is recommended that you use the default settings; this allows users to authenticate using the most common methods.

6. In *<SMTP Virtual Server> Properties*, on the **Access** tab, under **Secure communication**, click **Certificate** to configure a certificate (used for TLS encryption) that encrypts messages as they move from server to server. For more information about TLS encryption, see the Exchange online documentation.

7. On the **Access** tab, under **Connection control**, click **Connection** to allow or deny access to the server based on IP address. If you are using multiple SMTP virtual servers, and you want to deny access to specific hosts, you must perform the following procedure for each virtual server.
 - a. In **Connection**, click **All except the list below** for servers directly connected to the Internet.
 - b. To list only those hosts from which you do not want to receive mail, Click **Add** and then follow the instructions in the **Computer** dialog box. You can include any servers that you consider to be the source of unsolicited Internet e-mail
 - c. Click **OK** twice to apply the settings.

Setting Global Message Filters

You can use message filters to filter unsolicited commercial e-mails from known IP addresses, domains, or subnets.

To enable a message filter, you must perform the following two steps.

1. Create the message filter in **Global Settings**.
2. Apply the message filter to each SMTP virtual server you want to use the filter.

► **To create a message filter**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Global Settings**, right-click **Message Delivery**, and then click **Properties**.
3. Click the **Filtering** tab, and then click **Add**.
4. In **Add Sender**, type either a domain name or a combination of user and domain name. For example, if you want to filter all messages from the contoso domain, type **@contoso.com**. If you want to filter messages from a specific user in the contoso domain, type the user's entire e-mail address, for example, **sfine@contoso.com**.

After you create the message filter, you must apply it to each SMTP virtual server that you want to use the filter.

► **To apply a message filter**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand **<Server Name>**, expand **Protocols**, and then expand **SMTP**.
3. Right-click the SMTP virtual server you want, and then click **Properties**.
4. On the **General** tab, click **Advanced**.

5. In **Advanced**, under **IP Address**, click the IP address for which you want to apply the filter, and then click **Edit**. The **Identification** dialog box displays (Figure 25).




Figure 25 The **Apply Filter** check box in the **Identification** dialog box

6. Select the **Apply Filter** check box to apply the filter that you previously set.
7. If you have multiple virtual servers, repeat steps 3 through 6 for each virtual server on which you want to apply the filter.

Specifying Message Limits

On the **Messages** tab of the virtual server's properties, you can configure the default number of recipients per message. Reducing this number can mitigate the impact of unsolicited commercial e-mails by preventing the delivery of a single message to a large number of users. You can also reduce the maximum message size and the length of each session.

Note If your organization uses larger distribution lists that arrive through SMTP from Internet users, reducing the number of recipients per message can impact your users. However, Message Application Programming Interface (MAPI) recipients (such as Outlook users) are not affected.

► To specify message limits

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand *<Server Name>*, expand **Protocols**, and then expand **SMTP**.
3. Right-click the SMTP virtual server you want, and then click **Properties**.

4. Click the **Messages** tab to specify message limits for this server (Figure 26).

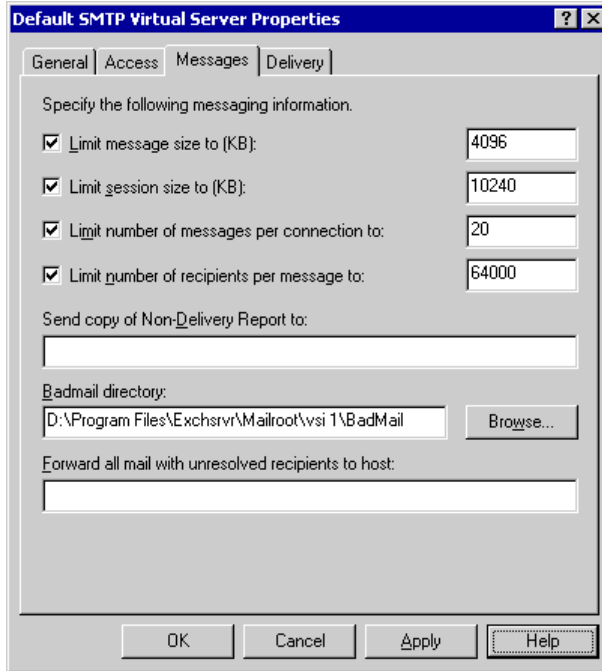


Figure 26 The **Messages** tab

5. Under **Specify the following message information**, select the **Limit message size to (KB)** check box to limit the maximum message size. To prevent users from sending large documents, type a small value in the corresponding box. However, if you do not limit the maximum message size at all, it can affect performance. It is recommended that you set the maximum message size, as is appropriate for your organization.

Note Documents expand in size approximately thirty three percent when sent outside the routing group or organization. For example, if you want to send documents up to 3 MB in size, set the maximum message size to 4,096 KB.

6. Select the **Limit session size to (KB)** check box, and type a value larger than the maximum message size. Sending a message requires network traffic that is greater than the message size.

7. Select the **Limit number of messages per connection to** check box to configure the system to drop the connection after it reaches the specified number of messages. This default setting optimizes message flow in most messaging topologies. Selecting this check box can lead to slight performance degradation if your system receives many messages from a single source.
8. Select the **Limit number of recipients per message to** check box to have Exchange return a non-delivery report (NDR) to senders whose messages exceed the maximum number of recipients. Selecting this check box allows you to keep users from sending an e-mail message to an excessive amount of recipients.

Configuring Advanced Outbound Settings

This section shows you how to configure advanced settings to control outgoing mail. Specifically, you will learn how to configure Internet mail message formats, outbound message limits, and advanced connector settings.

Configuring Internet Mail Message Formats

For each domain listed in Internet Message Formats, you can configure how you send Internet mail messages.

As a general rule, do not send mail exclusively in rich text format (RTF) because many non-Microsoft mail servers cannot read rich-text messages; instead, users receive an e-mail with a winmail.dat file attachment. To avoid this problem, ensure that your global message setting does not use the Exchange RTF exclusively.

- ▶ **To ensure that your Exchange server does not use RTF format exclusively**
 1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
 2. In the console tree, expand **Global Settings**, and then click **Internet Message Formats**.
 3. In the details pane, right-click the name you want, and then click **Properties**.
 4. Click the **Advanced** tab.

- Under **Exchange rich-text format**, ensure that either **Never use** or **Determined by individual user settings** is selected (Figure 27).

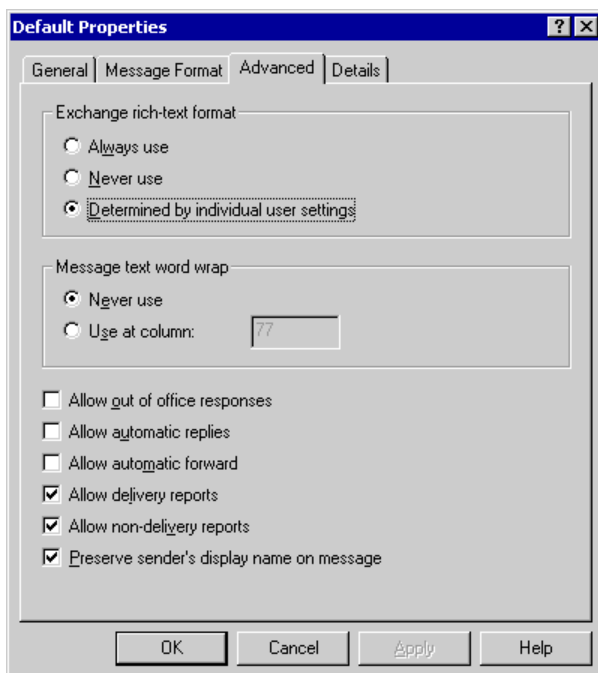


Figure 27 The **Advanced** tab for Internet Message Formats

Note Selecting **Always Use** can prevent users on non-Microsoft servers from reading your e-mails. They may receive an e-mail message with a winmail.dat file.

Configuring Outbound Message Limits

On your SMTP virtual server that handles outbound mail delivery, you can configure connection limits and time-out settings that the server uses with remote servers. Configure these limits to ensure that your server does not get overloaded.

► To set outbound limits on your SMTP virtual server

- Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
- In the console tree, expand **Servers**, expand <Server Name>, expand **Protocols**, and then expand **SMTP**.
- Right-click <Your Outgoing SMTP Virtual Server>, and then click **Properties**.
- Click the **Delivery** tab.

5. Under **Outbound**, you can modify the time in minutes for first, second, third, and subsequent retry attempts by entering the appropriate values for your organization (Figure 28).



Figure 28 Outbound settings on the **Delivery** tab

Note Setting the retry intervals too low may degrade performance, particularly when your Internet connection is down or the specified smart host is unavailable.

6. Under **Local**, set **Delay notification** and **Expiration timeout** for outbound messages by typing the values in the corresponding boxes, and then selecting the time in **Minutes**, **Hours**, or **Days**. It is recommended that you use the default settings.

Note Systems on the Internet may have different values for delay notification and expiration timeout. The values entered here refer to messages queued on this server.

7. Click **Outbound connections** to configure connection limits and timeout values that the server uses with remote servers. The **Outbound Connections** dialog box displays (Figure 29).

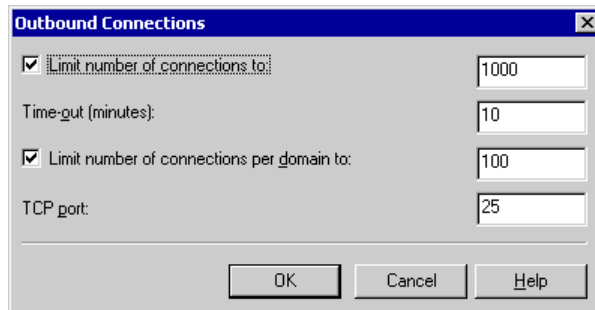


Figure 29 The **Outbound Connections** dialog box

8. Depending on your hardware, you can select the **Limit number of connections to** check box to limit connections to other servers and to reduce traffic. You can also select the **Limit number of connections per domain** check box. After you select the check boxes, enter the appropriate values for your organization.
9. Depending on your bandwidth and connection quality, you can change the **Time-out (minutes)** value.

Note Reducing the number of outbound connections and increasing the time-out period may cause all your outbound connections to wait for responses from remote servers. With such settings, e-mail remains in the queue for longer periods of time (potentially causing a delay in e-mail delivery), but network traffic is kept to a minimum.

Configuring Advanced Settings on the SMTP Connector

The SMTP connector offers several configuration options, which allows you to tailor your specifications for e-mail that is routed through this server. With the exception of message size limits, settings on the SMTP connector override settings on the SMTP virtual server. In this case, the lowest size limit is enforced.

