

Enterprise Vault™

Introduction and Planning

12.3

Enterprise Vault™: Introduction and Planning

Last updated: 2018-02-25.

Legal Notice

Copyright © 2018 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Enterprise Vault, Compliance Accelerator, and Discovery Accelerator are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on-premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<https://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Before you contact Technical Support, run the Veritas Quick Assist (VQA) tool to make sure that you have satisfied the system requirements that are listed in your product documentation. You can download VQA from the following article on the Veritas Support website:

https://www.veritas.com/support/en_US/vqa

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://www.veritas.com/docs/100040095>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

evdocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<https://www.veritas.com/community>

Contents

Chapter 1	About this guide	10
	Introducing this guide	10
	Where to get more information about Enterprise Vault	11
	Enterprise Vault training modules	13
Chapter 2	Introduction	14
	Overview of Enterprise Vault	14
	How archiving works	15
	How to access items in archives	16
	About classification	18
	About records management	20
	About Compliance Accelerator and Discovery Accelerator	21
	About Enterprise Vault add-ons	21
	How Enterprise Vault works	22
	About single instance storage	26
	About Enterprise Vault indexing	29
	About Index Server groups	33
	About Enterprise Vault Administration Console	36
	About Enterprise Vault sites, Directory, and Directory database	36
	About Enterprise Vault tasks	38
	About Enterprise Vault services	41
	About the Enterprise Vault Outlook Add-In	48
	About IMAP access to Enterprise Vault archives	50
	About the Enterprise Vault Client for Mac OS X	50
	About Microsoft Exchange forms	51
	About OWA Extensions	51
	About the Office Mail App for OWA 2013 and later and Outlook 2013 and later	52
	About Enterprise Vault extensions for Notes	52
	About Enterprise Vault Search	53
	About Enterprise Vault monitoring and reporting	54
	FIPS 140-2 compliance	54

Chapter 3	Enterprise Vault administration	55
	About Enterprise Vault administration	55
	Administration Console configuration of archiving	56
	Administration accounts and roles	56
	How to archive PST file contents	57
	How to archive NSF file contents	57
	How to export archived items	58
	Welcome message and other notifications	58
	About reporting and monitoring in Enterprise Vault	59
	Enterprise Vault Reporting	60
	Report mode	61
	Event and diagnostic logging	61
	System status in the Administration Console	62
	Enterprise Vault Operations Manager	62
	Automatic monitoring of events and performance	63
	Message queue monitoring	63
	Enterprise Vault auditing	64
	How to script management tasks	65
	Checklist of day-to-day management tasks	65
Chapter 4	Exchange Server archiving	67
	About Exchange Server archiving and user mailboxes	67
	Exchange Provisioning tasks	68
	Exchange Mailbox Archiving tasks	68
	Exchange archiving targets	69
	Exchange mailbox policies	70
	Exchange desktop policies	70
	Exchange archiving filters	71
	Exchange Server and journal mailbox archiving	72
	Exchange Server and journal filtering	73
	Compliance Accelerator and Exchange journaling	73
	Types of items to archive with Exchange Server archiving	73
Chapter 5	Exchange Public Folder archiving	75
	Exchange Public Folder tasks, Targets, and Policies	75
	How an Exchange Public Folder task archives	76
	User access to Exchange Public Folder archives	78
Chapter 6	File System Archiving	79
	About File System Archiving	80
	About File archiving policies	81

	About shortcut files with File System Archiving	81
	About setting up File System Archiving	85
	File System Archiving in a clustered environment	87
	The process of File System Archiving	89
	How File System Archiving handles older versions of archived files	91
	How File System Archiving synchronizes permissions	91
	File System Archiving reports	91
	How to restore files with File System Archiving	92
	About FSAUtility	93
	How to back up and scan shortcut files with File System Archiving	94
	Pass-through recall for placeholder shortcuts with File System Archiving	94
	Retention Folders and File System Archiving	95
	FSA Reporting	96
Chapter 7	Archiving Microsoft SharePoint servers	98
	About archiving Microsoft SharePoint servers	98
	How to configure SharePoint archiving	99
	SharePoint archiving tasks	99
	SharePoint archiving targets	100
	SharePoint archiving reports	101
	SharePoint archiving policies	101
	How to access archived SharePoint documents	102
	About Enterprise Vault shortcuts in SharePoint	102
Chapter 8	Domino mailbox archiving	103
	About Domino mailbox archiving and Enterprise Vault	103
	Domino provisioning groups	106
	Domino mailbox archiving tasks	107
	Domino mailbox archiving policies	108
	Domino mailbox archiving retention folders	108
	Domino mailbox archiving desktop policies	109
Chapter 9	Domino Journal archiving	111
	About Domino Journal archiving	111
	Domino Journal archiving policies	111
	Domino journal archiving database considerations	112
	How to set up Domino Journal Archiving	112
	Support for clustered Domino Journal databases	113

Chapter 10	SMTP Archiving	114
	About Enterprise Vault SMTP Archiving	114
	SMTP Archiving components	116
	Security and authentication	118
	Scalability and disaster recovery	119
Chapter 11	Skype for Business Archiving	120
	About Skype for Business Archiving	120
	About Skype for Business Archiving licensing	121
	Skype for Business Archiving overview	121
	Skype for Business Archiving components	122
Chapter 12	Enterprise Vault Accelerators	125
	About the Enterprise Vault Accelerators	125
	Differences between the Enterprise Vault Accelerators	126
	About Compliance Accelerator	126
	Compliance Accelerator components	128
	The Compliance Accelerator client application	129
	Compliance Accelerator configuration data	132
	About Discovery Accelerator	132
	The analytics facility in Discovery Accelerator	133
	Discovery Accelerator components	133
	Discovery Accelerator client application	136
	Discovery Accelerator configuration data	138
Chapter 13	Building in resilience	139
	About Enterprise Vault and VCS	139
	Supported VCS configurations and software	139
	About Enterprise Vault and the VCS GenericService agent	140
	Typical Enterprise Vault configuration in a VCS cluster	140
	About Enterprise Vault and Windows Server Failover Clustering	141
	Supported Windows Server Failover Clustering configurations	141
	Typical Enterprise Vault configuration in a Windows Server failover cluster	142
	About Enterprise Vault building blocks	143
	Building blocks and high availability	145

Chapter 14	Planning component installation	146
	About planning component installation	146
	About valid computer names for Enterprise Vault servers	147
	Prerequisites for Enterprise Vault components when planning installation	147
	Factors to consider when planning deployment of Enterprise Vault components	147
	Enterprise Vault Directory Service installation planning	148
	Where to set up the Enterprise Vault Services and Tasks	148
	How to plan installing Exchange Mailbox Archiving Tasks	153
	How to plan installing Exchange Journaling Tasks	154
	How to plan installing Exchange Public Folder Tasks	154
	How to plan installing Domino Journaling and Mailbox Archiving Tasks	154
	How to plan installing the Move Archive task	155
	How to plan installing the Storage Service	155
	How to plan installing the Indexing Service	155
	How to plan installing the Shopping Service	157
	How to plan installing File System Archiving	157
	How to plan installing SharePoint Archiving	157
	How to plan installing SMTP Archiving	157
	How to plan installing Accelerator Services	158
	Enterprise Vault databases and planning their installation	158
	Vault store groups and vault stores installation planning	160
	Administration Console installation	162
	Installation planning for client components	162
	Installation planning for Outlook Web App (OWA) and RPC over HTTP components	163
Chapter 15	Planning your archiving strategy	164
	About archiving strategies	165
	Where to define default settings for the Enterprise Vault Site	166
	How to allow users flexibility	167
	How to plan the types of items to archive	167
	How to define your archiving policy for user mailboxes	168
	How to plan enabling mailboxes	169
	How to plan controlling the appearance of the desktop	169
	How to plan the archiving policy for journal mailboxes	170
	How to plan the archiving strategy for Exchange public folders	170
	How to plan enabling public folders	171
	How to plan an archiving strategy for FSA	171
	How to plan a strategy for SharePoint archiving	173

How to plan settings for retention categories	173
About retention plans	174
How to plan the automatic deletion of archived items	176
How to plan PST migration	176
How to plan NSF migration	177
How to plan shared archives	177
How to plan vault stores and partitions	177
How to plan handling safety copies	178
How to choose a location for safety copies	179
How to plan single instance storage	179
About Enterprise Vault reports	180
Enterprise Vault Reporting feature	181

About this guide

This chapter includes the following topics:

- [Introducing this guide](#)
- [Where to get more information about Enterprise Vault](#)

Introducing this guide

This book gives an introduction to Veritas Enterprise Vault and associated products and explains how to plan your installation.

The book is designed to be read serially. The first part introduces the features and architecture of the various Enterprise Vault components. This gives you the background information you need to plan and set up your Enterprise Vault system. The second part describes the planning decisions that you need to make to set up Enterprise Vault.

To set up Enterprise Vault, you need a working knowledge of the following products:

- Microsoft Windows Server
- Microsoft SQL Server
- Microsoft Message Queue Server
- Microsoft Outlook
- Internet Information Services (IIS)

If you use Enterprise Vault with Microsoft Exchange Server or Microsoft SharePoint Portal Server, you need a working knowledge of those products.

Where to get more information about Enterprise Vault

[Table 1-1](#) lists the documentation that accompanies Enterprise Vault. This documentation is also available in PDF and HTML format in the [Veritas Documentation Library](#).

Table 1-1 Enterprise Vault documentation set

Document	Comments
Veritas Enterprise Vault Documentation Library	<p>Includes all the following documents in Windows Help (.chm) format so that you can search across them all. It also includes links to the guides in Acrobat (.pdf) format.</p> <p>You can access the library in several ways, including the following:</p> <ul style="list-style-type: none"> ■ In Windows Explorer, browse to the <code>Documentation\language\Administration Guides</code> subfolder of the Enterprise Vault installation folder, and then open the <code>EV_Help.chm</code> file. ■ On the Help menu in the Administration Console, click Help on Enterprise Vault.
<i>Introduction and Planning</i>	Provides an overview of Enterprise Vault functionality.
<i>Deployment Scanner</i>	Describes how to check the required software and settings before you install Enterprise Vault.
<i>Installing and Configuring</i>	Provides detailed information on setting up Enterprise Vault.
<i>Upgrade Instructions</i>	Describes how to upgrade an existing Enterprise Vault installation to the latest version.
<i>Setting up Domino Server Archiving</i>	Describes how to archive items from Domino mail files and journal databases.
<i>Setting up Exchange Server Archiving</i>	Describes how to archive items from Microsoft Exchange user mailboxes, journal mailboxes, and public folders.
<i>Setting up File System Archiving</i>	Describes how to archive files that are held on network file servers.
<i>Setting up IMAP</i>	Describes how to configure IMAP client access to Exchange archives and Internet Mail archives.
<i>Setting up SharePoint Server Archiving</i>	Describes how to archive documents from Microsoft SharePoint servers.

Table 1-1 Enterprise Vault documentation set (*continued*)

Document	Comments
<i>Setting up Skype for Business Archiving</i>	Describes how to archive Skype for Business sessions.
<i>Setting up SMTP Archiving</i>	Describes how to archive SMTP messages from other messaging servers.
<i>Classification using the Microsoft File Classification Infrastructure</i>	Describes how to use the classification engine that is built into recent Windows Server editions to classify all new and existing archived content.
<i>Classification using the Veritas Information Classifier</i>	Describes how to use the Veritas Information Classifier to evaluate all new and archived content against a comprehensive set of industry-standard classification policies. If you are new to classification with Enterprise Vault, we recommend that you use the Veritas Information Classifier rather than the older and less intuitive File Classification Infrastructure engine.
<i>Administrator's Guide</i>	Describes how to perform day-to-day administration procedures.
<i>PowerShell Cmdlets</i>	Describes how to perform various administrative tasks by running the Enterprise Vault PowerShell cmdlets.
<i>Auditing</i>	Describes how to collect auditing information for events on Enterprise Vault servers.
<i>Backup and Recovery</i>	Describes how to implement an effective backup strategy to prevent data loss, and how to provide a means for recovery in the event of a system failure.
<i>Reporting</i>	Describes how to implement Enterprise Vault Reporting, which provides reports on the status of Enterprise Vault servers, archives, and archived items. If you configure FSA Reporting, additional reports are available for file servers and their volumes.
<i>NSF Migration</i>	Describes how to import content from Domino and Notes NSF files into Enterprise Vault archives.
<i>PST Migration</i>	Describes how to migrate content from Outlook PST files into Enterprise Vault archives.
<i>Utilities</i>	Describes Enterprise Vault tools and utilities.

Table 1-1 Enterprise Vault documentation set (*continued*)

Document	Comments
<i>Registry Values</i>	A reference document that lists the registry values with which you can modify many aspects of Enterprise Vault behavior.
<i>Help for Administration Console</i>	The online Help for the Enterprise Vault Administration Console.
Help for Enterprise Vault Operations Manager	The online Help for Enterprise Vault Operations Manager.

For the latest information on supported devices and versions of software, see the Enterprise Vault [Compatibility Charts](#).

Enterprise Vault training modules

Veritas Education Services provides comprehensive training for Enterprise Vault, from basic administration to advanced topics and troubleshooting. Training is available in a variety of formats, including classroom-based and virtual training.

For more information on Enterprise Vault training, curriculum paths, and certification options, see <https://www.veritas.com/services/education-services>.

Introduction

This chapter includes the following topics:

- [Overview of Enterprise Vault](#)
- [How Enterprise Vault works](#)
- [FIPS 140-2 compliance](#)

Overview of Enterprise Vault

Enterprise Vault is a Windows application that enables an organization to store messaging and file system data automatically in centrally-held archives. Using Enterprise Vault clients, users can retrieve selected items easily and quickly when required.

