



Journaling with Exchange Server 2003

Product Version:
Reviewed By:
Latest Content:
Author:

Exchange Server 2003
Exchange Product Development
www.microsoft.com/exchange/library
John Speare



Microsoft



Journaling with Exchange Server 2003

John Speare

Published: May 2004

Applies to: Exchange Server 2003

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.

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.

©2004 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows NT, Windows Server, 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.

Acknowledgments

Project Editor: Cathy Anderson

Contributing Editors: Janet Lowen, Cathy Anderson

Technical Reviewers: Genevieve Orchard; Dan Longley; Max Ciccotosto; Elias Kaplan; Chandresh Jain; Ramon Infante; James Baker; Ross Smith IV; Bruce McKinstry

Graphic Design: Kristie Smith

Production: Sean Pohrilla

Table of Contents

- Introduction 1
 - What Is Covered in This Guide? 1
 - What Is Not Covered in This Guide?..... 1
 - Who Should Read This guide?..... 1
- What Is Journaling?..... 2
- Why Journal? 2
- Compliance Solution Framework 2
- Overview of Exchange Journaling..... 3
 - Types of Journaling 3
 - Where Journaling Does Not Work 4
 - Journal Data Format 4
- How Exchange Envelope Journaling Works 4
 - Example: Journaling for a Single Mailbox Database 5
 - Example: Journaling for Two Mailbox Databases..... 6
 - Example: Journaling for Two Mailbox Databases with Distribution List Expansion 7
 - Conclusions About How Envelope Journaling Works 10
- Planning an Exchange Journaling Deployment 11
 - Journaling Impact on User Mailbox Servers 11
 - Planning for Journal Recipient Mailbox Servers..... 12
 - Active Directory 13
- Deploying Journaling as Part of a Compliance Solution 13
 - Creating a Custom SMTP Recipient 14
 - Creating the Journal Recipient Mailboxes 15
 - Setting a Server-Side Rule for Journal Recipient Mailboxes 15
 - Configuring Mailbox Manager to Clean the Journal Recipient Mailbox 16
 - Configuring the SMTP Connector 17
 - Enabling Journaling..... 17
- Appendix 19
 - Resources Cited in This Book..... 19
 - Microsoft Knowledge Base Articles..... 19
 - Accessibility 19

Introduction

Regulatory compliance is an important issue for most companies in the financial and health care industries. For example, corporate officers in some financial sectors are responsible for the claims made by their employees to their customers. In the context of an enterprise messaging environment where large amounts of information are exchanged on a daily basis, supervising data is a complex but required task. However, e-mail is only one type of data that is regulated and is, therefore, one data source that a full-compliance solution has to manage. Microsoft® Exchange Server 2003 Service Pack 1 (SP1) (and Exchange 2000 Server SP3, with the envelope journaling software update) includes enhancements to the Journaling feature that help make messaging data more manageable for a broader compliance solution.

What Is Covered in This Guide?

This guide describes how journaling works in Exchange 2003 SP1 and in Exchange 2000 SP3 with the envelope journaling software update. This guide describes how to make messaging data available for a compliance solution, taking into account performance, scalability, and management. Specifically, this document answers the following questions:

- What is Exchange journaling?
 - Why is journaling required in some industries?
 - What is a compliance solution framework, and how does Exchange journaling fit in that framework?
 - What is Exchange envelope journaling, and how does it work?
 - How does enabling Exchange journaling affect my messaging environment?
 - How many more servers will I need to enable Exchange journaling in my environment?
 - How do I enable journaling?
 - How do I manage journaling in my messaging environment?
-

What Is Not Covered in This Guide?

This guide is intended to provide a broad introduction to the journaling capabilities of Exchange 2003 and how to enable journaling to fit in your overall compliance solution framework.

The following topics are not covered in this guide.

- Using Exchange journaling to provide an end-to-end journaling solution in small organizations. This guide is written primarily for enterprise deployments where a third-party compliance solution service manages journal data.
 - Developing supplementary software to create an end-to-end Exchange journaling solution for large enterprises.
 - Journaling across Exchange organizations.
 - Journaling for Exchange 2000 Instant Messaging Service.
 - Journaling for Microsoft Exchange 2000 Chat Service.
-

Who Should Read This guide?

This guide is written for the messaging architect and implementer. If you are responsible for implementing a data journaling solution in your organization, you should read this guide.

What Is Journaling?

Journaling is the ability to record all communications in an organization. E-mail communications are one of many different communication mechanisms that you may be required to journal. Therefore, journaling in Exchange has been developed to enable the messaging administrator to feed messaging data into a larger journaling solution, while using minimum overhead.

It is important to understand the difference between journaling and archiving. As stated earlier, journaling is the ability to record all communications; alternatively, archiving refers to reducing the strain of storing data by backing it up, removing it from its native environment, and storing it elsewhere. That said, you may use Exchange journaling as a tool in your e-mail retention or archival strategy.

Why Journal?

Because of new regulations, many organizations in the financial services, insurance, and healthcare industries must maintain records of communication that occur when employees perform daily business tasks.

Although journaling may not be required by a specific regulation, the terms of a regulation may force journaling as one way to comply. For example, corporate officers in some financial sectors are responsible for the claims made by their employees to their customers. To verify that the claims are accurate, the officer may set up a system where managers review some part of employee-to-client communications regularly. Every quarter, the managers, after verifying compliancy, approve their employees' conduct. After all managers report approval to the corporate officer, the corporate officer reports compliancy, on behalf of the company, to the regulating body. In this example, e-mail might be one of the employee-to-client communications that managers must review; therefore, all e-mail sent by client-facing employees is journalized. Other client communication mechanisms may include faxes and telephone conversations, which also must be recorded. Therefore, the ability to journal all classes of data in an enterprise is an important piece of the IT architecture.

The following is a list of some of the more well-known U.S. regulations with requirements that may rely on journaling technology. For more information about these regulations, see *Supporting Regulatory Compliance with Exchange Server 2003* (<http://go.microsoft.com/fwlink/?LinkId=29257>).

- Sarbanes-Oxley Act
 - SEC Rule 17A-4
 - NASD 3110 and 3111
 - Gramm-Leach-Bliley Act (Financial Institution Privacy Protection Act of 2001, Financial Institution Privacy Protection Act of 2003)
 - Healthcare Insurance Portability and Accountability Act of 1996 (HIPAA)
 - Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act of 2001 (Patriot Act)
-

Compliance Solution Framework

Making sure that e-mail in your company is journalized is a broad requirement. How you comply with this request is defined by the specific regulation that your company must follow and the advice of your legal counsel. For example, complying with a regulation that requires supervision of employee claims differs from a regulation that requires retention of all company officer communications. In either case, you must understand how your messaging data fits in your company's broader journaling solution.

