



Using Microsoft® Exchange 2000 Front-End Servers



KC Lemson
Michele Martin

Microsoft®

Using Microsoft® Exchange 2000 Front-End Servers



KC Lemson
Michele Martin

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, Windows, Windows NT, Active Directory, and Outlook 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: June 2000

Updated: October 2002

Applies To: Exchange 2000 Server SP3

Editors: Janet Lowen, Brendon Bennett

Technical Reviewers: KC Lemson, Allen Atwood, Bryan Atwood, Karim Batthish, Ron Mondri, Andrew Moss, Brendan Power

Artist: Kristie Smith

Production: Stephanie Schroeder

Table of Contents

Introduction

Assumed Knowledge	1
Overview of Front-End and Back-End Topologies	2
Advantages of a Front-End and Back-End Topology	3

Chapter 1

How Front-End and Back-End Topology Works	5
Integration with Internet Information Services (IIS).....	6
Dependency on DSAccess.....	6
DSAccess in Perimeter Networks	6
System Attendant on Front-End Servers	7
Supporting POP and IMAP Clients.....	8
Authentication for POP and IMAP Clients.....	9
IMAP Access to Public Folders.....	9
Running SMTP for POP and IMAP Clients.....	10
Supporting HTTP (Outlook Web Access and Web Folders).....	10
Finding User Mailboxes	11
Logging on to Outlook Web Access.....	12
Allowing Use of Outlook Web Access to Internal Clients Only.....	13
Simplifying the Outlook Web Access URL.....	13
Enabling the “Change Password” Feature	14
Finding Public Folders	15
Public Folder Referrals	16
The Default (MAPI) Public Folder Tree.....	18
General-Purpose Public Folder Trees	18
When Content Is Not Available on the Back-End Server.....	19
Back-End Server Downtime	19
Adding or Removing Back-End Servers	20

Authentication Issues for HTTP	21
Dual Authentication	22
Pass-Through Authentication	22
User Logon Information	23
Remote Procedure Calls (RPCs) in Front-End and Back-End Topology	23
Stopping RPC Traffic Across the Intranet Firewall	24

Chapter 2

Deployment Considerations	25
Recommended Server Configurations and Ratios.....	25
Load Balancing	26
Reducing Virtual Server Creation	27
Using Firewalls	27
Port Filtering.....	27
IP Filtering	28
Application Filtering	28
Securing Communication: Client to Front-End Server	29
Configuring SSL in a Front-End and Back-End Topology	29
SSL Accelerators.....	30
Securing Communication: Front-End to Other Servers.....	31
IP Security (IPSec).....	31
IPSec Protocols	32
IPSec Policy	32
IPSec with Firewalls and Filtering Routers	32
Service Packs: Upgrading Front-End and Back-End Servers.....	34
Upgrading Considerations for Outlook Web Access	34

Chapter 3

Scenarios	37
Standard Front-End and Back-End Topology Without a Firewall.....	38
Scenario	38
Setup Instructions	38
Discussion	38
Issues	39
Front-End Server Behind a Firewall	39
Scenario	40
Setup Instructions	40
Discussion	40

Web Farm With a Firewall.....	41
Scenario	41
Setup Instructions.....	41
Discussion	42
Issues	42
Front-End Server in a Perimeter Network.....	42
Scenario	42
Setup Instructions.....	43
Discussion	43
Issues	43
Advanced Firewall in a Perimeter Network	44
Scenario	45
Setup Instructions.....	46
Discussion	46
Issues	47

Chapter 4

Configuring a Front-End Server.....	49
Hosting Multiple Domains	49
Method One: Create Additional Virtual Servers	50
Method Two: Create Additional Virtual Directories	52
Creating HTTP Virtual Servers	52
Creating Virtual Directories	53
Configuring Authentication.....	54
Configuring The Front-End Server To Assume a Default Domain	55
Allowing the Use of an E-Mail Address as the Logon User Name	56
Disabling Unnecessary Services	56
Running the IIS Lockdown Wizard	57
Dismounting and Deleting Public and Mailbox Stores	58
Configuring Network Load Balancing	58
Configuring SSL.....	59
Configuring SMTP.....	59
Mail for Internal Domains.....	59
Mail for External Domains.....	60
Configuring DSAccess for Perimeter Networks.....	60
Disabling the NetLogon Check.....	61
Disabling the Directory Access Ping	61
Specifying Domain Controllers and Global Catalog Servers	62

Chapter 5

Configuring a Back-End Server	63
Creating and Configuring HTTP Virtual Servers on Back-End Servers	64
Method One: Additional Virtual Servers	64
Method Two: Additional Virtual Directories	65
Configuring Authentication on Back-End Servers	65

Chapter 6

Configuring Firewalls.....	67
Configuring an Internet Firewall.....	67
Configuring ISA Server.....	68
Configuring an Intranet Firewall.....	69
Basic Protocols	69
Active Directory Communication.....	70
Domain Name Service (DNS).....	70
IPSec.....	71
Remote Procedure Calls (RPCs).....	71
Filtering RPC Traffic with ISA Server.....	71
Stopping RPC Traffic.....	72
Restricting RPC Traffic.....	72

Chapter 7

Front-End and Back-End Topology Checklist	75
--	-----------

Chapter 8

Front-End and Back-End Topology Troubleshooting Steps	81
Troubleshooting Tools	81
General Troubleshooting Steps.....	82
Logon Failures.....	83
Troubleshooting Outlook Web Access	83

Additional Resources

Technical Papers.....	85
Microsoft Knowledge Base Articles	85
Other Useful Resources.....	87



Introduction

Microsoft® Exchange 2000 Server supports the deployment of Exchange in a manner that distributes server tasks among front-end and back-end servers. A front-end server accepts requests from clients and proxies them to the appropriate back-end server for processing. This book discusses how Exchange supports front-end and back-end server architecture, with an emphasis on HTTP for Outlook Web Access. This book also describes several front-end and back-end scenarios and provides configuration recommendations.

Important Most of the information in this book pertains to Exchange 2000 Service Pack 2 (SP2) or later versions. Therefore, if you are running Exchange 2000 Service Pack 1 (SP1) or earlier, you should upgrade to Exchange SP2 to take full advantage of the features described in this book.

Assumed Knowledge

You should have an understanding of Microsoft Outlook Web Access, Hypertext Transfer Protocol (HTTP), Simple Mail Transfer Protocol (SMTP), Post Office Protocol version 3 (POP3), and Internet Message Access Protocol (IMAP) version 4rev1 in a standard Exchange deployment, as well as basic Exchange 2000 and Microsoft Windows® 2000 Internet Information Services (IIS) concepts.

Overview of Front-End and Back-End Topologies

Figure 1 illustrates simple Exchange front-end and back-end topology. However, you can increase the security of your front-end and back-end topology by incorporating firewalls.

Figure 2 illustrates a recommended topology that uses an advanced firewall (such as Microsoft Internet Security and Acceleration [ISA] Server) between the Internet and the front-end server.

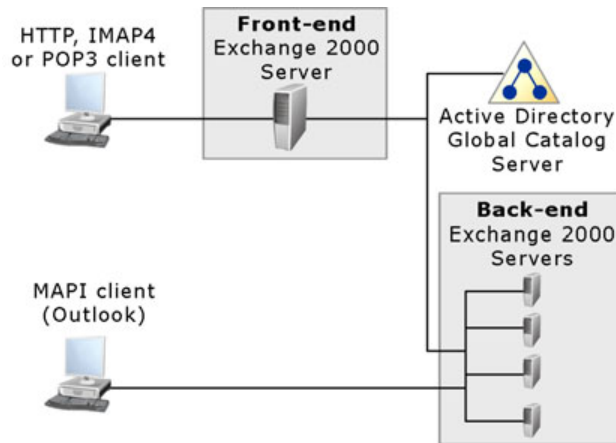


Figure 1 A simple Exchange 2000 front-end and back-end topology

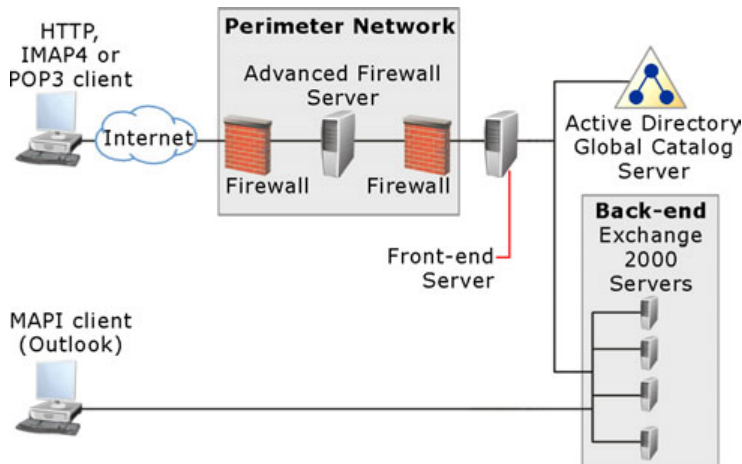


Figure 2 A secure Exchange 2000 front-end and back-end topology

Note Exchange 2000 Server can be used only as a back-end server in a front-end and back-end configuration. However, Exchange 2000 Enterprise Server can be used as a front-end server or a back-end server in a front-end and back-end configuration. For more information about the differences between Exchange 2000 Server and Exchange 2000 Enterprise Server, see Microsoft Knowledge Base article Q296614, “XADM: Differences Between Exchange 2000 Standard and Enterprise Versions” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=296614>).

Advantages of a Front-End and Back-End Topology

The front-end and back-end server topology is recommended for multiple-server organizations that use Microsoft Outlook Web Access (HTTP), POP, or IMAP and for organizations that want to provide HTTP, POP, or IMAP access to their employees over the Internet. After receiving a request, the front-end server uses Lightweight Directory Access Protocol (LDAP) to query the Microsoft Windows 2000 Active Directory® directory service and determine which back-end server holds the requested resource.

Note A front-end server is a specially configured server running Exchange 2000. A back-end server is a server with a standard configuration running Exchange 2000. There is no configuration option to designate a server as a back-end server. The term “back-end server” refers to all servers in an organization that are not front-end servers after a front-end server is introduced into the organization.

Using a front-end and back-end deployment has the following advantages:

Single namespace

The primary advantage of a front-end and back-end server architecture is the ability to expose a single, consistent namespace. You can define a single namespace for users to access their mailboxes (for example, <http://mail> for Outlook Web Access). Without a front-end server, each user must know the name of the server that stores their mailbox. This complicates administration and compromises flexibility, because every time your organization grows or changes and you move some or all mailboxes to another server, you must inform the users. With a single namespace, users can use the same URL or POP and IMAP client configuration, even if you add or remove servers or move mailboxes from server to server. In addition, creating a single namespace ensures that Outlook Web Access, POP, or IMAP access remains scalable as your organization grows.

Ability to balance processing tasks between servers

You can configure servers running Exchange 2000 to support Secure Sockets Layer (SSL) traffic between the client and the server to protect the traffic from third-party interception. However, encrypting and decrypting message traffic uses processor time. When SSL encryption is in use, front-end and back-end server architecture provides an advantage because the front-end servers can handle all encryption and decryption processing. In addition, you can use an SSL accelerator to further mitigate the impact encryption and decryption has on the server. An SSL accelerator improves performance by removing processing tasks from back-end servers, while still allowing data to be encrypted between the client and the server running Exchange.

Security

You can position the front-end server as the single point of access on or behind an Internet firewall that is configured to allow only traffic to the front-end from the Internet. Because the front-end server has no user information stored on it, it provides an additional layer of security for the organization. In addition, you can configure the front-end server to authenticate requests before proxying them, protecting the back-end servers from denial-of-service attacks.

Increased IMAP access to public folders

The IMAP protocol allows a server to refer a client to another server. Exchange 2000 supports this referral functionality in cases where a public folder store on a particular server does not contain the content requested and the client needs to be referred to another server. However, this requires a client that supports IMAP referrals, and most clients do not support referrals. (The University of Washington Pine client and toolkit is one example of a client that supports referrals.) When a non referral-enabled IMAP client connects through a front-end server, the client has access to the entire public folder hierarchy. When a front-end server proxies a command to a back-end server, it automatically handles any referral response that is passed back when attempting to access a folder that is not available on the back-end server. This makes the referral transparent to the client. For more information about nonreferral-enabled IMAP clients, see Request for Comments (RFC) 2221 and RFC 2193.

The front-end and back-end architecture supports HTTP, POP, and IMAP. You can also install SMTP on the front-end server, although there are essentially no differences between SMTP on a front-end server or back-end server.

Note The MAPI remote procedure call (RPC) protocol (used, for example, by Microsoft Outlook® 2002) is not supported by front-end and back-end architecture. MAPI clients have built-in support for handling situations where mailboxes are moved from one server to another or where content is not available on a server.

1

How Front-End and Back-End Topology Works

Although the general functionality of the front-end server is to proxy requests to the correct back-end servers on behalf of the client computers, the exact functionality of the front-end server depends on the protocol and the action being performed.

This chapter discusses the Windows and Exchange components that are essential to understanding how front-end and back-end topology works. Some of these components were modified in Exchange Service Pack 2 (SP2) to better support front-end and back-end topology. It is important that you understand how these components function in a front-end and back-end topology and assess whether the modifications will impact your organization.

This chapter also explains how front-end and back-end servers support the various client protocols.

Integration with Internet Information Services (IIS)

Exchange stores configuration information in Active Directory, whereas IIS stores configuration information in the metabase. The metabase is a local configuration database shared by the protocols that IIS supports. The Exchange System Attendant service replicates relevant configuration changes made in Active Directory through Exchange System Manager to the metabase at regular intervals. You can tell when the configuration replication has happened by looking for entries in Event Viewer from the metabase update service (MSEExchangeMU). To view MSEExchangeMU events, in the server's properties, on the **Diagnostics Logging** tab, set the MSEExchangeMU logging level to **Minimum** or greater.

Note If you have a perimeter network, and if allowing remote procedure calls (RPCs) across the internal firewall is an issue for your corporation, be aware that, by configuring your front-end server to authenticate requests, Internet Information Services (IIS) uses remote procedure calls (RPCs) to access Active Directory.

Dependency on DSAccess

DSAccess is a shared Microsoft Exchange 2000 Server component that accesses and stores directory information in cache. DSAccess dynamically detects the directory servers that other Exchange components should contact, based on criteria such as Active Directory site configuration and Active Directory server availability. Exchange 2000 front-end servers use DSAccess to determine which server contains a particular user's mailbox, the SMTP addresses that exist for a user object, the servers that contain public folder stores, and so on.

In Exchange 2000 SP1 and earlier versions, DSAccess used RPCs to connect to directory servers and discover the topology. In Exchange SP2 and later versions, DSAccess uses LDAP for most operations. However, DSAccess still uses RPCs to call the NetLogon service for each domain controller and global catalog server that it discovers.

DSAccess in Perimeter Networks

If you have a perimeter network in which you cannot allow RPC traffic between the perimeter network and the corporate network, the NetLogon RPC from DSAccess to domain controller and global catalog servers will fail. If this happens, DSAccess determines that RPC connectivity is simply blocked, and that the servers are still available. However, DSAccess continues to send the NetLogon RPC, which may affect performance.

To stop DSAccess from performing the NetLogon RPC check, you can create a registry key. For more information about optimizing DSAccess in a perimeter network, see “Configuring DSAccess for Perimeter Networks” in Chapter 4.

Note It is recommended that you use an advanced firewall server (such as ISA Server) rather than the front-end server in the perimeter network. For more information, see “Advanced Firewall in a Perimeter Network” in Chapter 3.

System Attendant on Front-End Servers

In Exchange SP2 and later versions, Exchange System Attendant no longer requires RPCs when it runs on a front-end server. Specifically, the components of System Attendant that use RPCs are no longer loaded on front-end servers; therefore, these components are disabled when you designate a server as a front-end server. The following list briefly describes these components:

DSProxy

The DSProxy service refers MAPI clients (such as Outlook 2002) to global catalog servers for global address list lookups. DSProxy also allows MAPI clients with older versions of Outlook to access Active Directory. DSProxy no longer runs on front-end servers; therefore, the front-end server can no longer determine which back-end server contains a MAPI client’s mailbox. As a result, you cannot point a MAPI client to the front-end server for the purpose of determining the user’s back-end server and routing the request to the appropriate server.