Enterprise Vault can archive any of the following types of data:

- Items in Microsoft Exchange user mailboxes
- Items in Microsoft Exchange journal mailboxes
- Microsoft Exchange Public Folder contents
- Items in Domino mail files
- Items in Domino journal databases
- Files held on network file servers
- Items held on Microsoft SharePoint servers
- Instant Messages and Bloomberg messages
- SMTP messages from other messaging servers

For a full list of the platforms and operating systems from which Enterprise Vault can archive, and the operating systems supported for client access of archived items, see the [Compatibility Charts](#).

How archiving works


At scheduled times, Enterprise Vault archiving processes check target servers for items to archive. Relevant items are then stored in Enterprise Vault archives.

Archived items are indexed by Enterprise Vault to enable fast searching and retrieval. The administrator can set the level of indexing required.

When an item is archived, it is automatically assigned a Retention Category, which defines how long it must be kept. The administrator can define different Retention Categories for different types of data. As Enterprise Vault monitors the archives, it can then delete items when the retention period expires.

Figure 2-1 describes how Enterprise Vault tasks and services store data in archives.

Figure 2-1 Enterprise Vault tasks and services store data in archives



The original items can be replaced with shortcuts to the archived copy. In this way primary storage on user computers and servers can be freed up. From the end users' point of view, they will still be able to access the items as before.

A wide range of configuration options allow the administrator to customize and control the archiving process.

Some examples of configuration options are as follows:

- The time and frequency of archiving runs.
- What is to be archived. Some of the attributes that can be used to define whether a file is to be archived are: where the file is located, file age, file type, size.
- Where particular archived items are to be stored.
- Available Retention Categories.
- Required level of indexing.
- Whether shortcuts are to be created and their contents.
- What access users are to be allowed to have to archived items.

Introduction to PST files and NSF files

In many organizations, PST files have been used to back up Exchange Server mailbox contents. Using Enterprise Vault to archive Exchange Server mailboxes removes the need for these PST files. To ensure that all legacy information is archived, Enterprise Vault includes PST migration tools for importing PST files into Enterprise Vault archives.

Enterprise Vault also includes migration tools to import content from Domino and Notes NSF files.

Benefits of archiving

Automatic archiving and archive management provide the following important benefits:

- It ensures that messages and documents are retained for the period of time required by compliance regulations or company policy
- The size of Exchange Server mailboxes and public folders and Domino mail files is easily controlled without the loss of data
- Primary disk space usage is reduced

How to access items in archives

One of the main drawbacks of conventional archiving applications is the high cost of finding and retrieving items from storage. In Enterprise Vault, item properties and content are indexed to enable fast searching of archived items.

Enterprise Vault can be configured to leave shortcuts to archived items in the original location (Exchange Server mailbox, public folder, PST file, Domino mail file, file system folder or SharePoint server). The shortcuts provide a text or HTML preview of the item. Users can double-click shortcuts to view the original item in its associated application or the contents of the shortcut, depending on how you configure Enterprise Vault. The user can save the item to their local computer or use options in an Enterprise Vault client to restore the item to its original location, depending on permissions, or to a specified location.

With SharePoint archiving, you can restore items from SharePoint document libraries but not social content items.

To enable users to access archives, search for items and manage archived items, the following Enterprise Vault client features are available:

- The Enterprise Vault Outlook Add-In and the Enterprise Vault Client for Mac OS X can be installed on user desktops to enable users to search mailbox archives and manage archived items.

The Enterprise Vault Outlook Add-In includes the Virtual Vault feature. With Virtual Vault enabled, users can access their archives in the Outlook Navigation Pane. To the user, the archive looks similar to their mailbox or personal folders.

- Users can access their Exchange archives and Internet mail archives from IMAP-enabled devices and applications.
- Enterprise Vault components for OWA can be set up on Exchange Servers to enable OWA users to manage archived items in mailboxes and public folders. OWA users do not require the Enterprise Vault Outlook Add-In to be installed on their desktop computers.
MAPI over HTTP and RPC over HTTP connections to Exchange Server mailboxes are also supported by Enterprise Vault. RPC over HTTP users do require the Enterprise Vault Outlook Add-In to be installed on their desktop computers.
- Enterprise Vault extensions for Notes and iNotes clients.
- Various browser-based components that users can launch from within the above mail clients or directly from a web browser, and that let the users browse their archives and search for archived items.

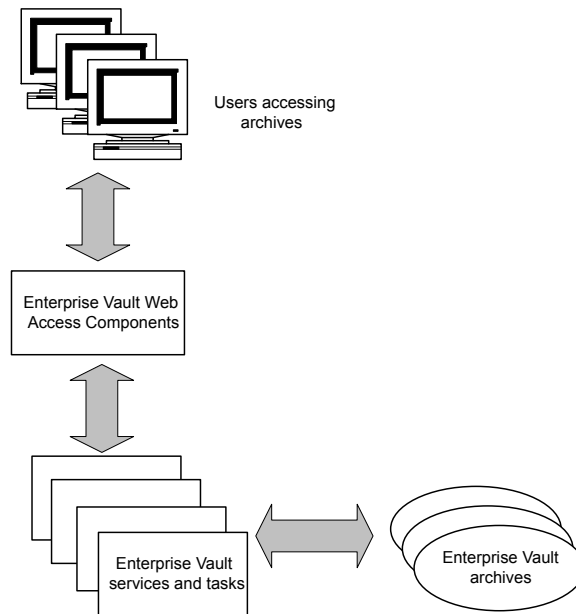
All the client features except the Enterprise Vault Client for Mac OS X are currently available in the following languages:

- | | | |
|------------------------|-------------|-----------|
| ■ Brazilian Portuguese | ■ French | ■ Korean |
| ■ Chinese, Simplified | ■ German | ■ Polish |
| ■ Chinese, Traditional | ■ Hebrew | ■ Russian |
| ■ Danish | ■ Hungarian | ■ Spanish |
| ■ Dutch | ■ Italian | ■ Swedish |
| ■ English | ■ Japanese | |

The Enterprise Vault Client for Mac OS X is currently available in English, French, German, Spanish, and Japanese only.

For a full list of the operating systems on which Enterprise Vault client features are supported, including opening shortcuts to archived items, see the [Enterprise Vault Compatibility Charts](#).

[Figure 2-2](#) gives an overview of the process involved when users access Enterprise Vault archives using the Enterprise Vault clients.

Figure 2-2 How users access stored items

For example, suppose that a user searches for an item stored in an archive. The request goes to the Enterprise Vault tasks and services, which perform the search and return a list of the search results to the user. The user then clicks the link to view the required item. Again, this request is passed to Enterprise Vault tasks and services, which return to the user an HTML version of the item or the item in its original format, depending on how the system has been configured.

In addition to searching and viewing items, users can restore items and, if permitted, delete items from archives.

The administrator can control which functionality is available to users.

About classification

A fundamental element of sound information governance plans, classification lets you analyze information and then make informed decisions about what to retain and what to delete. Through the classification feature in Enterprise Vault, you can do the following:

- Automatically classify all archived content.
- Assign unique retention to individual items.
- Tag items for faster search, discovery, or supervisory review.

- Reclassify an entire archive or specific archived items to adapt to changing retention requirements.

In this release of Enterprise Vault you can choose from two classification engines with which to classify your content:

- **Microsoft File Classification Infrastructure (FCI).** The FCI is a classification framework that is built into recent Windows Server editions and that you control through the File Server Resource Manager (FSRM) interface.

To define the criteria that an item must match to be awarded a specific classification value, you create rules with the FSRM. Enterprise Vault comes with a set of example classification rules, which you can import into the FSRM and use as a starting point to create your own set of rules.

For more information on this feature, see the *Classification using the Microsoft File Classification Infrastructure* guide.

- **Veritas Information Classifier.** This release includes the next generation of our Veritas Information Classifier engine for classifying all new and existing archived content. The new engine provides multifaceted classification through advanced proximity searches, regular expressions, default keywords and patterns, and checksum validation. It can also classify content by confidence level and hit rate.

The Veritas Information Classifier includes over 50 built-in classification policies, which cover many of the data protection regulations and corporate standards worldwide. For example, you can meet privacy regulations such as GDPR through policies to detect Personally Identifiable Information (PII). You can also create custom policies.

For more information on this feature, see the *Classification using the Veritas Information Classifier* guide.

You can use the Veritas Information Classifier in addition to or as an alternative to FCI classification. However, not only are the classification options in the Veritas Information Classifier more sophisticated than those in FCI classification but they are easier to implement. For this reason, the Veritas Information Classifier is the recommended way to classify archived content.

Choosing the classification options that you want to implement

Through the facilities that are available in both the Administration Console and PowerShell cmdlets, you can choose the classification options that you want to implement in your Enterprise Vault site. These options let you choose to do the following:

- Send items for classification and tag them with the results at the same time that Enterprise Vault indexes and archives them. This is also the case if you perform

an index rebuild of an archive or index volume, which causes Enterprise Vault to reclassify the associated items.

- Update the retention category of items when users manually delete them or Enterprise Vault automatically expires them—or optionally when Enterprise Vault indexes and archives the items.

You apply the selected classification options to one or more archives through retention plans.

About records management

The records management feature in Enterprise Vault lets you mark selected items as records, conduct searches for these records, and export them for transfer elsewhere. By doing this you can fulfil the requirements of Capstone, which is the approach to managing emails that the U.S. National Archives and Records Administration has advocated. However, the feature is sufficiently flexible to meet the needs of other records management systems.

By default, Enterprise Vault does not mark the items that it archives as records. Changing this behavior, however, is simply a matter of configuring the retention categories that Enterprise Vault assigns to the archived items. You can configure each retention category so that, at the same time that Enterprise Vault assigns the retention category to an item, it also marks the item as a record of some type. (For Capstone, this is either "permanent" or "temporary".) With a range of different retention categories to assign to items, you can mark each one as the appropriate type of record or non-record.

Enterprise Vault provides you with various ways to assign the required retention category to an item and thereby mark it as a record. For example, you can do the following:

- Create one or more retention plans, each of which is associated with a different retention category. Then you can set up multiple provisioning groups to apply the required retention plans to selected user mailboxes.
- By using facilities such as Enterprise Vault Policy Manager (EVPN), assign different retention categories to different folders in user mailboxes. Then you can allow users to mark their items as records by dragging the items from one mailbox folder and dropping them in another.
- Use the Enterprise Vault classification feature to assign the appropriate retention category to items that match certain criteria.

The facilities in Enterprise Vault Search make it easy to search an archive for the items marked as records. If you want to transfer the records for long-term retention in the U.S. National Archives or elsewhere, you can export them from the archives by running various PowerShell cmdlets.

See the *Administrator's Guide* for more details of the records management feature.

About Compliance Accelerator and Discovery Accelerator

Enterprise Vault also provides applications for compliance monitoring activities and data mining activities, such as legal discovery.

Legislation by regulatory bodies, such as SEC, NYSE and NASD, means that financial and legal organizations are obliged to retain an increasing amount of electronic correspondence and data. Outside the financial services sector, the Sarbanes-Oxley Act of 2002 extends the length of time that corporate data needs to be available for audit. Enterprise Vault enables organizations to capture, store and retrieve the relevant data. The Enterprise Vault Accelerator products, Compliance Accelerator and Discovery Accelerator, provide specialized web tools for capturing, searching and reviewing data that has been archived using Enterprise Vault.

Compliance Accelerator provides a capture and review system for monitoring employee messages to ensure compliance with industry regulations or company policy. A random sample of journaled messages (inbound and outbound) can be captured daily and reviewed by compliance officers using the Compliance Accelerator client. In addition, compliance administrators can run regular searches to find messages that meet certain criteria, for example, messages that contain unacceptable language.

Discovery Accelerator is specially tailored for performing enterprise-wide searches to find archived items relating to a legal case. Using the Discovery Accelerator client, case administrators can run searches for relevant data. Case reviewers can then review the items returned by the searches and produce the required items in a format suitable for presenting as evidence.

About Enterprise Vault add-ons

An increasing number of storage management applications are being integrated with Enterprise Vault to manage the collection and migration of archived data. For details of these integrations and instructions on what you need to configure to use the products with Enterprise Vault, see the Enterprise Vault [Compatibility Charts](#).

Enterprise Vault Extensions are partner solutions that extend the functionality of Enterprise Vault, and are displayed and managed in the Enterprise Vault Administration Console. Extension content providers are extensions that archive additional content such as UNIX file systems, text messages, social media content, and web pages. Extensions are provided by authorized partners in the Veritas Technology Ecosystem (VTE) program:

<https://www.veritas.com/vte>

Introduction to EnCase Ingest Connector

The EnCase® Ingest Connector is a utility that ingests files from forensic evidence files into archives, so that they can be searched using Discovery Accelerator, or to designated disk locations.

The EnCase Ingest Connector ingests files from Logical Evidence Files, EnCase images, and data definition (DD) images. Using the EnCase Ingest Connector, you can copy individual files or select sets of files based on metadata (for example, file name, file name extension, or creation, modification and last accessed dates).

Full documentation on how to install and configure the software is included with the software.

How Enterprise Vault works

This section introduces the Enterprise Vault components and gives an overview of the basic archiving and retrieval processes. Enterprise Vault is packaged as a number of components, which you can select at installation time.