Compliance solution framework refers to the IT infrastructure that your organization has implemented to comply with the specific regulations that apply to your industry. This framework includes journaling of specific data classes, such as document storage solutions, messaging systems, fax systems, and telephone communications. Another piece of this framework is storing and managing the data after it is captured. Finally, the framework may need to provide for searching, sorting, and otherwise manipulating the stored data.

Because of each company's unique criteria, including the existing IT infrastructure, the organization of the enterprise, and the regulation with which your company must comply, the compliance solution framework is a custom solution. For that reason, Exchange journaling is flexible in how the messaging data is delivered into the compliance solution framework.

There are many vendors and third-party partners that Microsoft works with to provide a full journaling solution. For more information about these vendors, see *Supporting Regulatory Compliance with Exchange Server 2003* (<http://go.microsoft.com/fwlink/?LinkId=29257>).

Overview of Exchange Journaling

As regulations for recording business communication have evolved, so has the journaling feature in Exchange. This section briefly describes the different types of Exchange journaling and explains some of the messaging data that is not journalized by Exchange. Journal data (message) format is also discussed.

Types of Journaling

There are three different types of journaling that you can enable in Exchange Server 2003.

- **Message-only journaling** Message-only journaling creates a copy of all messages and the corresponding P2 message header data to and from users on a mailbox database and sends the message copy to a specified mailbox. The P2 message header contains only the message recipient data that the sender declared to the recipients. If an external message is received from the Internet, Exchange journals the P1 message headers. The P1 message header is the address information that is used by message transfer agents (MTAs) to route mail. By default, when message-only journaling is enabled, Exchange does not account for blind carbon copy (Bcc) recipients, recipients from transport forwarding rules, or recipients from distribution group expansions.
- **Bcc journaling** Bcc journaling is message-only journaling with the added ability to capture the Bcc recipients. When Bcc journaling is enabled, Exchange captures all recipients (including Bcc recipients) that are known at the originating server. If this recipient list includes hidden distribution lists, query-based distribution lists, or distribution lists that are expanded on another server, the recipients for these lists will not be included in the journalized mail. This functionality is enabled by setting a registry key. For more information about setting this registry key, see Microsoft Knowledge Base article 810999, "XADM: Bcc Information Is Lost for Journalized Messages in Exchange 2000" (<http://go.microsoft.com/fwlink/?LinkId=3052&kbid=810999>).
- **Envelope journaling** Envelope journaling differs from message-only journaling and Bcc journaling because it permits you to archive transport envelope information (P1 message headers). This includes information about the recipients who actually received the message, including Bcc recipients and recipients from distribution groups. Envelope journaling delivers messages that are flagged to be archived by using an envelope message that contains a journal report together with the original message. The original message is delivered as an attachment. The body of the journal report contains the transport envelope data of the archived message.

Although three different journaling methods exist, the majority of regulations that require journaling will likely require envelope journaling for compliance. Therefore, unless specifically noted, all discussions about journaling in this guide refer to envelope journaling in an Exchange Server 2003 environment (or Exchange 2000 SP3 with the envelope journaling software update).

Where Journaling Does Not Work

Exchange does not journal the following scenarios and data-types:

- **Posts to public folders** Journaling cannot be enabled on public folder stores.
- **Mail sent to external contacts** External contacts (users or distribution lists) cannot be journalized. Exchange journals a record of the sender's mail that lists the external contacts, but if a distribution list external to the sending Exchange organization is listed, the recipients on the external distribution list will not be recorded as recipients.

Journal Data Format

Journaling is enabled at the mailbox store level. To enable journaling, you must enter a mailbox where the journalized messages are sent. When the message is delivered to the journal recipient mailbox and journalized, the format of the message is MAPI. Depending on the requirements of your compliance solution framework, MAPI format may be acceptable. However, most of the time, Multipurpose Internet Mail Extensions (MIME) is the preferred format, because it is standardized, widely understood, structured, and able to be streamed.

Sometimes, when a third party provides the storing and sorting functions of the compliance solution framework, sending Exchange data in the MIME format over Simple Mail Transfer Protocol (SMTP) is preferred. This is done by forwarding the messages from the journal recipient mailbox to the SMTP address by using a Microsoft Office Outlook® server-side rule.

In other cases, journalized messages can be retrieved from the journal recipient mailbox by using Post Office Protocol version 3 (POP3) or Internet Message Access Protocol version 4rev1 (IMAP4). This also provides a MIME format for the message.

The reason you cannot forward journalized messages directly from the mailbox database where journaling is enabled is because some of the envelope data (for example, Bcc recipients and expanded distribution list recipients) is added by the Exchange Information Store service upon delivery to the journaling mailbox. Therefore, if you journal all mail directly to SMTP, Bcc and expanded distribution list recipient data is lost.

How Exchange Envelope Journaling Works

As mentioned earlier, message-only journaling sends a copy of a message to the journaling mailbox every time a user in a journal-enabled mailbox database sends or receives a message. This is not suitable for most regulations that require journaling because of compliance.

Envelope journaling provides a much more useful service because it records data about all recipients that a message is delivered to. One way to understand how envelope journaling works is in the context of distribution groups. Most distribution lists change, and query-based distribution lists are specifically created based on the fact that lists change. Therefore, just knowing that a message was sent to a specific distribution list is not sufficient to comply with most of the regulations mentioned earlier in this guide. To comply, you must show who actually received a particular message, regardless of whether they were on the Bcc line, or were a member of a distribution group that has since changed.

To provide this level of data, Exchange provides the "envelope" or P1 headers for a particular message. The envelope data is the actual recipient list that Exchange generates as a message is in transit that enables a message to reach all recipients. Unlike the P2 headers (the recipients as declared by the sender in a particular piece of received mail), the envelope data changes as it travels to the many destinations; these changes occur because of distribution list expansion and hidden distribution list behavior.

To enable envelope journaling, you must enable message-only journaling and configure a setting in Active Directory® directory service. Message-only journaling is enabled by selecting a check box on the property page for the mailbox store that you want to journal and specifying a journal recipient mailbox for all journalized data. The setting in Active Directory is a global setting that turns on envelope journaling for the Exchange organization. After you configure the setting in Active Directory, whenever journaling is enabled on a mailbox database, it will be envelope journaling.

The following examples provide more detail about how journaling works in an Exchange organization.

Example: Journaling for a Single Mailbox Database

Figure 1 illustrates how journaling works for a single mailbox database.