Note If you want to enable DSProxy on the front-end server for routing MAPI client requests, install Exchange Service Pack 3 (SP3) and create the registry key described in Microsoft Knowledge Base article Q319175, “XADM: You Cannot Perform a Check Names Query Against a Front-End Exchange Computer” (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=319175>). Note that in order to receive these referrals, the client must have RPC access to the front-end server. In addition, the front-end server must have RPC access to domain controllers.

Recipient Update Service

The Recipient Update Service is responsible for updating recipients in the directory to match address lists or recipient proxy policies. The Recipient Update Service no longer runs on front-end servers. Ensure that none of your front-end servers are designated to run the Recipient Update Service. To do this, in Exchange System Manager, under **Recipients**, check the properties of each Recipient Update Service and ensure that no front-end servers are named in the **Exchange server** field.

Offline Address Book Generation (OABGen)

OABGen creates the off-line address book. Without the OABGen service, front-end servers no longer generate offline address books.

Group Polling

System Attendant uses group polling to ensure that the local computer remains a member of the Domain Exchange Servers group. System Attendant polls the Domain Exchange Servers group and adds the local computer back to the group if it is no longer a member. Front-end servers no longer perform this function.

Mailbox Management

The Mailbox Management service starts and stops the mailbox cleanup process according to the settings defined in Recipient Polices. Mailbox Management no longer runs on front-end servers.

Free/Busy (madfb.dll)

The free/busy service manages user schedules. This service no longer runs on front-end servers.

Supporting POP and IMAP Clients

When you use a front-end server, the names of the servers that host the mailboxes are hidden from the users. Client computers connect to one host name shared by the front-end servers. As a result, moving users between servers is transparent to the users and requires no reconfiguration of client computers.

To log on, a POP or IMAP client sends the front-end server a logon request that contains the name of the mailbox to be accessed. Unlike HTTP, the front-end server does not authenticate the user; it uses Active Directory to determine which back-end server contains the user's mailbox. The front-end server then proxies the logon request to the appropriate back-end server, where authentication is performed. The back-end server then sends the results of the logon operation back to the front-end server, which returns the results of the operation back to the client. Subsequent POP or IMAP commands are handled in a similar fashion.

SMTP must be available on the front-end server to allow POP and IMAP clients to submit e-mail. You can install SMTP on the front-end server or set up a separate SMTP server. E-mail submission through SMTP on the front-end server works the same way as it does on any other server running Exchange. For more information about how to configure SMTP on a front-end server, see Chapter 4, "Configuring a Front-End Server."

Authentication for POP and IMAP Clients

Unless you configure authentication, POP and IMAP e-mail clients send user and password information in clear text. If the front-end server is accessible from the Internet, you should configure SSL so that user authentication information and data is not passed over the Internet in clear text. You can also use an authentication method provided by the client e-mail application if one is available. For example, Outlook Express contains an option for using NT LAN Manager (NTLM) authentication.

IMAP Access to Public Folders

When a non referral-enabled IMAP client connects to a back-end server, it can access only public folders that have a replica on the user's home server. To access public folders that have replicas on other servers, an IMAP client must be referral-enabled. A referral-enabled client issues special commands to an IMAP server to create a list of the public folders available to the client. When the client computer requests a public folder that does not have a local replica, the server responds to the client request with a referral URL that contains the name of the server that has the public folder. The referral-enabled IMAP client computer then creates a new connection to that server to retrieve the appropriate information.

In a front-end and back-end topology, however, the front-end server acts as a referral-enabled client, so IMAP clients connecting to the front-end server do not need to support referrals; the front-end server handles referrals for them. It transparently maps non referral-enabled client requests to their referral counterparts, making the entire list of public folders available to a non referral-enabled client. When the front-end server receives a referral response from the back-end server, it does not pass this response back to the client. Instead it follows the referral for the client and makes a connection to the appropriate back-end server that has the data. The back-end server then responds with the requested item, which the front-end server then relays back to the client.

Note In Exchange 2000 SP1 and earlier versions, by default, IMAP4Svc and POP3Svc were dependent upon MExchangeIS. If your configuration did not allow RPC traffic across the internal firewall (between front-end servers and back-end servers), you were required to remove the IMAP and POP3 dependency on MExchangeIS so that you could stop the MExchangeIS and MExchangeSA services on the front-end server.

However, in Exchange 2000 SP2, IMAP and POP3 are no longer dependent upon the MExchangeIS service. Therefore, you do not need to remove this dependency.

Running SMTP for POP and IMAP Clients

POP and IMAP protocols are used only for receiving mail; you must configure SMTP on the front-end server so that POP and IMAP clients can submit mail.

Important To run SMTP on the front-end server and enable it to accept inbound mail (mail for your domains), you must mount a mailbox store on the front-end server. This mailbox store must not contain any mailboxes. You must mount a mailbox store on the front-end server because any non-delivery reports (NDRs) must be routed through the mailbox store for formatting.

To configure SMTP so that POP and IMAP clients can submit mail to external domains, you need to allow relaying on the front-end server.

By default, Exchange allows relaying only from authenticated clients. It is recommended that you keep this default. Clients such as Microsoft Outlook Express 5.0 and Microsoft Outlook 98, Microsoft Outlook 2000, and Outlook 2002 support SMTP authentication as well as Transport Layer Security (TLS) encryption.

You should not allow relaying in either of the following ways:

- You should not allow anonymous relaying to all IP addresses; if your front-end server is connected to the Internet, doing this allows anyone on the Internet to use your server to send mail.
- You should not allow relaying from specific client IP addresses. Even if you are familiar with the subnet from which clients send mail, the Internet environment makes it difficult to determine such a specific set of IP addresses.

Note If you want the front-end server to act as the bridgehead server between your company and the Internet, it is recommended that the server on the Internet that accepts mail for your domains has the ability to scan incoming messages for viruses.

Supporting HTTP (Outlook Web Access and Web Folders)

Whether generated by a browser or a specialized client, such as Windows 2000 Web Folders, HTTP requests from the client computer are sent to the front-end server. The front-end server uses Active Directory to determine which back-end server to proxy the request to.

After determining the appropriate back-end server, the front-end server forwards the request to the back-end server. Apart from specific header information that indicates the request was passed through a front-end server, the request is almost identical to the original request sent from the client. In particular, the HTTP host header, which matches the name of the front-end server to which the request was sent (meaning the hostname or fully qualified domain name that the user entered in the browser), remains unchanged. The front-end server contacts the back-end server using the hostname of the back-end server (for example, backend1), but in the HTTP headers of the request, the front-end server sends the host header used by the client, for example, www.adata.com. The host header setting ensures that the appropriate back-end Exchange virtual server handles the request. The virtual servers on the back-end must be configured to handle front-end server requests. For more information about configuring virtual servers on a back-end server, see Chapter 5, “Configuring a Back-End Server.”

For HTTP requests, the front-end server always contacts the back-end server over TCP port 80 (the default HTTP port), regardless of whether the client contacted the front-end server through port 80 or 443 (the SSL port). This means that:

- SSL encryption is never used between the front-end and back-end servers, although the client might use it in communication with the front-end server.
- HTTP virtual servers that differentiate themselves from other servers only by port number are not supported in a front-end and back-end topology. For example, if a back-end server has an HTTP virtual server listening on port 8080, a client can access that back-end server only if the client is pointed directly to the back-end server (for example, http://backend1:8080/data). A client connecting to the front-end server is not able to access this data.

The back-end server processes the HTTP request from the front-end normally, and the response is sent unchanged through the front-end server back to the client. In most cases, the back-end server handles the front-end server as if it were another HTTP client. The client, therefore, never needs to know that the request was not handled on the front-end server.

Finding User Mailboxes

To provide access to mailbox folders through HTTP, you must have a virtual directory on both the front-end and back-end servers that points to the mailboxes.

Note User mailboxes cannot be stored on the front-end server.

When you install Exchange, a virtual directory named “Exchange” is created in the default virtual server. This directory points to the default SMTP domain for the Exchange organization. When you configure the virtual directory on the front-end server through Exchange System Manager, you can select the SMTP domain name. Users connecting to that virtual server must have an e-mail address in their object in Active Directory with the same domain. In the dialog box in which the SMTP domain is selected, the list of domains is a list of all domains for which there are recipient policies. As a result, you might see duplicates in the list; it does not matter which one you choose.

When the front-end server detects a request to a location within the mailbox store (based on the setting of the virtual server or directory), it contacts an Active Directory global catalog server in the domain using LDAP and determines which back-end server contains the user’s mailbox.

Logging on to Outlook Web Access

Users can log on to Outlook Web Access through explicit logon or implicit logon.

Explicit logon

There are a few URLs that users can use to connect to Outlook Web Access. The usual URL is `http://<server>/exchange/<username>/`. Accessing Outlook Web Access using this URL is referred to as explicit logon.

Explicit logon must be used when the front-end server is not configured to authenticate users (for more information about authentication, see “Authentication Issues for HTTP” later in this chapter) or when a user is attempting to access a mailbox that is not his own but to which he has access.

When the front-end server receives an explicit logon request from a client, the user name is extracted from the URL and combined with the SMTP domain name associated with the virtual directory or virtual server to construct a fully qualified SMTP address. The front-end server looks up this address in Active Directory and determines which back-end server has the mailbox associated with the address. The front-end server then forwards the request to that back-end server, which processes the request as if it came directly from the client. The back-end server returns the response it generates to the front-end server, which returns it unchanged to the client.

Implicit logon

If the front-end server is configured to authenticate users, then users can access their mailboxes by omitting the username from the request, and pointing their browser to `http://<server>/exchange/`. After authenticating the user, the authentication information is used to look up the mailbox associated with the user in Active Directory and the back-end server on which the mailbox is located. The URL is updated with the user name and sent to the correct back-end server. This is known as implicit logon. Implicit logon is useful only for logging on to Outlook Web Access; specialized HTTP clients generally do not use implicit logon.

Allowing Use of Outlook Web Access to Internal Clients Only

You can deny access to Outlook Web Access to external clients who access your front-end server over the Internet while still allowing access to users within the corporate network. As mentioned earlier, users connecting to a virtual server must have an e-mail address in their user object in Active Directory with the same SMTP domain as the virtual server.

To allow only internal users access to Outlook Web Access, create a recipient policy with an SMTP domain name and apply the recipient policy to the user accounts for which you want to allow access. (Outlook Web Access users do not need to know this SMTP domain.) Then create a new virtual server on the front-end server that specifies the domain you used in the recipient policy. Users without an email address from that domain will not be able to log onto the virtual server. As long as you do not use this SMTP domain as the default domain, it will not appear in the **From** field when users send e-mail outside the organization. For more information, see Microsoft Knowledge Base article Q257891, “XWEB: ‘The Page Could Not Be Found’ Error Message When You Use OWA” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=257891>). You can also prevent specific users from accessing Outlook Web Access by disallowing HTTP and NNTP protocols for those users. To change a user’s protocol settings, in Active Directory Users and Computers, use the **Exchange Advanced** tab in a user’s properties.

Simplifying the Outlook Web Access URL

Users commonly request that a simpler URL be made available for accessing their mailbox. The following procedure provides a simple method for reducing the URL.

This procedure configures a request sent to the root of the Web server (<http://server/>) to redirect to the Exchange virtual directory. For example, a request to <http://mail/> is directed to <http://mail/exchange/>, which then triggers implicit logon.

► **To reduce the Outlook Web Access URL**

1. Using the Internet Services Manager, open the properties for the **Default Web Site**.
2. Click the **Home Directory** tab, and then select **A redirection to a URL**.
3. In **Redirect to**, type */<directory name>*, and then click **A directory below this one**. For example, if you want to redirect <http://mail/> requests to <http://mail/exchange/>, in **Redirect to**, you would type */exchange*.
4. If you want your users to use SSL to access their server, you can redirect client requests to <https://mail/<directory name>/>. To require users to use SSL, in **Redirect to**, type <https://mail/<directory name>>, and then click **The exact URL above**. This setting hard codes the name of the server; therefore if you redirect client requests to <https://mail/>, the client must be able to resolve the name **mail**. For information about another method for redirecting clients to SSL, see Microsoft Knowledge Base article Q279681, "How to Force SSL Encryption for an Outlook Web Access 2000 Client" (<http://go.microsoft.com/fwlink/?LinkId=3052&ID=279681>).

Note The procedure works only when the front-end server has authentication enabled (when implicit logon is possible). Users still must enter the full URL, including username, to access other mailboxes or content in other folders.

Enabling the "Change Password" Feature

If you are using Outlook Web Access, you can enable the Change Password feature in IIS to:

- Alert users when their passwords expire.
- Allow users to use the **Options** button in Outlook Web Access to change their passwords.

Keep in mind that if you want to use the Change Password feature, you must also use SSL between clients and the front-end server to secure the password during transmission. In addition, you must create a virtual directory named IISAdmPwd on the front-end server and back-end servers to handle the Change Password requests.

Note The only time you must require SSL on a back-end server is when you want users to be able to connect to the back-end server directly. Remember, however, that front-end servers cannot use SSL when connecting to back-end servers. Therefore, if you require SSL on the back-end server, ensure that you do not require SSL on the following directories so that front-end servers can still connect to them: Exchange, Public, ExchWeb, Exadmin, and any mailbox or public folder virtual roots.

For more information about how to configure the Change Password feature, see Microsoft Knowledge Base article Q327134, “XCCC: How the Change Password Feature Works in Outlook Web Access” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=327134>).

For more information about how to configure SSL, see “Securing Communication: Client to Front-End Server” in Chapter 2.

Finding Public Folders

Just as you must configure virtual directories for mailboxes, you must also configure virtual directories for each of the public folder trees that are to be accessible over HTTP through the front-end server.

When you install Exchange, a virtual directory called “public” is created under the default Exchange HTTP virtual server to allow access to the default (MAPI-accessible) public folder tree. When you create other public folder trees (for example, to host applications), you must also create virtual directories in Exchange System Manager to expose these trees over HTTP. Identical virtual directories must exist on each front-end server and on all back-end servers that host the public folder tree.

A request made to a URL in a public folder tree is handled differently when accessing the default (or MAPI-accessible) public folder tree than when accessing general-purpose public folders (also known as application Top-Level Hierarchies [TLH], or non-MAPI TLHs).

In both cases, however, the goal for accessing public folders is twofold:

Availability

If data exists in an Exchange 2000 public folder somewhere in the Exchange organization and is accessible over HTTP, it is available to the user.

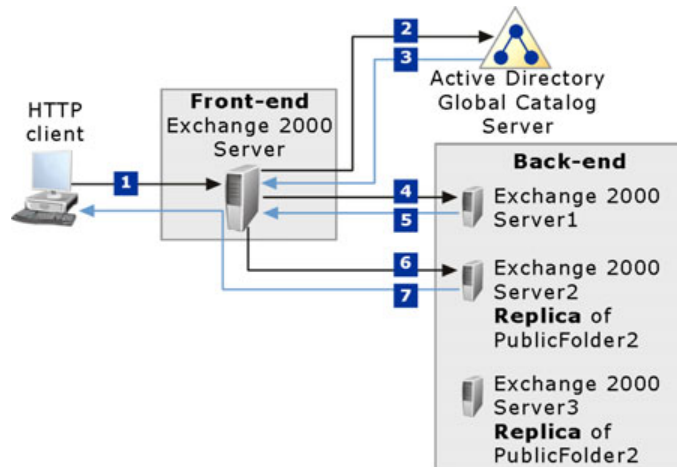
Consistency

As long as the server is available and the user is authenticated, the same public folder server services every request from that user. For authenticated users, ensuring that they go to the same public folder server means that they see the same data each time they access the public folder trees through the front-end server (including status of read and unread messages, which is stored on individual servers and is not replicated between public folder servers). The fact that users always reach the same back-end server is also important for server-based applications that maintain session state, such as some built with Active Server Pages (ASP).