The core Enterprise Vault components include the following:

- The Enterprise Vault Server component comprising services and tasks for performing the main archiving, indexing, storing and restoring functions.
- The Enterprise Vault Administration Console for configuring and managing services, tasks, indexes, and archives.
- Active Server Page (ASP) web access components for enabling users to access items in archives.

The following additional components are provided for Exchange Server archiving:

- Enterprise Vault Outlook Add-In for enabling users to access archived items from within their Outlook client.
- Enterprise Vault Client for Mac OS X for enabling Outlook for Mac 2011 or 2016 users to access archived items.
- Outlook Web App (OWA 2010) extensions for enabling users to access archived items from within OWA 2010 clients.
- Office Mail App for OWA 2013 and later, and Outlook 2013 and later, for enabling users with these clients to access archived items.

The following additional components are provided for Domino Server archiving:

- Enterprise Vault extensions for Notes and DWA clients.

The following, additional components are provided for file system archiving, SharePoint archiving and SMTP message archiving:

- The FSA Agent, which provides the FSA services on Windows file servers for the creation of placeholder shortcuts, and for FSA Reporting.
- Microsoft SharePoint components for archiving and restoring documents on SharePoint servers. Optional Enterprise Vault web parts provide archive search features for SharePoint users.
- SMTP archiving components for processing messages from third-party SMTP messaging servers

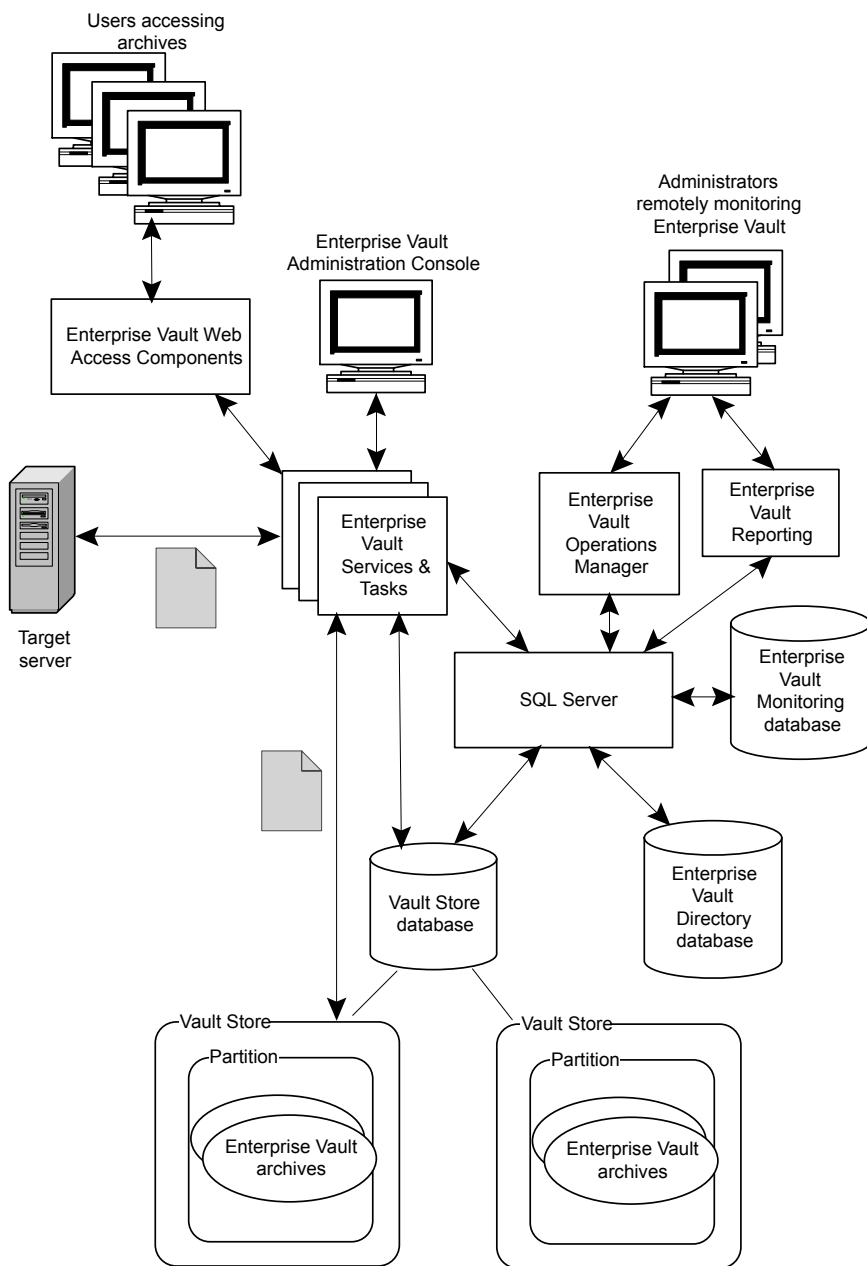
The following optional components provide enhanced management and reporting facilities:

- Enterprise Vault Operations Manager is a web application that makes remote monitoring of the status of Enterprise Vault possible from any computer on which Internet Explorer is installed.
- Enterprise Vault Reporting provides enterprise-level reporting, using Microsoft SQL Server Reporting Services.

Once installed and configured, the Enterprise Vault Server comprises a combination of Windows services and tasks, Microsoft SQL Server databases and Active Server Page (ASP) web access components. Services, tasks and archives are configured using the Enterprise Vault Administration Console, which is a snap-in to the Microsoft Management Console (MMC).

[Figure 2-3](#) shows the main components in an installed Enterprise Vault system. The target server in the diagram is a server from which items are to be archived. The illustration omits the components involved in single instance storage, which are described separately.

Figure 2-3 Illustration of an installed Enterprise Vault system



The Windows services and tasks perform background tasks such as scanning target servers for items to be archived, storing the items in archives, indexing item attributes and content and retrieving items from archives.

The Enterprise Vault Directory database and Vault Store database are SQL databases that hold Enterprise Vault configuration data and information about the archives.

The Enterprise Vault Monitoring database is a SQL database that holds monitoring data for use by the Enterprise Vault Operations Manager and Enterprise Vault Reporting components. A Monitoring agent on each Enterprise Vault server monitors the status of the Enterprise Vault services and archiving tasks, and the values of performance counters for vault stores, disk, memory, and processors. The agents collect data every few minutes and record it in the Enterprise Vault Monitoring database.

The first time that you configure a file server target for FSA Reporting, Enterprise Vault creates an FSA Reporting database (not shown in the figure). The FSA Reporting database holds the scan data that FSA Reporting gathers from the file server. When you configure another file server target for FSA Reporting, you can assign the file server to an existing FSA Reporting database, or create another database. Multiple FSA Reporting databases can provide scalability if you obtain FSA Reporting data for many file servers.

The Active Server Page web access components run on an IIS server and enable users to view, search and restore archived items using Enterprise Vault web client interfaces.

The physical organization of the components will depend on the requirements of your site. The various Enterprise Vault services and tasks can reside on one computer or be distributed over several computers. In a pilot system, for example, all the Enterprise Vault services, SQL server, IIS server and target server for archiving can, in most cases, reside on one computer.

The archives themselves can reside on your preferred storage system, for example, SAN, NAS, NTFS, WORM. You can also use certain storage devices that support the Enterprise Vault storage streamer API. Older archives can be moved off to more economic media for long-term storage. Enterprise Vault can migrate files from a vault store partition to a secondary storage location on the cloud such as Amazon Simple Storage Service, Microsoft Amazon Azure Blob Storage, and Google Cloud Storage.

The use of Hierarchical Storage Management (HSM) is also supported.

For details of supported software and storage devices, see the Enterprise Vault [Compatibility Charts](#).

Enterprise Vault organizes the archives in entities called *vault stores*. Each vault store can hold many archives, as follows:

- If a vault store is used for archiving from mailboxes, you can create archives as needed, and Enterprise Vault can create archives automatically when it enables mailboxes. This gives you control over where Enterprise Vault stores archived data.
- If a vault store is created for archiving from public folders, Enterprise Vault automatically creates one archive for each public folder that it archives.

A single vault store can be divided into a number of *partitions*, which can be on different disks or media. As a vault store grows, you can add partitions to extend the space available.

If you use the Enterprise Vault classification feature, you can archive different items to different partitions, depending on how the classification feature has tagged the items. For example, if you have configured the classification engine to detect and tag items that contain personally identifiable information (PII), you can choose to archive these items to one partition. Other types of items, such as bids and business proposals, can be archived to a different partition. These classification-related partitions are called *smart* partitions. They are identical to standard vault store partitions except in the following ways:

- Using the Vault Administration Console, you can associate a smart partition with one or more classification tags that you have defined in your chosen classification engine (Veritas Information Classifier or Microsoft File Classification Infrastructure). Only items to which the classification engine has assigned the chosen tags are archived to the smart partition.
- Multiple smart partitions can be open for archiving at the same time. This is not true of standard vault store partitions, which are limited to one open partition for each vault store.
- You can configure a standard vault store partition so that Enterprise Vault automatically rolls over to the next available partition when certain criteria are met. This rollover capability is not available for smart partitions.

On NTFS volumes, Enterprise Vault automatically uses NTFS file security. Although some elements of Enterprise Vault can be set up on FAT volumes (for example, the indexes) there will be no file security.

About single instance storage

Enterprise Vault's optimized single instance storage can provide a significant reduction in the storage space that is required for archived items. However, it can increase the network traffic between the Enterprise Vault servers and the storage devices that host the partitions.

Enterprise Vault single instance storage works on the following principles:

- Vault stores are grouped within vault store groups. A vault store group forms the outer boundary for sharing with single instance storage.
- Each vault store is assigned a sharing level, which is one of "Share within group", "Share within vault store" or "No sharing".
- Enterprise Vault archives an item using single instance storage if the target vault store has a sharing level of "Share within vault store" or "Share within group".
- A vault store's sharing level determines the vault store's sharing boundary, as follows:
 - If a vault store's sharing level is "Share within group", its sharing boundary includes all the vault stores in the group that have this sharing level.
 - If a vault store's sharing level is "Share within vault store", its sharing boundary contains only the vault store.
 - If a vault store's sharing level is "No sharing", the vault store has no sharing boundary. Enterprise Vault does not perform single instance storage for the vault store.
- Enterprise Vault identifies the parts of an item that are suitable for sharing, such as large message attachments. These parts are referred to as SIS parts. Enterprise Vault uses a minimum size threshold for SIS parts, to balance the likely storage savings against the resources that are required to create, archive, and retrieve them.
- Enterprise Vault stores each SIS part only once within the target vault store's sharing boundary. For each SIS part, Enterprise Vault accesses the vault store group's fingerprint database to determine whether a SIS part with the same fingerprint is already stored within the vault store's sharing boundary. A SIS part with the same fingerprint indicates an identical SIS part.
 - If an identical SIS part is not already stored within the sharing boundary, Enterprise Vault stores the SIS part and saves the SIS part's fingerprint information in the fingerprint database.
 - If an identical SIS part is already stored within the sharing boundary, Enterprise Vault references the stored SIS part. It does not store the SIS part again.
- Enterprise Vault stores the remainder of the item (the item minus any SIS parts) as the residual saveset file. The residual saveset file holds Enterprise Vault metadata about the item and unique information about it, such as the file name if it is a document or attachment, and follow up flags if it is a message.

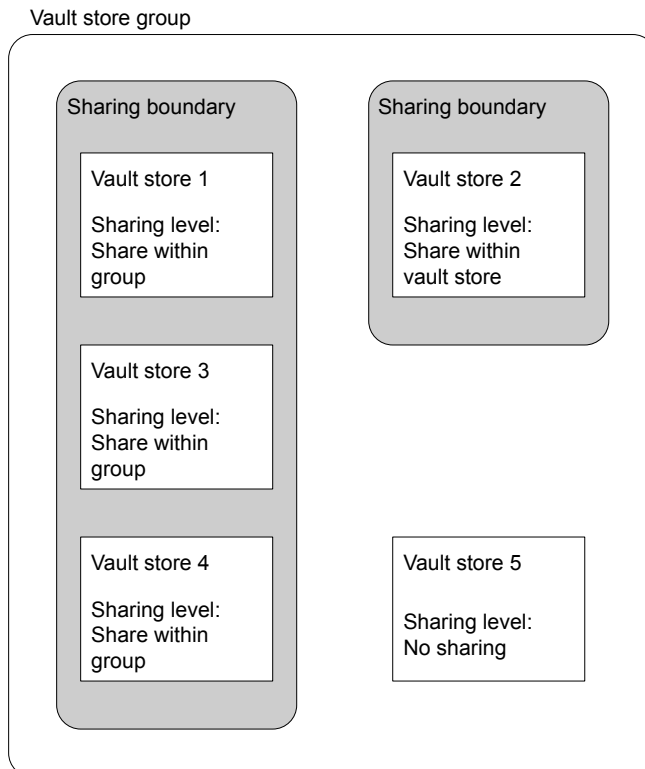
- When Enterprise Vault receives a request to restore an archived item, it reconstitutes the item from the item's residual saveset file and SIS part files.

Figure 2-4 shows an example vault store group that contains five vault stores:

- Vault stores 1, 3, and 4 all have the sharing level "Share within group". These vault stores are within the same sharing boundary. Enterprise Vault shares SIS parts across the three vault stores for items it archives to these vault stores.
- Vault store 2 has the sharing level "Share within vault store", so it has its own sharing boundary. Enterprise Vault shares SIS parts within the vault store for items it archives to the vault store.
- Vault store 5 has the sharing level "No sharing". The vault store is not included in any sharing boundary. Enterprise Vault does not perform Enterprise Vault single instance storage on the items that it archives to this vault store.