Specifying Delivery Restrictions

The default setting allows everyone in your organization to use this connector. In most situations this setting is sufficient, as you generally want your users to be able to send Internet mail. If you want to set more rigid restrictions, use the following procedures to enable the registry keys and set delivery restrictions.

You can use the **Delivery** tab to restrict the use of your connector. However, to enable these restrictions, you must also change certain registry key settings.

Important Be aware that restricting delivery is extremely process-intensive and can impact server performance.

A registry key on the Exchange 2000-based bridgehead server (which is the source for the connector that is being checked) controls the restriction checking functionality. If you need to configure a connector to restrict who can send data to the designated link, you must manually add the **restriction checking** registry value.

► **To enable the registry keys for delivery restrictions**

Warning If you use Registry Editor incorrectly, you may cause serious problems that may require you to reinstall your operating system. Microsoft cannot guarantee that you can solve problems that result from using Registry Editor incorrectly. Use Registry Editor at your own risk.

1. Start Registry Editor: From a command prompt, type **Regedt32.exe**.
2. Navigate to and select the following key in the registry:
HKEY_LOCAL_MACHINE/System/CurrentControlSet/Services/RESvc/Parameters/
3. On the **Edit** menu, click **Add Value**, and then add the following registry value:

```
Value Name: CheckConnectorRestrictions
Data Type: REG_DWORD
Date: 1
Radix: Decimal
```

4. Exit Registry Editor: On the **Registry** menu, click **Exit**.

After enabling the registry key, you can set delivery restrictions on the connector.

► **To set delivery restrictions on the SMTP connector**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, go to **Connectors** by performing one of the following steps:
 - Under the **Exchange organization**, expand **Connectors**.
 - If you do not have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, and then expand **Connectors**.
 - If you have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, expand **Routing Groups**, expand *<Routing Group Name>*, and then expand **Connectors**.
3. Right-click *<your SMTP connector>*, and then click **Properties**.

4. Click the **Delivery Restrictions** tab (Figure 30).




Figure 30 The **Delivery Restrictions** tab

5. To accept messages from everyone, but reject specified users:
 - a. Under **By default messages from everyone are**, verify that **Accepted** is selected.
 - b. Under **Reject messages from**, click **Add**, and then, in **Select Recipient**, type each user's name that you want to prevent from using the connector.
6. To reject messages from everyone but specified users:
 - a. Under **By default messages from everyone are**, click **Rejected**.
 - b. Under **Accept messages from**, click **Add**, and then, in **Select Recipient**, type each user's name that you want to allow to use the connector.

Setting a Connector Schedule for Connecting to a Network Service Provider

If you are using an SMTP connector to connect to a network service provider and download your Internet e-mail, you may want to schedule specific times for the connector to contact the network service provider's server. Alternatively, you can specify that a connector hold e-mail until a remote server triggers delivery.

► **To set a connector schedule**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, go to **Connectors** by performing one of the following steps:
 - Under the **Exchange organization**, expand **Connectors**.
 - If you do not have routing groups defined, expand **Administrative Groups**, expand <*Administrative Group Name*>, and then expand **Connectors**.
 - If you have routing groups defined, expand **Administrative Groups**, expand <*Administrative Group Name*>, expand **Routing Groups**, expand <*Routing Group Name*>, and then expand **Connectors**.
3. Right-click <*your SMTP connector*>, and then click **Properties**.
4. Click the **Delivery Options** tab (Figure 31).

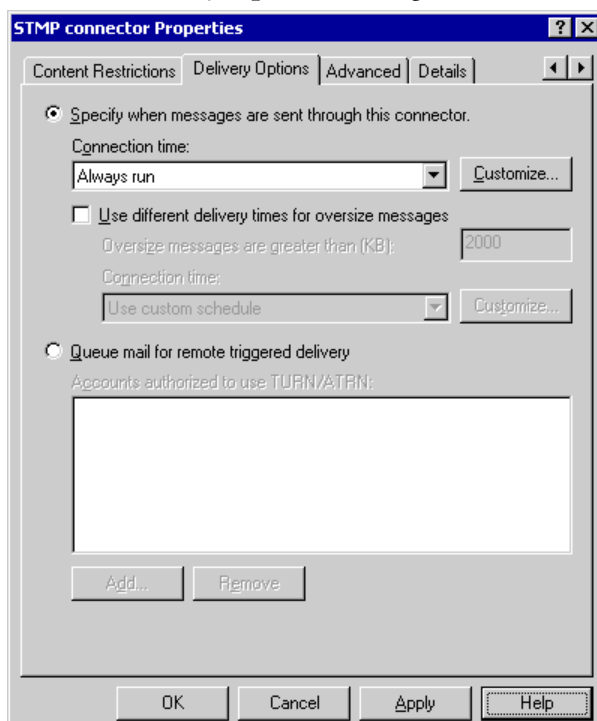


Figure 31 The **Delivery Options** tab

5. To specify a time when the connector runs, click **Specify when messages are sent through this connector**.
6. In the **Connection time** list, select a time or click **Customize** to create a custom schedule.

7. To schedule a different time for the connector to delivery oversize messages, select the **Use different delivery times for oversize messages** check box. If you select this check box, the following options appear:
 - **Oversize messages are greater than (KB)** In this box, type a threshold number that defines oversize messages.
 - **Connection time** In this list, select a time or click **Customize** to create a custom schedule.
8. To hold e-mail until a remote server triggers delivery, click **Queue mail for remote triggered delivery**, and then click **Add** to add authorized accounts that can trigger remote delivery.

Setting Content Restrictions

You can restrict the type of messages delivered through a connector.

► **To set content restrictions on an SMTP connector**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, go to **Connectors** by performing one of the following steps:
 - Under the **Exchange organization**, expand **Connectors**.
 - If you do not have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, and then expand **Connectors**.
 - If you have routing groups defined, expand **Administrative Groups**, expand *<Administrative Group Name>*, expand **Routing Groups**, expand *<Routing Group Name>*, and then expand **Connectors**.
3. Right-click *<your SMTP connector>*, and then click **Properties**.

4. Click the **Content Restrictions** tab (Figure 32).

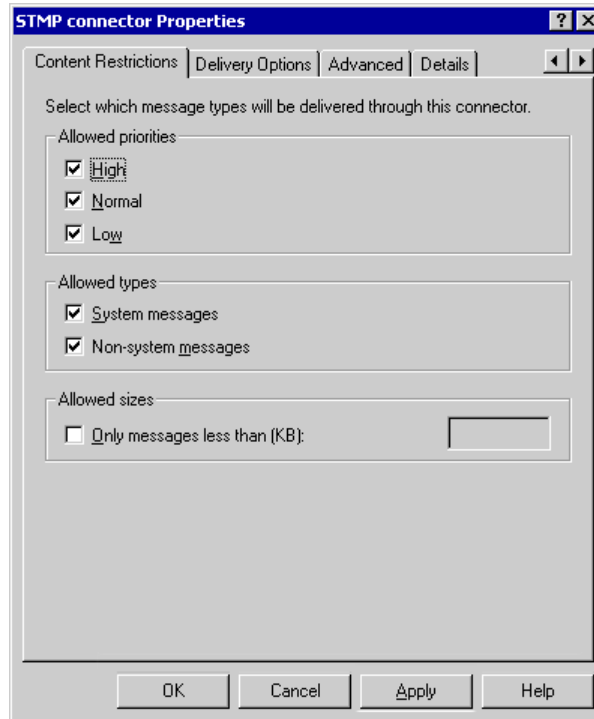


Figure 32 The **Content Restrictions** tab

5. Under **Allowed priorities**, select each type of priority messages you want to send through this connector.
6. Under **Allowed types**, select each type of messages, system or non-system, that you want to send through the connector.
7. Under **Allowed sizes**, if you want to set a size restriction, select the **Only messages less than (KB)** check box, and then type a size limit.

Configuring Notification of Delivery Reports

Use the following procedure to control how undeliverable mail is handled on a specific virtual server. You can always use the postmaster account to handle all NDRs for an organization. If you are sharing a namespace with another mail system such as UNIX, and you want to accept mail for these users and forward this mail to the other system by designating it as a smart host, specifying undeliverable mail handling on a virtual server can be useful.

► **To specify how undeliverable mail is managed**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand <*Server Name*>, expand **Protocols**, and then expand **SMTP**.
3. Right-click the SMTP virtual server you want, and then click **Properties**.
4. Click the **Messages** tab (Figure 33).

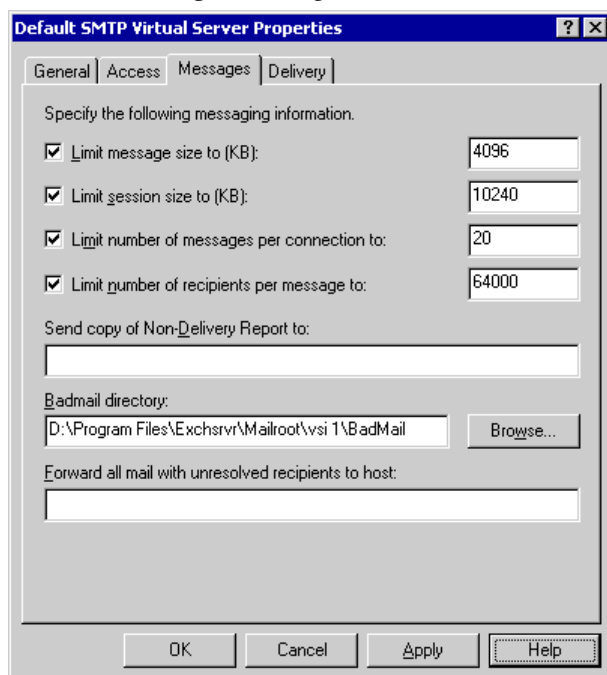


Figure 33 The **Messages** tab in the **Default SMTP Virtual Server Properties** dialog box

5. In the **Send copy of Non-Delivery Reports to** box, type the SMTP address of the Exchange administrator who you want to receive copies of NDRs. You can use the NDRs to help you diagnose user problems. For information about examining NDRs, see “Understanding Non-Delivery Reports” in Chapter 7.

Note NDRs often occur because users type the wrong e-mail address. You may want to disable this feature until you experience problems and need to investigate NDRs.

6. In the **Badmail directory** box, you can modify the location of the messages that are misrouted and cannot be delivered. It is recommended that you keep the default location.

Caution Moving the Badmail directory to a disk separate from the queuing directory may degrade performance and make it difficult to track bad messages.

7. In the **Forward all mail with unresolved recipients to host** box, you can specify an alternate host to which undeliverable messages are forwarded. This is useful if you are sharing a namespace with another mail system—specifically if there are mail recipients with your domain name who do not belong to the Exchange organization. For example, *exchange.user@contoso.com* resides in the Exchange organization, and *unix.user@contoso.com* resides outside the Exchange organization. In this example, users at *exchange.user@contoso.com* can send mail to users at *unix.user@contoso.com*, and Exchange forwards the message to the specified alternate host.