Public Folder Referrals

In Exchange, you can configure public folder replication on a per-folder basis. The actual public folder tree hierarchy is available on all Exchange servers in the organization, but the contents of each folder may not. This information is not stored in Active Directory but is maintained as a property on each folder in the public folder store. Therefore, special handling is required when the back-end server selected by the front-end server does not contain the contents of the folder requested by the client.

Figure 3 illustrates how public folder referrals travel through a front-end server.



1. An HTTP client authenticates against the front-end server and requests /public/PublicFolder2.
2. The front-end server authenticates the user against Active Directory and requests the location of the user's default public folder store.
3. Active Directory tells the front-end server that the user's default public folder store is on Server1.
4. The front-end server sends the client request to Server1.
5. Server1 tells the front-end server that it does not have the contents of /public/PublicFolder2, but Server2 and Server3 do.
6. The front-end server performs a hashing algorithm against the list of servers with the content (in this case, Server2 and Server3). The results of the hash in this case turn out to be Server2, so the front-end server forwards the request to Server2.

Note A hashing algorithm applies a given number (in this case, the user's security token) and uses it to generate a position in a list so that the distribution of all possible inputs is even over the list.

7. Server2 returns the contents of /public/PublicFolder2 to the front-end server, which then sends the contents to the HTTP client.

Figure 3 Public folder referral through a front-end server

The Default (MAPI) Public Folder Tree

When a client accesses the default public folder tree in Outlook Web Access, an attempt is made to maintain parity with MAPI clients such as Outlook 2002. Each mailbox store is associated with a particular public folder store somewhere in the organization (sometimes on the same server as the mailbox store, sometimes on a dedicated public folder server). The public folder store associated with the user's mailbox store is the public folder store that displays the public folder hierarchy (tree) in Outlook 2002.

When a user requests a public folder in the default public folder tree through HTTP, the front-end server authenticates the user and looks up the user in Active Directory to see which public store is associated with that user's mailbox store. The front-end server then forwards the request to the user's public folder server.

If an authenticated user does not have a mailbox associated with it (for example, when users with accounts in a Microsoft Windows NT® 4.0 domain have access to Exchange 2000 mailboxes), the front-end server treats accesses to the default public folder tree the same way as it treats access to the general-purpose public folder trees. Also note that if the front-end server is not configured to authenticate users, requests for public folders are not load balanced.

General-Purpose Public Folder Trees

Default public folder tree servers have an association with mailbox stores because of their MAPI heritage; general-purpose public folder trees do not have such an association. As a result, requests for folders in general-purpose public folder trees are handled slightly differently than requests for folders in the default public folder tree.

When a client makes a request to access a general-purpose public folder tree, the front-end server first contacts Active Directory to find a list of all servers running Exchange 2000 in the organization that have a replica of the particular general-purpose public folder tree that the client is attempting to access. This list is filtered to remove servers running Exchange 5.5, because Exchange 5.5 does not support the HTTP extensions (Web Distributed Authoring and Versioning) that Exchange 2000 supports.

The front-end server then uses the user's authentication token in a hashing algorithm against the list of servers to ensure that:

- Users are load-balanced across the available servers.
- Individual user requests are always processed by the same back-end server, regardless of the HTTP client used.

Note If you add or remove a back-end server, the output from the hashing algorithm changes, and users may be redirected to a different server from then on. For more information, see “Adding or Removing Back-End Servers” later in this chapter.

When Content Is Not Available on the Back-End Server

The front-end and back-end topology has special handling for times when the back-end server receives a request for a public folder for which it does not have a replica. This handling happens for folders in the default public folder store as well as folders in general-purpose public folder trees.

When a back-end server receives such a request, it returns a list of the servers that have the contents of the requested folder. This is the only case in which the back-end server does not process requests from a front-end server in the same way it processes requests directly from clients. The front-end server does not pass this information back to the client, but runs the same hashing algorithm against the new list of servers again, to ensure load balancing. As a result, in organizations that use partial replicas of public folder trees, the front-end server may have to perform two HTTP requests to satisfy the client’s single request. However, in processing the client’s request, the front-end server caches information about which servers have the content, allowing the front-end server to avoid extra requests when data in the same folder is accessed in the future.

The caches maintained by the front-end server substantially reduce the number of queries sent to Active Directory and back-end servers for both public and private folder accesses. Cache information expires after ten minutes and is also reset when changes in server configuration are detected.

Back-End Server Downtime

If a back-end server is down for maintenance or is otherwise inaccessible over HTTP, the front-end server cannot connect to it. The front-end server marks that server “unavailable” for a period of 10 minutes and sends the request to a different server if there are other servers available; the request fails if no other servers are available. While the back-end server is unavailable, the front-end server automatically directs requests to other servers. Therefore, after a back-end server returns to production, it might be inaccessible through the front-end server for as long as 10 minutes, because the front-end server might still have that back-end server marked as unavailable.

Adding or Removing Back-End Servers

The goal of the hashing algorithm is load balancing; however, a condition of the algorithm is that the distribution of users across servers is dependent on the number of servers.

Therefore, if the list of servers hosting the content for a public folder changes because of the addition or removal of a server, the result of the hashing algorithm may direct the user to a new server for future requests. Typically, when the server processing a user's request changes the user cannot tell that anything physical changed, with the exception of the following:

- Users may observe that the read or unread state of messages is reset.
- Users of Web-based applications (running in general-purpose public folder trees) that maintain session state may need to restart their application session if they use the application during the transition period.

Therefore, it is recommended that administrators inform their users before adding or removing public folder servers.

At a high level, the hashing algorithm works as follows: Two back-end servers, numbered 1 and 2, with a particular public folder tree are deployed. Then, if six different users—A, B, C, D, E, and F—try to access data, the front-end server distributes their requests over the two servers as follows:

- Users A, B, and C get their data from server 1.
- Users D, E, and F get their data from server 2.

Note This load balancing is done transparently—users do not know which back-end server is actually handling the request.

Then another server is added, server 3. Now the users are distributed as follows:

- Users A and B get the data from server 1.
- Users C and D get the data from server 2.
- Users E and F get the data from server 3.

In this example, users A, B, and D did not change servers, but users C, E, and F did.

Authentication Issues for HTTP

The front-end server handles authentication in two ways: either the front-end server authenticates the user itself, or it forwards the request anonymously to the back-end server. Either way, the back-end server also performs authentication.

It is recommended that you use dual authentication, in which you configure both front-end and back-end servers to authenticate users, particularly when the front-end server is exposed to the Internet. For more information about dual authentication, see “Dual Authentication” later in this chapter.

If the front-end server is in a locked-down perimeter network in which RPCs to back-end servers are blocked, and the front-end server cannot authenticate users, you would have to use pass-through authentication. However, it is important to understand the risks of using pass-through authentication. For more information about pass-through authentication, see “Pass-Through Authentication” later in this Chapter.

Exchange HTTP front-end servers support only HTTP 1.1 basic authentication between client computers and front-end servers, as well as between front-end and back-end servers. Basic authentication is a simple authentication mechanism defined by the HTTP specification that lightly encodes the user’s user name and password before sending it to the server. To achieve real password security in a front-end and back-end topology, you should use SSL encryption between the client and the front-end server.

Note Front-end servers do not support integrated Windows authentication (which includes both NTLM and Kerberos authentication) or HTTP 1.1 Digest authentication.

Basic authentication does not support single sign on. Single sign on is when a user logs on to a computer running Windows, the user authenticates against a domain, and then the user can access all resources and applications in the domain without re-entering his or her credentials. Microsoft Internet Explorer versions 4.0 and later allow single sign on for Web applications, including Outlook Web Access, if the server being accessed has Integrated Windows authentication enabled. Because front-end servers do not support Integrated Windows authentication, when users access HTTP applications, the front-end server always prompts them for authentication and they must re-enter their credentials, even if they already used Windows to log on. Users only have to enter credentials once per browser session, however, because their credentials are cached in the browser process.

Important When using a kiosk, be aware that caching credentials can pose a security risk if you cannot close the browser and end the browser process between sessions. This risk occurs because a user's credentials remain in the cache when the next user accesses the kiosk. If you want to enable the use of Outlook Web Access on a kiosk, ensure that you can close the browser between sessions and end browser processes. Otherwise, consider using a third-party product that incorporates two-factor authentication, in which the user must present a physical token along with a password to use Outlook Web Access on the kiosk.

Dual Authentication

In dual authentication, both front-end and back-end servers are configured to authenticate users with HTTP basic authentication. Basic authentication is the only form of authentication that front-end servers support. Back-end servers support both basic authentication and Integrated Windows Authentication. To enable connections through a front-end server, basic authentication must be enabled on both the front-end server and back-end servers.

HTTP basic authentication carries the authentication information in every request from the client to the front-end server; the front-end server then passes the information to the back-end server. The authentication information in the request is sufficient for both servers; therefore, after the front-end server requests authentication information from the user, the back-end server receives the same information and does not need to request it from the user again.

You should configure front-end servers to perform authentication whenever possible. If you cannot enable authentication on the front-end server, implicit log on does not work, and you cannot load balance public folder requests. You can use explicit logon can to gain access, regardless of how authentication is configured.

Note Exchange relies on IIS to perform the authentication for HTTP requests. IIS uses RPCs to directory servers to perform authentication. If RPCs are not allowed between the front-end server and the directory server, you must use pass-through authentication. For more information about how to enable pass-through authentication and the risks of doing so, see "Pass-Through Authentication" later in this chapter.

Pass-Through Authentication

In pass-through authentication, the front-end server is configured with anonymous authentication, so it does not ask the user for an authorization header. The front-end server forwards the user's request to the back-end server, which asks the user for authentication. The back-end server's request for authentication and the user's response are routed unchanged through the front-end server.

Warning When you use pass-through authentication, anonymous HTTP requests go directly to the back-end server where they are authenticated. You should use pass-through authentication only if the front-end server cannot authenticate—for example, in a locked down perimeter network where RPCs are not allowed across the intranet firewall. Because pass-through authentication allows requests from any source, valid or invalid, to be passed to your back-end servers, you may want to re-evaluate policies that restrict RPC traffic across the intranet firewall. It may be more secure to allow RPCs than to allow anonymous requests to reach back-end servers. For more information, see Chapter 3, “Scenarios.”

When pass-through authentication is used, the front-end server is not able to load balance public folder requests, because it does not have the authentication token on which to perform a hashing algorithm. Additionally, implicit logon will not work. Users have to enter the full URL including their user name to log on.

User Logon Information

When authenticating against a front-end server, by default, the user must enter his or her user name in the following format: *domain\username*. You can configure the front-end server to assume a default domain so that users do not need to remember their domain.

An additional option for authentication is to configure a user principal name (UPN) for users. This allows users to enter their e-mail address as their user name. For more information, see Chapter 4, “Configuring a Front-End Server.”

Note If you want to set a default domain or configure a UPN, you should upgrade to Exchange 2000 SP1 or later. For more information, see Microsoft Knowledge Base article Q267936, “XIMS: Directory Service to Metabase Service May Not Replicate the Default Logon Domain for Virtual Servers” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=267936>).

Remote Procedure Calls (RPCs) in Front-End and Back-End Topology

Exchange 2000 SP1 and earlier versions used RPCs to perform Active Directory service discovery (which is performed by DSAccess) and to authenticate users (which is performed by IIS). In Exchange 2000 SP2, DSAccess no longer uses RPCs to perform Active Directory service discovery. However, IIS still uses RPCs to authenticate requests on the front-end server. Therefore, if you enable basic authentication on the front-end server, you must open certain RPC ports. For more information about opening RPC ports on the intranet firewall, see “Configuring an Intranet Firewall” in Chapter 6.

Stopping RPC Traffic Across the Intranet Firewall

Corporations that have perimeter networks often restrict the type of traffic that passes from the perimeter network into the corporate intranet.

Without RPC access to Active Directory servers, the front-end server cannot authenticate clients. Therefore, features that require authentication on the front-end server (such as implicit logon and public folder tree load balancing) will not work. Public folder access is possible, but the front-end server cannot load balance the requests because the front-end server cannot determine the identity of the user. Without the user's authentication token, the front-end server cannot perform the load balancing hashing algorithm. As a result, all anonymous requests for a public folder are routed to the same back-end server.

Note If you do not allow RPC traffic across the internal firewall, you cannot support IMAP and POP clients because these protocols require that SMTP run on the front-end server for sending e-mail. When RPC traffic is blocked, you cannot run MExchangeIS or MExchangeSA on the front-end server, and as a result, you cannot run SMTP on the front-end server.

If open RPC ports are not allowed between the perimeter network and the corporate intranet, you must use pass-through authentication. With pass-through authentication, the front-end server passes requests to the back-end anonymously, and then the back-end server performs the authentication.

Note It is recommended that you use an advanced firewall server (such as ISA Server) rather than the front-end server in the perimeter network. For more information, see "Advanced Firewall in a Perimeter Network" in Chapter 3.

2

Deployment Considerations

When deploying a front-end and back-end topology, you must take into account several factors: expected load, hardware needs, administrative overhead, load balancing and security to name a few. The following sections goes over these considerations in more detail.

Recommended Server Configurations and Ratios

Server configuration depends on many factors, including the number of users for each back-end server, the protocols used, and the expected load on the system. The configuration of particular models of servers should be done in consultation with a hardware vendor or consultant. For more information about server sizing, see the capacity and topology calculator at (<http://go.microsoft.com/fwlink/?Linkid=1716>).

As a general rule, one front-end server is reasonable for every four back-end servers. However, this number is provided only as a suggested ratio, not as a rule. Use this information only as a starting point. Front-end servers do not need large or particularly fast disk storage, but should have fast CPUs and a large amount of memory. There is no need to back up the disks on the front-end server unless you choose to enable SMTP, because SMTP commits queued mail to the local disk. For POP, IMAP, and HTTP, no user data is stored on the drive.

For more information about hardware requirements for front-end and back-end servers, see the following technical papers:

- *Microsoft Exchange 2000 Front-End Server and SMTP Gateway Hardware Scalability Guide*
(<http://go.microsoft.com/fwlink/?Linkid=1713>)
- *Microsoft Exchange 2000 Server Back-End Mailbox Scalability*
(<http://go.microsoft.com/fwlink/?Linkid=1711>)

Load Balancing

In a corporate or hosting environment, you may want to load balance the front-end servers. Windows 2000 provides load balancing through Network Load Balancing (NLB). Other load-balancing mechanisms include domain name service (DNS) round robin and a hardware load-balancing solution.

Windows NLB works by clustering two or more servers that represent the cluster with a single IP address. Each computer receives traffic to its own unique IP address and the shared IP address. Each member of the cluster performs a hashing algorithm to map incoming clients to one of the members of the cluster based on the client IP address, port, and other information. When a packet arrives, all servers or hosts perform the same hashing algorithm, and the output is one of the hosts. That host then responds to the packet. The mapping does not change unless the number of hosts in the NLB cluster changes. The configuration of every server in the NLB cluster must be same, otherwise clients may experience different behavior depending on which server they are routed to.

Note NLB has no health monitoring; if the World Wide Web Publishing Service on a front-end server is not running, for example, NLB continues to send requests to that server. You can run Microsoft Application Center 2000 on a front-end server to set up NLB and monitor the health of load-balanced servers. (However, you cannot manage Exchange resources or replicate Exchange configuration information through Application Center.) For more information about Application Center, see the Microsoft Application Center Web site (<http://go.microsoft.com/fwlink/?Linkid=591>).

Although it is not a requirement, you should ensure that each user is always sent to the same front-end server for the duration of a session. This makes use of the caching and connection state information already maintained on the front-end server. In NLB, this is referred to as “client affinity.” Many hardware solutions also have this ability.