Note that a vault store group can have only one sharing boundary that contains multiple vault stores.

Figure 2-4 Sharing boundaries in a vault store group



Single instance storage can save storage space in a number of ways:

- If you use separate vault stores for journaling and mailbox archiving, Enterprise Vault can share the SIS parts between the vault stores, provided that they are in the same sharing boundary.
- If a number of separate messages with the same large attachment are sent to multiple recipients, Enterprise Vault stores the attachment only once within a sharing boundary.
- Enterprise Vault identifies a SIS part from the content, not the file name. If two messages both have the same large file attachment, Enterprise Vault can share the attachments, even if they have different file names.
- Enterprise Vault can share the identical SIS parts that result from different types of archiving, such as an Exchange message attachment that is also stored as a file on a file server.

Note the following:

- **Partitions on Dell EMC Centera devices.** Enterprise Vault single instance storage is not performed when items are stored to partitions that are hosted on Dell EMC Centera devices. Enterprise Vault provides a separate device-level sharing option to take advantage of the sharing capabilities of Centera devices.
- **Smart partitions.** Enterprise Vault shares the SIS parts between items in the same smart partition, but it does not share the SIS parts between a smart partition and other partitions.

For example, suppose that two employees receive the same email, which has an attachment. For compliance reasons, Enterprise Vault archives one employee's emails to a smart partition; it archives the other employee's emails to the standard vault store partition. If Enterprise Vault first archives the email and its attachment to the standard partition then, when it subsequently archives the email to the smart partition, it should not normally archive the attachment again. This would mean that the data on the smart partition is not fully compliant, however, so in this case Enterprise Vault archives both the email and the attachment again.

About Enterprise Vault indexing

To provide fast and efficient searching of archived data, Enterprise Vault indexes items as they are archived. An index is created for each archive. As items are added to or deleted from the archive, associated index documents are added to or deleted from the index. When a user performs a search for an item in an archive, Enterprise Vault searches the index, not the actual archive.

In Enterprise Vault you can set the required level of indexing. If required, different levels can be set for different groups of users. Two levels of indexing are available: brief and full:

- Brief indexing. This level enables users to search on attributes of an archived item such as author, subject, recipients, created date, file extension, retention category and so on. With brief indexing, the content of the item is not indexed.
- Full indexing. This level enables users to search as for brief indexing, and also provides content searching.

The more information that is indexed, the more disk space is required for the index. [Table 2-1](#) shows the estimated size of an index as a percentage of the size of the unarchived item for the different indexing levels.

Table 2-1 Estimated size of index data

Indexing level	Estimated size of index
Brief	4%
Full	12%

At Enterprise Vault 10.0 a new 64-bit search engine was introduced. A 64-bit index is created for the items that are indexed using Enterprise Vault 10.0 or later releases. The indexes that were created using earlier releases of Enterprise Vault were 32-bit indexes. When a search is performed on an archive that has both 32-bit and 64-bit indexes, Enterprise Vault automatically searches across both 32-bit and 64-bit indexes. An update tool is provided in the Enterprise Vault Administration Console to upgrade 32-bit indexes to 64-bit, if required. For details of this tool, see *Managing indexes* in the *Administrator's Guide*.

The Enterprise Vault Indexing service manages the various tasks that create, update, and search indexes. It interoperates closely with the Storage service to index items as they are stored, or later, and find items for the Storage service to retrieve. The Storage service converts items to HTML or text, if possible, and this converted content is then used to index the item. As Enterprise Vault does not index the content of items that cannot be converted to text or HTML, it is not possible to search on the content of such items. For example, the content of some binary file types cannot be converted, or searched. However, Enterprise Vault does index the attributes of items it cannot convert, so that items can still be found in the archive.

You do not have to install the Indexing Service on every Enterprise Vault server. For example, in larger deployments of Enterprise Vault the Indexing and Storage services can be located on more powerful computers to optimize search and retrieve performance. Associated Storage and Indexing services can reside on different

computers. To ensure good performance, the connection between such computers must be fast.

Enterprise Vault servers that perform indexing can be grouped together in Index Server groups to provide scalability and load balancing.

See [“About Index Server groups”](#) on page 33.

About index volumes

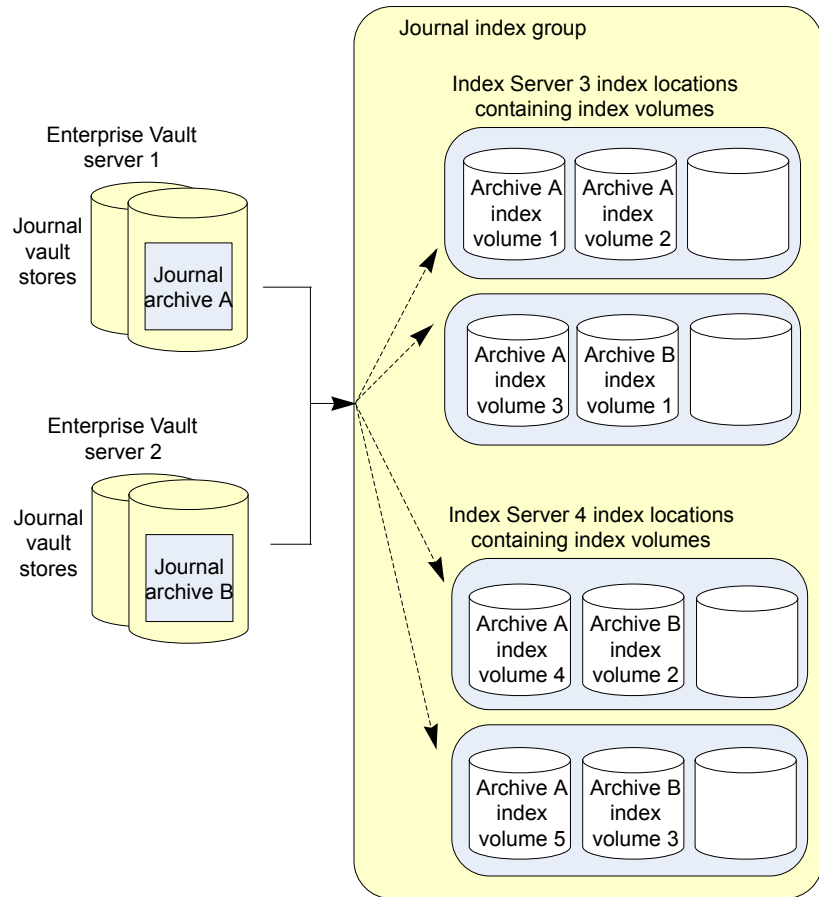
The index for an archive comprises one or more sequential index volumes. Each index volume contains index documents for the items that are stored in the archive.

Typically, a user mailbox index requires only one or two volumes. Indexes for larger archives, such as file system archives and journal archives, are likely to have multiple volumes.

When an index volume is full, the Index Server automatically creates a new index volume. The location that is selected for new index volumes depends on whether the vault store is managed by a single Index Server or an Index Server group. If the vault store is indexed by an Index Server group, the index volumes for an archive may be distributed across multiple Index Servers and locations, as shown in [Figure 2-5](#). In general, when the index for a mailbox rolls over, the existing Index Server is used, if possible. The index location is selected using a round robin algorithm, to balance the load across all index locations. For larger archives, such as journal archives, new volumes are distributed evenly across the Index Servers and index locations.

The Enterprise Vault Administration Console includes tools for managing index volumes.

See [“About managing indexes and index volumes”](#) on page 32.

Figure 2-5 Spread of index volumes in an Index Server group

About managing indexes and index volumes

You use the Enterprise Vault Administration Console to manage Index Servers, Index Server groups, and index locations. The Administration Console also includes a range of tools for managing indexes and index volumes. Tasks that you can perform using indexing tools include the following:

- Browse and manage the index volumes that are associated with an archive.
- Verify the integrity of indexes, and report any archived items that are not indexed or not deleted from the index.
- Synchronize indexes to ensure that they are up to date.

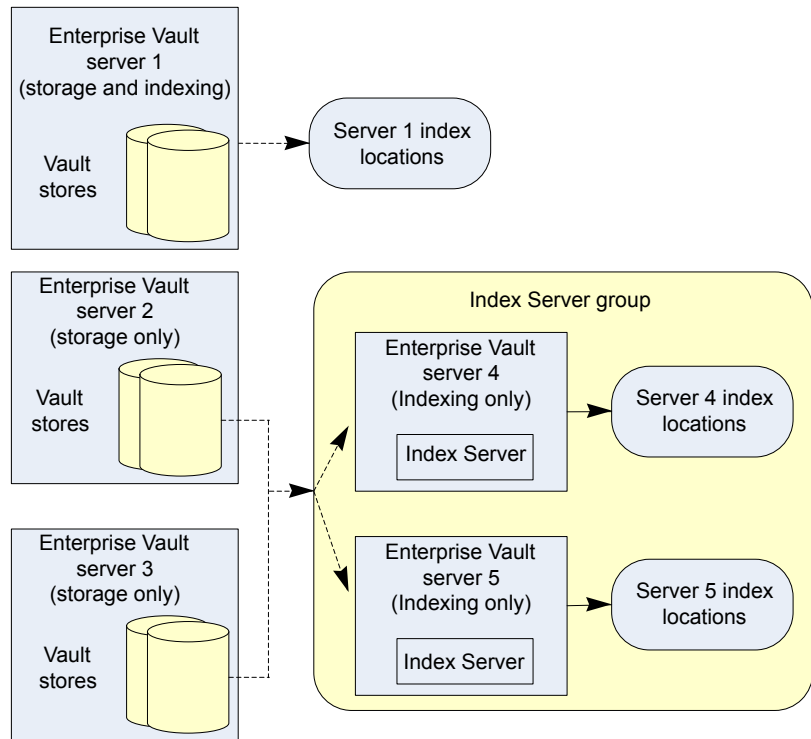
- Rebuild indexes from scratch.
- Change the location of index volumes.
- Upgrade archive indexes from 32-bit to 64-bit indexes.
- Monitor and manage the various indexing tasks.

Additional management functions are provided by PowerShell cmdlets. For example, there are PowerShell cmdlets that enable you to list the Index Servers that are assigned to index locations, and manage the backup mode for a vault store or index location.


For details of the tools, see *Managing indexes* in the *Administrator's Guide*.

About Index Server groups

An Index Server is an Enterprise Vault server that has the Enterprise Vault Indexing service installed. An Index Server can be added to an Index Server group, or *ungrouped*. In [Figure 2-6](#) The Index Server group includes Enterprise Vault servers 4 and 5. Enterprise Vault server 1 is an ungrouped Index Server. An Index Server can only be ungrouped if the server has both the Storage and Indexing services installed.

Figure 2-6 Ungrouped Index Server

Index Server groups provide load-balanced indexing services for large or distributed Enterprise Vault environments. In a distributed environment, some Enterprise Vault servers may host Storage services, while others host Indexing services. An example of this environment is shown in [Figure 2-7](#).

Figure 2-7 Index Server groups in a distributed environment

For archives to be indexed, the vault store to which they belong must be associated with an Index Server or an Index Server group. An Index Server creates indexes in index locations, and each Index Server has one or more index locations assigned to it. The Index Servers in a group share the task of indexing the items that are stored in the archives in the vault stores assigned to the group.

You use the Enterprise Vault Administration Console to perform the following tasks:

- Create Index Server groups, and add Index Servers to each group.
- Assign physical index locations to each Index Server.

- Assign each vault store to either a standalone Index Server, or an Index Server group.

About Enterprise Vault Administration Console

The Enterprise Vault Directory and all the various entities in Enterprise Vault sites are configured using the Enterprise Vault Administration Console, which is a snap-in to the Microsoft Management Console (MMC).

The Enterprise Vault Administration Console is currently available in English, Japanese and Simplified Chinese.

In the left-hand pane, a tree structure displays site entities including the following:

- Archives
- Vault stores and partitions
- Enterprise Vault servers and the Enterprise Vault services and tasks running on each server
- Index Servers and Index Server groups
- Targets for archiving (for example, Exchange Servers and mailboxes, and Domino Servers and mail files)
- Policies for defining which items are to be archived and the archiving actions to be taken
- Retention Categories and Retention Plans
- Classification policies
- Target computers and files for PST migration using the Locate and Migrate feature

You can create new entities and view or change the properties of existing entities using right-click options in both the tree and the right-hand pane.

Using the wide range of options in the Enterprise Vault Administration Console, you can configure and manage archiving within an enterprise and define the functionality that is available to users in the various client interfaces.

About Enterprise Vault sites, Directory, and Directory database

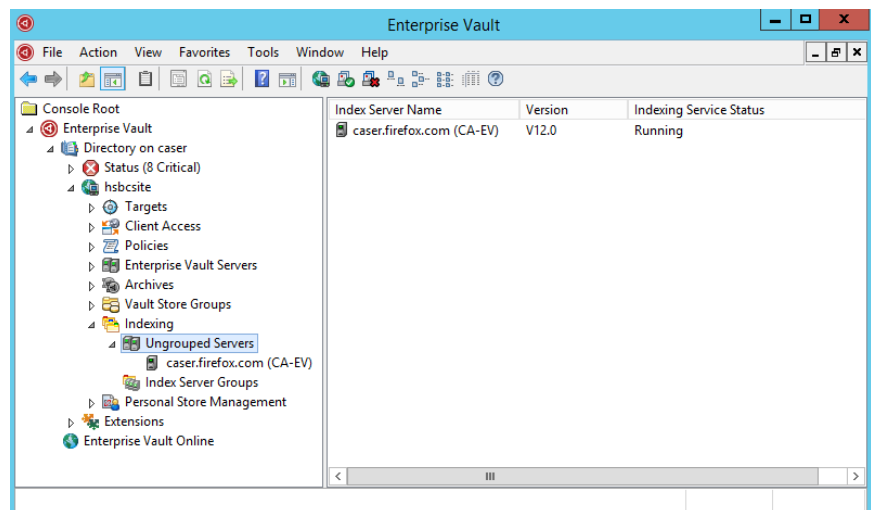
During its initial configuration, each Enterprise Vault server must join an Enterprise Vault site. A site comprises one or more Enterprise Vault servers running one or more Enterprise Vault services and tasks to archive items from specified targets: for example, Microsoft Exchange Servers, Domino mail servers, Microsoft SharePoint Servers, and file servers.