Using Distribution Lists in Multi-Domain Environments

To expand distribution lists into individual recipients, Exchange contacts a global catalog server. The global catalog server has a copy of all local and universal groups in its domain, but it does not have a copy of global groups from other domains. This becomes important in multi-domain environments. In a multi-domain environment, if a message is destined for a global distribution group in a domain separate from the global catalog server, Exchange cannot expand the distribution group on that message. Because the global catalog server does not have copies of global groups for domains outside its own, it does not contain any information about the distribution list; therefore, the categorizer cannot expand the distribution list. To avoid this problem, you should always use universal distribution groups in multi-domain environments. Use global groups within single domains only.

6

Security Considerations

Network attacks are more common than ever, and that trend will only continue. Therefore, after configuring SMTP in your Exchange organization, it's crucial that you take measures to secure it.

Messages that are routed to and from Exchange servers and other external systems also travel across your local network and over the Internet. To prevent malicious Internet users from intercepting your organization's mail and attacking your servers, it's important to secure your Internet connections.

There are three types of Internet connectivity (generally, each of these requires a different level of security):

- Using connectors over the Internet to have e-mail connectivity between your organization and other external systems
- Using connectors to connect Exchange routing groups within your organization over the Internet
- Allowing Exchange clients to use Internet mail protocols or Outlook Web Access to access Exchange mailboxes in your organization

To effectively implement security on your mail system, you must understand how to secure your infrastructure and your servers. The remainder of this chapter focuses on how to secure both of these.

For additional information about securing Exchange, see the following technical papers:

- *Security Operations Guide for Exchange 2000 Server*
(<http://go.microsoft.com/fwlink/?LinkId=11906>)
- *Securing Exchange Server*
(<http://go.microsoft.com/fwlink/?LinkId=11923>)

Securing Your Infrastructure

This section focuses on some important infrastructure components that you can implement for greater security.

IIS Lockdown Wizard

As discussed in “Internet Information Services” in Chapter 3, Microsoft provides IIS Lockdown Wizard as security tool. This wizard turns off unnecessary IIS services, thereby reducing your exposure to attack through these services. To provide defense against attackers, IIS Lockdown Wizard integrates URLscan with customized templates for Exchange servers. IIS Lockdown Wizard is designed primarily to secure Outlook Web Access servers and front-end servers; however, it is also useful for checking the security configuration on any Exchange server

For optimal security, run IIS Lockdown Wizard on each Exchange server and domain controller in your organization. You can download IIS Lockdown Wizard from the Microsoft Download Center (<http://go.microsoft.com/fwlink/?LinkId=12281>).

For more information about IIS Lockdown Wizard, see Microsoft Knowledge Base article Q309508, “XCCC: IIS Lockdown and URLscan Configurations in an Exchange Environment” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=309508>).

Firewalls

A firewall prevents unauthorized access to data on servers that reside behind the firewall. Whether your organization has an existing network or is setting up a new one, firewall planning is extremely important.

With software such as Microsoft Internet Security and Acceleration (ISA) Server, you can route all Internet traffic through a single location. Although this requires more setup and planning than a simple direct Internet connection, it provides increased security for the servers in your organization.

You can use a firewall to allow only essential Internet traffic through ports that you specify; for example, you can configure your network to allow only SMTP (port 25) traffic to pass through your firewall, thereby preventing connections on all other ports.

For Exchange to operate properly in a firewall environment, specifically in regard to remote clients, there are certain requirements necessary to maintain Internet connectivity. For instance, firewalls can filter certain TCP ports or block them entirely. Therefore, for remote clients and servers to communicate through a firewall, you cannot change or block the port assignments for the various protocols that Exchange supports. For more information about the ports that Exchange requires, see “Common Ports Used by Exchange” in Chapter 8 and Microsoft Knowledge Base article Q278339, “XGEN: TCP/UDP Ports Used By Exchange 2000 Server” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=278339>).

If you need a simple SMTP server in the perimeter network of a firewall, often a Windows 2000 SMTP service computer is all that is required. Exchange 2000 Enterprise Server, Windows 2000 Network Address Translation (NAT), Microsoft ISA Server, or any solution that buffers the Internet from the internal LAN can add additional security.

If you do not implement a firewall connection to the Internet, you must consider how security will be affected. All Exchange servers within a network that have a direct connection to the Internet are exposed to the Internet. For more information about firewalls, see the following resources:

- Chapter 13, “System Security” of the *Microsoft Exchange 2000 Server Planning and Installation* guide
- *Security Operations Guide for Exchange 2000 Server* (<http://go.microsoft.com/fwlink/?LinkId=11906>)
- *Securing Exchange Server* (<http://go.microsoft.com/fwlink/?LinkId=11923>)

Virtual Private Networks

The Windows 2000 Routing and Remote Access Service (RRAS) is an open, extensible platform for routing and internetworking. RRAS offers remote access over the Internet and to organizations in LAN and WAN environments by using secure virtual private network (VPN) connections. VPNs are secure, authenticated links across public or private networks, such as the Internet.

The Windows 2000 Remote Access Service (RAS) and RRAS utilities offer options that remote users can use for dial-up Internet access. To function properly, these access services require two things:

- A remote connection method called Point-to-Point Tunneling Protocol (PPTP)
- An Internet connection to create a VPN

PPTP is designed to support VPNs. Because of Digital Subscriber Line (DSL) and cable modem Internet connections, VPNs are less expensive to start up and support than traditional WANs. A VPN eliminates long-distance telephone charges, while offering secure connections, mutual authentication, and packet filtering.

After a PPTP server authenticates a remote client, the VPN connection opens. The PPTP session acts as a tunnel through which network packets flow. The packets are first encrypted when sent. The packets then travel over the tunnel and are decrypted upon receipt. For example, an organization can allow remote clients to connect to a corporate network across the Internet using a VPN. Although a broadband connection is not required for a VPN, a broadband VPN connection can benefit remote VPN users. With a broadband VPN connection, users can connect to a corporate network over the Internet and then use the corporate network as if they were directly logged on.

Securing Your Exchange Server

This section focuses on ways that you can secure your Exchange server.

Disabling Open Relaying on All SMTP Virtual Servers

As explained in “Relay Restrictions” in Chapter 2, it is essential that you do not allow anonymous or open relaying on your SMTP virtual servers. Relaying is when a user uses your Exchange server to send mail to an external domain.

In its default configuration, Exchange allows only authenticated users to relay mail—in other words, only authenticated users can use Exchange to send mail to an external domain. If you modify the default relay settings to allow unauthenticated users to relay, or if you allow open relaying to a domain through a connector, then unauthorized users can use your Exchange server to send unsolicited commercial e-mail. As a result, your server may be blacklisted and thereby be prevented from sending e-mail to legitimate remote servers. To prevent unauthorized users from using your Exchange server to relay mail, you should always use the default relay restrictions.

Note Relaying is often confused with unsolicited commercial e-mail. Relay control does not block unsolicited commercial e-mail. If you are receiving unsolicited e-mail, consider using an event sink or a third-party product that filters and protects against unsolicited commercial e-mail.

Implementing Sender Filtering for Your SMTP Mail Domain on Inbound Gateway Servers

The most secure SMTP configuration prevents internal spoofing and allows only authenticated users to relay to the Internet. Internal spoofing occurs when an external user forges an internal recipient address in an attempt to impersonate a legitimate internal user.

If you use dedicated SMTP virtual servers for your IMAP and POP clients that connect remotely, you can require authentication to prevent spoofing. However, this functionality is limited on bridgehead servers that accept inbound Internet mail. Generally, you cannot configure these bridgehead servers to require authentication because, to accept Internet mail, they must allow anonymous access.

However, to protect your organization from unauthorized users, you can implement sender filtering on your Internet gateway servers. Specifically, you can configure your gateway servers to reject inbound Internet e-mail that has the same recipient address as your SMTP mail domain. As a result, unauthorized users are prevented from successfully impersonating internal users. In this situation, you would assume that all e-mail coming from your SMTP mail domain should originate internally; therefore, if inbound Internet e-mail comes from your SMTP mail domain, it is probably a forged address.

You can also set up other filtering or scanning methods to determine what constitutes valid e-mail on the inbound bridgehead server. In particular, you can set restrictions by IP or domain. You can configure these restrictions in the SMTP virtual server properties.

Preventing Anonymous Access on Internal SMTP Virtual Servers and Dedicated SMTP Virtual Servers for IMAP and POP Clients

For increased security, you can prevent anonymous access on your internal SMTP virtual servers and on any SMTP virtual servers dedicated to accept incoming mail from remote IMAP and POP users. When sending internal mail, Exchange servers automatically authenticate; therefore, by preventing anonymous access on your internal servers, mail flow is not disrupted, and an extra layer of security is provided on your internal SMTP virtual server.

Similarly, IMAP and POP clients authenticate before sending mail to SMTP virtual servers. So, if you use dedicated SMTP virtual servers for your IMAP and POP clients, you can configure these servers to allow only authenticated access. To prevent anonymous access, on the **Access** tab in the SMTP virtual server properties, click **Authentication**, and then clear the **Anonymous Access** check box. For step-by-step instructions about how to prevent anonymous access, see "Configuring Access Controls and Security Settings" in Chapter 5.

Important Do not disable anonymous access on your Internet bridgehead SMTP virtual servers. SMTP virtual servers that accept mail from the Internet must allow anonymous access.

Controlling Unsolicited Commercial E-Mail

Controlling unsolicited commercial e-mail is a difficult task, but there are several prevention methods you can use:

- Implement message filters to prevent specific types of e-mail from being delivered. The following section, “Using Message Filters,” discusses this topic in more detail.
- Educate your users not to respond to or forward unsolicited e-mail. You should also instruct users not to click any “remove” links included in the mail, as they are often used to verify addresses.
- Write SMTP event sinks to prevent unsolicited mail and spoofing. Third-party products can also provide additional anti-spoofing and spamming protection. For more information about event sinks, see “Event Sinks” in Chapter 8.

The following Web sites offer additional information about unsolicited commercial e-mail. This list is not intended to be a comprehensive list of all of the sites available, nor does it imply an endorsement by Microsoft:

Important The third-party contact information included in this book is provided to help you find the technical support you need. This contact information is subject to change without notice. Microsoft in no way guarantees the accuracy of this third-party contact information.

- Network Abuse Clearinghouse (<http://www.abuse.net>)
- Sam Spade (<http://www.samspace.org>)
- Mail Abuse Prevention System (MAPS) (<http://www.mailabuse.org/>)
- Ultradesign Xperimental Network (UXN) Spam Combat (<http://combat.uxn.com/>)

For more information about how to control relaying, see Microsoft Knowledge Base article Q304897, “XIMS: Microsoft SMTP Servers May Seem to Accept and Relay E-Mail Messages in Third-Party Tests” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=304897>).