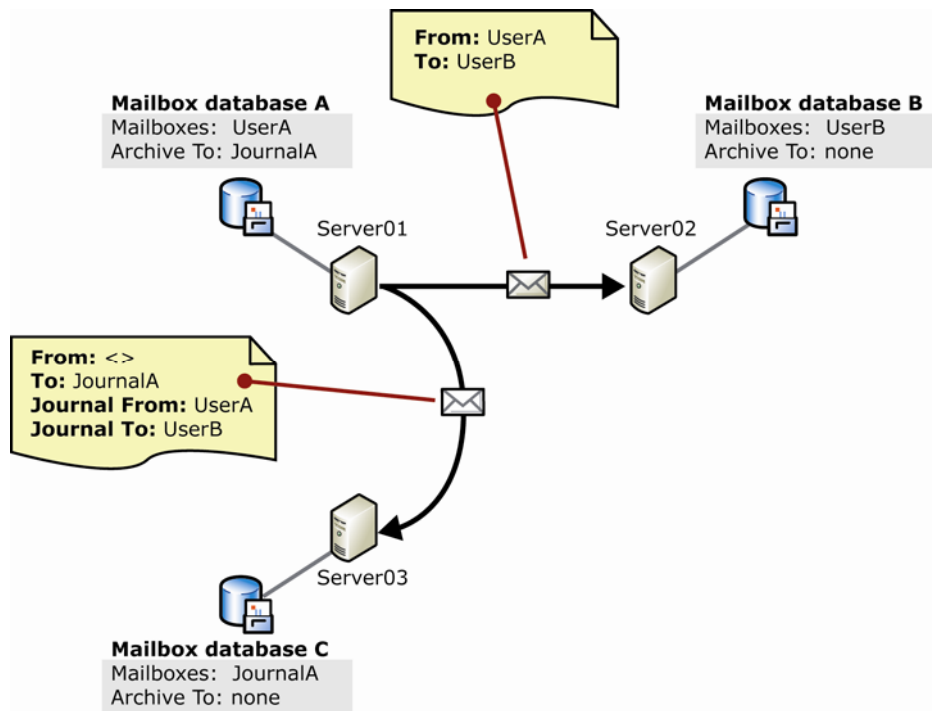


Figure 1 Example of journaling for a single mailbox database example

In this example, UserA has a mailbox on Server01 in Mailbox database A, where journaling is enabled. UserA sends a message to UserB. UserB's mailbox is on Server02, in Mailbox database B, where journaling is not enabled. Server03 is a mailbox server that hosts only one mailbox: JournalA. Server03 is a dedicated journal recipient mailbox server.

Exchange performs several actions before a message leaves the sender's Exchange server. First, Exchange sets a journaling property on the message that identifies it as a journalized message. This property travels with the message to its various destinations in the Exchange organization.

Note

In envelope journaling, all types of messages except journal messages themselves are journalized. This includes delivery status notifications, read receipts, meeting requests, and out of office replies. Message-only journaling does not journal data source names or read receipts.

Next, Exchange adds a list of journal recipient mailboxes to the journaling property. So, in this example, the journaling property contains an entry that specifies the JournalA recipient mailbox with the corresponding recipient, UserA.

Finally, as a normal part of the send operation, Exchange looks up the recipient's address information. Exchange specifically uses the recipient's "home" mailbox database attribute as the destination. If the home mailbox database attribute includes a journaling mailbox address, the recipient also must be journalized, and the recipient's journal mailbox and the recipient information are added to the journal property on the message. In this example, because UserA is the only user who is marked for journaling, no other recipient mailboxes or recipients are added to the journaling property.

When the message is sent to UserB, a message is also sent to JournalA for journaling. When Server02 receives the message, it recognizes the mail as a message that requires journaling, reads the journaling property (which specifies that UserA journal message has been recorded and sent), and delivers the message to UserB's mailbox. When Server03 receives the message, it creates an envelope message that specifies UserA as the sender and UserB as the recipient, attaches the original message to the envelope message, and delivers it to JournalA's mailbox.

Example: Journaling for Two Mailbox Databases

When journaling is enabled for more than one mailbox database and more than one journal recipient mailbox server is used, the message flow becomes more complicated (Figure 2).

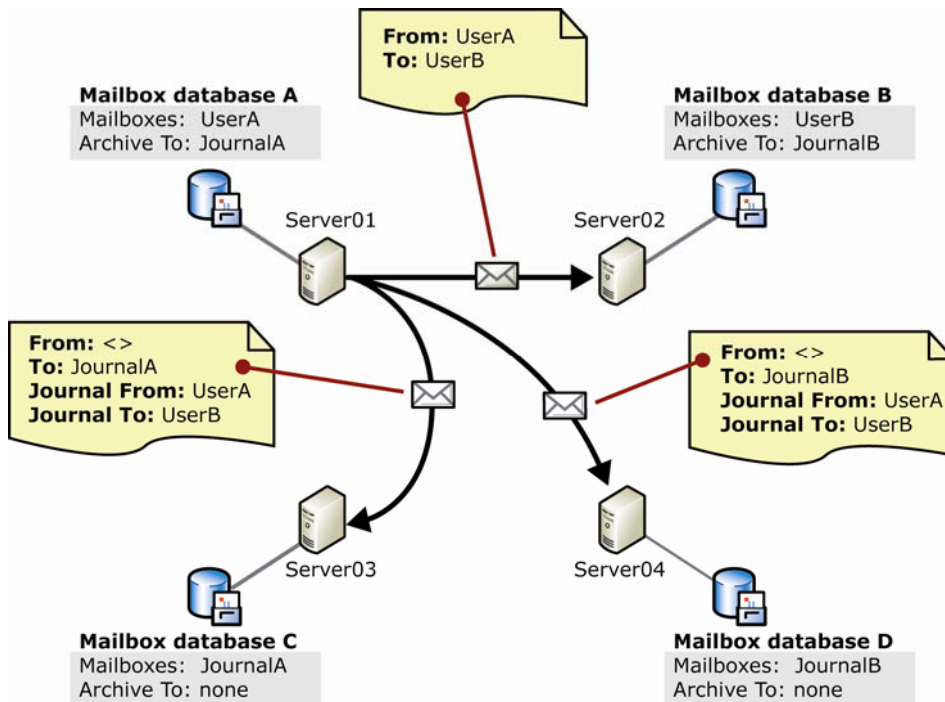


Figure 2 Example of journaling for two mailbox databases

In this example, the mailbox database on Server02 is enabled for journaling. However, the journal recipient mailbox for Mailbox database B is on a new Exchange server, Server04.

Before Exchange sends UserA's message to UserB, it looks up UserB's recipient information in Active Directory. UserB's Mailbox database B is marked as enabled for journaling. Therefore, Exchange lists UserB's journal recipient mailbox in the journaling property on the message together with UserA's journal recipient mailbox.

Therefore, when UserA sends a message to User B, Server01 sends three messages:

- **A message to JournalA on Server03** The journal property on the message specifies UserA as the sender of the attached message to UserB. Server03 attaches this message to the envelope message that it creates for submission to the JournalA mailbox.
- **A message to JournalB on Server04** The journal property on the message specifies UserB as the recipient of the attached message from UserA. Server04 attaches this message to the envelope message that it creates for submission to the JournalB mailbox.
- **A message to UserB on Server02** The message is marked with the journaling property specifying that a journal message has been sent for UserB; therefore, Server02 does not send a journal message for UserB.