A site also contains a collection of vault stores, index servers, archiving policies that define how and when items are to be archived, and Retention Categories and Retention Plans that define how long items are to be stored before being deleted. A site may also include a list of target computers for the automatic importing of PST files.

An Enterprise Vault site is located in an Enterprise Vault Directory on an Enterprise Vault server computer. An Enterprise Vault Directory can contain one or more sites. The hierarchy of the sites, and Enterprise Vault servers belonging to those sites is shown in the Enterprise Vault Administration Console.

[Figure 2-8](#) illustrates how the Enterprise Vault Administration Console shows the contents of an Enterprise Vault Directory.

Figure 2-8 Enterprise Vault Administration Console showing the contents of an Enterprise Vault Directory



When you configure an Enterprise Vault server for the first time you can either create a Directory and site on the computer you are configuring, or join a site in a Directory on another Enterprise Vault server computer.

The Enterprise Vault Directory is accessed by a Directory Service. The other Enterprise Vault services and tasks use this service to access the configuration information in the Enterprise Vault Directory database.

See [“Introduction to the Directory Service”](#) on page 42.

The Enterprise Vault Directory database holds configuration information for each site in the Directory. (In general, configuration information is not shared across

Enterprise Vault sites.) This SQL database can be located on a computer separate from the Enterprise Vault servers.

The Enterprise Vault Directory, sites, index server groups, targets, policies and Retention Categories and Plans are all configured using the Enterprise Vault Administration Console. You can use Microsoft SQL management tools to help with managing the SQL database.

About Enterprise Vault tasks

Archiving and restoring activities are performed by Enterprise Vault tasks. Different tasks are used depending on the type of data that is being archived. Enterprise Vault tasks are managed by the Task Controller Service.

Tasks are also used to configure Exchange Server mailboxes or Domino mail files for archiving, and to locate and import PST files automatically.

Introduction to provisioning and archiving tasks

The following kinds of tasks can be created and managed using the Enterprise Vault Administration Console:

- Exchange Provisioning task. This task processes Exchange provisioning groups to configure the user mailboxes for archiving. Every target Exchange user mailbox must be in a provisioning group.
- Exchange Mailbox Archiving task. This task archives items from user mailboxes. It updates mailbox configuration information with archiving policy changes, and synchronizes permissions on the mailbox with those on the associated archive. It also updates the location and the retention category of archived items whose shortcuts have been moved or copied to a different folder.
- Exchange Public Folder task. This task archives items from public folders.
- Exchange Journaling task. This task archives items from Exchange Server journal mailboxes.
- Domino Provisioning task. This task configures Domino mail files for archiving, updates mail file configuration information with archiving policy changes, and synchronizes permissions on the mail file with those on the associated archive. Every target Domino user mail file must be in a provisioning group.
- Domino Mailbox Archiving task. This task archives items from Domino mail files.
- Domino Journaling task. This task archives messages from Domino journal databases.
- SharePoint task. This task archives documents from document libraries on SharePoint servers.

- **File System Archiving task.** This task archives files from file systems, including NTFS file systems, NetApp® Filer devices, and Dell EMC Celerra/VNX file servers.
- **SMTP Archiving task.** The Enterprise Vault SMTP service places the SMTP messages in the SMTP holding folder. The SMTP Archiving task then examines SMTP message files in the holding folder, and archives message files that contain target SMTP addresses.
- **SMTP Provisioning task.** This task assigns SMTP policies and archives to the target users in the provisioning groups.
- **Move Archive task.** This task manages Move Archive operations initiated in the Move Archive wizard.
- **Client Access Provisioning task.** This task applies Enterprise Vault Search policies and IMAP policies to the target users in selected provisioning groups.

When an archiving task is configured, it is assigned a set of targets. Each target has an archiving policy assigned to it. The targets define the location of the items to be archived and the policy defines how and when the items are to be archived. A single task can archive several targets on different servers.

You can schedule most tasks. At the times that you schedule, the archiving task scans the configured target for items that are ready for archiving, that is, those items that satisfy the archiving policy. This is automatic or background archiving. With Exchange mailbox and public folder archiving and Domino mail file archiving, users can also store specific items in the archive using the Store in archive option in their mail client. This is manual archiving.

For each type of archiving, the archiving task collects the items that are to be archived and passes them on to the Storage Service. When the Storage Service has safely stored an item, the archiving task can delete the original, and create a shortcut to the archived item. You can configure whether the task deletes the original and also whether it creates a shortcut. Some archiving tasks, such as the journaling and SMTP Archiving tasks, do not create shortcuts.

You can configure the archiving task to leave the original item as a safety copy until the vault store containing the archived copy is backed up.

When Enterprise Vault archives an item, the Enterprise Vault Directory database is accessed to find out where the required archive is. In addition, information is written to the Vault Store database.

If an archiving task creates shortcuts, the task is also responsible for the automatic deletion of shortcuts.

The automatic deletion of shortcuts could be as follows:

- When an archived item is deleted automatically at the end of the retention period. This deletion is optional.
- When an archived item is deleted explicitly by users who have delete access to them.
- When the archive itself is deleted.

The item in the archive is deleted by the Storage Service, but shortcuts are deleted by the archiving task.

For some types of archiving, you can configure automatic shortcut deletion separately from item deletion. For example, you could set all shortcuts to be deleted after a year. The archived items would remain in the archives, still available if required.

Introduction to retrieval processes

A hidden retrieval process is associated with each archiving task. This process is used when retrieving items from archives for viewing or restoring to a location specified by the user.

The retrieval process receives requests from a number of sources, such as the following:

- The Enterprise Vault Outlook Add-In and Enterprise Vault Client for Mac OS X, when a user views or restores an item from a shortcut in a mailbox.
- The Enterprise Vault extensions for Notes, when a user views or restores an item from a shortcut in a mail file.
- Enterprise Vault Search, when a user views or restores an item in an archive
- The Shopping Service, when a user uses the browser-based search facilities to select items to be restored.
- FSA shortcuts, when a user double-clicks the shortcut to open the file.
- SharePoint shortcuts or archived version links, when a user clicks the shortcut or link to open the document.

The retrieval process responds to these requests by instructing the Storage Service to retrieve the items from the archives. If the requested item has been stored offline, there may be a delay in retrieving it.

Introduction to PST migration tasks

Enterprise Vault provides the following tools for migrating (importing) the contents of PST files to archives:

- **Locate and Migrate** — This locates PST files on users' computers, copies them to a central location, and then migrates them. Unless you have only a few PST files to migrate, Locate and Migrate is likely to require least effort on your part.
- **Client-driven migration** — This is similar to Locate and Migrate, but finding PST files and sending them to a PST collection area is performed automatically by the user's computer, as opposed to by the Enterprise Vault Server Tasks.
- **Scripted migration using Policy Manager** — This is ideal for performing bulk migrations of PST files, but you need to collect the PST files in a central location.
- **PST Migrator wizard-assisted migration** — If you have a small number of PST files, this provides a quick and easy way of migrating them to Enterprise Vault.

Locate and Migrate partially automates the process of migrating the contents of PST files into Enterprise Vault. It can automatically search for PST files on users' computers and move them to a central holding area, from which they can be automatically migrated.

Locate and Migrate comprises the following Enterprise Vault task types:

- **A Locator task.** This searches your network for computers and PST files. You can configure specific paths to include or exclude during PST searching. There can be only one Locator task in your Enterprise Vault site.
- **A Collector task.** This moves PST files that the Locator task has found to a central holding folder, ready for them to be migrated. There can be many Collector tasks in your Enterprise Vault site.
- **A Migrator task.** This migrates the contents of PST files that are in the holding folder to Enterprise Vault archives. There can be many Migrator tasks in your Enterprise Vault site.

The tasks are configured and scheduled using the Enterprise Vault Administration Console.

It is possible to configure the Enterprise Vault Outlook Add-In so that users can perform their own PST migrations. The underlying mechanism that is used is still Locate and Migrate, but users can queue their own PST files for migration.

About Enterprise Vault services

In addition to tasks, the Enterprise Vault Server has the following core Windows services:

- **Directory Service**
- **Storage Service**
- **Indexing Service**

- Shopping Service
- Task Controller Service
- Admin Service

With the exception of the Directory Service and Admin Service, the services are listed and configured in the Enterprise Vault Administration Console. The Directory Service and Admin Service are managed using the Windows Services Management Console.

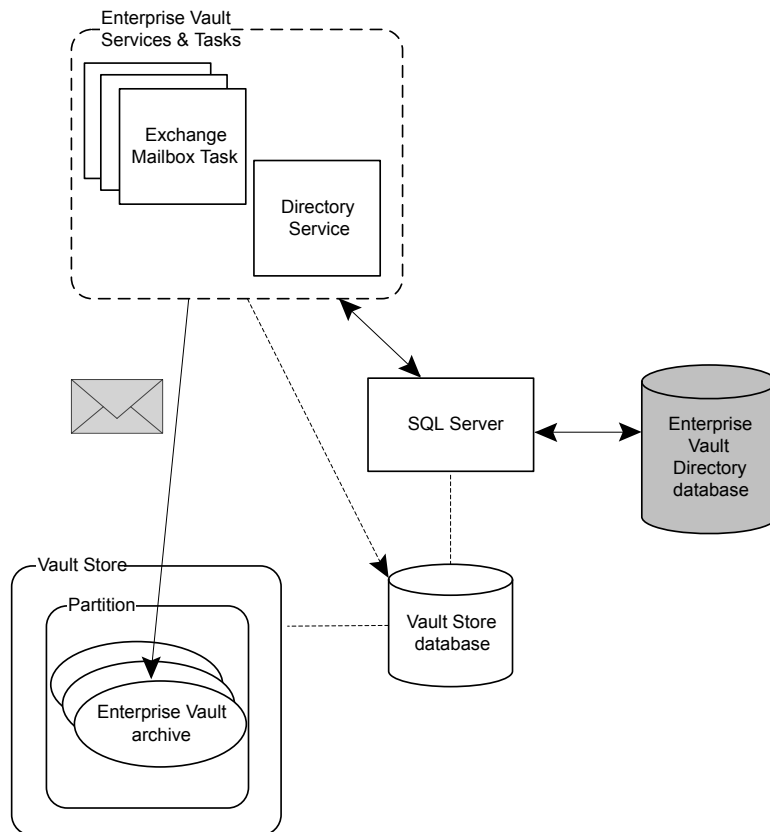
If you also install the Enterprise Vault SMTP Archiving components, the Enterprise Vault SMTP Service is created. This service is displayed in the Windows Services Management Console, but not in the Enterprise Vault Administration Console. The SMTP Service is described later in this guide.

See [“About Enterprise Vault SMTP Archiving”](#) on page 114.

Introduction to the Directory Service

The Enterprise Vault Directory has a Directory Service that the other Enterprise Vault services and tasks use to access the configuration information for their site. The Directory Service uses Microsoft SQL Server to access the configuration information in the Enterprise Vault Directory database.

[Figure 2-9](#) illustrates what happens when an Exchange Mailbox Archiving task archives an item.

Figure 2-9 Accessing the Enterprise Vault Directory database

The Exchange Mailbox Archiving task requires information such as the location of the archive, permissions on the archive and the Indexing Service to contact, and instructs the Directory Service to retrieve this information from the Enterprise Vault Directory database. (The task also updates information in the Vault Store database.)

One Directory Service can provide access to configuration information for one or more Enterprise Vault sites. Because of this, the Directory Service is considered to be outside the Enterprise Vault site rather than a part of it.

The Directory Service is a Windows Service and is listed in the Windows Services Management Console. There is no requirement for the Directory Service to be on the same computer as its database.

A single computer can never run more than one Directory Service.

Introduction to the Storage service

The Storage service manages the vault stores and archives on the computer where it is running.

The role of the Storage service can be summarized as follows:

- The Storage service accepts items for archiving from the archiving tasks. If possible, it generates a text or HTML version of each item, which the Indexing service uses to compile indexing data for the item. The Storage service compresses and stores the items (and the text or HTML versions) in the appropriate archives.

The default content converters in Enterprise Vault cannot convert some file types to text or HTML. As Enterprise Vault supports Windows IFilter, you may be able to procure a suitable IFilter for file types in your environment that are not supported by the default content converters. Also, in the default configuration, very large files (larger than 50 MB) are not converted.

Information about each item that is archived is stored in the Vault Store database.

- The Storage service responds to requests from the retrieval tasks to restore items.
- The Storage service monitors open partitions to identify those that have met their partition rollover criteria.
- The Storage service responds to requests to view archived items. It can also provide an HTML preview of the item (if a preview of the item is available).
- The Storage service deletes archived items. This can be a manual deletion by a user or an automatic deletion when the retention period on an item expires. The archiving task deletes shortcuts.