Using Message Filters

Message filters allow you to filter e-mail from specific domains. Message filters are useful if you want to block unsolicited commercial e-mail (also known as “spam”) from a specific domain. For step-by-step instructions about how to create message filters, see “Setting Global Message Filters” in Chapter 5.

Identifying Spoofed Mail

You can educate your users on how to identify spoofed mail. By default, Exchange resolves incoming e-mails to their display name stored in Active Directory—Exchange does not verify the e-mail address. However, you can modify the default behavior of Exchange so that mail originating from outside the Exchange organization does not resolve to its display name. In this configuration, when mail is sent from a forged address, Exchange does not resolve the sender's e-mail address to its display name.

For example, if your Exchange server has an internal user named Suzan Fine, and she sends mail internally from your domain example.com, the e-mail shows her sending address as **Suzan Fine**, which is the display name configured in Active Directory for sfine@example.com. (This is because when Suzan Fine sends mail, she is an authenticated user.) Exchange then verifies that Suzan Fine has “send as” permissions under her credentials and then resolves her e-mail address to her display name in Active Directory. Spoofing occurs when an unauthorized user pretends to be Suzan Fine by forging this address and then sending mail to another user in your domain.

If you configure Exchange not to resolve e-mail addresses that originate externally, Exchange will not resolve the sending address in the **From** line to its display name. Instead, the **sfine@example.com** will appear in the **From** line of the e-mail. If your users understand this difference, they can at least identify spoofed mail.

► To configure Exchange to not resolve e-mail addresses that originate externally

Caution This section contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Microsoft Knowledge Base article Q256986, “Description of the Microsoft Windows Registry” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=256986>).

1. Start Registry Editor: Click **Start**, click **Run**, type **regedt32**, and then click **OK**.
2. Locate or create the following key in the registry (where one *1* is the SMTP virtual server number):

```
HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services/
MsExchangeTransport/Parameters/1
```

Note You may need to create both the **Parameters** key and the **1** key.

3. On the **Edit** menu, click **Add Value**, and then add the following registry value:

```
Value name: ResolveP2
Data type: REG_DWORD
```

4. Use the following flags to determine which value to use:

Field	Value
-----	-----
FROM:	2
TO: and CC:	16
REPLY TO:	32

5. To determine the value that you want to use, add the values for all of the elements that you want to be resolved. For example, to resolve all of the fields except the sender, type **48** ($16+32=48$). To resolve only the recipients, type only **16**. By default, Exchange 2000 resolves everything (you can specify this behavior either by removing the key or by setting the value with this formula: $2+16+32=50$).
6. Quit Registry Editor.
7. Restart the SMTP virtual server.

Be cautious when you select the servers on which you want to enable this setting. If you change the behavior on the default SMTP virtual server, and there are multiple servers in your organization, all internal mail that originates on other Exchange 2000 servers is also affected. Therefore, because Exchange 2000 uses SMTP to route internal mail between servers, you may want to create a new SMTP virtual server, or perhaps apply this setting only on an incoming SMTP bridgehead server.

7

Troubleshooting Mail Flow

Even after you've successfully configured SMTP in your Exchange organization and taken every measure to secure it, there is still the possibility you will experience mail flow problems. But don't worry—this chapter covers many of the common problems you may encounter.

Specifically, you will learn how to:

- Use Telnet.
- Understand NDR reports.
- Use the SMTP and X.400 queues.
- Use message tracking.
- Use Event Viewer.
- Configure diagnostic logging for SMTP.

However, before considering the troubleshooting recommendations in this chapter, first ensure that Exchange is configured correctly to send and receive mail. The lists below briefly summarize the requirements necessary for inbound and outbound mail to flow properly.

For incoming Internet mail to flow correctly:

- Your recipient policies must be configured correctly.
- Your SMTP virtual server that accepts Internet mail must be configured on port 25 and allow anonymous connections.
- Your DNS server must contain an MX record pointing to your external or Internet domain of your mail server.
- Your Internet mail server must be accessible to remote servers on the Internet.

For outgoing Internet mail to flow correctly:

- Your SMTP virtual server that sends Internet mail must be configured to use port 25.
- If you are using SMTP connectors, at least one connector must contain an address space of *, which specifies all external domains.
- Your Exchange server must be able to resolve external DNS names. You can resolve external DNS names in the following ways:
 - Use an internal DNS server that forwards mail to an external DNS server
 - Configure your SMTP virtual server to use a specific DNS server
 - Route mail to a smart host that performs DNS resolution

For more information about how to configure Exchange to send and receive e-mail, see Chapter 5, “Configuring Exchange to Send and Receive E-Mail.”

Using Telnet

Telnet is an extremely useful tool for troubleshooting issues regarding SMTP and mail flow. For example, you can use telnet to:

- Verify that SMTP is installed properly, with all the necessary commands.
- Ensure that your server is accessible over the Internet.
- Attempt mail delivery directly over the TCP port.
- Determine that all servers are accepting connections.
- Determine if a firewall is blocking a connection.
- Ensure that a single user can receive mail.
- Ensure that a specific domain can receive mail.
- Ensure that a specific user or domain can send mail to your domain.

► To use telnet to test SMTP communication

Note The following procedure shows you how to test the process of an internal user sending mail to a remote user when basic authentication is required for relaying mail outside your organization

1. Open a telnet session: From a command prompt, type **telnet**, and then press ENTER.
2. Type **set local_echo** on a Windows 2000 computer or **SET LOCALECHO** on a Windows XP computer, and then press ENTER. This command allows you to view the responses to the commands.

Note For a list of available telnet commands, type **set ?**.

3. Type **o <your mail server domain> 25**
4. Type **EHLO <your mail server domain>** and then press ENTER.

5. Type **AUTH LOGIN**. The server responds with an encrypted prompt for your user name.
6. Enter your username encrypted in base64. You can use one of several utilities available to encode your user name.
7. The server responds with an encrypted base64 prompt for your password. Enter your password encrypted in base64.
8. Type **MAIL FROM:<sender@domain.com>** and then press ENTER. If the sender is not permitted to send mail, the SMTP server returns an error
9. Type **RCPT TO:<recipient@remotedomain.com>** and then press ENTER. If the recipient is not a valid recipient or the server does not accept mail for this domain, the SMTP server returns an error.
10. Type **DATA**.
11. If desired, type message text, press ENTER, type a period (.), and then press ENTER again.
12. If mail is working properly, you should see a response similar to following indicating that mail is queued for delivery:

```
250 2.6.0 <INET-IMC-01UWr81nn9000fbad8@mail11.example.com>  
Queued mail for delivery
```

Example 7 shows a telnet test sending mail from example.com to a remote domain with a successful result. User input appears in bold.

Example 7 Telnet session sending mail to a remote domain

```
ehlo example.com  
250-mail11.example.com Hello [172.16.0.0]  
250-TURN  
250-ATRN  
250-SIZE 5242880  
250-ETRN  
250-PIPELINING  
250-DSN  
250-ENHANCEDSTATUSCODES  
250-8bitmime  
250-BINARYMIME  
250-CHUNKING
```

```
250-VERFY
250-X-EXPS GSSAPI NTLM
250-AUTH GSSAPI NTLM
250-X-LINK2STATE
250-XEXCH50
250 OK
auth login
334 VXNlcm5hbWU6
c2ZpbmVz
334 UGFzc3dvcmQ6
cGFzc3dvcmQ=
235 2.7.0 Authentication successful.
mail from:sfine@example.com
250 2.1.0 sfine@example.com...Sender OK
rcpt to:tbremmer@contoso.com
250 2.1.5 tbremmer@contoso.com
data
354 Start mail input; end with <CRLF>.<CRLF>
This is a test mail sent on SMTP port 25 to verify test my SMTP
server
.
250 2.6.0 <INET-IMC-01UWr81nn9000fbad8@mail1.example.com> Queued
mail for delivery
```

Understanding Non-Delivery Reports

Non-delivery reports (NDRs) are a type of delivery status notification. NDRs are generated whenever a message cannot be delivered. If a server detects the reason for the delivery failure, it associates the reason to a status code and a corresponding error message is printed. Understanding various NDR status codes can assist you in troubleshooting mail flow issues. NDRs with a numerical code of 5.x.x indicate a permanent failure, while 4.x.x codes indicate transient conditions. Also, the server that reports the error is listed prior to the numerical code. However, sometimes the server that reports the problem is not the server experiencing the issue.

Table 2 lists the most common NDR numerical codes and corresponding error conditions.

Table 2 NDR numerical codes and corresponding error conditions

NDR Code	Possible Cause	Troubleshooting
4.3.1	An out of memory error occurred. A resource problem, such as a full disk, can cause this problem.	Ensure that your Exchange server has enough disk storage. If possible, move your mail queues to an NTFS disk partition.
4.3.2	Available in Exchange 2000 SP1 and later. This NDR is generated when a queue has been frozen.	Unfreeze the queue.
4.4.1	A host is not responding. Transient network conditions can cause this error. The Exchange server automatically tries to connect to the server again and deliver the mail. If delivery fails after multiple attempts, an NDR with a permanent failure code is generated.	
4.4.2	A connection dropped between the servers. Transient network conditions or unavailable servers can cause this error. The server attempts to deliver the message for a specific time period, and then generates further status reports.	

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
4.4.6	The maximum hop count was exceeded for the message. A looping situation between the sending and receiving servers in different organizations can cause this error. The message simply bounces between the servers until the hop count is exceeded.	The maximum hop count property is set per virtual server, and you can manually override the default setting of 15 . You should also check for situations that might cause looping between servers.
4.4.7	The message in the queue has expired. The sending server tried to relay or deliver the message, but the action was not completed before the message expiration time occurred. This message can also indicate that a message header limit has been reached on a remote server, or some other protocol timeout occurred while communicating with the remote server.	This message usually indicates an issue on the receiving server. Check the validity of the recipient address and determine if the receiving server is configured correctly to receive messages. You may have to reduce the number of recipients in the message header for the host about which you are receiving this error. If you resend the message, it is placed in the queue again. If the receiving server is available, the message is delivered.