Example: Journaling for Two Mailbox Databases with Distribution List Expansion

This example builds on the two earlier examples by adding a distribution list that is expanded on a different server than the server where the message originates (Figure 3). This example also explains other details about how journaling in Exchange works.

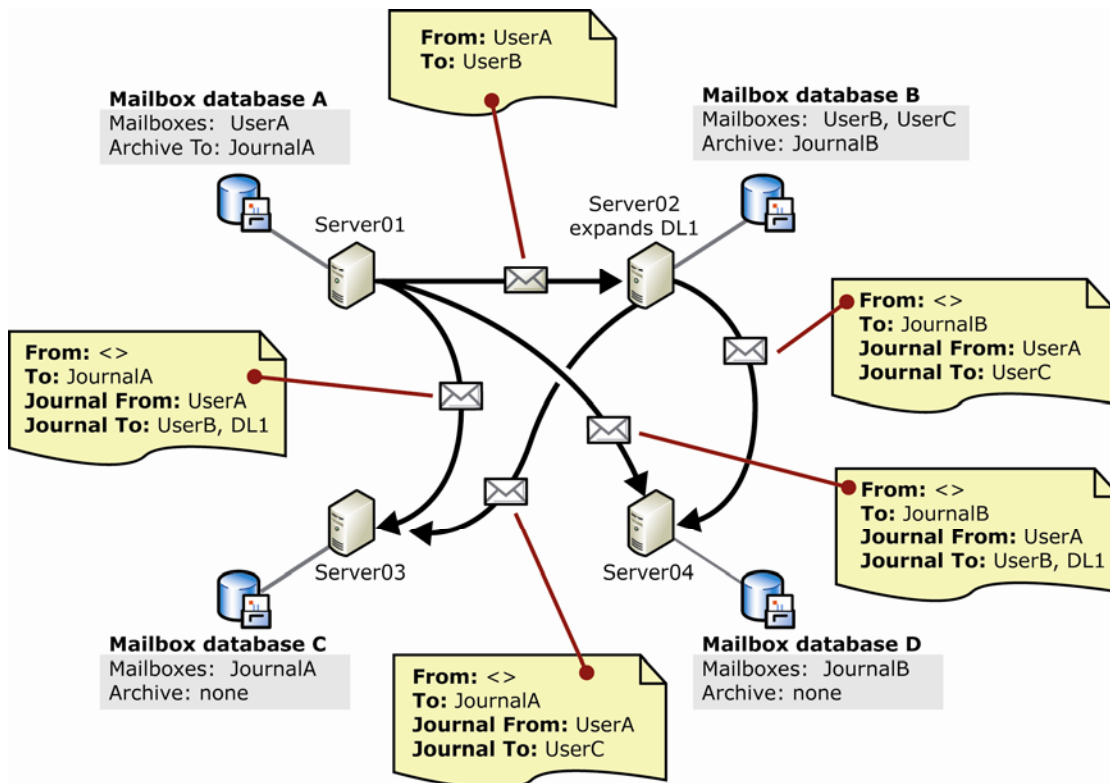


Figure 3 Example of journaling for two mailbox databases with distribution list expansion

In this example, UserA sends a message to UserB and the distribution list, DL1. DL1 contains one member, UserC. All users have mailboxes that are journalized. The expansion server for DL1 is on Server02.

As in the earlier examples, Exchange looks up recipient information in Active Directory before sending the message. UserB is recognized and marked as a journalized recipient in the message journaling property. Because the distribution list is expanded on Server02, Server01 cannot determine who is on the distribution list and therefore cannot determine whether the recipients on the list are journalized.

Server01 sends three messages:

- **A message to JournalA on Server03** The journal property on the message specifies UserA as the sender of the attached message to UserB. Server03 attaches this message to the envelope message that it creates for submission to the JournalA mailbox.
- **A message to JournalB on Server04** The journal property on the message specifies UserB as the recipient of the attached message from UserA. Server04 attaches this message to the envelope message that it creates for submission to the JournalB mailbox.
- **A message to UserB on Server02** The message is marked with the journaling property specifying that a journal message has been sent for UserA and UserB; therefore, Server02 does not send a journal message for UserB.

When Server02 receives the message, the journaling property on the message will indicate that it is a journalized message. As the expansion server for the distribution list, DL1, Server02 sends a message to the JournalA recipient mailbox; the message journal property specifies UserC as a recipient. Additionally, because UserC is on Mailbox database B, which is also journalized, Server02 also sends a message to the JournalB recipient mailbox, where Server04 creates an envelope message that specifies UserC as the recipient of a message from UserA.

This example shows that multiple journal messages are frequently sent for a single message. Each of these envelope messages has a different message ID, but the attached message (the original message being journalized) has the same message ID, which is also added to the body of the envelope message. Additionally, if UserB and UserC were on different mailbox stores on Server02, with different journal recipient mailboxes, Server02 would send two messages to the two different journal recipient mailboxes.

Note

Because the sender is identified in the journal properties on the message, Server02 does not have to calculate this as it would in the case where a message is delivered to a hidden distribution list. (Hidden distribution list expansion causes the P1 sender to be changed every time the message is received by a different server.)

Envelope Message

The journal messages that are sent to the journaling mailboxes are MAPI messages (the body of which is the original message sent by the sender). These MAPI messages contain the journaling property, which also contains all the recipient information for the message. The journaling mailbox requests that the Exchange Store service package an envelope message that contains all this data. The resulting message (Figure 4) is similar to a non-delivery report (NDR) message. The body contains Exchange-generated data (in this case the sender, the original message ID, and a list of recipients), and the original message is attached.

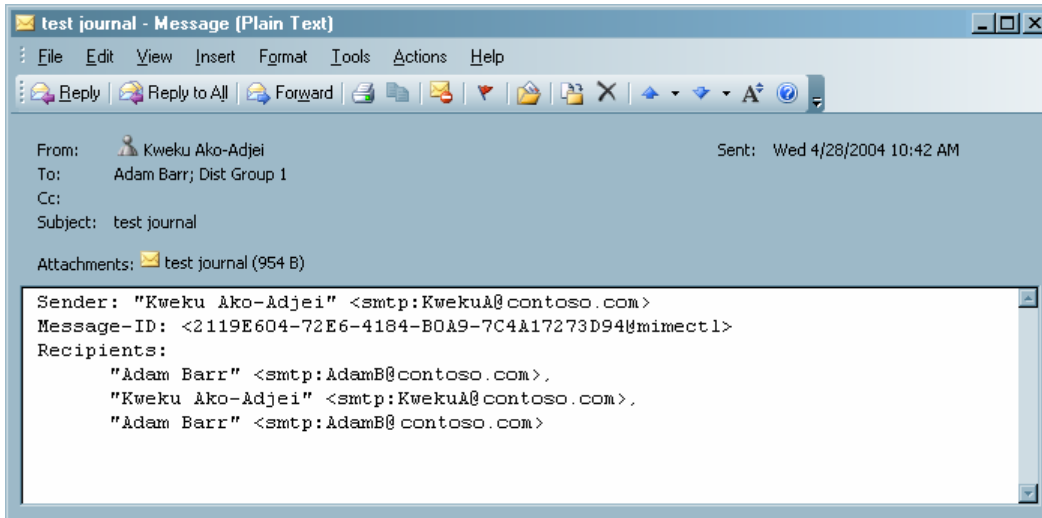


Figure 4 Sample journal message

Until the journal message is created by the Exchange server that hosts the journal recipient mailbox, envelope data is attached as a property to the message instead of copied as embedded content. Attaching the envelope data as a property is done so that existing mailbox servers do not have to manage the additional transaction volume that the Exchange Store service creates when constructing journal reports locally at the point of data gathering. Instead, the mailbox servers that host the journal mailboxes assume the responsibility for constructing journal reports at delivery time.