Vault stores have been designed so that they are suitable for being managed by storage management software, such as Veritas NetBackup. If it is available on the system, the storage management software manages the migration of files to a secondary storage location and retrieves files from the secondary storage to the vault store on behalf of the Storage service. The secondary storage can be offline devices, such as optical disks or tapes, or cloud storage providers such as Amazon Simple Storage Service, Microsoft Amazon Azure Blob Storage, and Google Cloud Storage.

Introduction to the Indexing Service

The Indexing Service manages the indexes of archived data to enable users to search for archived items that they want to retrieve. There are two levels of indexing: brief, and full. With brief indexing, only information about the item such as the subject and author, can be searched. With full indexing you can also search the content of each item, including phrase searches.

When users search the archives to which they have access, the index files are searched. The more information that is indexed about an item, the easier it is for a user to find.

When you set up Enterprise Vault, you specify where indexes are to be stored. You also specify the default level of indexing that you want to apply across the site, although you can override this for groups or archives.

The role of the Indexing Service can be summarized as follows:

- On instruction from the Storage Service, the Indexing Service indexes items as they are archived.
The locations of index files are specified in the Indexing Service properties, and also in the Index Server properties.
- If an index is out of date, the Indexing Service automatically updates the index.
- In response to search requests by users, the Indexing Service searches these indexes and returns information about the archived items that match the search criteria.

Introduction to the Shopping Service

When a user chooses to restore one or more items through the browser-based facilities for browsing and searching archives, the Shopping Service manages the process.

With Exchange Server archives, the selected items are placed in containers that are called *shopping baskets*. The Shopping Service is responsible for managing these shopping baskets. It instructs the retrieval process to retrieve the contents of the shopping baskets from storage, if required. With Domino archives, the Shopping Service still manages the restoring of selected items, but there are no shopping baskets.

Introduction to the Task Controller Service

Enterprise Vault archiving tasks are controlled by this service. If a task is configured as automatic, it will start when the Task Controller Service is started.

See [“About Enterprise Vault tasks”](#) on page 38.

Introduction to the Admin Service

The Admin Service has two main functions: it installs new Enterprise Vault license keys, and it provides a general monitoring service that runs automatically whenever any other Enterprise Vault task or service starts. It is installed automatically when you install any of the other Enterprise Vault services.

The Enterprise Vault Admin Service monitors the following:

- Free space on local hard disks. By default, the Admin Service monitors all local hard disks, but you can restrict it to specific disks if required.
- The amount of available virtual memory.
- The number of items on system message queues.

The Admin Service has a warning threshold and a critical threshold for each type of check. When the warning threshold is reached, the Admin Service writes a warning message to the Windows Application Log. When the critical threshold is reached, the Admin Service shuts down all Enterprise Vault tasks and services running on the same computer.

By shutting down Enterprise Vault before problems arise, this behavior helps to maintain the stability of Enterprise Vault.

The online help for the Enterprise Vault Administration Console describes how to control the behavior of the Admin Service.

About the Archive Discovery Search Service

The Archive Discovery Search Service provides a simple web service API that enables Veritas partners to develop search client applications for performing discovery-type searches across multiple vault stores and archives in an Enterprise Vault installation. The methods that this service provides lets the client applications do the following:

- Create and submit searches of Enterprise Vault archives.
- Check the status of searches.
- Retrieve the results of searches.
- Cancel and resubmit searches.
- Close searches to discard the search results and reclaim disk space.

You can install the Archive Discovery Search Service on any Enterprise Vault server. You require a license for the service to submit searches.


Note the following:

- The Archive Discovery Search Service is optimized for simultaneous searches across a large number of archives. It is not designed for a large number of simultaneous searches on a small number of archives.
- Veritas does not support the use of this service in Enterprise Vault building blocks environments. However, you can use the service in both Veritas Cluster Server (VCS) and Windows Server Failover Clustering (MSCS) environments.

About the Archive Discovery Search Service components

Figure 2-10 shows the components in a typical Archive Discovery Search Service environment.

Figure 2-10 Components in an Archive Discovery Search Service environment



A client application submits its search requests to an Enterprise Vault server that hosts the *Request Web Service* of the Archive Discovery Search Service. The Request Web Service is a web application that is hosted in Microsoft Internet Information Services (IIS), and it is the point of entry for the Archive Discovery Search Service system.

The Request Web Service is a lightweight component that forwards all calls to the *Middle Tier*, which manages the searches and aggregates their results. The *Middle Tier* always runs on the same Enterprise Vault server as the Request Web Service. If you have multiple Enterprise Vault servers, you need to install these components on just one of them.

One of the functions of the Middle Tier is to interact with a *Search Broker* that is running on an Enterprise Vault index server. In Archive Discovery Search Service

environments, each index server has a dedicated Search Broker. The Search Broker is responsible for retrieving the results of searches and storing them as XML files on disk. Each index server also hosts a *Result Web Service* from where the client application can retrieve the results of a search.

After you have installed the Enterprise Vault Operations Manager on any Enterprise Vault server, you can use it to monitor the Archive Discovery Search Service. For Archive Discovery Search Service purposes, you must install Operations Manager but you do not need to enable data collection with Operations Manager.

About the Enterprise Vault Outlook Add-In

The Enterprise Vault Outlook Add-In provides Enterprise Vault functionality in Outlook clients so that users can access archives and manage items that have been archived from Exchange Servers. It also supports RPC over HTTP connections to Exchange Server mailboxes.

The Enterprise Vault Outlook Add-In is supplied as a Windows Installer (.msi) package on the Enterprise Vault release media.

You can also enable the Office Mail App for Outlook 2013 and later users, as an alternative to the Outlook Add-In or in addition to it.

For users who access their Exchange Server mailbox using IMAP or POP3 clients, you can configure customized shortcuts using the Enterprise Vault Administration Console.

Introduction to the Outlook Add-In

The Outlook Add-In can be installed on users' computers to give access to Enterprise Vault mailbox archives. Enterprise Vault buttons and menu options are added to Outlook. The Outlook Add-In also adds Enterprise Vault online help to Outlook.

The Outlook Add-In includes the Virtual Vault feature. With Virtual Vault enabled, users can access their archives in the Outlook Navigation Pane, in a similar way to other mailbox folders and personal folders.

The Outlook Add-In lets users access and manage their mailbox archives in a number of ways, including the following:

- Manually archive items. Users can select the destination archive and a Retention Category for the item.
- Delete items (if permitted).
- Use shortcuts to view items.
- Restore items.
- Search archive.

- Access Enterprise Vault Search.

The administrator can control which functionality is available to users using policy settings in the Enterprise Vault Administration Console. In particular, the administrator can set the Outlook Add-In to work in full mode or light mode. In full mode, there are no functional restrictions on the behavior of the Outlook Add-In. In light mode, the following restrictions apply:

- Users have no access to the Enterprise Vault properties of folders.
- When users archive items manually, they cannot specify the destination archive and retention category.
- When users restore archived items, they cannot choose the destination folder. The Outlook Add-In only restores items to the folders where the shortcuts are.

Enterprise Vault can perform automatic archiving without the Outlook Add-In being installed, but, depending on shortcut configuration, users may not be able to use shortcuts or modify archiving settings. However, users can still access archived items by using the Enterprise Vault Search facilities in stand-alone browser sessions.

Introduction to offline users

The Enterprise Vault Outlook Add-In can maintain a personal, offline Vault Cache for offline workers.

The Vault Cache has the following advantages:

- It provides instant access to archived items, even when the user is not connected to the corporate network.
- It operates in addition to the normal online archive, not instead of it.
- It is useful to laptop computer users who work offline. They are often used to synchronizing their offline folders and online folders.
- It may be useful in normal offices if you need to conserve bandwidth or improve performance. The retrieval of an archived item takes place on the local computer.

When an offline user starts Outlook, the Enterprise Vault client scans the user's synchronized folders, looking for the following:

- Items that will be archived from the mailbox fairly soon — These items are copied into the Vault Cache so that they will already be there when the items become shortcuts in the user's mailbox. These items have already been downloaded as part of the Outlook synchronize, so the copy takes place on the user's computer with no further download required
- Enterprise Vault shortcuts — If the corresponding items are not in the Vault Cache the client adds them to its download list.

Users can access and manage items in their Vault Cache using facilities such as Virtual Vault and shortcuts. They can also use Windows Desktop Search to search for items in their Vault Cache.

About IMAP access to Enterprise Vault archives

Enterprise Vault's IMAP feature provides IMAP client access to existing Exchange mailbox archives, and to new Internet mail archives for users of other mail services.

Once users have been provisioned for IMAP access, and have configured their IMAP devices and clients, they can access and manage the contents of their archives as far as the Enterprise Vault policies permit. For example, users can archive new content both manually and automatically using their own client rules.

In the case of Exchange archives, the existing regime for archiving and retention remains in place. For new archives that are created for Internet mail users, retention is dictated by the IMAP policies you create and apply to the users when you provision and enable them for IMAP access.

For information about setting up IMAP access and making archives accessible to users from IMAP devices and clients, see the *Setting up IMAP* guide.

About the Enterprise Vault Client for Mac OS X

The Enterprise Vault Client for Mac OS X provides Enterprise Vault functionality to Microsoft Outlook for Mac 2011 or 2016 users.

The installer kit for the Enterprise Vault Client for Mac OS X is available as a disk image (.dmg) file on the Enterprise Vault distribution media. After the user installs the client, a **Veritas Enterprise Vault Client** menu appears in the menu bar. By default, an Enterprise Vault toolbar also appears whenever the user starts the email client. This toolbar is optional, and users can choose to hide or show it by selecting a command on the **Veritas Enterprise Vault Client** menu.

The options on the menu and toolbar enable users to interact with their archives in the following ways:

- Download items from the archive and open them in their original form.
- Reply to and forward items in the archive.
- Manually archive items.
- Restore items from the archive.
- Delete items from the archive (if permitted).
- Search the archive.

The administrator can control which functionality is available to users by using configuration settings in the Exchange Desktop Policy in the Enterprise Vault Administration Console.

About Microsoft Exchange forms

For Exchange Server archiving, the Enterprise Vault Exchange forms extend the Microsoft Exchange Forms Library to include the forms needed by Enterprise Vault.

The forms are currently available in the following languages:

- | | | |
|------------------------|-------------|-----------|
| ■ Brazilian Portuguese | ■ French | ■ Korean |
| ■ Chinese, Simplified | ■ German | ■ Polish |
| ■ Chinese, Traditional | ■ Hebrew | ■ Russian |
| ■ Danish | ■ Hungarian | ■ Spanish |
| ■ Dutch | ■ Italian | ■ Swedish |
| ■ English | ■ Japanese | |

When the Microsoft Exchange forms are installed on Microsoft Exchange Server computers, users whose mailboxes are on those computers and who have the Enterprise Vault Outlook Add-In installed on their computers can use the features of Enterprise Vault.

The Microsoft Exchange forms are copied to the Enterprise Vault server computer when you install the Enterprise Vault Server component. You can either configure the Outlook Add-In to copy the forms automatically to the Outlook Personal Forms Library or you can install the forms in the Organization Forms Library on the Exchange Server.

About OWA Extensions

To enable users to access archives and manage archived items from within OWA 2010 clients, Enterprise Vault OWA Extensions must be installed on Exchange Server 2010 CAS computers. To provide Enterprise Vault access in OWA 2013 and later clients, the Enterprise Vault Office Mail App is used instead of OWA Extensions.

In OWA clients, archives can be searched and items can be archived, viewed, restored and deleted, if permitted. Enterprise Vault buttons and menu items are added to the client. In the Enterprise Vault Administration Console, the administrator can configure what functionality to make available to OWA users.

OWA users do not require the Enterprise Vault Outlook Add-Ins to be installed on their desktop computers.

About the Office Mail App for OWA 2013 and later and Outlook 2013 and later

In OWA 2013 and later, an Enterprise Vault Office Mail App provides access to archived items. You can also enable the Office Mail App for Outlook users, as an alternative to the Outlook Add-In or in addition to it. Depending on how you configure the Office Mail App, it can make the following Enterprise Vault functions available to users:

- View an archived item from its shortcut
- Archive an item manually
- Restore an archived item
- Delete an archived item
- Cancel an action
- Open Enterprise Vault Search

For more information on the Office Mail App, see the *Setting up Exchange Server Archiving* guide.

About Enterprise Vault extensions for Notes

These extensions provide users with Enterprise Vault functionality in Notes mail clients and Domino Web Access clients. The extensions consist of template and database files, which you install on target Domino mail servers. They are currently available in the following languages:

- | | | |
|------------------------|-------------|-----------|
| ■ Brazilian Portuguese | ■ French | ■ Korean |
| ■ Chinese, Simplified | ■ German | ■ Polish |
| ■ Chinese, Traditional | ■ Hebrew | ■ Russian |
| ■ Danish | ■ Hungarian | ■ Spanish |
| ■ Dutch | ■ Italian | ■ Swedish |
| ■ English | ■ Japanese | |

These extensions give users full management access to mail file archives. Users can do the following:

- Manually archive items.
- Delete items (if permitted).
- Use shortcuts to view items.
- Restore items.

- Search archive.

About Enterprise Vault Search

Enterprise Vault Search provides client users with the ability to browse and search their archives and conduct simple or advanced searches. Users can also do the following:

- Save frequently-used mailbox folders and search queries as favorites for one-click access and reuse.
- Preview archived items and their attachments.
- Copy and move archived items out of an archive, within an archive, and from one archive to another.
- Restore archived items.
- Export archived items to a .nsf, .pst, or .zip file, depending on archive type and search policy.
- Delete archived items.
- Change the retention categories of items.