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.0.0	<p>Note Prior to Exchange 2000 SP1, the following numerical codes all appeared under the 5.0.0. code:</p> <ul style="list-style-type: none"> ● 4.3.2 ● 5.4.0 ● 5.4.4 ● 5.5.0 <p>Possible causes include:</p> <ul style="list-style-type: none"> ● There is no route for the given address space; for example, an SMTP connector is configured, but this address does not match. ● DNS returned an authoritative host that was not found for the domain. ● The routing group does not have a connector defined; mail from one server in one routing group does not have a route to another routing group. ● An SMTP error occurred. 	<p>On one or more SMTP connectors, add an asterisk (*) value as the SMTP address space; verify that DNS is working; ensure that routing groups have connectors connecting them. If necessary, apply the latest Exchange service pack (SP3) to isolate the issue.</p>

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.1.0	<p>This NDR is caused by a general categorizer-based failure (bad address failure). An e-mail address or another attribute could not be found in Active Directory. Contact entries without the targetAddress attribute set can cause this problem.</p> <p>Another possible cause could be that the categorizer is unable to determine the homeMDB attribute of a user. The homeMDB attribute corresponds to the Exchange server on which the user's mailbox resides.</p> <p>Another common cause of this NDR is if you used Outlook to save your e-mail as a file, and then someone opened the message offline and replied to the message. The message property only preserves the legacyExchangeDN attribute when Outlook delivers the message, and therefore the lookup could fail.</p>	<p>Either the recipient address is incorrectly formatted, or the categorizer was not able to resolve the recipient properly. The first step in resolving this error is to check the recipient address and resend the message.</p>

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.1.1	<p>The e-mail account does not exist in the organization where the message was sent. This can occur when users move to new locations within a site. For instance, if a former Administrative_Group_1 user moves to Administrative_Group_2, and then replies to an old mail or does not recreate an Outlook profile, an old Administrative Group style LegacyDN address will be used, and this NDR is issued. Likewise, sending mail to obsolete personal address book entries will result in this error.</p> <p>Also, if you configured your SMTP contact with invalid SMTP characters (as per RFC821), the categorizer will reject the delivery with this diagnostic code.</p>	<p>Either the recipient address is formatted incorrectly, or the categorizer was not able to resolve the recipient properly. The first step in resolving this error is to check the recipient address and resend the message.</p>
5.1.3	<p>This NDR is caused by incorrect address syntax. For example, a contact that was configured in Active Directory with a targetAddress attribute but without an address would result in this error.</p>	<p>Either the recipient address is formatted incorrectly or the categorizer was not able to resolve the recipient properly. The first step in resolving this error is to check the recipient address and resend the message.</p>

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.1.4	Two objects have the same (proxy) address, and mail is sent to that address. This issue can also occur if the recipient does not exist on the remote server.	Check the recipient address and resend the message.
5.1.6	Available in Exchange 2000 SP2 and later. One possible cause of this NDR is that the user directory attributes such as homeMDB (the user's home mailbox store) or msExchHomeServerName (the server on which the user's mailbox resides) is missing or corrupted.	Check the user directory attribute's integrity and rerun the Recipient Update Service to ensure the validity of the attributes that are required for transport.
5.1.7	Available in Exchange 2000 SP2 and later. The sender has a malformed or missing mail attribute in the directory service. The categorizer cannot deliver the mail item without a valid mail attribute.	Check the sender directory structure and determine if the mail attribute exists.
5.2.1	Local mail is refused because the message is too large. A missing Master Account Security ID (SID) number on the recipient can also cause this error.	Check access permissions as well as the message size. Check if the recipient has a SID in Active Directory.
5.2.2	Available in Exchange 2000 SP3. This NDR is generated when the recipient's mailbox exceeds its storage limit.	Check the mailbox storage or the queue storage quota limit.

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.2.3	The message is too large and the local quota is exceeded. For example, a remote Exchange user might have a restriction on the maximum size of an incoming message.	Resend the message without attachments, or set the server or the client side limit to allow a larger message size limit.
5.3.3	When the Exchange 2000 remote server runs of disk storage to hold mail, it could respond with this NDR. This error usually occurs when the sending server is sending mail with an ESMTP BDAT command. This error also indicates a possible SMTP error.	Ensure that the remote server has enough storage to hold mail. Check the SMTP log.
5.3.5	A mail-looping situation is detected. This means that the server is configured to loop mail back to itself. If you have multiple SMTP virtual servers configured on your Exchange server, ensure that they are serving unique incoming ports. Also, to avoid looping between local SMTP virtual servers, ensure that the outgoing SMTP port configuration is valid.	Check the configuration of the server's connectors for loops. If there are multiple virtual servers, ensure that none are set to "All Unassigned."

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.4.0	<p>Available on Exchange SP1 and later.</p> <p>Possible causes include:</p> <ul style="list-style-type: none"> ● Authoritative host not found in DNS ● Smarthost entry is incorrect ● FQDN name in HOSTS file (fixed in Windows 2000 SP3) ● DNS failure occurred, or you configured an invalid IP address as your smart host ● SMTP virtual server does not have a valid FQDN or lookup of your SMTP virtual server ● A contact's SMTP domain does not resolve to any SMTP address spaces 	<p>Use Nslookup to check the DNS configuration. Verify that the IP address is in IPv4 literal format. Verify the valid DNS entry for the server/computer name in question. If you rely on an FQDN in a HOSTS file, ignore and update entry in Exchange System Manager with valid IP address or correct name.</p>
5.4.4	<p>Available in Exchange 2000 SP1 and later versions.</p> <p>This NDR occurs if no route exists for message delivery, or if the categorizer could not determine the next hop destination.</p> <p>You set up a routing group topology, but no routing group connector exists between the routing groups.</p>	<p>Add or configure your routing group connector between routing groups.</p>

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.4.6	<p>A categorizer forward loop is detected.</p> <p>The targetAddress attribute is set on a mailbox-enabled user.</p> <p>This common hosting configuration problem occurs when someone creates a contact in one organizational unit, and then uses the provisioning tool to create a user in another organizational unit with the same e-mail address.</p>	<p>This happens when <i>contact A</i> has an alternate recipient that points to <i>contact B</i>, which then has an alternate recipient that points back to <i>contact A</i>. Check the contact's alternate recipient. Check and remove targetAddress attribute from mailbox-enabled users.</p> <p>For hosting, where you want to send mail from one user in one company in an organizational unit to another company in a separate organizational unit, you should configure the following two related objects:</p> <p>User: SMTP proxy: user@example.com</p> <p>Contact: targetAddress: user@example.com; SMTP proxy: contact@contoso.com, where contoso.com is the name of the second company.</p>

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.4.8	Available in Exchange 2000 SP1 and later versions. This message warns of a looping condition. One possibility is that one of the recipient policies includes a local domain that matches the FQDN of an Exchange server in the organization. When the categorizer is processing mail that is destined for a domain matching an Exchange server's FQDN, it will return this NDR.	Check your recipient policies. If a recipient policy contains an Exchange server's FQDN, you must remove that entry. Your recipient policy should not contain the FQDN of your server; instead, it should contain the mail domain only—for example, instead of server1.example.com, you would enter example.com.
5.5.0	Available in Exchange 2000 SP1 and later versions. A generic protocol error or an SMTP error causes this NDR. The remote SMTP server responds to a sending server's identifying EHLO with a 500-level error. The sending system will then terminate the connection and deliver an NDR indicating that the remote SMTP server cannot handle the protocol. For example, if a Microsoft Hotmail® e-mail account is no longer active, a 550 SMTP error will occur.	Run the SMTP Log or Netmon trace to see why the remote SMTP server rejects the protocol request.

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.5.2	<p>A generic SMTP error occurs when SMTP commands are sent out of sequence. For example, a server attempts to send an AUTH (authorization) command before identifying itself with an EHLO command.</p> <p>It is possible that this error can also occur when the system disk is full.</p>	<p>Run the SMTP Log or Netmon trace, and ensure there is enough disk storage and virtual memory for SMTP to operate.</p>
5.5.3	<p>Too many recipients on a message can cause this NDR.</p>	<p>The recipient limit is a configurable setting. To resolve this issue, either increase the recipient limit or break up the message into multiple messages to fit the server limit.</p> <p>Note The default recipient limit on an SMTP message is 5,000. To change this limit, start Exchange System Manager, expand Global Settings, right-click Message Delivery, click Properties, and then use the Defaults tab. This can also be a per-user setting in Active Directory.</p>

Table 2 NDR numerical codes and corresponding error conditions (continued)

NDR Code	Possible Cause	Troubleshooting
5.7.1	Possible causes include: <ul style="list-style-type: none"> ● General access denied, sender access denied—the sender of the message does not have the required permissions necessary to complete delivery. ● You are trying to relay your mail through another SMTP server, and the server does not permit you to relay. ● The recipient may have mailbox delivery restrictions enabled. For example, if a recipient's mailbox delivery restriction is set to receive mail from a distribution list only, non-member's mail will be rejected and produce this error. 	Check system privileges and attributes for the contact, and try sending the message again. Also, to resolve other potential issues, ensure that you are running Exchange 2000 SP1 or later.

Using the SMTP and X.400 Queues

SMTP uses the SMTP queues to deliver mail internally and externally; Exchange 5.5 servers, MAPI clients (such as Microsoft Outlook), and other mail connectors (such as Microsoft Exchange Connector for Lotus Notes and Microsoft Exchange Connector for Novell Groupwise) use the X.400 queues to send mail to and receive mail from Exchange. The following sections explain how to use both the SMTP and X.400 queues to troubleshoot message flow.

Understanding the SMTP Queues

During message categorization and delivery, the advanced queuing engine sends all mail through the SMTP queues of an SMTP virtual server. If there is a problem delivering the message at any point in the process, the message remains in the queue where the problem occurred.

Use the SMTP queues to isolate possible causes of mail flow issues. If a queue is in a “Retry” status, you should check the properties of the queue to determine the cause. For example, if the queue properties display a message similar to “An SMTP error has occurred,” you should review your server’s event logs to locate any SMTP errors. If there are no events in the log, you should increase the SMTP Protocol logging level. For more information about how to increase the SMTP Protocol logging level, see “Using Event Viewer” and “Configuring Diagnostic Logging for the SMTP Protocol” later in this chapter.

Table 3 lists the SMTP queues, including their descriptions and troubleshooting information for message accumulation in each queue.

Table 3 SMTP queues and probable causes for message accumulation

SMTP Queue	Description	Troubleshooting
[Local domain name] (Local Delivery)	Contains messages that are queued on the Exchange server for local delivery to an Exchange mailbox.	Messages can accumulate in this queue if the Exchange server is not accepting messages for local delivery. Slow or sporadic message delivery can indicate a looping message or a performance problem. This queue is affected by the Exchange store. Increase diagnostic logging for the Exchange store as described in “Configuring Diagnostic Logging for the SMTP Protocol” later in this chapter.

Table 3 SMTP queues and probable causes for message accumulation (continued)

SMTP Queue	Description	Troubleshooting
Messages awaiting directory lookup	Contains messages to recipients who have not yet been resolved against Active Directory. Messages are also held here while distribution lists are expanded.	Generally, messages accumulate in this queue because the advanced queuing engine is unable to categorize the message. The advanced queuing engine may not be able to access the global catalog servers and access recipient information, or the global catalog servers are unreachable or performing slowly. The categorizer affects this queue. Increase diagnostic logging for the categorizer as described in “Configuring Diagnostic Logging for the SMTP Protocol” later in this chapter.
Messages waiting to be routed	Holds messages until their next-destination server is determined, and then moves them to their respective link queues.	Messages accumulate in this queue if Exchange routing problems exist. Message routing may be backed up. Increase diagnostic logging for routing as described in “Configuring Diagnostic Logging for the SMTP Protocol” later in this chapter.