Recipient Data Gathering

To understand envelope journaling, you must understand how recipient data is gathered and the types of recipients that are gathered. When envelope journaling is enabled, Exchange gathers the P1 recipient data for every recipient that a message will be delivered to. This includes data for:

- All recipients that are added to a message because of a distribution list expansion or another recipient forwarding rule.
- A distribution list recipient that is being forwarded to another server for expansion, including the distribution list recipient itself.
- Recipients that are forwarded to an SMTP virtual server configured as an alternative server. (Recipients that are forwarded to an alternative server are those that would otherwise cause an NDR because of having addresses that are not found in the directory but fall under an authoritative domain from the recipient policy configuration.)

This does not include data for:

- Recipients that cause an NDR.
- Distribution list recipients that were expanded locally and subsequently removed from the recipient list.
- Journal recipients.

Included recipients

- Mailbox and public folder recipients
- Contacts
- One-offs: recipients that are not included in the global address list
- Alternative recipients
- Distribution lists and query-based distribution list members
- Unexpanded distribution lists and query-based distribution lists being forwarded for expansion
- Alternative server recipients

Excluded recipients

- Distribution list and query-based distribution list names. Although excluded from the actual envelope report, if distribution list and query-based distribution list names were entered on the To or Cc line of the original mail, they will be included in the message that is attached to the envelope message. However, if the distribution list or query-based distribution list is configured to expand on a separate server from the sender's server, the envelope message derived from the sender's server will include the name of the distribution list or the query-based distribution list. In this case, a separate envelope message is derived from the expansion server, as described earlier in this guide.
- Journal recipients.

When multiple address types are available, the address is selected in preferential order by type as follows:

1. SMTP
2. X.400
3. Other, or Custom Address
4. LDAP distinguished name
5. Legacy DN

Conclusions About How Envelope Journaling Works

Based on the behavior explained in this guide, there are two significant pieces of functionality that journaling provides, as long as the message is sent within an organization:

- The journal recipient for the sender's mailbox database will receive a complete record of all recipients who receive a copy of a message sent by that sender. The record may arrive through multiple journal reports, but all recipients will be covered. This is because the sender is always at the root of the distribution tree.
- The journal recipient for a recipient's mailbox database will receive at least a partial record that mentions every recipient in that mailbox database for all messages received by recipients in that mailbox database. That record may arrive through multiple separate journal reports.

Because of these pieces of functionality, the archive can answer the following questions:

- **For a particular piece of mail sent by someone in the organization, who received it?**
By checking the journal mailbox for the sender's mailbox database for all journal reports containing that message, you can obtain the complete record of all recipients by aggregating the reported recipients from those reports. This covers all recipients in the organization, in addition to contacts and one-offs outside the organization.
- **For a particular piece of mail received by someone in the organization, who else received it?**
If the message was sent from in the organization, it is quickest to check the sender's mailbox database archive because it will always contain the complete recipient record for sent mail. For mail received from an external sender, the only way to uncover all internal recipients is through an exhaustive search of all journal mailboxes for mailbox databases in the organization. If any recipient received the message, there will be a journal report reflecting that in his or her archive mailbox.

Planning an Exchange Journaling Deployment

Widespread journaling will have an impact on the performance of Exchange. Therefore, it is likely that you will have to deploy more hardware to maintain the current level of messaging service in your organization. Because every organization is different and compliance solution frameworks vary greatly, this section provides only high-level recommendations about performance analysis for server and topology sizing. It is recommended that final decisions be made from information gathered in a lab environment that is as similar as possible to your production environment.

Most of the CPU and disk input/output (I/O) consumed by journaling is the result of an increased burden on the store process on the journalized mailbox databases and on the servers hosting the journal recipient mailboxes.

Therefore, from a performance and server load perspective, there are two classes of server you have to size: the user mailbox servers that you want to journalize and the mailbox servers that host the journal mailboxes.

The following sections about planning discuss how to size these two classes of servers. First however, a brief discussion of how the surrounding infrastructure will be affected is necessary.

The main impact on infrastructure is the increase in network traffic because of enabling journaling. The impact is directly proportional to the number of users who are being journalized and how much mail is sent between mailbox databases in the organization. For example, if only a small part of users are journalized to satisfy the requirements of the Sarbanes-Oxley regulation, network impact will be much less than if an entire division is journalized to satisfy the requirements of the Sec 17a regulation. As for the impact on mail sent between mailbox databases, mail sent from a user to another user on the same journalized mailbox database generates only a single message to a journal recipient mailbox. However, if a journalized user sends mail to a journalized user on a different mailbox database with a different journal recipient, two messages are sent to two different journal recipient mailboxes. Therefore, mailing habits and how you group users on the mailbox databases in your organization affects the overall network load after journaling is enabled.

Journaling Impact on User Mailbox Servers

When journaling is enabled in a mailbox database and a user with a mailbox on that mailbox database sends a message, the server generates two messages: one for the recipients and one for the journal recipient. When a message is submitted to a journalized mailbox database, the mailbox database processes the message as it normally would to deliver it, but it also creates a message for the journaling recipient. When a journalized mailbox database receives a message, most of the time, the message has been journalized already. In the receive case, extra processing (beyond reading the journaling property) is required only when the receiving server is the expansion server for the distribution list or when the distribution list is hidden or query-based.

Therefore, you can estimate the impact of journaling on a mailbox database by assuming that the enabled mailbox database can process approximately half of the messages being sent, as long as all other conditions, such as CPU power, bandwidth, storage space, and disk speed, remain constant. This result is approximately a 15 to 35 percent performance degradation.

Note

This approximation is just a starting figure for planning purposes. Only complete testing in a lab environment that closely resembles your production environment can approximate a more accurate evaluation.

Planning for Journal Recipient Mailbox Servers

For most enterprises that must comply with regulation, hundreds of thousands of messages will be sent to the journal recipient mailboxes on a daily basis. These messages are then likely forwarded, or downloaded to a third-party storage solution offsite. Therefore, it is highly recommended that you house the journal recipient mailboxes on servers that are separate from the regular user mailboxes. This section discusses how to organize the journaling mailboxes and other standard considerations for server sizing.