Enterprise Vault Search is designed as a search tool for end users and does not include the features provided by an eDiscovery application like Veritas Discovery Accelerator. For example, you cannot use Enterprise Vault Search for case management, review, legal hold, reporting, or auditing.

About Enterprise Vault Search Mobile edition

Designed for use on Android, iOS, and Windows Mobile devices, Enterprise Vault Search Mobile edition lets users access their archives through the web browsers on their smartphones. The search facilities that are available in Enterprise Vault Search Mobile edition closely resemble those in the desktop and tablet implementations of Enterprise Vault Search. Some facilities that are available in Enterprise Vault Search on the desktop and tablet are not currently available in Enterprise Vault Search Mobile edition. For example, mobile users cannot do the following:

- Access archives of types other than Exchange Mailbox, Exchange Public Folder, and Internet Mail.
- Search in more than one archive at a time.
- Save search queries for reuse.
- Export archived items.

- Restore archived items to any mailbox folder other than the Restored Items folder.

Enterprise Vault Search Mobile edition is a browser-based application that you deploy for intranet or Internet access using Microsoft Internet Information Services (IIS). If you want to give your users Internet access without exposing your Enterprise Vault server to unnecessary security risks, you can install the required components on a proxy server. The proxy server acts as an intermediary for requests from Enterprise Vault Search clients that are seeking resources from the Enterprise Vault server.

About Enterprise Vault monitoring and reporting

Enterprise Vault Operations Manager and Enterprise Vault Reporting are two optionally installable features that provide remote browser-based monitoring and reporting of your Enterprise Vault servers.

Note that you must set up an Enterprise Vault Monitoring database even if you do not install these features.

See “[Enterprise Vault Reporting](#)” on page 60.

See “[Enterprise Vault Operations Manager](#)” on page 62.

FIPS 140-2 compliance

Federal Information Processing Standards (FIPS) 140-2 is a standard for cryptographic modules in computer systems.

For more information about Enterprise Vault and compliance with the FIPS 140-2 standard, and about using Enterprise Vault in a FIPS 140-2-compliant environment, see the following article on the Veritas Enterprise Support website:

https://www.veritas.com/content/support/en_US/doc/ev_12_FIPS_00

Enterprise Vault administration

This chapter includes the following topics:

- [About Enterprise Vault administration](#)
- [Administration Console configuration of archiving](#)
- [Administration accounts and roles](#)
- [How to archive PST file contents](#)
- [How to archive NSF file contents](#)
- [How to export archived items](#)
- [Welcome message and other notifications](#)
- [About reporting and monitoring in Enterprise Vault](#)
- [How to script management tasks](#)
- [Checklist of day-to-day management tasks](#)

About Enterprise Vault administration

This section gives an introduction to Enterprise Vault management.

For the majority of Enterprise Vault administration you use the Enterprise Vault Administration Console, which is a Microsoft Management Console (MMC) snap-in.

For some management tasks you can use Enterprise Vault's PowerShell cmdlets. For a complete list of the available Enterprise Vault cmdlets, see the *PowerShell cmdlets* guide.

You can also use other standard Windows tools to perform general management tasks, such as granting permissions and viewing event logs. You use Microsoft SQL Enterprise Manager to maintain the SQL databases.

A number of tools are provided for monitoring and reporting, including the Operations Manager, and the Enterprise Vault Reporting feature.

Administration Console configuration of archiving

Within the Administration Console, Enterprise Vault configuration of archiving is, broadly speaking, broken down into the following:

- **Tasks** — A task is a job of archiving work that Enterprise Vault is to perform. For example, archiving from mailboxes on a particular Exchange Server computer.
- **Policies** — A policy specifies how the task is to be carried out. For example, when archiving from a mailbox the policy specifies the age at which items are to be archived and whether to create shortcuts to items that are archived.
- **Targets** — A target is the object on which a task acts. For example, depending on the type of task the target could be a Domino mail file, a disk volume, or an Exchange Server public folder.

Additionally, there are settings that affect the general behavior of Enterprise Vault and all tasks. These settings are editable in the Enterprise Vault Directory and site properties.

Administration accounts and roles

The most important Enterprise Vault account is the Vault Service account. You must set up the Vault Service account and give it suitable permissions before you install Enterprise Vault. This account must be used to run the Enterprise Vault configuration wizard when you are setting up Enterprise Vault.

Enterprise Vault services and tasks use the Vault Service account when accessing Enterprise Vault databases. In Exchange Server archiving, the Vault Service account is used by Enterprise Vault tasks when connecting to the Microsoft Exchange Server.

Enterprise Vault tasks can run under the Vault Service account or you can, if required, specify different accounts for individual tasks.

See the *Installing and Configuring* guide for more information on creating the Vault Service account.

Enterprise Vault also provides administration roles that can be assigned to other Enterprise Vault administrators. These roles provide limited privileges to allow the

users to perform given management tasks using the Administration Console. A number of predefined roles are provided for specific management tasks.

How to archive PST file contents

You can migrate (import) the contents of Outlook PST files (Personal Folder files) into Enterprise Vault. You also have the option of creating shortcuts in users' mailboxes, which they can open to go directly to the archived items.

Enterprise Vault provides administrators with the following ways to migrate the contents of PST files into Enterprise Vault:

- Wizard-assisted migration — If you have a small number of PST files, this provides a quick and easy way of migrating them to Enterprise Vault.
- Scripted migration — This scripted migration facility is provided by Enterprise Vault Policy Manager. It is ideal for performing bulk migrations of PST files. It offers more flexibility than is available when using the migration wizard.
- Locate and Migrate. This locates PST files on users' computers, copies them to a central location, and then migrates them. You can configure specific paths to include or exclude during PST searching. Locate and Migrate is designed to minimize the difficulties of collecting PST files from users' computers and is likely to require the least effort on your part.
- Client-driven migration — It is possible for you to configure the Enterprise Vault Outlook Add-In so that users can perform their own PST migrations. The underlying mechanism that is used is Locate and Migrate, but users are able to queue their own PST files for migration. This can be useful if, for example, there are users with laptop computers who are in the office only one or two days a week, thus making it difficult to obtain their PST files by other methods.

To aid PST migration you can configure desktop clients so that, when a user starts Outlook, the client writes a marker into each PST file that is listed in the mail profile. When a marked PST file is subsequently imported, the marker indicates the owning mailbox.

PST migration is described fully in the *PST Migration* guide.

How to archive NSF file contents

You can migrate (import) the contents of Domino and Notes NSF files into Enterprise Vault. You also have the option of creating shortcuts in users' mail files, which they can open to go directly to the archived items.

Enterprise Vault provides administrators with the following ways to migrate the contents of NSF files into Enterprise Vault:

- Wizard-assisted migration — If you have a small number of NSF files, this provides a quick and easy way of migrating them to Enterprise Vault.
- Scripted migration — This scripted migration facility is provided by Enterprise Vault Policy Manager. It is ideal for performing bulk migrations of NSF files. It offers more flexibility than is available when using the migration wizard.

NSF migration is described in the *NSF Migration* guide.

How to export archived items

Enterprise Vault provides Administrators with a wizard to export archived items.

You cannot use the wizard to export Enterprise Vault Domino archives. If you want to export items from Domino archives, see the description of Domino Archive Exporter in the *Utilities* guide.

The wizard enables you to export the following:

- Archives to PST files.
- Archives to their original Exchange Server mailboxes.
- A single archive to any Exchange Server mailbox.

You can export the following archive types:

- Exchange Server mailbox and journal archives.
- File system archives.
- Shared archives. (Shared archives are special archives that you can create and enable several users to access. These archives do not contain folders.)

You can filter the output by date and by Retention Category. So, for example, you could export items less than a year old that were archived with a Retention Category of "Business".

When you export archives to PST files you can then import them back into Enterprise Vault. This is useful if, for example, you are moving a mailbox to a different Enterprise Vault site and want to move its archived items too.

Welcome message and other notifications

When you enable archiving for an Exchange Server mailbox or Domino mail file, Enterprise Vault automatically sends a Welcome message to the user. The message contains instructions for the users, describing what they have to do to start using Enterprise Vault.

Just what the users have to do depends on how you have set up Enterprise Vault. Consequently, you must edit the supplied template message before it is sent out so that it gives users appropriate information about how you have set up Enterprise Vault.

There is a template Welcome message for each Enterprise Vault client language.

Similarly, Enterprise Vault automatically sends a Goodbye message when you disable archiving for an Exchange Server mailbox or Domino mail file. Once again there is a supplied template that you need to edit appropriately.

There are also notification messages sent when PST and NSF files are migrated, and when the size of users' Exchange Server mailbox archives are approaching their limit, if archive size limits are set. You may want to edit these.

About reporting and monitoring in Enterprise Vault

Enterprise Vault provides a host of facilities with which you can report on and monitor its operation. These facilities include the following:

- An Enterprise Vault Reporting feature, which provides reports on the status of Enterprise Vault servers, archives, and archived items. If you configure FSA Reporting, additional reports are available for file servers and their volumes. See [“Enterprise Vault Reporting”](#) on page 60.
- The option to run Enterprise Vault tasks in report mode. This mode lets you gauge the usage of Enterprise Vault when it runs with particular settings, but without archiving any items. See [“Report mode”](#) on page 61.
- The facility to view critical and informational Enterprise Vault events in three Windows event logs. See [“Event and diagnostic logging”](#) on page 61.
- A System Status feature with which you can monitor the health of your Enterprise Vault system from within the Administration Console. See [“System status in the Administration Console”](#) on page 62.
- A browser-based Enterprise Vault Operations Manager application, which lets you monitor Enterprise Vault remotely from any computer on which Internet Explorer is installed. See [“Enterprise Vault Operations Manager”](#) on page 62.
- Facilities to monitor Enterprise Vault events and performance automatically, including the option to use Microsoft Operations Manager (MOM) or Microsoft

System Center Operations Manager (SCOM) to monitor critical Enterprise Vault events and alerts.

See [“Automatic monitoring of events and performance”](#) on page 63.

- The facility to monitor the performance of the Microsoft Message Queue (MSMQ) queues, which Enterprise Vault uses to transfer information between components. See [“Message queue monitoring”](#) on page 63.
- Options to enable auditing for a number of different types of events for individual Enterprise Vault servers. See [“Enterprise Vault auditing”](#) on page 64.

Enterprise Vault Reporting

The Enterprise Vault Reporting feature provides enterprise-level reporting for Enterprise Vault. It uses Microsoft SQL Server Reporting Services as the reporting mechanism. Administrators manage report content and view reports using the SQL Server Reporting Services Report Manager web application.

The reports include information on the following topics:

- Exchange and Domino mailbox archiving status, health, and trends.
- The volume of items that are archived per vault store.
- The status of the Enterprise Vault archives.
- Archive quota usage per user.
- Archived item access trends.
- Vault Store usage by archive, or by billing account.
- The storage reduction that is due to Enterprise Vault single instance storage.
- The number of savesets and collections.
- Content providers' licensed capacity usage and ingest history.

Administrators can do the following:

- Customize report content using the parameters provided by the report.
- Choose from a number of report export formats, including PDF, XLS, HTML, and TIFF.
- Schedule reports to be emailed to a configured email address, or saved to a shared folder.

The Enterprise Vault Reporting feature generates some of its Enterprise Vault server reports using data obtained from the Enterprise Vault Monitoring database.

If you configure FSA Reporting for a file server, you can also generate FSA Reporting reports for that file server and its volumes.

See [“FSA Reporting”](#) on page 96.

Report mode

To gauge the usage of Enterprise Vault when it runs with particular settings, you can run tasks in report mode. In this mode, Enterprise Vault does not archive items but produces a report about what would be archived on a normal run.

The reports are logged in a file in the `Reports` folder, which is a subfolder of the Enterprise Vault installation folder. The fields in the file are tab-separated, so you can easily import the file contents into a spreadsheet program such as Microsoft Excel.

The following are examples of the types of reports that are generated:

- Exchange mailbox preparation. When processing Provisioning Groups, the Provisioning task reports on mailboxes that have been prepared for archiving and the policies assigned.
- PST migration. In the case of Locate and Migrate there are separate reports for each phase of the migration: Location, Collection, and Migration.
- File System Archiving actions.
- SharePoint archiving actions.
- In the Administration Console, Vault Store Usage Reporter can generate different reports detailing vault store usage.

Event and diagnostic logging

Enterprise Vault logs events to three Windows event logs. You can use the Windows Event Viewer to view these logs. Additionally, you may have your own or various third-party tools that you can use to monitor the Enterprise Vault log entries.

The logs that Enterprise Vault uses are as follows:

- Windows Application Event Log — This is used for events that are deemed to be critical. Service start-up and shutdown events are logged here and also events arising from the integrated monitoring that is on Site Properties. These include, for example, warnings if databases have not been backed up recently or if there is a backlog of items in journal mailboxes with a status of archive pending.

When Enterprise Vault logs an event from the integrated monitoring it also shows the event in the Status pane of the Administration Console. The Status pane

provides the quickest way for you to check the health of the Enterprise Vault system.

- Enterprise Vault Log — This is used for all events that are not deemed to be critical. For example, events relating to the progress of mailbox or public folder archiving. Additionally, events that are placed in the Windows Application Event Log are also placed in the Enterprise Vault Log, thus ensuring that the Enterprise Vault Log contains a complete record of all events.
- Enterprise Vault Converters Log — This contains events arising from document conversions.

For each of the Enterprise Vault services, you can select the level of diagnostics that is reported. The diagnostic reports are logged in the Enterprise Vault Log.