Table 3 SMTP queues and probable causes for message accumulation (continued)

SMTP Queue	Description	Troubleshooting
Remote delivery [Connector name Server name Remote domain]	Holds messages destined for a remote delivery. The name of the queue matches the remote delivery destination. It may be a connector, a server, or a domain.	If messages accumulate in this queue, you must first identify the status of the queue. If the queue is in "Retry," check the queue properties to determine the reason it is in this state. For DNS issues, use Nslookup and telnet to troubleshoot. If the host is unreachable, use telnet to ensure that the remote server is responding.
Final destination currently unreachable	The final destination server for these messages cannot be reached. For example, Exchange cannot determine a network path to the final destination.	Messages can accumulate in this queue if no route exists for delivery. Additionally, any time a connector or a remote delivery queue is unavailable or in "Retry" for a period of time, and no alternate route exists to the connector or remote destination, new messages queue here. This allows an administrator to fix the problem or define an alternate route. To get new messages to flow to their remote destination queue so you can force a connection and get a Netmon trace, simply restart the SMTP virtual server.
Pre-submission	Holds messages that have been acknowledged and accepted by the SMTP service. The processing of these messages has not begun.	Messages that are accumulating constantly may indicate a performance problem. Occasional peaks in performance can cause messages to appear in this queue intermittently.

Understanding the X.400 Queues

The X.400 queues are used by Exchange 2000 to submit mail to and receive mail from Exchange 5.5 servers and send mail through connectors to other mail servers. If you are experiencing mail flow problems when sending mail to an Exchange 5.5 or earlier server, or to another mail system, you should check the X.400 queues on Exchange 2000. If you are experiencing mail flow problems when sending mail to servers that are running earlier versions of Exchange 2000, you should also check the MTA queues on those servers.

Table 4 lists the X.400 queues, including their descriptions and troubleshooting information for message accumulation in each queue.

Table 4 X.400 queues and probable causes for message accumulation

X.400 Queue	Description	Troubleshooting
PendingRerouteQ)	Contains messages that are waiting to be rerouted after a temporary link outage.	Messages can accumulate in this queue if a route to a connector, to a different mail system, or to an Exchange 5.5 server is unavailable.
Next hop MTA	Contains messages destined to one of the following: <ul style="list-style-type: none"> ● Another gateway, such as a connector for Lotus Notes or Novell Groupwise ● X.400 link to an Exchange 5.5 site or a destination outside of the organization ● An Exchange MTA over the LAN—for example, destined to an Exchange 5.5 server in a mixed mode environment 	Messages can accumulate in this queue when Exchange 2000 experiences problems sending to another mail system, to an Exchange 5.5 server, or through an X.400 link. Increase diagnostic logging for the X.400 service as described in “Configuring Diagnostic Logging for the X.400 Service (MSExchangeMTA)” later in this chapter.

Viewing the Properties of a Queue

To view a queue's status (for example, "Ready" or "Retry"), as well as the explanation of why it is in this state, you must open the queue's properties.

► To check the properties of a queue

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand <*Server Name*>, expand **Protocols**, and then perform one of the following tasks:
 - To access the SMTP queues, expand **SMTP**, expand the SMTP virtual server you want, and then click **Queues**.
 - To access the X.400 queues, expand **X.400**, and then click **Queues**.
3. In the details pane, right-click the queue you want, and then click **Properties** to view additional information about the queue's state.
4. In the queue **Properties**, you can see all the information listed in the column beside the queue in the detail pane, as well additional information about its state under **Status** (Figure 34).

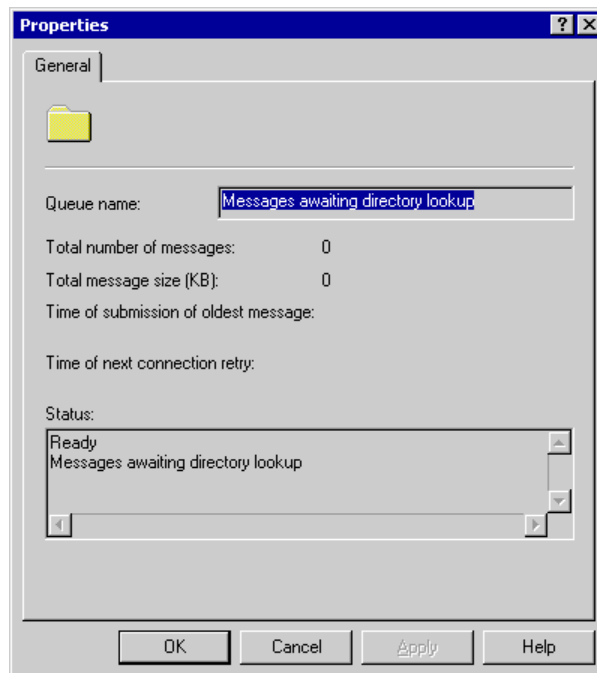


Figure 34 The SMTP queue **Properties** dialog box

Viewing the Messages in a Queue

If you experience mail flow problems, it is important to determine if you are having global problems or problems with individual recipients or domains. Viewing the messages in a queue can help you figure this out.

► **To view messages in a queue**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, expand *<Server Name>*, expand **Protocols**, and then perform one of the following tasks:
 - To access the SMTP queues, expand **SMTP**, expand the SMTP virtual server you want, and then click **Queues**.
 - To access the X.400 queues, expand **X.400**, and then click **Queues**.
3. In the details pane, right-click the queue you want, and then click **Enumerate 100 messages** to view the first one hundred messages. If any messages are in the queue, the first one hundred will display in the details pane.
4. To view the properties of an individual message, in the details pane, right click the message you want, and then click **Properties**. The message's **Properties** dialog box displays the sender's name, the recipient's name, the message size, and other details about the message.

Note For more information about using the SMTP queues, see Microsoft Knowledge Base article Q268163, "XCON: How To Configure a Virtual Server Part 2" (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=268163>).

Using Message Tracking Center

To log information about messages that are sent over your messaging system, you can use Message Tracking Center in Exchange. Message Tracking Center logs information about the sender, the mail message, and the message recipients. Specifically, you can determine statistics such as the time the message was sent or received, the message size and priority, and the list of message recipients. You can also log the subject line of e-mail messages. Message Tracking Center searches for all types of messages, including system messages, public folder messages, and e-mail messages.

You must enable Message Tracking Center on each server for which you want to track messages. When enabled, all messages routed through a server are added to the message tracking logs.

► **To enable Message Tracking Center on a server**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, right-click the server on which you want to enable message tracking, and then click **Properties**.
3. On the **General** tab, select the **Enable message tracking** check box.
4. To record the subject of any message sent to, from, or through the server, select the **Enable subject logging and display** check box.

Note Enabling subject logging causes some performance degradation.

5. Under **Log file maintenance**, you can prevent the removal of log files or modify the length of time the log files are kept. The default period that tracking logs are kept is seven days.

Note On servers that process large quantities of mail, the tracking logs grow quickly. Ensure that you have adequate disk space for the log files and for other services or applications that use this disk.

6. Click **OK** or **Apply**. You do not need to restart services for this change to take effect.

For more information about how to use Message Tracking Center, see Microsoft Knowledge Base article Q262162, “XADM: Using the Message Tracking Center to Track a Message” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=262162>).

Using Event Viewer

In Event Viewer, both the Application Log and the System Log contain errors, warnings, and informational events related to the operation of Exchange, the SMTP service, and other applications. To help you identify the cause of message flow issues, carefully review the data contained in the Application Log and System Log.

Viewing the Application Log

Use the following procedure to view errors, warnings, and informational events in the Application Log.

► **To view the Application Log in Event Viewer**

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Event Viewer**.
2. In the console tree, click **Application Log**.

3. To sort the log alphabetically and quickly locate an entry for an Exchange service, in the details pane, click **Source**.
4. Double-click a log entry to open an event's properties page.
5. To filter the log to list entries for a specific type of Exchange-related event, from the **View** menu, click **Filter**.
6. In **Application Log Properties**, use the **Event source** list to select an Exchange-related event source. For example:
 - **MSExchangeTransport** Select this event source to view events recorded when SMTP is used to route messages.
 - **IMAP4Svc** Select this event source to view events related to the service that allows users to access mailboxes and public folders through IMAP4.
 - **MSExchangeAL** Select this event source to view events related to the service that addresses e-mail through address lists.
 - **MSExchangeIS** Select this event source to view events related to the service that allows access to the Exchange Information Store service.
 - **MSExchangeMTA** Select this event source to view events related to the service that allows X.400 connectors to use the message transfer agent (MTA).
 - **MSExchangeMU** Select this event source to view events related to the metabase update service, a component that reads information from Active Directory and transposes it to the local IIS metabase.
 - **MSExchangeSA** Select this event source to view events recorded when Exchange uses Active Directory to store and share directory information.
 - **MSExchangeSRS** Select this event source to view events recorded whenever Site Replication Service (SRS) is used to replicate computers running Exchange 2000 with computers running Exchange 5.5.
 - **POP3Svc** Select this event source to view events recorded whenever POP3 is used to access e-mail.
7. In the **Category** list, select a specific set of events or, to view all events for that event source, leave the default setting at **All**.
8. Click **OK**.

Viewing the System Log

Use the following procedure to view errors, warnings, and informational events in the System Log for the SMTP service.

► **To view the System Log in Event Viewer**

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Event Viewer**.
2. In the console tree, click **System Log**.
3. To sort the log alphabetically and quickly locate an entry for an Exchange service, in the details pane, click **Source**.
4. Double-click a log entry to open an event's properties page.
5. To filter the log to list entries for a specific type of SMTP service events, from the **View** menu, click **Filter**.
6. In **System Log Properties**, in the **Event source** list, select **SMTPSVC**.
7. In the **Category** list, select a specific set of events or, to view all events for the SMTP service, leave the default setting at **All**.
8. Click **OK**.

Configuring Diagnostic Logging for the SMTP Protocol

To help you determine the root of a transport issue, view events that relate to `MSExchangeTransport`. If you experience problems with Exchange message flow, immediately increase the logging levels relating to `MSExchangeTransport`. Logging levels control the amount of data that is logged in the Application Log. The more events logged, the more transport-related events you can view in the Application Log; therefore, you have a better chance in determining the cause of the message flow problem. The SMTP log file is located in the `Exchsrvr\ Server_name .log` folder.