Organizing Journal Recipient Mailboxes

When it comes to organizing the journal recipient mailboxes, you must answer the following questions:

- How many journal recipient mailboxes are required?
- How much storage is required for each journal recipient mailbox?
- How should the journal recipient mailboxes be configured on the hard disks?

As with all sizing estimates, test estimates in a lab before you implement solutions in your production environment. To estimate how many journaling mailboxes you will require, look at the load on the mailbox servers. In some organizations, mailbox servers are clustered and run at high resource usage levels (80 percent or more). In other organizations, mailbox servers are run below 30 percent load. If your resource usage levels run high, in a lab setting, add a journal recipient mailbox for every three to five mailbox servers. If your resource usage levels run on the lower end, start with one journaling recipient mailbox for every seven to nine mailbox servers.

From an organizational perspective, it may be easier to manage fewer journal recipient mailboxes if you do not plan to hire a third-party to store and organize the data for you. Generally, however, minimizing the number of journal recipient mailboxes is a good practice for reasons already mentioned in this guide (bandwidth, management, performance).

For storage size, there are a number of factors to consider. The list below discusses these factors, but the biggest factor is how much mail is used in the organization. For a simple place to start testing, assume that for each mailbox database that is journalized, you need two to three times the storage space on the corresponding journal recipient mailbox. However, because the messages in the journal mailboxes are largely "transient," meaning that the messages that reside there are quickly forwarded or downloaded and then deleted, you may be able to run closer to a one-to-one storage ratio with the user mailboxes.

Some factors that will affect the storage requirements and performance characteristics of the journal recipient mailbox include:

- **Envelope message overhead** There is a small overhead for each journalized message, because the envelope journal message includes a plain text report with the original message attached. Unless there are thousands of recipients listed for a particular message, the plain text report adds less than 1 kilobyte (KB) of overhead, in addition to the original message.
- **Multiple journalized messages** If expansion servers are used for a distribution list or alternative recipients are added as a message flows through your system, there will be multiple instances of the same journal message that each report different recipients. Calculating the overhead that this produces is difficult because it depends on whether your organization uses expansion servers or hidden distribution lists and the habits users have for sending their mail.
- **Storage group configuration** As much as possible, minimize the number of storage groups per mailbox database.

As for configuration of the hard disk, follow the general recommendations specified in the *Exchange Server 2003 Performance and Scalability Guide* (<http://go.microsoft.com/fwlink/?LinkId=28660>). Additionally, the *Exchange Server 2003 High Availability Guide* (<http://go.microsoft.com/fwlink/?LinkId=21277>) provides detailed server sizing recommendations.

Again, only complete testing in a lab and careful monitoring in production will provide accurate sizing data.

Active Directory

Although there is no performance impact or special considerations about the impact of journaling on Active Directory, it is important to understand how journaling uses Active Directory. As with almost all user and configuration data on which Exchange relies, journaling configuration data is stored in Active Directory. There are two relevant Active Directory attributes:

- **heuristic** This attribute enables envelope journaling at the Exchange organizational level.
- **msExchMessageJournalRecipient** This attribute specifies the journal mailbox for the particular mailbox store.

Both of these attributes are read and cached by Directory Service Access (DSAccess), which is the local directory cache on each Exchange server. DSAccess is updated every 15 minutes; therefore, any configuration change you make to these attributes takes no more than 15 minutes to update on the local Exchange computer (you must also take into account replication for multiple domains).

Heuristic Attribute

The heuristic attribute is an attribute on the Exchange organization name object (Configuration\Services\Microsoft Exchange*Organization Name*). Two values are recognized by Exchange journaling for this attribute: null, and the integer value, 512. If this attribute is not set (null), mailbox databases that are journal-enabled perform message-only journaling. After this value is set (integer value, 512), all mailbox databases in the Exchange organization that are journal-enabled perform envelope journaling. For more information about setting this attribute, see "Enabling Journaling" later in this guide.

msExchMessageJournalRecipient Attribute

The msExchMessageJournalRecipient attribute is set when journaling is enabled on a mailbox database. This attribute is a mailbox database attribute and holds the distinguished name of the journaling mailbox. If the value is null, journaling is not enabled on that mailbox database. For information about enabling a mailbox database for journaling, see "Enabling Journaling" later in this guide.

Deploying Journaling as Part of a Compliance Solution

Until now, this guide discussed what journaling is, how Exchange journaling works, and how to plan for and enable journaling in your organization. This section discusses how to configure other pieces of the messaging system to provide Exchange journalized data in an overall compliance solution framework.

This section assumes that your organization pushes Exchange journalized data to a custom-built or third-party compliance solution framework as shown in Figure 5. It is assumed that the Exchange data is provided to a data storage solution in a structured format, such as MIME.

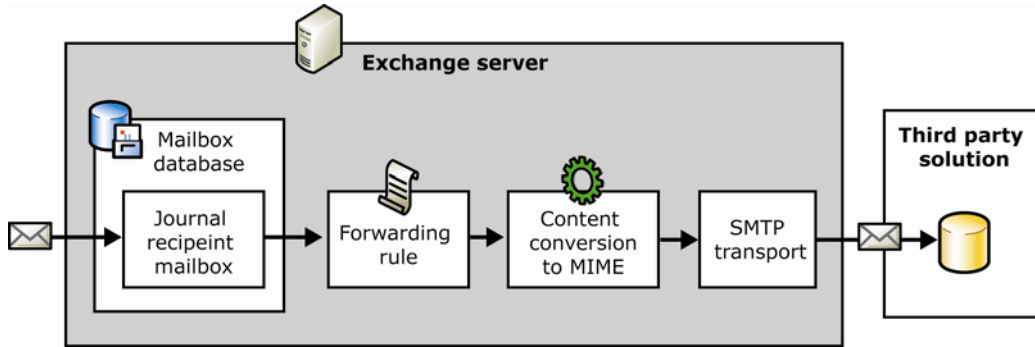


Figure 5 A journaled message is pushed to a third-party compliance storage solution

For large compliance solutions where hundreds of thousands of messages (or more) are journaled daily, it is highly recommended that you hire a third-party service to manage the massive volume of journaling data. A service that specializes in archiving journalized data for the purposes of regulatory compliance will be prepared to receive, parse, sort, and index the journal data by using an SMTP connection from your enterprise.

Providing Exchange journaling data to a third-party service requires extra management and specific configurations. Specifically, you must configure the journaling mailboxes to be sent to the compliance storage solution and then deleted before the local journaling mailbox is full. Special configuration of the server-side forwarding rule and the SMTP connector that provides the transport path to the compliance storage solution is also required to maintain fidelity of the messages.