Reducing Virtual Server Creation

In some circumstances, it might be important to reduce the number of virtual servers created on the back-end servers. You should not reduce the number of virtual servers unless you fully understand how HTTP virtual servers work. You can reduce virtual server creation by either of two methods.

Make an analysis of the users and data on each back-end server to determine if users will ever be directed to that particular server. If a back-end server contains mailboxes for only adatum.com, then there is no need for that back-end server to have a virtual server for contoso.com. If users from contoso.com are later added to that back-end server, however, then an administrator will need to create a virtual server for contoso.com.

Similarly, you only need to create virtual directories for resources your users will need access to. On a server that has no public store, the public virtual directory is not necessary.

Using Firewalls

If your network is exposed to the Internet, it is highly recommended that you use either a software or hardware firewall solution. Firewalls control traffic to the network by using such methods as port filtering, IP filtering, and, in advanced firewall solutions, application filtering.

There are several options for incorporating a firewall into a front-end and back-end topology; The “Scenarios” section in this document describes these various options. In general, it is recommended that you use an advanced firewall server in your topology (for more information about using an advanced firewall, see “Advanced Firewall in the Perimeter Network” in Chapter 3).

Port Filtering

At a minimum, any firewall you use to protect servers from the Internet must use port filtering. Port filtering restricts the type of network traffic that comes through the firewall by allowing access only to information sent to specific ports. For example, you may configure the firewall facing the Internet to accept only HTTPS traffic by opening TCP port 443.

The following two sections describe two important concepts related to TCP connections: source port versus destination port, and direction of the TCP connection.

Source Port Versus Destination Port

When computer A opens a TCP connection to computer B, two ports are used: the source port (on computer A), and the destination port (on computer B). The network stack on the computer that initiates the connection usually selects source ports at random. Destination ports are the ports on which the specified service is listening (for example, port 443 for HTTPS). In this document, any reference to a port used by a specific service refers to the destination port.

Direction of the TCP Connection

When you open firewall ports, most firewalls require you to specify the direction of the connection. For example, in order to allow a front-end server to contact back-end servers, you must open port 80 for HTTP traffic. However, back-end servers never initiate new TCP connections to the front-end server; they only respond to requests that were initiated by the front-end. Therefore, on your firewall, you only need to allow HTTP port 80 connections from the front-end to the back-end. In this document, such connections are referred to as “inbound” (in other words, the connections are inbound to the corporate network).

IP Filtering

Many firewall solutions also support IP filtering. IP filtering improves the reliability of the firewall by allowing you to restrict traffic through the firewall to specific servers. For example, in a perimeter network, you may want to configure DSAccess to use specific domain controllers and global catalog servers, and then use IP filtering to ensure that the front-end servers connect to only those domain controllers and global catalog servers.

Application Filtering

Advanced firewalls such as ISA Server can provide advanced inspection at the application protocol level. This inspection allows the firewall to perform such functions as filtering RPC interfaces and validating HTTP request syntax. Application filtering is the main reason why using an advanced firewall in your topology provides the most security.

Securing Communication: Client to Front-End Server

To secure data transmitted between the client and the front-end server, it is highly recommended that the front-end server be SSL-enabled. In addition, to ensure that user data is always secure, access to the front-end server without SSL should be disabled (this is an option in the SSL configuration). When using basic authentication, it is critical to protect the network traffic by using SSL to protect user passwords from network packet sniffing.

Warning If you do not use SSL between clients and the front-end server, HTTP data transmission to your front-end server will not be secure. It is highly recommended that you configure the front-end server to require SSL.

It is recommended that you obtain an SSL certificate by purchasing a certificate from a number of third-party certification authorities. Purchasing a certificate from a certification authority is the preferred method because the majority of browsers already trust many of these certification authorities.

Alternately, you can use Microsoft Certificate Server to install your own certification authorities. Although installing your own certificate authority may be less expensive, browsers will not trust your certificate, and users will receive a warning message indicating that the certificate is not trusted.

For more information, see Microsoft Knowledge Base article Q320291, “XCCC: Turning On SSL for Exchange 2000 Server Outlook Web Access” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=320291>).

Configuring SSL in a Front-End and Back-End Topology

You do not need to configure SSL on back-end servers when using a front-end server, because the front-end server does not support using SSL to communicate with back-end servers. You can configure SSL on the back-end servers for use by clients that are directly accessing them.

Back-end servers sometimes need to generate absolute URLs, such as a list of URLs for the messages in a user’s inbox. If SSL is used between the client and the front-end server, the back-end server needs to know this, so it can formulate URLs using HTTPS instead of HTTP. If the SSL decryption is done on the front-end server, the front-end server knows SSL was used, and it notifies the back-end server of this by passing an HTTP header that says, “Front-End-Https: on” in all requests to the back-end server.

If there is a separate server between the client and the front-end server that performs the SSL decryption, that server must be able to pass the “Front-End-Https: on” header to the front-end server, which then passes it to the back-end server. As an alternative, you can configure SSL between the SSL decryption server and the front-end server. However, if you added that separate server to offload the additional traffic caused by SSL encryption and decryption, this method defeats that purpose. This method allows that separate server to filter the traffic.

Tip Ensure that the Windows 2000 License Logging Service is running on the front-end server. IIS does not allow more than ten simultaneous SSL connections unless this service is running.

SSL Accelerators

There is a decrease in performance involved in setting up and tearing down SSL connections, so you may want to investigate adding an SSL accelerator to your front-end and back-end topology.

SSL accelerators generally come in two forms:

- A card you can put on each front-end server. One example is the Compaq AXL200.
- A separate device or computer you place between the clients and the front-end servers. Examples are the F5 BigIP SSL accelerator and the Microsoft Internet Security and Acceleration Server 2000 Service Pack 1 (ISA). These accelerators support adding the “Front-End-Https: On” header for Outlook Web Access.

Note For information about configuring the ISA server for Outlook Web Access, see Microsoft Knowledge Base article Q307347, “Secure OWA Publishing Behind ISA Server May Require Custom HTTP Header” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=307347>).

Accelerator cards are generally used directly on the front-end server, and they offload the encryption and decryption overhead; this increases the throughput of each connection and decreases the amount of work the software on the server needs to accomplish.

External accelerator devices sit between the clients and the front-end servers. Traffic coming from the client is decrypted on the accelerator device and sent to the front-end server unencrypted. Likewise, traffic from the front-end server is sent to the accelerator device unencrypted, and then it is encrypted for transmission to the client.

The most important factor to consider when choosing what type of SSL accelerator to use is the number of front-end servers in your topology. If you have a small number of front-end servers, adding SSL accelerator cards to each of them is a simple, cost-effective way to offload SSL duties. Because the SSL decryption is done on the front-end server, there is no need for extra configuration of the “Front-End-Https: on” header for Outlook Web Access.

For a large number of front-end servers, the cost of additional accelerator cards and the administrative cost of storing and configuring SSL certificates on each server eventually ceases to be cost effective. In this case, a separate SSL accelerator device may be a more cost effective option for your topology because it needs to be configured only once, regardless of the number of front-end servers. These devices generally cost more than an accelerator card, so weigh the options in your own topology to determine which to use. Keep in mind that for Outlook Web Access, an external SSL device must have the ability to notify the front-end server that SSL was used with the “Front-End-Https: on” header.

Securing Communication: Front-End to Other Servers

HTTP, POP, and IMAP communication between the front-end server and any server with which the front-end server communicates (such as back-end servers, domain controllers, and global catalog servers) is not encrypted. When the front-end and back-end servers are in a trusted physical or switched network, this is not a concern. However, if front-end and back-end servers are kept in separate subnets, network traffic may pass over unsecured areas of the network. The security risk increases when there is greater physical distance between the front-end and back-end servers. In this case, it is recommended that this traffic be encrypted to protect passwords and data.

IP Security (IPSec)

Windows 2000 supports IPSec, which is an Internet standard that allows a server to encrypt any IP traffic, except traffic that uses broadcast or multicast IP addresses. Generally, you use IPSec to encrypt HTTP traffic; however, you can also use IPSec to encrypt LDAP, RPC, POP, and IMAP traffic.

With IPSec you can:

- Configure two servers running Windows 2000 to require trusted network access.
- Exchange data that is protected from modification (using a cryptographic checksum on every packet).
- Encrypt any traffic between the two servers at the IP layer.

In a front-end and back-end topology, you can use IPSec to encrypt traffic between the front-end and back-end servers that would otherwise not be encrypted.

IPSec Protocols

The method in which data is secured using IPSec depends on which protocol is used: Authentication Header (AH) or Encapsulating Security Payload (ESP). With AH, the packets are not encrypted; AH adds a checksum to the IP packet. AH guarantees that the packet came from the expected host, was not impersonated, and was not modified in transit. AH uses IP protocol 51. ESP, which uses IP protocol 50, encrypts the entire contents of the IP packet. Both forms of IPSec provide a reliable and trusted communication channel that an attacker cannot easily insert data into or interrupt.

IPSec encryption affects the performance on both the front-end and back-end servers; the precise extent to which it affects performance, however, depends on the type of encryption used.

IPSec Policy

The IPSec policy filter you configure on the back-end servers should be designed to apply only to traffic initiated by the front-end server that is sent to TCP port 80 on a remote server. You should configure IPSec on the back-end servers so that they respond appropriately when they receive a request for IPSec communication. However, the back-end servers should not require that all communication from all clients be encrypted using IPSec.

Windows 2000 has three IPSec policies installed by default. Select the “Client (respond only)” policy for the back-end server. With this policy enabled on the back-end server, the front-end server can use IPSec to communicate safely with the back-end server, while other clients (including MAPI clients like Outlook 2002) and servers can communicate with the back-end server without needing to use IPSec.

IPSec with Firewalls and Filtering Routers

When a firewall or filtering router is used between the front-end and back-end servers, the filters must allow IPSec to pass through it.

Note IPSec does not work if there is a Network Address Translation (NAT) server between the perimeter network and the corporate network.

When using IPSec, configure the ports as follows:

- **HTTP (TCP port 80)** HTTP (TCP port 80) is no longer required and should be blocked.
- **500/UDP** Open the IPSec negotiation port at 500/UDP, which is required. 500/UDP is the well-known Internet assigned port for the Internet Key Exchange (IKE) Standard, which is defined in RFC 2409. This standard uses UDP and not TCP to avoid the security weaknesses of relying on TCP connections. IKE provides strong security for the negotiation data packets. It establishes and maintains the IPSec connections, called security associations.
- **IP protocol 50 or 51** Allow either IP protocol 50 (AH) or IP protocol 51 (ESP), depending on the protocol you are using.
- **UDP port 88 and TCP port 88** Open UDP port 88 and TCP port 88 to allow Kerberos traffic. Kerberos is the core Windows 2000 security protocol typically used for IPSec authentication. Kerberos traffic uses a UDP/TCP protocol source and destination port 88. Kerberos is itself a security protocol that does not need to be secured by IPSec. Therefore, TCP or UDP packets sent to or from port 88 are automatically permitted. For more information about configuring IPSec with firewalls, see Microsoft Knowledge Base article Q233256, “How to Enable IPSec Traffic Through a Firewall” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=233256>).
- In a perimeter network, you may want the increased security of IPSec without losing the ability to filter traffic for security reasons. In this case, you can set up an IPSec tunnel from the front-end server to the internal firewall, and then a separate tunnel from the internal firewall to the back-end server. This approach secures the information on the wire while allowing you to use filtering or intrusion detection software or techniques on the traffic before it reaches your internal network.

Note IPSec encryption occurs after the application (in this case, Exchange) passes the request to Windows to send to the server. So, as far as Exchange is concerned, the request is made over HTTP using TCP port 80. However, before the traffic leaves the server, it is intercepted, possibly encrypted if ESP is configured, and sent over a separate channel (IP protocols 50 or 51). Thus, the encryption is transparent to the Exchange applications running on each server, and the fact that the data never used port 80 is not an issue to these applications.

Service Packs: Upgrading Front-End and Back-End Servers

When new service packs are released, upgrading your front-end servers is a straightforward process. It is simpler than upgrading a back-end server, because there is no user data stored on the front-end server. As long as there are multiple front-end servers, taking one server offline does not mean an interruption in service for the user. If you have multiple front-end servers that are load balanced, you can remove the front-end server you want to upgrade from the load-balancing cluster, upgrade it, and add it back to the load-balancing cluster with no interruption in service to users.

Upgrading Considerations for Outlook Web Access

If Outlook Web Access is deployed in your organization, you should upgrade every front-end server in the organization before you upgrade any back-end servers. This is because of the way Outlook Web Access is built. At a high level, Outlook Web Access data is comprised of two parts: templates and controls. Templates are things such as forms (e-mail messages, calendar items, and other forms). Controls are files such as DHTML behaviors, JScript files, style sheets, and images that are stored in the /exchweb virtual directory in IIS.

When a user accesses Outlook Web Access through the front-end server, the templates actually come from the back-end server, and the controls come from the front-end server. The templates come from the back-end for performance reasons. The controls come from the front-end server because there is no mechanism to proxy a request for data that is not in the Exchange store from a front-end server to a back-end server. This is not a problem as long as the front-end and back-end servers run the same version and service pack of Exchange.

A problem arises when the servers run different versions of service packs. If a template on a back-end server references a control on a front-end server, and the front-end server is running a previous service pack, the control the back-end server is referencing might not exist. As a result, Outlook Web Access does not work for users whose mailboxes are on the upgraded back-end server.

The reverse situation does not have the same problem. If you upgrade a front-end server first, then users see templates from the back-end server. These templates reference the previous versions of the controls, which still exist on the front-end server, because the files are versioned and not removed in an upgrade. To address this issue, it is recommended that you upgrade all front-end servers before upgrading any back-end servers.

In addition, ensure that the required services are running before you upgrade. For Exchange setup to run, you must install and enable (but not necessarily start) Network News Transfer Protocol (NNTP), SMTP, w3SVC and IIS Admin. If the MExchangeMTA, IMAP4, POP3, and MExchangeIS services are disabled, Setup still runs; however, Setup will enable these services after it starts. After setup is complete, you can disable unnecessary services.

3

Scenarios

This chapter discusses common scenarios where Exchange front-end and back-end topology is deployed. The scenarios can be broadly divided into intranet and extranet scenarios, with the intranet scenarios focused on performance and scalability and the extranet scenarios focused on security.

In each scenario, the following topics are discussed:

Scenario

What is the scenario, and when does it apply?

Setup instructions

How to set up the scenario, in general terms. (Specific configuration instructions appear later in this book.)

Discussion

What is special about this scenario? How does it work? What additional information is needed to make decisions about this scenario?

Issues

Caveats or limitations of this scenario.

Of the following five scenarios, four require a firewall. You can use software and hardware solutions as a firewall. Port filtering is the minimum requirement for a firewall that protects the server from the Internet.

Standard Front-End and Back-End Topology Without a Firewall

Figure 4 illustrates a basic front-end and back-end topology.

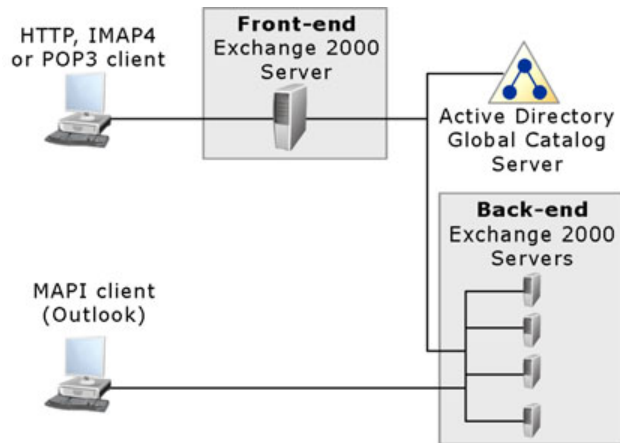


Figure 4 A standard front-end and back-end topology

Scenario

A company wants to maintain a single namespace for their e-mail servers but cannot fit all of their users on a single server.