As discussed in “Understanding the SMTP Queues” earlier in this chapter, issues with specific routing and transport components can cause messages to accumulate in a queue. If you are having problems with a specific queue, turn up logging for the component affecting the queue.

Modifying Logging Settings

The following procedure explains how to modify diagnostic logging related to `MSExchangeTransport`.

► **To modify logging settings for `MSExchangeTransport`**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, right-click *<Server Name>*, and then click **Properties**.
3. Click the **Diagnostics Logging** tab.

4. Under **Services**, click **MSExchangeTransport**.
5. Under **Categories**, click the category for which you want to configure the logging level:
 - Select **Routing Engine/Service** to troubleshoot routing issues. Increase the logging level for this component if messages are accumulating in the **Messages waiting to be routed** SMTP queue.
 - Select **Categorizer** to troubleshoot problems with address resolution in Active Directory, distribution list expansion, and other categorizer issues. Increase the logging level for this component if messages are accumulating in the **Messages waiting to be routed** SMTP queue.
 - Select **Connection Manager** to troubleshoot issues with dial-up and virtual private network connectivity through Connection Manager.
 - Select **Queuing Engine** to troubleshoot problems with the queuing engine. Increase the logging level for this component if you are experiencing mail flow problems and mail is not accumulating in any of the queues.
 - Select **Exchange Store Driver** to troubleshoot issues with the Exchange store driver. Increase the logging level for this component if messages are accumulating in the local delivery SMTP queue, the X.400 queues, or if you have problems receiving mail from Exchange 5.x servers or other mail systems.
 - Select **SMTP Protocol** to troubleshoot general SMTP issues. Increase the logging level for this component if messages are accumulating in the **Remote delivery** SMTP queue to determine if SMTP errors are causing the bottleneck.
 - Select **NTFS store driver** to troubleshoot issues with the NTFS store driver. Increase the logging level for this category if messages are accumulating in the local delivery SMTP queue.
6. Under **Logging level**, click **None**, **Minimum**, **Medium**, or **Maximum**. Click **Maximum** for troubleshooting purposes.

Caution If you increase the logging levels for Exchange services, you will experience some performance degradation. It is recommended that you increase the size of the Application Log to contain all the data produced. If you do not increase the size of the Application Log, you will receive frequent reminders that the Application Log is full.

Enabling Debugging Level Logging

If you are experiencing mail flow issues and want to view all events, you can modify a registry key to set logging to the highest level (level 7).

Caution This procedure contains information about modifying the registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore the registry if a problem occurs. For information about how to back up, restore, and edit the registry, see Microsoft Knowledge Base article Q256986, "Description of the Microsoft Windows Registry" (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=256986>).

► **To enable logging at the debugging level**

1. Start Registry Editor: From the **Start** menu, click **Run**, and then type **regedt32**.
2. In Registry Editor, locate and click the following registry key:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\
MSExchangeTransport\Diagnostics\SMTP Protocol
3. Set the value to 7.

Configuring Diagnostic Logging for the X.400 Service (MSExchangeMTA)

This section explains how to configure diagnostic logging for the X.400 service (MSExchangeMTA) on your Exchange 2000 server. If you have to troubleshoot mail flow problems for servers running Exchange 5.5 and earlier, for other mail systems, or for X.400 connectors, it is very useful to increase the logging level for MSExchangeMTA .

► **To configure logging for MSExchangeMTA**

1. Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In the console tree, expand **Servers**, right-click <*Server Name*>, and then click **Properties**.
3. Click the **Diagnostics Logging** tab.
4. Under **Services**, click **MSExchangeMTA**
5. Under **Categories**, click **X.400 Service** to troubleshoot delivery problems to servers running Exchange 5.5 and earlier and other systems.
6. Under **Logging level**, click **None**, **Minimum**, **Medium**, or **Maximum**. Click **Maximum** for troubleshooting purposes.

8

Reference

This chapter contains reference material about the following topics:

- SMTP commands and definitions
- Internal SMTP transport mechanisms
- SMTP event sinks
- Common ports used by Exchange

SMTP Commands and Definitions

Table 5 lists the SMTP commands and functions provided by Windows SMTP service.

Table 5 SMTP commands and functions

SMTP Command	Command Function
HELO	Sent by a client to identify itself, usually with a domain name.
EHLO	Enables the server to identify its support for ESMTP commands.
MAIL FROM	Identifies the sender of the message; usually used in the form MAIL FROM:.
RCPT TO	Identifies the message recipients; used in the form RCPT TO:.

Table 5 SMTP commands and functions (continued)

SMTP Command	Command Function
TURN	Allows the client and server to switch roles and send mail in the reverse direction without having to establish a new connection.
ATRN (Authenticated TURN)	The ATRN command optionally takes one or more domains as a parameter. The ATRN command must be rejected if the session has not been authenticated
SIZE	Provides a mechanism by which the SMTP server can indicate the maximum size message supported. Compliant servers must provide size extension to indicate the maximum size message that can be accepted. Clients should not send messages larger than the size indicated by the server.
ETRN	An extension of SMTP. ETRN is sent by an SMTP server to request that another server send any e-mail messages it has.
PIPELINING	Provides the ability to send a stream of commands without waiting for a response after each command.
CHUNKING	An ESMTP command that replaces the DATA command. So the SMTP host does not have to continuously scan for the end of the data, this command sends a BDAT command with an argument that contains the total number of bytes in a message. The receiving server counts the bytes in the message and, when the message size equals the value sent by the BDAT command, the server assumes it has received all of the message data.
DATA	Sent by a client to initiate the transfer of message content.

Table 5 SMTP commands and functions (continued)

SMTP Command	Command Function
DSN	An ESMTP command that enables delivery status notifications.
RSET	Nullifies the entire message transaction and resets the buffer.
VRFY	Verifies that a mailbox is available for message delivery; for example, vrify sfine verifies that a mailbox for sfine resides on the local server.
HELP	Returns a list of commands supported by the SMTP service.
NOOP	Verifies that the SMTP service is still operating on commands. If the service is still running when this command is sent, a 250 "OK" code is returned.
QUIT	Terminates the session.

Table 6 lists the extended SMTP commands that Exchange makes available to the SMTP service.

Table 6 Extended SMTP commands

SMTP Command	Command Function
X-EXPS GSSAPI	A method used by Exchange 2000 servers to authenticate.
X-EXPS=LOGIN	A method used by Exchange 2000 servers to authenticate.
X-EXCH50	Provides the ability to propagate message properties during server-to-server communication.
X-LIN2STATE	Adds support for link state routing in Exchange.

Understanding the Internal SMTP Transport Mechanisms

This section provides detailed descriptions of the components involved in sending and receiving mail in SMTP. The following are two important SMTP components involved in mail transport:

Advanced queuing engine

The advanced queuing engine is responsible for several aspects of message delivery. Specifically, the advanced queuing engine retrieves messages from SMTP or the Exchange store driver, categorizes them, determines each message's destination, and then provides an interface to the multiple queues to which a message can be assigned while awaiting delivery.

Message categorizer

The message categorizer is a component of the advanced queuing engine that sends lightweight directory access protocol (LDAP) queries to the global catalog server to perform directory lookups. These queries retrieve the following information:

- The recipient e-mail addresses
- The mailbox store on which a recipient mailbox resides
- The Exchange server hosting that mailbox store

Figure 35 illustrates the transport components involved in mail flow. The shaded areas depict SMTP transport components.

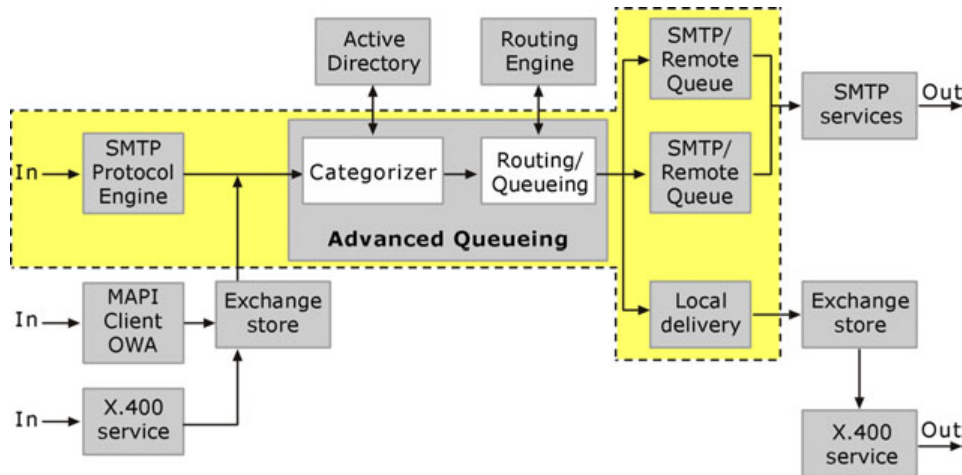


Figure 35 Message flow through internal transport components

Receiving Internet Mail

Exchange receives incoming Internet mail in the following manner.

1. The SMTP server uses DNS to query for the preferred MX (mail exchanger) record of the destination domain or target server. DNS returns a list of A (host) records, which resolve to an IP address or addresses of the server.
2. The sending SMTP server initiates a connection to port 25 of the destination SMTP server. The destination SMTP server is the SMTP virtual server located on the physical gateway server that is configured to accept incoming Internet mail for the domain to which the mail is addressed.
3. Ideally, the inbound SMTP server only accepts the incoming message if it is destined for an SMTP mail domain defined in a recipient policy (unless the server is open to relay, which is strongly discouraged).
4. When the message is accepted, the SMTP virtual server creates an envelope for the message—this message structure is called IMAILMSG. IMAILMSG contains all of the properties of the message, including the sender and recipient names.
5. The message categorizer (a component of the advanced queuing engine) makes an LDAP query to the global catalog server to find the homeMdb attribute of recipient. The message categorizer then stamps the FQDN of this Exchange server on the IMAILMSG object. The homeMdb attribute is the user's home mailbox server; this is the location where the user's mailbox store and mailbox reside.
6. The categorizer marks the message for local delivery, and the advanced queuing engine transfers the message to the Exchange store driver. The Exchange store driver then delivers the message to the mailbox store.
7. If the user's mailbox store is not located on that Exchange server, the message categorizer transfers the message to the advanced queuing engine. The advanced queuing engine then calls the routing engine to determine the best way to send the message to the server (based on link state routing), and determines the next destination or hop in the route to the user's home server.
8. Finally, complete with destination information from the message categorizer and routing information from the routing engine, the advanced queuing engine sends the message to its final destination in one of the following ways:
 - If the destination is a local domain, the message is delivered to the SMTP virtual server located on the Exchange server where the user's mailbox resides. If the user's mailbox is in a remote routing group, the message may have to be sent through other connectors.
 - If the destination is remote, the message is delivered to the SMTP server for remote domains in a different remote queue. An incoming message will be sent to a remote domain only if one of the following configurations is applied:

- The Exchange server is open for relay.
- The user sending the message is authorized to relay.
- Another connector is configured that allows relaying to these domains.
- If the destination is a connector to another system or to an earlier version of Exchange, the Exchange store driver submits the mail to the MTA.