Note

This section assumes that your organization plans to "push" journalized data to a third-party compliance solution provider over an SMTP connection. Some service providers recommend a "pull" model where POP3, IMAP4, or even a proprietary solution built on CDOEX is used. Both the "push" and "pull" solutions are valid methods; the SMTP "push" method is explained here to illustrate an end-to-end solution.

The following sections discuss the steps that you must take to push journalized data to a compliance storage solution. Perform the following steps after you plan how you will support journaling in your organization:

1. Create a custom SMTP recipient.
2. Create the journal recipient mailboxes.
3. Set a server-side forwarding rule for the journal recipient mailboxes.
4. Configure Exchange Mailbox Manager to clean out the journal recipient mailboxes.
5. Configure an SMTP connector to transmit messages to the compliance storage solution destination.
6. Enable journaling in your organization.

Creating a Custom SMTP Recipient

To send MIME-formatted versions of the journalized messages to a compliance storage solution, you must create at least one contact in Active Directory as a recipient for the journalized data, this contact is set up with an SMTP address and is referred to as the "custom SMTP recipient." This is the recipient that you specify when you set up a server-side rule for the journalized mailboxes in your organization.

Important

It is critical that you do not specify the custom SMTP recipient as the journaling mailbox when you enable journaling on a mailbox database. To record all recipient information, the journaling mailbox must be hosted on an Exchange server. This is because there is relevant journaling information stored in MAPI properties that are discarded when the message is sent over SMTP. The only way to make sure an accurate journal report is sent over SMTP is to forward the journal message from the journal mailbox that is stored on the Exchange server.

Depending on your requirements, you may have to create multiple recipients for different journal mailboxes.

To create a custom SMTP recipient

1. Open **Active Directory Users and Computers** Microsoft Management Console (MMC) snap-in, and then connect to the domain in which you want to create the custom SMTP recipient.
2. Right-click the organizational unit in which you want to create the custom SMTP recipient, point to **New**, and then click **Contact**.
3. In the first **New Object - Contact** dialog box, enter a name and a display name for the custom SMTP recipient, and then click **Next**.
4. In the second **New Object - Contact** dialog box, verify that the **Create an Exchange e-mail address** check box is selected, enter an e-mail alias in the **Alias** box, and then click **Modify**.
5. In the **New E-mail Address** dialog box, select **SMTP Address**, and then click **OK**.
6. In the **Internet Address Properties** dialog box, on the **General** tab, enter the e-mail address of the storage solution where your journaling messages will be stored, click **OK**, and then click **Next**.
7. On the last **New Object - Contact** dialog box, click **Finish**.

Creating the Journal Recipient Mailboxes

After you decide how many journal recipient mailboxes you have to create, you have to set up the journaling server or servers and create the journal mailboxes. For more information about planning for mailboxes, see "Planning for Journal Recipient Mailbox Servers" earlier in this guide. Except for how you organize your journaling mailboxes, there are no other special considerations when setting up the mailboxes on the journaling recipient mailbox servers.

Setting a Server-Side Rule for Journal Recipient Mailboxes

For each journal mailbox that you create, you must create a server-side rule that forwards all messages as they arrive to the third-party storage solution as a MIME message. To do this, you can use Outlook to create a rule that forwards all messages to the custom SMTP recipient that you created. Because the message travels over SMTP, Exchange automatically converts it to MIME for delivery. Alternatively, when you set up the rule, you may decide to redirect the message, which is essentially forwarding the mail, but the redirected mail does not contain the headers of the original mail.

After a message has been forwarded to the third-party storage solution, it must be deleted from the local mailbox store to reduce storage bloat. To do this, you must add an action to the rule that moves the message to the Delete Items folder. After the messages are moved to the Deleted Items folder, you must set up an Exchange recipient policy that Mailbox Manager can run to permanently delete the messages.

To create a server-side rule for journal recipient mailboxes

1. Create an Outlook profile for the journal recipient mailbox for which you will be creating the forward rule.
2. Open Outlook 2003 by using the profile created in Step 1.
3. In Outlook, click **Tools**, and then click **Rules and Alerts**.
4. In the **Rules and Alerts** dialog box, on the **E-mail Rules** tab, click **New Rule**.
5. On the first page of the **Rules Wizard**, select **Start from a blank rule**, and then click **Next**.
6. On the second page of the **Rules Wizard**, do not select any conditions in the **Select condition(s)** section. Instead, click **Next**. A message will prompt you to verify that the rule you are creating is for all messages that are received in this mailbox. Click **Yes**.
7. On the third page of the **Rules Wizard**, in the **Select action(s)** section, select **forward it to people or distribution list**. In the **Edit** section, click **people or distribution list**.

8. In the **Rule Address** dialog box, select the custom SMTP recipient you created earlier, click **To**, and then click **OK**.
9. On the third page of the **Rules Wizard**, in the **Select action(s)** section, select **move it to the specified folder**. In the **Edit** section, select **specified folder**, locate the **Deleted Items** folder, and then click **OK**.
10. On the third page of the **Rules Wizard**, click **Finish**.

Configuring Mailbox Manager to Clean the Journal Recipient Mailbox

After you configure the journal recipient mailbox to forward messages to the third-party storage solution, and then move them to the Deleted Items folder, you have to configure Mailbox Manager to permanently delete the messages.

This section explains how to use an Exchange Recipient Policy together with Mailbox Manager to permanently delete the messages after they have been sent to the third-party storage solution.

It is recommended that you verify through testing and monitoring how long you want messages to remain in the Deleted Items folder before they are permanently deleted. For example, you may want to allow two days before messages are permanently deleted so, if there are any problems with delivering the messages, or any problems receiving or storing them at the destination, you can resend all messages from the last two days.

Note

The age you specify for deleted mail in the recipient policy begins when the message is moved to the Deleted Items folder, not when it was received.

To configure messages to be permanently deleted from the journal recipient mailbox, you must create a recipient policy that identifies the journal recipient mailbox and permanently deletes old items in the Deleted Items folder. Then, you must set Mailbox Manager to run the policy at an appropriate interval.

To create a recipient policy that permanently deletes messages from the Deleted Items folder in the journal recipient mailbox

1. Open **Exchange System Manager** in the Exchange organization where the journaling mailboxes reside.
2. Expand the **Recipients** folder, right-click the **Recipient Policies** folder, point to **New**, and then click **Recipient Policy**.
3. In the **New Policy** dialog box, select **Mailbox Manager Settings**, and then click **OK**.
4. In the **Properties** dialog box, on the **General** tab, enter a **Name** for the policy. Under **Filter rules**, click **Modify**.
5. In the **Find Exchange Recipients** dialog box, construct a search query that will return the journal mailboxes that you have created. When you finish building the query, click **OK**.

Note

There are a number of ways to construct the search query; in large part, the right query depends on how many journaling mailboxes or journaling servers you have configured. The search criteria you enter here will create an LDAP query that will run when Mailbox Manager runs this policy. Make sure to test this query to verify that it does not display other mailboxes when run.