Setup Instructions

1. Set up a standard collection of servers running Exchange.
2. Set up a single server running Exchange configured as a front-end server.
3. Direct HTTP, POP, and IMAP users to this server, not to their back-end servers.
4. Ensure that all virtual directories and servers are configured identically on all front-end and back-end servers.

Discussion

This is the default configuration. You do not need to perform any steps other than the standard front-end and back-end configuration steps.

Issues

If the network permits connections between the client and the back-end servers, there is nothing to prevent users from circumventing the front-end server and connecting directly to the back-end server. If this is undesirable, you must change the network routing configuration or the back-end server configuration to prevent direct connections between a client and a back-end server.

- ▶ **To configure the back-end server to only allow connections to a virtual server from a certain group of IP addresses**
 1. In Internet Services Manager, right-click the virtual server you want, and then click **Properties**.
 2. Click the **Directory Security** tab. In **IP address and domain name restrictions**, click **Edit**. Select **Denied Access**, and then click **Add** to allow specific IP addresses or a range of IP addresses (using a network ID and a subnet mask) to access that virtual server.

Front-End Server Behind a Firewall

Figure 5 illustrates a front-end and back-end topology where the front-end server is behind the firewall.

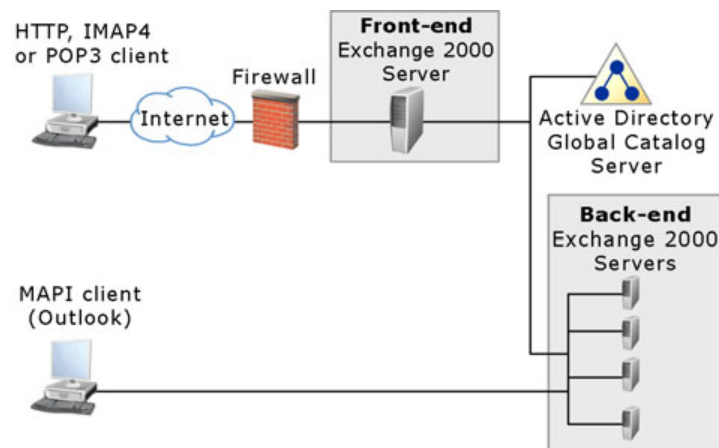


Figure 5 A simple Exchange firewall topology

Scenario

To achieve security and still provide access to Outlook Web Access, POP, or IMAP from the Internet, a corporation wants to place the Exchange system behind the corporate firewall.

Setup Instructions

1. Set up a standard Exchange front-end and back-end environment in the corporation.
2. Configure a firewall between the front-end server and the Internet. For more information about how to configure an Internet firewall for use with a front-end server running Exchange, see Chapter 6, “Configuring Firewalls.”

Discussion

Because the entire configuration is inside the firewall, Exchange does not require any special configuration. After a request comes through the firewall to the front-end server, the front-end server returns a response without any configuration changes.

IP address filtering is highly recommended to limit requests through the firewall to only those going to the front-end server (or servers) running Exchange and block requests through the firewall to other servers in the organization.

Web Farm With a Firewall

Figure 6 illustrates a Web farm scenario.

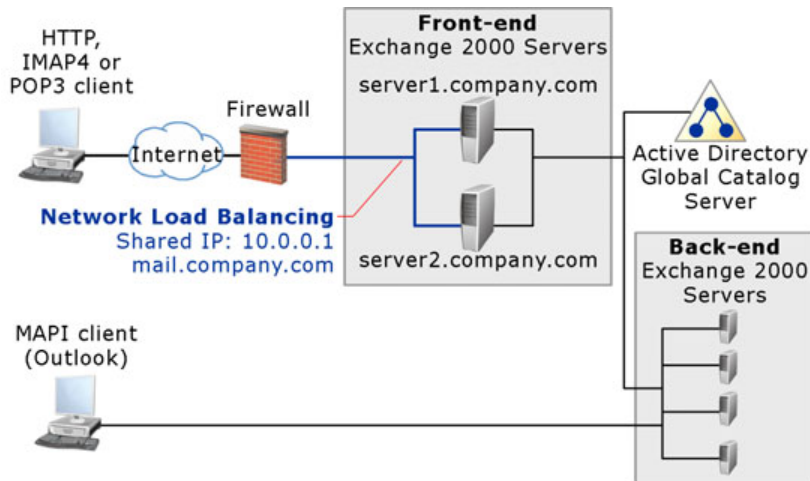


Figure 6 Front-end and back-end topology in a Web farm

Scenario

A corporation is rolling out Outlook Web Access to 200,000 users. The goal is to have a single namespace (for example, `http://mail`) in which users can reach their mailboxes. In addition, for performance reasons, the corporation wants to avoid having a bottleneck at the front-end server or a single point-of-failure, so they want to spread the load over multiple front-end servers by using Network Load Balancing (NLB). This scenario is referred to as a “Web Farm.”

Note Although this is the only scenario that depicts NLB, you can use NLB to distribute load among front-end servers in any of the scenarios described in this book.

Setup Instructions

1. Set up a group of servers as back-end servers in the same forest.
2. Distribute users over the servers.
3. Set up another group servers as front-end servers and configure Network Load Balancing on all these servers.

Discussion

For information about how to set up Network Load Balancing, see the Windows online documentation. Configuring Exchange on the front-end servers does not require any special steps.

Issues

As mentioned earlier, the load-balancing solution you use should ensure that each user is always sent to the same front-end server for the duration of a session.

Front-End Server in a Perimeter Network

Figure 7 illustrates an Exchange 2000 front-end server in a perimeter network.

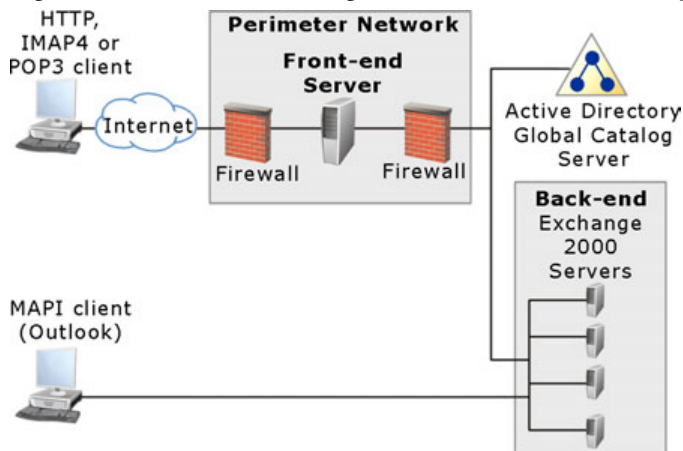


Figure 7 Exchange 2000 front-end server in a perimeter network

Scenario

In Figure 7, the corporation places the front-end server between two separated firewalls. The first firewall separates the front-end server from the Internet and allows requests only to that front-end server. The second firewall separates the front-end server from the internal network. The systems between the two firewalls lie in what is known as a perimeter network. This is also referred to as a demilitarized zone (DMZ). A perimeter network configuration provides more security because if the front-end server is compromised, there is still another barrier between the intruder and the rest of the network.

Note Placing front-end servers inside the perimeter network is one approach to deploying front-end and back-end topology within a perimeter network. However, the recommended approach is depicted in the next scenario, “Advanced Firewall in a Perimeter Network.” This approach involves placing the front-end and back-end servers inside the intranet and placing an advanced firewall (such as ISA Server) in the perimeter network. The advanced firewall can provide application protocol filtering and perform additional authentication on requests before it proxies them to the internal network.

Setup Instructions

1. Configure the outer (Internet) firewall for a firewall in this environment, limiting access to only the ports required and to only the designated front-end server.
2. Configure the inner (intranet) firewall to have certain ports open to support authentication, DNS, and Active Directory access. The exact list depends on the balance of security and features that each corporation chooses.

For information about how to configure Internet and intranet firewalls, see Chapter 6, “Configuring Firewalls.”

Discussion

Typically, corporations that have deployed and standardized the use of a perimeter network have restrictions on the type of network traffic allowed through the intranet firewall by limiting the network ports that are enabled on the intranet firewall. However, the front-end server requires certain ports to operate fully.

Issues

Some corporations that have deployed perimeter network topologies for other services have policies that restrict computers located within the perimeter network from initiating connections with servers inside the corporate intranet. A front-end server running Exchange is not supported in this configuration because it must initiate connections.

Additionally, the front-end server must be a member of the same domain as the back-end servers. Some corporations do not allow member servers in the perimeter network; for these corporations, deploying a front-end server in the perimeter network is not an option.

It is recommended that you completely configure the front-end server before the intranet firewall is put in place or locked down. Configuring settings on the front-end server in Exchange System Manager requires the System Attendant (MSEExchangeSA) service to be running so that the configuration information can replicate to the metabase. The MSEExchangeSA service requires RPC access to the back-end servers, and RPCs often are not allowed across an intranet firewall in a perimeter network.

The DSAccess component in Exchange 2000 SP2 is redesigned to provide better support for perimeter networks in which RPC traffic is not allowed across the internal firewall. However, there are two additional registry keys that you should set on the front-end server to disable NetLogon and the Directory Access ping:

NetLogon

DSAccess connects to Active Directory servers to check available disk space, time synchronization, and replication participation by using NetLogon service with RPC. If you do not allow RPC traffic across the internal firewall, you should stop the NetLogon check by creating the DisableNetlogonCheck key on the front-end server.

Directory Access ping

By default, Directory Access uses Internet Control Message Protocol (ICMP) to ping each server to determine whether the server is available. However, in a perimeter network in which there is no ICMP connectivity between the server running Exchange and the domain controllers, Directory Access determines that every domain controller is unavailable. Directory Access then discards old topologies and performs new topology discoveries, which impact server performance. To avoid these performance issues, you should turn off the Directory Access ping on the front-end server by creating the LdapKeepAliveSecs registry key for the Windows implementation of LDAP (wLDAP).

For information about how to set these registry keys, see “Configuring DSAccess for Perimeter Networks” in Chapter 4.

Advanced Firewall in a Perimeter Network

Figure 8 illustrates an advanced firewall scenario, in which an advanced firewall is placed inside the perimeter network, between the Internet firewall and the internal firewall. Front-end and back-end servers are placed in the same network behind the internal firewall. This is the recommended topology for the following reasons:

- It provides security by isolating intruders from the rest of the network.
- It provides application protocol filtering.
- It performs additional authentication on requests before it proxies them to the internal network.

Note As an alternative to placing the advanced server within a perimeter network behind a separate Internet firewall, the advanced server itself can function as the Internet firewall.

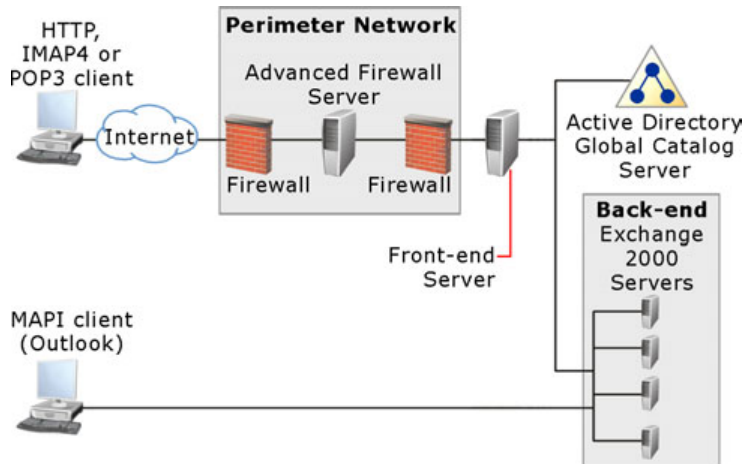


Figure 8 Proxy server between two firewalls

Scenario

A corporation places an advanced firewall such as ISA Server between two separated firewalls. The corporation's decision to set up this advanced firewall topology is based on the following benefits:

- Advanced firewalls provide additional security to the network by protecting against unauthorized access, inspecting traffic, and alerting the network administrator to attacks.
- Advanced firewalls allow you to use such methods as port filtering and IP filtering to control traffic.
- Advanced firewalls allow you to restrict access by users and groups, application type, time of day, content type, and destination sets.

Setup Instructions

The following steps are based on using ISA Server as the proxy server software.

1. Configure the outer (Internet) firewall for a firewall in this environment, limiting access to only the ports required and to only the designated front-end server.
2. Configure the inner (intranet) firewall to have certain ports open to support authentication, DNS, and Active Directory access. The exact list depends on the balance of security and features that each corporation selects.
3. Configure the advanced firewall. The following are general guidelines you should follow when deploying an ISA Server in a front-end and back-end topology. (For detailed information about how to configure ISA Server, see the ISA Server product documentation.)
 - a. Configure a listener for SSL.
 - b. Create a destination set that contains the external IP address of the ISA Server. This destination set will be used in the Web publishing rule.
 - c. Create a Web publishing rule that redirects requests to the internal front-end server.
 - d. Create protocol rules to open ports in ISA Server for outgoing traffic.
 - e. Configure the ISA server for Outlook Web Access. (For information about how to configure an ISA server for Outlook Web Access, see Microsoft Knowledge Base article Q307347, "Secure OWA Publishing Behind ISA Server May Require Custom HTTP Header" (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=307347>).

For more information about ISA Server, including product information and technical resources, see the ISA Server Web site (<http://www.microsoft.com/isaserver>).

Discussion

ISA Server contains two types of rules:

Server publishing rules

These rules, which can apply to any protocol, inspect incoming requests at the receiving port. If an incoming request is allowed, the protocol rule forwards it from the receiving port to an internal IP address.

Web publishing rules

These rules apply to HTTP or HTTPS (80/443) requests only. You can set up Web publishing rules to filter incoming requests based on the service type, port, source computer name, and destination computer name. You can also allow only specific servers or deny high-risk servers.

If you are supporting HTTP clients, create a Web publishing rule to handle HTTP or HTTPS traffic. If you are supporting POP or IMAP clients, create server publishing rules to handle these protocols.

Unlike the perimeter network scenario, the ISA server in the perimeter network does not have to be a member server unless you configure ISA Server to authenticate requests. Generally, if you have configured authentication on the front-end server, you do not need to configure ISA Server to authenticate users. However, if you want to restrict incoming requests to those that originate from specific users, you must create a Web publishing rule that specifies the users and enables authentication on the ISA Server. In this case, the ISA server must be a member of the Windows domain. In addition, ISA Server does not delegate the user's credentials to back-end servers. Therefore, although ISA Server can authenticate users and restrict access to the network, users cannot be pre-authenticated for Outlook Web Access.

Issues

In the advanced firewall scenario, there is no need for RPC access to the internal network. This is often regarded as an advantage because fewer ports must be open on the internal firewall; however, regardless of the number of open ports, the potential for a security breach exists. To avoid this security risk, ensure that the appropriate filters are set up for each open port.

In the advanced firewall scenario, you can configure SSL in one of two ways:

- Between the client and ISA server only.
- Between the client and ISA server, and between the ISA server and the front-end server.

The second option is usually used if customer policy dictates that e-mail traffic within the perimeter network is encrypted. After ISA Server receives the SSL request from the client, it terminates the session and reopens a new SSL session with a new certificate to contact the front-end server. The name on each certificate is important. The certificate name on the incoming request must match the name the user typed into the URL. Furthermore, The certificate name on the request to the front-end server must match the name or IP address of the front-end server.

To configure SSL in ISA Server, use the **Bridging** tab in the Web publishing server rule to direct SSL traffic. If you are hosting multiple domains and want to use SSL, you must set up a listener and a different IP address for each domain; this is because the certificates must be named so that they match the destination names or IP addresses.

4

Configuring a Front-End Server

A front-end server is an ordinary Exchange server until it is configured as a front-end server. A front-end server must not host any users or public folders.

A front-end server must be a member of the same Exchange organization as the back-end servers (therefore, a member of the same Windows 2000 forest).

► **To designate a front-end server**

1. Install the server running Exchange in the organization.

Note Only servers running Exchange 2000 Server Enterprise Edition can be configured as front-end servers.