Sending Internet Mail

Internet mail is sent through Exchange in the following manner.

1. An internal user sends a message to a remote domain. The message is submitted on the Exchange server on which the user's mailbox resides.
2. The message is submitted to the advanced queuing engine in one of two ways:
 - If the message was sent using an Outlook Web Access or Outlook (MAPI) client, the Exchange store submits the message to the advanced queuing engine through the store driver.
 - If the message was sent using a Post Office Protocol (POP) or an Internet Mail Access Protocol (IMAP) client, the Exchange store submits the message to SMTP, which in turn passes it to the advanced queuing engine.
3. The message categorizer then queries the global catalog server with the recipient address to find the user. If the recipient address is not in a recipient policy, or if a matching recipient with a proxy address does not exist, the recipient address will not be stored in Active Directory; as a result, the message categorizer determines that the message is bound for a remote domain.
4. The SMTP virtual server located on the Exchange server that performs categorization then uses its metabase information to locate the SMTP virtual server or connector that handles remote domains for “*” (all addresses) or to locate another SMTP connector that contains an address space that more closely matches the remote domain.
5. With this information, the server determines whether to send the message, to route to the smart host, or to use an SMTP connector with the remote address space.
6. If there are multiple connectors or virtual servers that handle outbound mail, the advanced queuing engine determines the virtual server or connector with the address space that most closely matches the address space of the remote domain.

7. The message is routed to the connector or to the outbound SMTP virtual server that is responsible for delivery.
8. The connector or SMTP virtual server on the Exchange server then performs one of two tasks:
 - Uses DNS to look up the IP address for the target domain, and then attempts delivery of the message
 - Forwards the message to a smart host that assumes responsibility for the DNS resolution and delivery

Event Sinks

Event sinks can be used to extend and modify the behavior of the Windows 2000 SMTP service. In fact, the reason Exchange 2000 requires the Windows 2000 SMTP service to function is that most of the transport functionality in Exchange 2000 is accomplished with this architecture. Therefore, after you reinstall IIS or the Windows 2000 SMTP service, you must also reinstall Exchange or use the SP2 SMTP Reinstall tool.

An SMTP service event is the occurrence of some activity within the SMTP service, such as the transmission or arrival of an SMTP command or the submission of a message into the SMTP service transport component. When a particular event occurs, the SMTP service uses an event dispatcher to notify registered event sinks of the event. When notifying event sinks, the SMTP service passes information to the sink in the form of COM object references.

SMTP service events can be broken down into two general categories:

Protocol Events

Protocol events occur when SMTP commands are either received or transmitted over the network. These events occur when:

- A client SMTP service or mail user agent uses SMTP to transmit messages for delivery to the local service.
- The SMTP service relays messages to other SMTP services.

Note Protocol events must be written in C++ using Component Object Model (COM) or Active Template Library (ATL).

Transport Events

Transport events occur when the SMTP service receives a message, and that message passes through the SMTP core transport. During the passage through the transport, the message is categorized (examined and placed into categories), and then either delivered to a local storage location or, if it is not local, relayed to another destination.

The default Windows 2000 protocol and transport events are only accessible by writing Component Object Model (COM) objects in Microsoft Visual C++®. These events are fast, require no extra processing, and offer access to the lowest level message properties; however, these events are more complex to write. For smaller jobs that don't require high performance, you can use the CDO_OnArrival event, which can be written using Microsoft Visual Basic®, Scripting Edition (VBScript).

If you would like to write one of these event sinks, you can download the Platform SDK (<http://go.microsoft.com/fwlink/?LinkId=12059>) or view the MSDN online document "Microsoft SMTP Server Events for Windows 2000" (<http://go.microsoft.com/fwlink/?LinkId=12279>).

Common Ports Used by Exchange

Table 7 lists the common ports used by Exchange. For detailed information about which ports need to be opened internally or externally, see the Exchange online book, *Using Microsoft Exchange 2000 Front-End Servers* (<http://go.microsoft.com/fwlink/?LinkId=12055>).



Table 7 Ports used by Exchange

Protocol	Port	Description
SMTP	TCP: 25	The SMTP service uses TCP port 25.
DNS	TCP/UDP: 53	DNS listens on port 53. Domain controllers use this port.
SMTP/LSA	TCP: 691	The Exchange Routing Engine (also known as RESvc) listens for routing link state information on this port.
LDAP	TCP/UPD 389	LDAP used by Active Directory, Active Directory Connector, and the Microsoft Exchange Server 5.5 directory uses this port.
LDAP/SSL	TCP/UDP 636	LDAP over Secure Sockets Layer (SSL) uses this port.
LDAP.	TCP/UDP: 379	The Site Replication Service (SRS) uses this port.
LDAP.	TCP/UDP: 390	This is the recommended alternate port to configure the Exchange Server 5.5 LDAP protocol when Exchange Server 5.5 is running on an Active Directory domain controller.

Table 7 Ports used by Exchange (continued)

Protocol	Port	Description
LDAP	TCP: 3268	Global catalog. The Windows 2000 Active Directory global catalog (a domain controller “role”) listens on TCP port 3268.
LDAP/SSLPort	TCP: 3269	Global catalog over SSL. Applications that connect to TCP port 3269 of a global catalog server can transmit and receive SSL encrypted data.
IMAP4	TCP: 143	IMAP uses this port.
IMAP4/SSL.	TCP: 993	IMAP4 over SSL uses this port.
POP3	TCP: 110	POP3 uses this port.
POP3/SSL	TCP: 995	POP3 over SSL uses this port.
NNTP	TCP: 119	Network News Transfer Protocol (NNTP) uses this port.
NNTP/SSL	TCP: 563	NNTP over SSL uses this port.
HTTP	TCP: 80	HTTP uses this port.
HTTP/SSL	TCP: 443	HTTP over SSL uses port 443.

Appendix



A

Additional Resources

The following resources provide valuable information regarding Exchange and SMTP.

Technical Papers

Configuring Microsoft Exchange 2000 Server for the Internet

(<http://go.microsoft.com/fwlink/?LinkId=12056>)

Configuring and Securing Exchange 2000 Server and Clients

(<http://go.microsoft.com/fwlink/?LinkId=10733>)

Securing Exchange Server

(<http://go.microsoft.com/fwlink/?LinkId=11923>)

Security Operations Guide for Exchange 2000 Server

(<http://go.microsoft.com/fwlink/?LinkId=11906>)

Microsoft Knowledge Base Articles

The following Microsoft Knowledge Base articles are available on the Web at

<http://support.microsoft.com/>:

Q288175, “XCON: Recipient Policy Cannot Match the FQDN of Any Server in the Organization, 5.4.8 NDRs”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=288175>)

Q140933, “XFOR: SMTP Proxy Address Generated Incorrectly”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=140933>)

Q262162, “XADM: Using the Message Tracking Center to Track a Message”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=262162>)

- Q265293, “XFOR: How to Configure the SMTP Connector in Exchange 2000”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=265293>)
- Q276388, “XIMS: How to Configure Exchange 2000 Behind Proxy Server 2.0”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=276388>)
- Q260973, “XCON: Setting Up SMTP Domains for Inbound and Relay E-Mail in Exchange 2000 Server”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=260973>)
- Q284204, “XCON: Delivery Status Notifications in Exchange 2000 Server”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=284204>)
- Q275596, “XADM: MAPI Messages Stack Up in Send Queue to the Host Specified in Forward Unresolved Recipients”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=275596>)
- Q278339, “XGEN: TCP/UDP Ports Used By Exchange 2000 Server”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=278339>)
- Q288635, “XIMS: ResolveP2 Functionality in Exchange 2000 Server”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=288635>)
- Q280794, “XIMS: Message Cannot Be Sent to Domains with MX Record Pointing to CNAME Record”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=280794>)
- Q298448, “Windows 2000 DNS and Active Directory Information and Technical Resources”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=298448>)
- Q309508, “XCCC: IIS Lockdown and URLscan Configurations in an Exchange Environment”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=309508>)
- Q315591, “XCON: Authoritative and Non-Authoritative Domains in Exchange 2000”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=315591>)
- Q321721, “XCON: Sharing SMTP Address Spaces in Exchange 2000”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=321721>)
- Q266686, “XADM: How to Configure a SMTP Virtual Server Part I”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=266686>)
- Q268163, “XADM: How to Configure a SMTP Virtual Server Part II”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=268163>)
- Q265293, “XFOR: How to Configure the SMTP Connector in Exchange 2000”
(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=265293>)

Q251700, “XCON: Dial-up Feature of Internet Mail Service Not Available in Exchange 2000”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=251700>)

Q253108, “XADM: Domain-Specific Information Not Transferred When You Upgrade Internet Mail Service to Exchange 2000”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=253108>)

Q246739, “XADM: Exchange Front-end/Back-end Terminology and Implementation”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=246739>)

Q284204, “XCON: Delivery Status Notifications in Exchange 2000 Server”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=284204>)

Q293800, “XCON: How to Set Up Windows 2000 as a SMTP Relay Server or Smart Host”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=293800>)

Q319759, “XADM: How to Configure Exchange 2000 Server to Forward Messages to a Foreign Messaging System That Shares the Same SMTP Domain Name Space”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=319759>)

Q304897, “XIMS: Microsoft SMTP Servers May Seem to Accept and Relay E-Mail Messages in Third-Party Tests”

(<http://go.microsoft.com/fwlink/?LinkId=3052&ID=304897>)

Other Useful Resources

Microsoft Exchange 2000 Server Resource Kit

(<http://go.microsoft.com/fwlink/?LinkId=12058>)

Exchange 2000 Service Pack 3 Deployment Guide

(<http://go.microsoft.com/fwlink/?LinkId=12337>)

The Microsoft Platform SDK

(<http://go.microsoft.com/fwlink/?LinkId=12059>)

Microsoft SMTP Server Events For Windows 2000

(<http://go.microsoft.com/fwlink/?LinkId=12279>)

The IIS Lockdown Wizard

(<http://go.microsoft.com/fwlink/?LinkId=12281>)



Does this book help you? Give us your feedback. On a scale of 1 (poor) to 5 (excellent), how do you rate this book?

<mailto:exchdocs@microsoft.com?subject=Feedback: Configuring SMTP in Microsoft Exchange 2000 Server>.

For more information about Exchange, see (<http://www.microsoft.com/exchange/>).

To download a self-extracting executable of all Exchange Product Team technical articles and online books, see (<http://go.microsoft.com/fwlink/?LinkId=10687>).