6. In the **Properties** dialog box, on the **Mailbox Manager Settings (Policy)** tab, select **Delete Immediately** from the **When processing a mailbox** drop-down menu.
7. On **Mailbox Manager Settings (Policy)** tab, in the **Folder** list, clear all folders except the **Deleted Items** folder. Select **Deleted Items**, and then click **Edit**.
8. In the **Folder Retention Settings** dialog box, clear the **Message Size (KB)** check box, enter an appropriate age limit in the **Age Limit (Days)** box, and then click **OK**. Click **OK** again to close the dialog box.

9. In the **Recipient Policies** details pane, right-click the policy that you have just created, and then click **Apply this policy now**. A message prompts you to confirm the update; click **Yes**.

After you create a policy, you must schedule when the policy will run on the journaling mailbox server. You must configure this schedule for each server that hosts journal recipient mailboxes.

To schedule Mailbox Manager to run recipient policies

1. Open **Exchange System Manager** and locate the server that is hosting your journal recipient mailbox.
2. Right-click the journal recipient mailbox server and then click **Properties**.
3. In the server **Properties** dialog box, click the **Mailbox Management** tab, use the drop-down menu to set a mailbox management schedule, and then click **OK**.

Configuring the SMTP Connector

The connection between your organization and the third-party compliance solution provider is an SMTP connection. Although there are no settings strictly recommended for configuring the SMTP connector, this section provides some considerations for setting up the SMTP connection.

Specifically, it is recommended that you have a dedicated connector for the target custom recipient addresses because this allows better control of the mail flow from your organization to the third-party service. Additionally, if you plan a large journaling deployment, it is recommended that you have a dedicated connector server. If two target bridgehead servers are used for the SMTP connector, load-balancing with failover is also recommended for maximum reliability and performance. For more information about setting up and load-balancing an SMTP connector, see the *Exchange Server 2003 Transport and Routing Guide* (<http://go.microsoft.com/fwlink/?LinkId=26041>).

Enabling Journaling

By default, envelope journaling is disabled. Enabling envelope journaling involves two steps:

1. Enable standard journaling in Exchange System Manager.
2. Enable envelope journaling.

Note

If you are running Exchange 2000 and you want to enable envelope journaling, you must install Service Pack 3 and the Exchange 2000 envelope journaling software update.

For more information about installing Exchange 2000 SP3, see *Service Pack 3 for Exchange Server 2000 and Exchange Server 2000 Enterprise Server* (<http://go.microsoft.com/fwlink/?LinkId=17058>).

For more information about the Exchange 2000 envelope journaling software update, see Microsoft Knowledge Base article 834634, "A hotfix is available to enable the Envelope Journaling feature in Exchange 2000 Server" (<http://go.microsoft.com/fwlink/?LinkId=3052&kbid=834634>).

Enabling Standard Journaling

Before you enable envelope journaling, you must enable standard journaling on each mailbox store in your organization for which you want envelope journaling enabled. It is recommended that you designate a dedicated Exchange server as the journaling server. Additionally, if you use a dedicated journaling server, you do not have to enable standard journaling on the server. Enable journaling only on those servers with mailbox stores for which you want to journal.

To enable standard journaling

1. In **Exchange System Manager**, expand **Servers**, expand <your Exchange server>, expand <storage group>, and then right-click the mailbox store.

2. On the **General** tab, select **Archive all messages sent or received by mailboxes on this store**, and then click **Browse** to specify a mailbox as the journaling mailbox. All journalized messages for senders on this mailbox store are sent to the mailbox you specify.

Enabling Envelope Journaling

You can enable envelope journaling by using the `exejcfg` tool or by manually setting the last bit on the Exchange organization heuristic objects. You can run the tool from any server with access to Active Directory, but it is recommended that you run the tool from a domain controller.

The `exejcfg` tool is available in the Exchange Server 2003 SP1 download in the `i386\RTW` directory and can be used in Exchange 2000 or Exchange 2003 environments.

To enable envelope journaling by using the `exejcfg` tool

1. Download and unzip the `exejcfg` tool to a directory of your choice.
2. Open a command prompt.
3. Go to the directory where you installed `exejcfg`.
4. Type the following command to enable envelope journaling:

```
exejcfg -e
```

You can use `exejcfg` with any of the following parameters (Table 1).

Table 1 Parameters for `exejcfg` tool

Parameter	Description
-e	Enables envelope journaling
-d	Disables envelope journaling
-l	Lists the envelope journaling setting—whether envelope journaling is enabled or disabled.
/?	Provides the list of options and a short help. If you run the command with no options, it defaults to this option.

To manually enable envelope journaling

1. Use a directory modification tool of your choice, such as LDP (`ldp.exe`) or ADSIEdit (`AdsiEdit.msc`), to access the domain controller.
2. Browse to the Exchange organization object under Configuration, Services, Microsoft Exchange, Organization name.
3. Set the `heuristics` attribute on the Organization name by adding 512 to the existing value.

Appendix

Resources Cited in This Book

- *Exchange Server 2003 Performance and Scalability Guide*
(<http://go.microsoft.com/fwlink/?LinkId=28660>)
 - *Exchange Server 2003 High Availability Guide*
(<http://go.microsoft.com/fwlink/?LinkId=21277>)
 - *Exchange Server 2003 Transport and Routing Guide*
(<http://go.microsoft.com/fwlink/?LinkId=26041>)
 - Supporting Regulatory Compliance with Exchange Server 2003
(<http://go.microsoft.com/fwlink/?LinkId=29257>)
 - Service Pack 3 for Exchange Server 2000 and Exchange Server 2000 Enterprise Server
(<http://go.microsoft.com/fwlink/?LinkId=17058>)
-

Microsoft Knowledge Base Articles

- 834634, "A hotfix is available to enable the Envelope Journaling feature in Exchange 2000 Server"
(<http://go.microsoft.com/fwlink/?LinkId=3052&kbid=834634>)
 - 810999, "XADM: Bcc Information Is Lost for Journaled Messages in Exchange 2000"
(<http://go.microsoft.com/fwlink/?LinkId=3052&kbid=810999>)
-

Accessibility

For information about accessibility for people with disabilities, see the Microsoft Accessibility Web site (<http://go.microsoft.com/fwlink/?LinkId=21487>).

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

Mail feedback to exchdocs@microsoft.com.

For the latest information about Exchange, see the following Web sites:

- Exchange Product Team technical articles and books
<http://go.microsoft.com/fwlink/?LinkId=21277>
- Exchange Tools and Updates
<http://go.microsoft.com/fwlink/?LinkId=21030>
- Exchange Server Community
<http://go.microsoft.com/fwlink/?LinkId=14927>
- Self-extracting executable containing all Exchange Product Team technical articles and books
<http://go.microsoft.com/fwlink/?LinkId=10687>