2. Use **Exchange System Manager** to go to the server object, right-click the server object, and then click **Properties**.
3. Select **This is a front-end server**, and then close the page.
4. To begin using the front-end server do one of the following:
 - Restart the computer.
 - Stop and restart the HTTP, POP, and IMAP services.

Hosting Multiple Domains

In the simplest topology, you do not need to configure the front-end server beyond designating it as a front-end server. However, if you are hosting multiple domains, organizations, or public folder trees, you will need to create additional virtual servers or directories.

For each domain you host, you must have a unique virtual server or directory so that users with e-mail addresses from those domains can log on to the appropriate virtual server or directory. There are two methods you can use to configure the HTTP virtual servers when hosting additional domains:

- Method one: Create additional virtual servers
- Method two: Create additional virtual directories

In the following methods, your default Exchange domain is microsoft.com, and you are hosting mailboxes for adatum.com and public folders for contoso.com. The methods describe how to configure your front-end and back-end servers for these hosted companies.

Method One: Create Additional Virtual Servers

In this method, you create a virtual server for each domain you host. For example, you have three HTTP virtual servers on each server running Exchange: one is the default Exchange virtual server, one is the virtual server for adatum.com's mailboxes, and one is the virtual server for contoso.com's public folder tree. This method allows for maximum flexibility in determining the resources available to each domain.

Each HTTP virtual server must have a unique combination of IP address, host header, and port. The host header is the DNS name of your Web site. For example, if users access your Web site by typing **http://www.contoso.com/example** into a browser, the host header is **www.contoso.com**. By specifying the host header for the virtual server, you can host multiple Web sites using the same port and IP address combination because the host header ensures uniqueness. However, if you try to create a virtual server that has the same combination of these settings as another server, the new virtual server will not start.

When additional virtual servers are set up and sessions are non-encrypted, client requests are routed to the correct virtual server as follows:

1. The client connects to the front-end server.
2. The server receives a packet containing the IP address, requested port, and host header.
3. The server uses this information to find the appropriate virtual server to handle the request.

When a user's request is routed to the correct virtual server, there is no guarantee that the user will be able to access Outlook Web Access successfully. To log on to the virtual server, the user must also have an e-mail address from the SMTP domain that is associated with the virtual server. Each virtual server and each virtual directory that points to mailboxes is associated with a specific SMTP domain. Therefore, when you create additional virtual servers that point to mailboxes, you must associate the virtual server with the appropriate domain. (Exchange Setup associates the default Exchange HTTP virtual server with the default Exchange domain; this default cannot be changed.)

After the user's request is routed to the correct virtual server, the process continues as follows:

1. The virtual server uses the user's authentication credentials to look up the user in Active Directory.
2. If the user has an e-mail address from the SMTP domain associated with this virtual server, the virtual server allows access. Otherwise, the virtual server denies access.

Note Using the authentication credentials to look up the user's SMTP proxy address works only if authentication is enabled on the front-end server. If authentication is not enabled, the server uses the alias specified in the URL (because accessing Outlook Web Access through a front-end server with pass-through authentication requires the user to enter an explicit logon in the format **http://server/exchange/user**). Then, to form the proxy address, the server combines the alias with the SMTP domain on the virtual server properties. .

Note When SSL is used, the host header is contained in the encrypted portion of the packet; only the IP address and port are available in the unencrypted portion. To determine which SSL certificate to use to decrypt the data, the server must be able to determine the appropriate virtual server with a unique combination of IP address and port. Therefore, when SSL is used (such as when the front-end server is deployed on the Internet), the IP address must be specifically associated with the appropriate virtual server.

After you create the additional virtual server, you need to create the Exchange and Public virtual directories underneath that virtual server. For instructions about how to create virtual directories, see "Creating Virtual Directories" later in this chapter.

Note If you want to enable the Change Password feature for your users, you must enable the feature on all virtual servers on the front-end. Each virtual server must contain a virtual directory named IISAdmPwd.

Method Two: Create Additional Virtual Directories

The second way of configuring multiple hosted domains is to add a virtual directory for each domain. For access to mailbox stores, you must specify the domain in the properties of the virtual directory. For access to public stores, you must specify the root public folder. For your first hosted company, *adatum.com*, add a virtual directory under the default Exchange HTTP virtual server, with a name that uniquely identifies the hosted company, such as *Adatum*. In the properties of the virtual directory, click **Modify**, and then select *adatum.com* as the SMTP domain, just as you did for the virtual server. Users from *adatum.com* will now be able to go to `http://mail.microsoft.com/adatum/` to access their mailboxes.

Add another virtual directory for *Contoso*, and this time select **Public folder** and specify the public folder root for *Contoso*. Users from *contoso.com* will now be able to go to `http://mail.microsoft.com/contoso/` to access their public folders.

The main advantage to this method has to do with SSL. SSL certificates are issued for a specific host or domain name—in this case, *mail.microsoft.com*. When you have multiple virtual servers with different domain names (for example, *mail.microsoft.com* and *mail.adatum.com*), you need one SSL certificate for each domain, and that costs money. Because, in method two, clients access their data through a single domain—`http://mail.microsoft.com`—you save on the cost of the certificates as well as the additional step of configuring SSL on each virtual server.

Creating HTTP Virtual Servers

You must use Exchange System Manager, not Internet Services Manager when you create virtual servers. When you create virtual servers in Exchange System Manager, you do not need to simplify the URL; after you create a virtual server that points to mailboxes and set the host headers, users can type `http://<virtual server name>/` without having to type `/exchange`.

► To create a virtual server

1. In **Exchange System Manager**, in the **HTTP Protocols** container for the front-end server, right-click **HTTP**, and then select **New Virtual Server**.

Note For a name, it is recommended that you use something following the form of “*adatum.com (front-end)*.” Consistent naming of the new virtual servers ensures that each virtual server’s purpose and associated domain can be easily determined. The name of the virtual server is used only for identification purposes and does not affect its operation.

2. Click **Mailboxes for** or **Public folder**, click **Modify**, and then do one of the following:

- If the virtual server points to mailboxes, select the domain.

Note The list of domains in **Select SMTP Domain** is pulled from the domains of the SMTP addresses in the Exchange organization's recipient policies, so if you have more than one recipient policy for the same domain, you will see duplicates in this dialog box. It does not matter which one you choose.

- If the virtual server points to a public folder, select the appropriate public folder to act as the root public folder for this virtual server.
3. Click **Advanced**, and then add host headers that define all the names a client might use to contact this front-end server.

Note If a front-end server is used internally and externally, it is recommended that you list both a hostname and a fully qualified domain name.

You can leave the IP address field at **All unassigned** (the default) or restrict it to the particular IP address assigned to the server. If you know SSL will be used to connect to this front-end server, you may want to configure a specific IP address for the virtual server.

Creating Virtual Directories

► To create virtual directories

1. In **Exchange System Manager**, right-click the HTTP virtual server, click **New**, and then click **Virtual Directory**.
2. Name the virtual directory.

Note The name you use is important, because this is what the client must enter into their browser (for example, <http://mail.microsoft.com/virtualdirectoryname>). Do not use the following character sequences in the directory name because IIS blocks them:

- Period (.)
 - Double period (..)
 - Period and forward slash(/)
 - Backslash (\)
 - Percent sign (%)
 - Ampersand (&)
-

3. Right click the virtual directory you just created, and then click **Properties**. Specify whether or not the directory is for access to a mailbox or a public store. If the directory is for access to a mailbox store, the SMTP domain must be specified. If the virtual directory is for access to a public folder store, then select the appropriate public folder to act as the root public folder for this virtual directory.

If you are hosting multiple domains, as explained in the first method for adding virtual servers, where a virtual server is created for each domain, it is recommended that you use the standard virtual directory names—“Exchange” for mailbox access (make sure to specify the domain again) and “public” for public folder store access. Do not create an “exadmin” virtual directory on any additional virtual servers; this is used only by System Manager and is not proxied by the front-end server.

Configuring Authentication

It is highly recommended that you use dual-authentication, in which both front-end and back-end servers are configured to authenticate users.

However, if you have a locked down perimeter network in which RPCs are not allowed across the intranet firewall, and it is impossible for the front-end server to authenticate users, you can use pass-through authentication.

Warning When you use pass-through authentication, anonymous HTTP requests go directly to the back-end server where they are authenticated. You should use pass-through authentication only if the front-end server cannot authenticate users. For more information, see Chapter 3, “Scenarios.”

► To configure authentication

1. Start **Exchange System Manager**: Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. Go to the “Exchange” or “Public” virtual directory.
3. Right-click the “Exchange” or “Public” virtual directory, and then click **Properties**.
4. Click the **Access** tab, and then click **Authentication**.
5. Do one of the following:
 - To configure the front-end server to authenticate users (as in dual authentication), select the **Basic authentication** check box.
 - To configure pass-through authentication, select the **Anonymous access** check box, and then clear the **Basic authentication** check box.

Configuring The Front-End Server To Assume a Default Domain

You can configure the front-end server to assume a default domain so that users do not need to remember their domain.

► **To configure the front-end server to assume a default domain**

1. Start **Exchange System Manager** on the front-end server: Click **Start**, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. Right-click the “Exchange” virtual directory, and then click **Properties**.
3. Click the **Access** tab, and then click **Authentication**.
4. Select the **Basic authentication** check box, and then type the domain in the **Default domain** box. (See Figure 9.)
5. Repeat this procedure for the “Public” virtual directory.

Note Ensure that the System Attendant service is running so that the configuration settings replicate from the directory to the IIS metabase (you can also restart Exchange System Attendant to force replication).

After replication is complete, users can log on with just their user name and password; no domain is required.

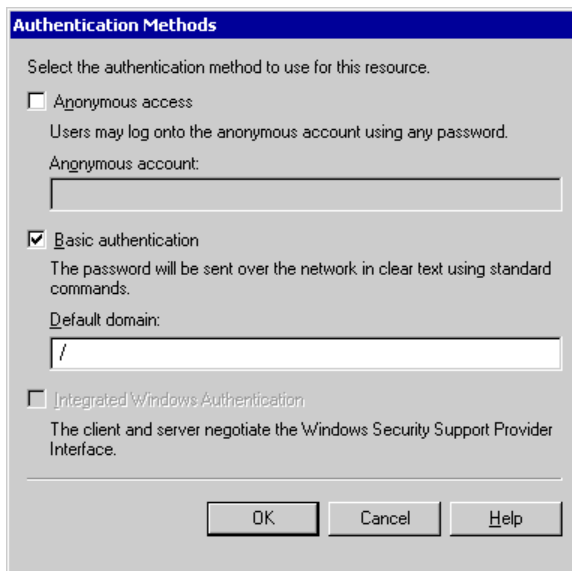


Figure 9 Configure the front-end server to assume a domain

Allowing the Use of an E-Mail Address as the Logon User Name

An additional option for authentication is to configure a user principal name (UPN) for users. This allows users to enter their e-mail address as their user name. To configure UPN, in the **Authentication Methods** dialog box (see Figure 9), in the **Default domain** box, type a backslash (\). You must configure the UPN on all virtual servers and all virtual directories on each front-end and back-end server. Use ISM to configure the default virtual server, and use System Manager to configure any additional virtual servers. After you configure a UPN in the properties of the virtual servers and directories on all front-end and back-end servers, users can authenticate by using *user@domain.com* as their user name. In addition, users can still use the *domain\username* logon.

Note If you want to set a default domain or configure UPN, you should upgrade to Exchange 2000 SP1 or later. For more information, see Microsoft Knowledge Base article Q267936, "Directory Service to Metabase Service May Not Replicate the Default Logon Domain for Virtual Servers" (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=267936>).

Disabling Unnecessary Services

Not all Exchange services are required on a front-end server, depending on the protocols being exposed and whether or not you will be making configuration changes after the initial setup. To stop and disable services, use the Services snap-in in the Microsoft Management Console (MMC). The following list shows the Exchange services required; stop and disable all other Exchange services:

- **HTTP** No Exchange-specific services required. However, the World Wide Web Publishing service (w3svc) must be running in order for Outlook Web Access to work.
- **POP** Exchange POP (POP3Svc) and Microsoft Exchange System Attendant Service (MSEExchangeSA).
- **IMAP** Exchange IMAP (IMAP4Svc) and MSEExchangeSA.
- **SMTP** MSEExchangeIS and MSEExchangeSA.
- **Exchange System Manager** MSEExchangeSA.
- **Routing Engine** Microsoft Exchange Routing Engine Service (RESvc). The routing engine must be running on all Exchange servers.
- **NNTP** NNTP must be enabled on a server during upgrade; however, if you are not offering NNTP to your users, you can disable it.

Running the IIS Lockdown Wizard

It is important that you secure IIS before you expose servers to the Internet. To secure IIS, turn off all features and services except those that are absolutely necessary. The IIS Lockdown Wizard helps you disable unnecessary IIS features and services based on the type of server software you are running. To provide multiple layers of protection against attackers, the IIS Lockdown Wizard also contains URLScan, which analyzes HTTP requests as IIS receives them and rejects any suspicious requests.

IIS Lockdown Wizard also contains a configuration template for Exchange that turns off unwanted features and services. To use this configuration template, run IIS Lockdown Wizard, select the Exchange template you want, and then change or accept the default configuration options.

After you run IIS Lockdown Wizard, the URLScan application is installed in the folder *<drive:>/<Windows directory>/system32/inetsrv/urlscan*.

Unless you configure the following settings in the **Urlscan.ini** file, after you run the wizard, you may experience some problems with Outlook Web Access functionality:

- **Allow Dot In Path** Make sure that this setting is set to "1" to ensure that Outlook Web Access attachments can be accessed and that earlier-version browsers can use Outlook Web Access.
- **File Extensions** By default, .htr files are disabled. If this file type is disabled, the Outlook Web Access Change Password feature does not function.
- **Deny Url Sequences** In the [DenyUrlSequences] section, sequences that are explicitly blocked can potentially affect access to Outlook Web Access. Any mail item subject or mail folder name that contains any of the following character sequences is denied access:
 - Period (.)
 - Double period (..)
 - Period and forward slash(/)
 - Backslash (\)
 - Percent sign (%)
 - Ampersand (&)

If you experience additional problems when you attempt Outlook Web Access requests with URLScan enabled, check the UrlScan.log file for the list of requests that are being rejected. For more information, see Microsoft Knowledge Base article Q309677, “XADM: Known Issues and Fine Tuning When You Use the IIS Lockdown Wizard in an Exchange 2000 Environment” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=309677>).

For information about how to install and use IIS Lockdown Wizard, see Microsoft Knowledge Base article Q325864, “HOW TO: Install and Use the IIS Lockdown Wizard” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=325864>).

The IIS Lockdown Wizard is available at <http://go.microsoft.com/fwlink/?LinkId=10194>.

Note To maximize the security of your Exchange servers, apply all of the necessary hotfixes before and after applying IIS Lockdown Wizard. The hotfixes ensure that servers remain protected against known security vulnerabilities.

Dismounting and Deleting Public and Mailbox Stores

You must dismount and delete the public folder stores on the front-end server before you place the server in production.

If you are not using SMTP on the front-end server, you should also dismount and delete the mailbox stores. (If you are using SMTP, a mailbox store is required, but the mailbox store must not contain any mailboxes. For more information, see “Running SMTP for POP and IMAP Clients” in Chapter 1.

The Information Store service will not start unless there is at least one storage group defined. If you need to leave the Information Store service running, do not delete the storage group on the front-end server after deleting the private and public stores.

Note If the MExchangeIS service is not running, you will not be able to make configuration changes using Internet Services Manager (ISM). If you must make configuration changes using ISM (for example, configuring SSL encryption), be sure to complete these steps before you remove the mailbox and public folder stores.

Configuring Network Load Balancing

You may want to load balance the front-end servers. Do not use Windows Clustering Services to load balance the front-end servers because there is no data stored on the front-end servers; each front-end is essentially a copy of every other, so failing over is not useful. Use Network Load Balancing to evenly spread client requests across multiple front-end servers.

See the Windows 2000 online documentation for information about configuring Network Load Balancing.

Configuring SSL

The steps to configure SSL for POP, IMAP, and SMTP differ from HTTP.

► **To configure SSL for POP, IMAP, and SMTP**

1. In **Exchange System Manager**, right-click the virtual server, and then click **Properties**.
2. Click the **Access** tab, and then click **Certificate**.
3. Follow the instructions in the wizard, using the guidelines in the online help to request and install the SSL certificate.

► **To configure SSL for HTTP**

1. In **Internet Services Manager**, for the default Exchange HTTP virtual server, right-click **Default Web Site**, and then click **Properties**.
2. Click the **Directory Security** tab, and then select **Server Certificate**.
3. Follow the instructions in the wizard, using the IIS documentation as a reference, to request and install the SSL certificate. HTTP SSL is configured individually for each virtual server or Web site, so if you create additional HTTP virtual servers in System Manager, you will need to configure SSL for each server (or Web site as it is referred to in Internet Services Manager).

Configuring SMTP

SMTP must be available on the front-end server to allow POP and IMAP clients to submit e-mail. You can install SMTP on the front-end server or set up a separate SMTP server. If you want to install SMTP on the front-end server, configure SMTP for internal and external domains, as described in the following two sections.

Mail for Internal Domains

If you want the front-end server to accept mail inbound from the Internet, the front-end server needs to know the domains for which it should accept mail. Adding recipient policies for each of your domains tells all servers in the Exchange organization to accept mail for those domains. Additionally, you must enable anonymous access for other SMTP servers on the Internet to successfully route mail to your organization (this is the default setting).

Mail for External Domains

In the default configuration, any SMTP mail that is submitted to your server and addressed to external domains is denied. This occurs because relaying is turned off for all anonymous access (however, authenticated users can still send mail to any external domain). Users who try to anonymously submit mail to external domains get an error, such as “550 5.7.1 Unable to relay for suzan@adatum.com.” The clients must be configured to use SMTP authentication.

Note Although you can allow relaying for anonymous access, it is not recommended and should never be necessary. Allowing unauthenticated relaying allows anyone on the Internet to use your server to send mail.

If you choose to require SMTP authentication for mail submitted by your users from the Internet, you should also require SSL for clear text or basic authentication, so that corporate usernames and passwords are not sent out unencrypted. Configure SSL for basic authentication in the **Properties** of the SMTP virtual server: On the **Access** tab, click **Authentication**, and then select **Requires TLS encryption**.

Configuring DSAccess for Perimeter Networks

The DSAccess component in Exchange 2000 SP2 was redesigned to provide better support for perimeter networks in which RPC traffic is not allowed across the internal firewall. However, to prevent performance problems, there are two additional registry keys that you should set on the front-end server to disable NetLogon and the Directory Access ping. In addition, you can configure DSAccess so that your front-end servers contact specific domain controllers and global catalog servers. The following sections describe how to configure these settings.

Important This section contains information about editing the registry. Before you edit the registry, make sure you understand how to restore it if a problem occurs. For information about how to do this, view the "Restore the Registry" Help topic in Regedit.exe or Regedt32.exe.

Warning Using Registry Editor incorrectly can cause serious problems that may require you to reinstall your operating system. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. For information about how to edit the registry, view the "Change Keys and Values" Help topic in Registry Editor (Regedit.exe) or the "Add and Delete Information in the Registry" and "Edit Registry Information" Help topics in Regedt32.exe. Note that you should back up the registry before you edit it. If you are running Windows NT or Windows 2000, you should also update your Emergency Repair Disk (ERD).

Disabling the NetLogon Check

DSAccess connects to Active Directory servers to check available disk space, time synchronization, and replication participation by using the NetLogon service with RPC. If you do not allow RPC traffic across the internal firewall, you should stop the NetLogon check by creating the DisableNetlogonCheck registry key on the front-end server.

► **To create the DisableNetlogonCheck registry key**

1. Start Registry Editor (Regedt32.exe).
2. Locate and select the following key in the registry:

```
HKEY_LOCAL_MACHINE\  
  System\  
    CurrentControlSet\  
      Services\  
        MExchangeDSAccess\  
          DisableNetlogonCheck
```

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

```
Value name: DisableNetlogonCheck  
Data type: REG_DWORD  
Value data: 1
```

Disabling the Directory Access Ping

In a perimeter network, you must also create a registry key on the front-end server to prevent Directory Access from pinging domain controllers. Create the LdapKeepAliveSecs registry key on the front-end server.

► **To create the LdapKeepAliveSecs registry key**

1. Start Registry Editor (Regedt32.exe).
2. Locate and select the following key in the registry:

```
HKEY_LOCAL_MACHINE\  
    System\  
        CurrentControlSet\  
            Services\  
                MExchangeDSAccess \
```

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

```
Value name: LdapKeepAliveSecs  
Data type: REG_DWORD  
Value data: 0
```

Specifying Domain Controllers and Global Catalog Servers

In a perimeter network, you may want to configure DSAccess to use specific domain controllers and global catalog servers, and then use IP filtering to ensure that the front-end servers connect to only those domain controllers and global catalog servers. To specify domain controllers and global catalog servers, use the **Directory Access** tab in the *<server name> Properties* dialog box. Specifying servers on the **Directory Access** tab sets keys in the registry. The **Directory Access** tab was not available in earlier versions of Exchange; if you have previously set registry keys to specify domain controllers and global catalog servers, these registry keys will still work.

5

Configuring a Back-End Server

Exchange configuration is stored in Active Directory on a per-forest basis, which means that all front-end and back-end servers must be in the same forest. Back-end servers can be accessed directly if required, with no effect on the behavior of the front-end and back-end configuration.

If you did not configure any extra virtual servers or directories on any front-end servers, then you do not need to configure any on the back-end server. If you created additional virtual servers or directories on any front-end servers, however, you must add matching virtual servers and directories on the back-end servers.

The specific back-end server configuration steps depend on whether the back-end server is hosted on a cluster or not. The steps below describe configuration for non-clustered back-end servers. For information about configuring clustered back-end servers, see the technical paper *Deploying Exchange 2000 Server Clusters* (<http://go.microsoft.com/fwlink/?LinkId=10193>).

Note When back-end servers are clustered, you cannot access the cluster HTTP virtual directory using the name of the actual physical clustered computer. You must use the virtual group name.

Creating and Configuring HTTP Virtual Servers on Back-End Servers

In the “Configuring a Front-End Server” section, two methods for configuring your front-end and back-end topology when hosting multiple domains were discussed:

- Method one involves setting up a virtual server for every hosted domain.
- Method two involves setting up a virtual directory for every hosted domain.

Back-end configuration differs slightly depending on which method you chose. Additionally, if you create additional virtual servers on the front-end servers for other reasons (such as to host Web applications), you must add similarly configured virtual servers to any back-end servers on which the Web applications will exist.

Method One: Additional Virtual Servers

If you chose to create an additional virtual server for each additional domain (with an Exchange virtual directory beneath each additional domain), you will need to create a matching virtual server on the back-end server. The steps are slightly different from those for configuring the front-end server.

► **To configure additional virtual servers on a back-end server**

1. Give the virtual server a consistent name, such as “adatum.com (back-end)”.
2. Click **Modify**, and then select the appropriate domain from the list (adatum.com in this case).
3. Click **Advanced**, and then add the appropriate host headers (mail.adatum.com in this case). The address through which the client browser accesses the front-end server is forwarded by the front-end server to the back-end server, so the back-end server must be aware of every name a client might use to reach the front-end server (for example, http://mail, http://mail.adatum.com, and so on).

Note On the back-end server, the TCP port must be 80. This is the only port used for HTTP communication between front-end and back-end servers, regardless of the port used by the client to communicate with the front-end server. You can leave the SSL port at the default setting.

Method Two: Additional Virtual Directories

If you chose to create an additional virtual directory (under the default domain) for each additional domain, configure the virtual directory structure to match the virtual directory structure on the front-end servers. As with the front-end servers, ensure that you specify the appropriate SMTP domain for virtual directories associated with mailbox stores.

Note The directory name must not contain the following character sequences in the directory name because IIS blocks them:

- Period (.)
 - Double period (..)
 - Period and forward slash(/)
 - Backslash (\)
 - Percent sign (%)
 - Ampersand (&)
-

Configuring Authentication on Back-End Servers

By default, HTTP virtual servers on the back end are configured to allow both basic authentication and Integrated Windows Authentication. You should use this default configuration.

Basic authentication passes the user name and password across the network in a lightly encoded (not encrypted) format. Integrated Windows Authentication refers to a package of authentication mechanisms (such as NTLM and Kerberos) that are more secure and do not send the password across the network in clear text. Only Internet Explorer supports Integrated Windows Authentication.

If you configure the front-end HTTP virtual servers to authenticate requests, the front-end server requests authentication information from the user. The user sends authentication information to the front-end server, which authenticates the user and then passes on the information to the back-end server. The back-end server then authenticates the user, but it does not need to request authentication information from the user again.

6

Configuring Firewalls

This chapter discusses how to configure firewalls for use with an Exchange front-end and back-end topology. This chapter also provides additional configuration information for the front-end servers in these environments.

Configuring an Internet Firewall

In Internet scenarios, a firewall is usually placed between the corporate network and the Internet. This firewall controls the connections that are allowed between computers on the Internet and computers in the corporate network. When you configure this firewall, it is important to consider the direction of traffic. . For a detailed discussion of port direction, see “Port Filtering” in Chapter 2.

You must configure the firewall to allow requests to certain IP addresses and over certain TCP/IP ports. Table 1 lists the ports required for different services. These ports are specific to inbound traffic (from the Internet to the front-end server).

Table 1 Ports required to be open on the Internet firewall

Internet Firewall Mail Protocols	
Destination Port number/transport	Protocol
443/TCP inbound	HTTPS (SSL-secured HTTP)
993/TCP inbound	SSL-secured IMAP
995/TCP inbound	SSL-secured POP
25/TCP inbound	SMTP

Note In Table 1, “Inbound” means that you should configure the firewall to allow computers outside (on the Internet, in this case) to initiate connections to the front-end server. The front-end server never has to initiate connections to the computers on the Internet; the front-end server only responds to connections initiated by computers on the Internet.

Configuring ISA Server

If you are using ISA Server, you must configure ISA Server as follows. (These are general guidelines; for detailed information about how to configure ISA Server, see the ISA Server product documentation.)

1. Configure a listener for SSL.
2. Create a destination set that contains the external IP address of the ISA server. This destination set will be used in the Web publishing rule.
3. Create a Web publishing rule that redirects requests to the internal front-end server.
4. Create protocol rules to open ports in ISA Server for outgoing traffic.
5. Configure the ISA server for Outlook Web Access. For information about how to configure an ISA Server for Outlook Web Access, see Microsoft Knowledge Base article Q307347, “Secure OWA Publishing Behind ISA Server May Require Custom HTTP Header” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=307347>).

Configuring an Intranet Firewall

When you set up a perimeter network, you must configure a firewall between the front-end server and the corporate intranet. In a perimeter network, the front-end server must be able to initiate connections to back-end servers and Active Directory servers. Therefore, you would configure the internal firewall with a rule that allows inbound port 80 traffic from the perimeter network into the corporate network. This rule will not allow outbound port 80 traffic from inside the corporate network to the front-end server. All of the port discussions that follow refer to inbound ports carrying traffic from the server in the perimeter network (whether it is a front-end Exchange server or an advanced firewall server) to the back-end servers.

Basic Protocols

In all cases, all the supported protocol ports must be open on the inner firewall. The SSL ports do not need to be open, because SSL is not used in communication between the front-end server and the back-end servers. Table 2 lists the ports required for the intranet firewall. These ports are specific to inbound traffic (from the front-end server to the back-end servers).

Table 2 Protocol ports required for the intranet firewall

Intranet Firewall—Mail Protocols	
Port number/transport	Protocol
80/TCP inbound	HTTP
143/TCP inbound	IMAP
110/TCP inbound	POP
25/TCP inbound 691/TCP	SMTP Link State Algorithm routing

Note In Table 2, “Inbound” means that the firewall should be configured to allow computers in the perimeter network, such as the advanced firewall server, to initiate connections to the front-end server on the corporate network. The front-end server never has to initiate connections to the computers in the perimeter network; the front-end server only responds to connections initiated by the computers in the perimeter network.

Active Directory Communication

To communicate with Active Directory, the Exchange front-end server requires LDAP ports to be open. Both TCP and UDP are required: Windows 2000 on the front-end server will send a 389/UDP LDAP request to a domain controller to check if it is available for use; the LDAP traffic after that uses TCP. Windows 2000 Kerberos authentication is also used; therefore, the Kerberos ports must also be open. Both TCP and UDP are required for Kerberos as well: Windows uses UDP/88 by default, but when the data is larger than the maximum packet size for UDP, it uses TCP. Table 3 lists the ports required for communicating with Active Directory.

Table 3 Ports required for Active Directory communication and Kerberos

Intranet Firewall—Active Directory Communication	
Port number/transport	Protocol
389/TCP	LDAP to Directory Service
389/UDP	
3268/TCP	LDAP to Global Catalog Server
88/TCP	Kerberos authentication
88/UDP	

There are two sets of optional ports that can be opened in the firewall. The decision to open them depends on the policies of the corporation. Each decision involves tradeoffs in the areas of security, ease of administration, and functionality.

Domain Name Service (DNS)

The front-end server needs access to a DNS server to correctly look up server names (for example, to convert server names to IP addresses) Table 4 lists the ports required for access.

If you do not want to open these ports, you must install a DNS server on the front-end server and enter the appropriate name to IP mappings for all of the servers it might need to contact. If you choose to install a DNS server, it is essential to keep these mappings up to date when changes are made to the organization.

Table 4 Ports required for access to DNS server

Intranet Firewall—DNS	
Port number/transport	Protocol
53/TCP	DNS Lookup
53/UDP	

Note Most services use UDP for DNS lookups and only use TCP when the query is larger than the maximum packet size. The Exchange 2000 SMTP service, however, uses TCP by default for DNS lookups. For more information, see Microsoft Knowledge Base article Q263237, "XCON: Windows 2000 and Exchange 2000 SMTP Use TCP DNS Queries" (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=263237>).

IPSec

Table 5 lists the necessary requirements for allowing IPSec traffic across the intranet firewall. You only need to allow the port that applies to the protocol you configure; for example, if you choose to use ESP, it is only necessary to allow IP protocol 50 across the firewall.

Table 5 Ports required for IPSec

Intranet Firewall—IPSec	
Port number/transport	Protocol
IP protocol 51	Authentication Header (AH)
IP protocol 50	Encapsulating Security Payload (ESP)
500/UDP	Internet Key Exchange (IKE)
88/TCP	Kerberos ()
88/UDP	

Remote Procedure Calls (RPCs)

DSAccess no longer uses RPCs to perform Active Directory service discovery. However, if your front-end server is configured to authenticate requests, IIS must still have RPC access to Active Directory in order to authenticate the requests. Therefore, you must open the RPC ports that are listed in Table 6.

Filtering RPC Traffic with ISA Server

If you have a perimeter network, you can use ISA Server as the internal firewall and configure it to filter RPC traffic. You can configure ISA Server to authorize access to only specific RPC interfaces and block all others. This feature allows you to open only a narrow range of RPC ports.

Stopping RPC Traffic

If you have a locked-down perimeter network in which it is impossible for the front-end server to authenticate users, you may not be allowed to open the RPC ports that are listed in Table 6. Without these RPC ports, the front-end server cannot perform authentication. You can configure the front-end server to allow anonymous access, but you should understand the risks of doing so. For more information, see “Authentication Issues for HTTP” in Chapter 1.

Rather than stopping all RPC traffic, it is recommended that you restrict RPC traffic by opening one port (as described in the next section).

Restricting RPC Traffic

If you want the features that require RPCs, such as authentication or implicit logon, but do not want to open the wide range of ports above 1024, you can configure your domain controllers, global catalog servers, and all other back-end servers to use a single known port for all RPC traffic. For more information about how to restrict RPC traffic, see Microsoft Knowledge Base article Q224196, “Restricting Active Directory Replication Traffic to a Specific Port” (<http://go.microsoft.com/fwlink/?LinkID=3052&ID=224196>).

In order to authenticate clients, the registry key (described in the above Knowledge Base article and listed below) must be set on any server that the front-end server may contact with RPCs such as a global catalog server. Set the following registry key to a specific port, such as 1600:

HKEY_LOCAL_MACHINE\CurrentControlSet\Services\NTDS\Parameters

Registry Value: TCP/IP Port

Value Type: REG_DWORD

Value Data: (available port)

On the firewall between the perimeter network and your intranet, you need to open only two ports for RPC communication—the RPC portmapper (135) and the port you specify (port 1600, as listed in Table 6). The front-end server first attempts to contact back-end servers with RPCs over port 135, and the back-end server responds with the RPC port it is actually using.

Note Exchange System Administrator uses RPCs to administer Exchange servers. It is recommended that you do not use Exchange System Administrator on a front-end server to administer back-end servers because this requires configuring RPC access from the front-end to each back-end server. Instead, you should use Exchange System Administrator from an Exchange client computer or a back-end server to administer back-end servers. You can still use Exchange System Administrator on the front-end server to administer the front-end server itself.

Table 6 RPC ports needed for authentication

Intranet Firewall—RPCs: Service Discovery and Authentication	
Port number/transport	Protocol
135/TCP	RPC port endpoint mapper
1024+/TCP or 1600/TCP	All service ports (Example) RPC service port, if restricted

7

Front-End and Back-End Topology Checklist

The following checklist summarizes the steps required to configure front-end servers, back-end servers, and firewalls.

Warning The following procedures contain information about editing your registry. Using Registry Editor incorrectly can cause serious problems that may require you to reinstall your operating system. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. For information about how to edit the registry, view the "Change Keys and Values" Help topic in Registry Editor (Regedit.exe) or the "Add and Delete Information in the Registry" and "Edit Registry Information" Help topics in Regedt32.exe. Note that you should back up the registry before you edit it. You should also update your Emergency Repair Disk (ERD).

Table 7 Front-end and back-end topology checklist

Configure Front-End Server(s)	
<input checked="" type="checkbox"/>	Step 1. Install Exchange: Install Exchange on the front-end server. Designate the server as front-end in the server properties. Restart the server, or stop and restart the necessary services

Table 7 Front-end and back-end topology checklist (continued)

Configure Front-End Server(s)	
<input checked="" type="checkbox"/>	<p>Step 2. Configure HTTP virtual servers or directories on the front-end server for access to mailbox and public stores as necessary:</p> <p>For additional virtual servers, specify the SMTP domain, IP address, and host headers or ports. Leave the Basic authentication check box selected.</p> <p>For additional virtual directories for public stores, specify the appropriate public store root.</p> <p>For additional virtual directories for mailbox stores, specify the SMTP domain.</p>
<input checked="" type="checkbox"/>	<p>Step 3. Disable unnecessary services:</p> <p>Stop any services that are not required for the protocols you are using.</p>
<input checked="" type="checkbox"/>	<p>Step 4. Important Dismount and delete the public folder store.</p>
<input checked="" type="checkbox"/>	<p>Step 5. Dismount and delete the mailbox store if necessary:</p> <p>If you are not running SMTP, dismount and delete all mailbox stores.</p> <p>If you are running SMTP, leave a mailbox store mounted, but make sure the mailbox store does not contain any mailboxes.</p>
<input checked="" type="checkbox"/>	<p>Step 6. Set up front-end server load balancing if necessary:</p> <p>Install load balancing on all front-end servers.</p> <p>(Recommended) Enable client affinity.</p>
<input checked="" type="checkbox"/>	<p>Step 7. Configure SSL (recommended):</p> <p>Option 1: Configure SSL on the front-end server.</p> <p>Option 2: Set up a server between the client and the front-end server to offload SSL decryption.</p>

Table 7 Front-end and back-end topology checklist (continued)

Configure Front-End Server(s)	
<input checked="" type="checkbox"/>	<p>Step 8. If you use a perimeter network and do not want to allow RPCs across the intranet firewall: Disable authentication on the front-end server.</p> <p>Note If you disable authentication on the front-end server, you allow anonymous requests to reach your back-end servers.</p> <p>Create the DisableNetlogonCheck registry key and set the REG DWORD value to 1</p> <p>Create the LdapKeepAliveSecs registry key and set the REG DWORD value to 0</p> <p>Note It is recommended that you use an advanced firewall server (such as ISA Server) rather than the front-end server in the perimeter network. For more information, see “Advanced Firewall in a Perimeter Network” in Chapter 3.</p>
<input checked="" type="checkbox"/>	Step 9. If required, create an IPSec policy on the front-end servers.
Configure Back-End Servers	
<input checked="" type="checkbox"/>	<p>Create and configure HTTP virtual servers or directories to match the front-end: For additional virtual servers, set the host headers and IP addresses as appropriate. The TCP port must be left at 80. Make sure the Basic authentication and Integrated Windows Authentication check boxes are both selected.</p> <p>For additional virtual directories for public folder stores, specify the appropriate public folder store root, to match the root configured on the front-end server.</p> <p>For additional virtual directories for mailbox stores, specify the SMTP domain.</p>
Configure Firewalls	
<input checked="" type="checkbox"/>	<p>Step 1. Configure the Internet firewall (between the Internet and the front-end servers):</p> <p>Open TCP ports on the Internet firewall for the mail protocols:</p> <p>443 for HTTPS</p> <p>993 for SSL-enabled IMAP</p> <p>995 for SSL-enabled POP</p> <p>25 for SMTP (including TLS)</p>

Table 7 Front-end and back-end topology checklist (continued)

Configure Firewalls (continued)	
<input checked="" type="checkbox"/>	<p>Step 2. (continued) If using ISA Server, configure as follows:</p> <p>Configure a listener for SSL.</p> <p>Create a destination set that contains the external IP address of the ISA server. This destination set will be used in the Web publishing rule.</p> <p>Create a Web publishing rule that redirects requests to the internal front-end server.</p> <p>Create protocol rules to open ports in ISA Server for outgoing traffic.</p> <p>Configure the ISA server for Outlook Web Access (for more information about how to configure an ISA server for Outlook Web Access, see Microsoft Knowledge Base article Q307347, "Secure OWA Publishing Behind ISA Server May Require Custom HTTP Header" (http://go.microsoft.com/fwlink/?LinkID=3052&ID=307347)).</p>
<input checked="" type="checkbox"/>	<p>Step 3. If using a perimeter network, configure the intranet firewall:</p> <p>Open TCP ports on the intranet firewall for the protocols you are using:</p> <p>80 for HTTP</p> <p>143 for IMAP</p> <p>110 for POP</p> <p>25 for SMTP</p> <p>691 for Link State Algorithm routing protocol</p> <p>Open ports for Active Directory Communication:</p> <p>TCP port 389 for LDAP to Directory Service</p> <p>UDP port 389 for LDAP to Directory Service</p> <p>TCP port 3268 for LDAP to Global Catalog Server</p> <p>TCP port 88 for Kerberos authentication</p> <p>UDP port 88 for Kerberos authentication</p> <p>Open the ports required for access to the DNS server:</p> <p>TCP port 53</p> <p>UDP port 53</p> <p>Open the appropriate ports for RPC communication:</p> <p>TCP port 135 - RPC endpoint mapper</p> <p>TCP ports 1024+ - RPC service ports</p> <p>If you use IPSec between the front-end and back-end, open the appropriate ports. If the policy you configure only uses AH, you do not need to allow ESP, and vice versa.</p> <p>UDP port 500 - IKE</p>

Table 7 Front-end and back-end topology checklist (continued)

Configure Firewalls (continued)	
<input checked="" type="checkbox"/>	<p>IP protocol 51 – AH IP protocol 50 – ESP UDP port 88 and TCP port 88 – Kerberos (Optional) If you want to limit RPCs across the intranet firewall, edit the registry on servers in the intranet to limit RPC traffic to a specific port. Then open the appropriate ports on the internal firewall: TCP port 135 – RPC endpoint mapper TCP port 1600 (example) – RPC service port</p>

8

Front-End and Back-End Topology Troubleshooting Steps

Problems experienced with front-end and back-end architectures are often caused by the inability of network traffic to flow from the front-end server to the correct back-end servers because of incorrect configurations on the server or the network routers. In all cases, event log entries may help troubleshoot the particular issue. When you troubleshoot reported or observed problems with a front-end and back-end topology, step through the issues below to see if they might apply to your problem.

Troubleshooting Tools

When troubleshooting problems in a front-end and back-end topology, the following tools can help you:

Network Monitor

Use Network Monitor to monitor the traffic and determine exactly what is happening between the front-end and the other servers. Set up a client to connect to the front-end server and monitor the traffic between the front-end servers and the intranet servers. You can also use Network Monitor to monitor between the client and the front-end server if SSL is not being used.

Event Viewer

Check the event logs on the front-end and back-end servers and any other involved servers (DNS, global catalog, and other servers). There may be entries that provide a clue as to what the problem is.

RPC Ping

To test RPC connectivity between the front-end server and a global catalog or back-end server, use the Rpsings.exe tool. It is available in the support directory of the Exchange CD.

Telnet

Use telnet.exe to attempt to connect directly to the user's back-end server using the port that the mail protocol uses. For example, if Outlook Web Access is not working when you connect to the front-end server, try using telnet from the front-end server to port 80 on the back-end server.

General Troubleshooting Steps

- Make sure that all of the appropriate services are started on the front-end and back-end servers. This includes the relevant Exchange services as well as the World Wide Web Publishing service and SMTP service, if applicable.
- If you have a perimeter network, make sure that the appropriate ports are open on the internal firewall as described in Chapter 6, "Configuring Firewalls."
- Ensure that the front-end server can successfully connect to the global catalog servers and DNS server. This is particularly important when the front-end server is in a perimeter network. Telnet from the front-end server to the appropriate ports on the servers in the intranet—389, 3268, 53, and other ports.

Note Windows 2000 telnet uses TCP/IP and cannot be used to connect to UDP ports.

- If you are unable to connect to the back-end server from the front-end server using the hostname with any protocol, try using the IP address. If this works, then verify you can connect to the DNS server the front-end server is using. Also verify that the name to IP mapping is correct in DNS.
- If the front-end server is configured with the list of domain controllers and global catalog servers in the registry, verify that the front-end can reach each of those servers exactly as specified in the registry entry.
- Make sure the combination of IP address and host header is unique for each virtual server.
- If you have a load balancing solution for the front-end servers, ensure that the shared IP can be reached from client computers.

- Administration: If you want to use Exchange System Manager, ensure that the System Attendant service is running. Also recall that you cannot use the Internet Services Manager after deleting the stores on the front-end server.
- If users complain that the state of read and unread messages in public folders fluctuates, consider the following:
 - Was a back-end public folder server added or removed?
 - Is authentication enabled on the front-end?
 - Are any back-ends that host the folder down?

Logon Failures

If your users have problems logging on to POP, IMAP or Outlook Web Access, consider the following common problems:

- Is the user entering the username in the proper format—*domain\username*, *username@domain.com*, *username*?
- If UPN or a default domain is configured and the user is entering the username in the proper format, verify that the default domain setting is correct on all virtual servers and virtual directories in Exchange System Manager. Verify the same setting in Internet Services Manager. If the domain is correct in Exchange System Manager but not in Internet Services Manager, there is most likely a problem replicating settings from Exchange System Manager to Internet Services Manager. Try restarting the MExchangeSA service to fix this.
- Verify that the host headers for the HTTP virtual server match exactly what the client browser is using to connect to the server. Verify that the host headers are correct and there are no typing mistakes on the back-end and front-end virtual servers and directories.
- If you have multiple virtual servers for multiple domains, make sure that the SMTP domain is properly configured.
- Ensure that the user attempting to log on has an e-mail address for the domain configured on the virtual server the user is accessing.

Troubleshooting Outlook Web Access

For detailed information about troubleshooting Outlook Web Access, see the technical paper *Troubleshooting Outlook Web Access in Microsoft Exchange 2000 Server* (<http://go.microsoft.com/fwlink/?linkid=9634&clcid=0x409>).

Additional Resources

The following technical papers, Microsoft Knowledge Base articles, and other resources provide valuable information regarding Exchange 2000 front-end and back-end topology.

Technical Papers

Front-end and back-end scalability

Microsoft Exchange 2000 Front-End Server and SMTP Gateway Hardware Scalability Guide

(<http://go.microsoft.com/fwlink/?Linkid=1713>)

Microsoft Exchange 2000 Server Back-End Mailbox Scalability

(<http://go.microsoft.com/fwlink/?Linkid=1711>)

IPSec

IP Security for Windows 2000 Server

(<http://go.microsoft.com/fwlink/?LinkId=7197>)

Exchange 2000 clustering

Deploying Exchange 2000 Server Clusters

(<http://go.microsoft.com/fwlink/?LinkId=10193>)

Microsoft Knowledge Base Articles

The following Microsoft Knowledge Base articles are available on the Web at (<http://support.microsoft.com/>):

Differences between Exchange 2000 Server and Exchange 2000 Enterprise Server

Q296614, "XADM: Differences Between Exchange 2000 Standard and Enterprise Versions"

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=296614>)

Configuring ISA Server

Q307347, “Secure OWA Publishing Behind ISA Server May Require Custom HTTP Header”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=307347>)

Configuring UPN logons

Q267936, “XIMS: Directory Service to Metabase Service May Not Replicate the Default Logon Domain for Virtual Servers”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=267936>)

DNS lookups with Exchange 2000 SMTP service

Q263237, “XCON: Windows 2000 and Exchange 2000 SMTP Use TCP DNS Queries”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=263237>)

Directory lookup information

Q250570, “XADM: Directory Service Server Detection and DSAccess Usage”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=250570>)

Restricting RPC traffic to a specific port

Q224196, “Restricting Active Directory Replication Traffic to a Specific Port”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=224196>)

DSAccess registry key

Q320529, “XADM: Using DSAccess in a Perimeter Network Firewall Scenario Requires a Registry Key Setting”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=320529>)

Disabling the NetLogonCheck registry key

Q320228, “The ‘DisableNetLogonCheck’ Registry Value and How to Use It”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=320228>)

Enabling DSProxy on the front-end server

Q319175, “XADM: You Cannot Perform a Check Names Query Against a Front-End Exchange Computer”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=319175>)

Allowing only internal users to access Outlook Web Access

Q257891, “XWEB: ‘The Page Could Not Be Found’ Error Message When You Use OWA”

(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=257891>)

Redirecting clients to use SSL

Q279681, “How to Force SSL Encryption for an Outlook Web Access 2000 Client”
(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=279681>)

Configuring the Change Password feature

Q327134, “XCCC: How the Change Password Feature Works in Outlook Web Access”
(<http://go.microsoft.com/fwlink/?LinkID=3052&ID=327134>)

Other Useful Resources

Microsoft Exchange Server Web Site

(<http://www.microsoft.com/exchange/>)

Security Services Web site

(<http://go.microsoft.com/fwlink/?Linkid=1706>)

Internet Standard Documents

HTTP 1.1: RFC 2616 (www.ietf.org/rfc/rfc2616.txt)

POP: RFC 1939 (www.ietf.org/rfc/rfc1939.txt)

IMAP: RFC 2060 (www.ietf.org/rfc/rfc2060.txt)

IMAP Referrals: RFC 2221 (<http://www.ietf.org/rfc/rfc2221.txt>) and RFC 2193
(<http://www.ietf.org/rfc/rfc2193>)

Server Sizing Calculator

(<http://go.microsoft.com/fwlink/?Linkid=1716>)

Microsoft Application Center Server 2000

(<http://go.microsoft.com/fwlink/?Linkid=591>)

For more information: (<http://www.microsoft.com/exchange/>).

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: Microsoft Exchange Server Front-End and Back-End Topology>