System status in the Administration Console

Enterprise Vault automatically runs checks to monitor the health of the Enterprise Vault system. If any check finds a problem Enterprise Vault displays an alert in the **Status** pane of the Administration Console.

You can also run the checks from the **Status** pane without waiting for the scheduled checks.

You can configure each of the checks on the **Monitoring** tab of **Site Properties**. For each check you can set the following:

- The threshold level. Enterprise Vault displays an alert when this threshold is exceeded.
- The frequency, which is how often you want the check to run.
- The start time. You can specify any required time. If you do not select a time, the statistics are collected when the monitoring process starts and then according to the frequency that you have defined for each check.

The **Status** pane provides the quickest way for you to check the health of the Enterprise Vault system.

Enterprise Vault Operations Manager

Enterprise Vault Operations Manager is a web application that makes remote monitoring of Enterprise Vault possible from any computer on which Internet Explorer is installed.

Enterprise Vault Operations Manager lets the administrator monitor the following:

- The status of Enterprise Vault services.
- The status of Enterprise Vault archiving tasks.

- Performance counters for vault stores, disk, memory, and processors.
- Exchange Server journal mailbox target archiving parameters, including message counts for Inbox, Archive Pending, and failed operations such as Failed DL Expansion.
- Domino Server journaling target archiving parameters, including message counts for Inbox, Archive Pending and failed operations.
- The status of the Archive Discovery Search Service components and of the searches that you have conducted with them.

A Monitoring agent on each Enterprise Vault server collects data at scheduled intervals, typically every few minutes. This data is then stored in the Monitoring database. The Enterprise Vault Operations Manager webpages display the data from when the system was last monitored. Summary pages provide at-a-glance status assessment, while detailed data can help identify problems and bottlenecks.

Automatic monitoring of events and performance

Enterprise Vault has the following further mechanisms that you can use for automatic monitoring:

- In the Administration Console, the Monitoring tab in Site Properties enables you to turn on performance monitoring of important aspects of Enterprise Vault. When a monitored item reaches its threshold settings, an alert is logged in the following places:
 - The Status pane in the Administration Console
 - The Enterprise Vault event log
 - The Windows Application event logIf you have other tools to monitor the log you can then, if necessary, be notified when such messages are logged.
- If you have Microsoft Operations Manager (MOM) or Microsoft System Center Operations Manager (SCOM) then you can use the supplied Enterprise Vault Management Pack to monitor Enterprise Vault operations and performance and to take appropriate actions as required.

Message queue monitoring

Enterprise Vault uses Microsoft Message Queue (MSMQ) Server to transfer information between Enterprise Vault components. It is important that you monitor MSMQ queues so that you can quickly spot any problems that may occur.

You can use the Windows Performance Monitor to monitor the performance of the queues. You may find it useful to have the Windows Performance Monitor running continuously, showing the number of messages on all the queues.

You will quickly become used to the normal behavior of the queues and will notice excessive backlogs. Investigate the cause of any such backlogs promptly.

More information on the queues used by Enterprise Vault processes is given in the *Administrator's Guide*.

Enterprise Vault auditing

Enterprise Vault includes flexible auditing that you can enable for individual Enterprise Vault servers. The auditing events are written to a SQL Server database—you can have a single auditing database for all Enterprise Vault Servers in a site.

For example, the audit events record the following:

- The time an event occurred
- The account that initiated the event
- The archive in which an item was archived
- The category of the event, such as View, Archive, or Delete

You can enable auditing for a number of different types of event, showing for example, details of the following:

- Actions taken using the Administration Console
- Searches
- Viewing an item
- Deletions

For most types of event you can specify detail levels of Summary or Details, or both:

- Summary gives information about the event, such as the date and time, account used, vault used.
- Details lists more information, such as extracts from the content of a message, for example Subject, Mailbox Owner, and Folder.

Note that there will be a slight reduction in performance when you enable auditing.

By default, auditing is disabled.

For information on how to set up auditing, see the *Auditing* guide.

How to script management tasks

Enterprise Vault Policy Manager provides a scripted method of modifying and controlling Exchange Server mailboxes and archives so that they conform to your Enterprise Vault archiving policies.

Additionally, you can use Policy Manager to migrate the contents of PST files and NSF files to Enterprise Vault.

Policy Manager enables you to apply settings to individual mailboxes in a much more specific manner than you can when using the Administration Console.

For example, you could write a script to do the following:

- Define a filter that archives all items older than 1 month.
- Create a folder called Personal Archive in all mailboxes and apply the new filter to the folder.
- Apply the Personal Retention Category to the new Personal Archive folder.

Policy Manager runs in Command Prompt window and uses an initialization file of settings to apply to mailboxes, archives, and to PST and NSF file migrations.

An additional Provisioning API is available. This can be used to provide the auto-enabling of mailboxes from a web page, for example. The API can also be used in conjunction with Enterprise Vault Policy Manager.

Checklist of day-to-day management tasks

The following provides a checklist of the main day-to-day administration tasks required to maintain optimal performance of your Enterprise Vault system:

- Checking the system status in the Administration Console
- Checking logs
- Monitoring Enterprise Vault tasks and services
- Starting or stopping tasks or services
- Monitoring Exchange Server journal mailboxes and Domino journal databases
- Monitoring disk usage
- Monitoring MSMQ queues
- Maintaining SQL databases
- Backing up vault stores

- Enable archiving for new Microsoft Exchange Server mailboxes or Domino mail files
- Importing PST files (Personal Folder files)
- Importing NSF files
- Monitoring licenses
- Modifying the list of who has access to an archive that is being shared by a number of users

Details of how to perform these tasks are given in the *Administrator's Guide*, and in the *Backup and Recovery*, *NSF Migration*, and *PST Migration* guides.

There are a number of utilities available for performing a variety of tasks, such as recreating FSA Placeholder shortcuts on a file server, moving archived data from an NTFS device to a Dell EMC Centera device, and managing FSA archive points.

See the *Utilities Guide*.

Exchange Server archiving

This chapter includes the following topics:

- [About Exchange Server archiving and user mailboxes](#)
- [Exchange Server and journal mailbox archiving](#)
- [Types of items to archive with Exchange Server archiving](#)

About Exchange Server archiving and user mailboxes

User mailboxes hold many types of information, for example, messages, documents, spreadsheets, graphics, and voice mail. You can specify the types of items that Enterprise Vault archives (message classes) in the properties of the Enterprise Vault Directory or in the Exchange mailbox policy.

See [“Types of items to archive with Exchange Server archiving”](#) on page 73.

In Enterprise Vault, you create an Exchange Mailbox task to archive items from user mailboxes. The user mailboxes that the task is to archive are defined using Targets. Enterprise Vault automatically creates an archive for each user mailbox to be archived. How the mailboxes are to be archived is defined in Exchange Mailbox Policies.

A single Enterprise Vault site can serve more than one Exchange domain (Exchange Organization in Active Directory). Using policies, you can apply the same archiving strategy to all users in a domain or you can configure different archiving strategies for different groups of users within the domain.

If you use a database availability group (DAG) in your Exchange environment, you must set up archiving for all members of the DAG.

Mailbox archiving does not automatically archive information held in PST files stored on the user's computer. However, the administrator can use Enterprise Vault PST migration tools to copy items from PST files into user mailbox archives.

See [“How to archive PST file contents”](#) on page 57.

Exchange Provisioning tasks

You use provisioning groups to group the user mailboxes that are to be archived using the same archiving policy.

You can select the mailboxes to be associated with a provisioning group using any of the following:

- Windows group
- Windows user
- Distribution Group (the Active Directory Group type, Distribution)
- Organizational Unit
- LDAP query
- Whole Exchange Server organization

The provisioning groups are then processed by the Exchange Provisioning task. This task assigns the correct policy settings to each mailbox.

The Provisioning task can also be used to relink mailboxes to the associated archives, if, for example, the mailboxes are moved to a different Exchange Server.

Exchange Mailbox Archiving tasks

In the Administration Console, on the required Enterprise Vault server under Enterprise Vault Servers, you create an Exchange Mailbox Archiving task for each Exchange Server with user mailboxes to be archived. These tasks are controlled by the Task Controller Service.

The Exchange Mailbox Archiving task is responsible for the following:

- Enabling mailboxes that have been processed by the Exchange Provisioning task. The Exchange Mailbox Archiving task creates an archive for the mailbox and enables the mailbox for archiving.
When new mailboxes are created, these can be enabled manually or automatically, according to settings for the provisioning group.
- Accessing each mailbox and archiving items according to the policy set for the mailbox. The task works in cooperation with the Indexing Service, which converts

and indexes items, and the Storage Service, which compresses and stores the items in the associated archive.

You can specify the Indexing Service and indexing level to use at various places in the Enterprise Vault Administration Console tree.

- The mailbox folders are replicated in the associated mailbox archive, as are the permissions set on the folders. The archiving task synchronizes folder permissions in Outlook with folder permissions in the archive. This means that if other users have been given access to an Outlook folder, they will also have access to that folder in the mailbox archive.

If required, you can set up shared archives that several users can access. For example, the mailboxes of legal staff could all be set up to have a folder for a case and, in each mailbox, the folder configured to be archived to a shared archive. Unlike mailbox archives, shared archives do not contain folders.

To obtain an estimate of the number of items that will be archived, without actually archiving anything, you can run the task in Report Mode.

Exchange Mailbox Archiving tasks run automatically according to the schedule defined for the Enterprise Vault site.

Exchange archiving targets

The mailboxes to be archived are defined by targets that you create in the Administration Console under Targets > Exchange. A target can be a whole Exchange domain, defined in Active Directory as an Exchange Organization, or a group of users within the Exchange Organization, or an individual user. Within an Exchange Organization in Enterprise Vault, there may be several Exchange Servers. Exchange Server mailbox targets are defined using provisioning groups.

For a provisioning group, you can configure settings including the following:

- The Exchange mailbox policy and desktop policy to be used for the target group
- The default Retention Category or Retention Plan to be used for items being archived
- The vault store to store archives in
- The Indexing Service to use
- Whether mailbox archives are to be enabled automatically
- A policy for importing PST file contents to archives

An Exchange mailbox policy and desktop policy are associated with each provisioning group. Together, these policies define how that group of user mailboxes is archived, and the user desktop experience in terms of the Enterprise Vault features and functions available. If you want to use different settings for different groups of

users, you will need to create, in the Administration Console, a provisioning group for each group of users and a suitable Exchange mailbox policy and desktop policy for each group.

Exchange mailbox policies

In Enterprise Vault Administration Console you create mailbox archiving policies under Policies > Exchange > Mailbox.

An Exchange mailbox policy supplies information for the archiving task to use when processing the target mailboxes, including the following:

- The indexing level to use.
- The archiving strategy. You can base the archiving strategy for an Exchange mailbox policy on one of the following:
 - Age: items are archived when they have not been modified for the time that you specify.
 - Quota: archiving keeps a percentage of each user's Exchange mailbox storage limit free.
 - Age and quota: Enterprise Vault performs age-based archiving first. If age-based archiving does not make the required percentage of mailbox storage limit free, quota-based archiving continues until the required percentage is reached.
- Archiving actions, such as deleting the original item or creating shortcuts after archiving an item.
- Whether shortcuts are created and what they contain.
- Whether items are archived from Exchange managed folders, and whether Enterprise Vault uses retention settings that are based on Exchange managed content settings.
- The types of items (Message Classes) that are archived. (You can set the default list of Message Classes in the Directory properties in the Administration Console.)

You can lock policy settings to prevent users from being able to change them in their Outlook client. This can be done from the Archiving Actions tab of the mailbox policy properties.

Exchange desktop policies

An Exchange desktop policy defines the end users' experience when using the Enterprise Vault Exchange clients. It contains the settings that control the Enterprise Vault features and functionality available with these clients.

The desktop policy settings include the following options:

- Show or hide Enterprise Vault buttons and menu options, such as Search Vaults, Store in Vault, Restore from Vault, and Delete from Vault.
- Customize deletion behavior when the user deletes a shortcut.
- Add all servers to the local intranet zone of users' web browsers, so that the users are not prompted for their logon details when they search their archives or view or restore archived items.
- Control the availability, the maximum size, and the available features of Vault Cache.
- Control the availability and the behavior of Virtual Vault.
- Change the method of deploying Exchange forms for Enterprise Vault.

In Enterprise Vault Administration Console you create Exchange desktop policies under **Policies > Exchange > Desktop**. When you create a provisioning group you assign a desktop policy to it. You can create multiple desktop policies if you want different provisioning groups to use different policy settings.

Exchange archiving filters

An advanced filtering feature enables you to customize the way certain messages are processed by the archiving task. Messages can be filtered on a variety of attributes, such as sender, recipients, subject, message direction or custom MAPI properties added by a third party application. Using XML configuration files, you can define the required action for messages that meet the filter rules. For example, you may want messages from a certain domain to be given a different Retention Category and archived in a different archive.

The custom filtering feature is particularly useful if your enterprise adds custom MAPI properties to items, and you want to be able to search archived items using these custom properties. Using the XML configuration files, you can specify custom properties that are to be indexed by Enterprise Vault when the message is archived.

You define in the XML configuration files which properties are available in your search interface. If you have a proprietary search application, you can also define how the search options for these properties are displayed. Alternatively, Enterprise Vault Search lets you search on custom properties.

An API is also available to enable you to add custom property searching to proprietary archive search applications.

Custom filtering can be configured for particular types of Exchange Server archiving (mailbox, public folder or journal archiving). If required, custom filtering can be restricted to particular mailboxes.

How to configure custom filters and properties is described in the *Setting up Exchange Server Archiving* guide.

If you want a customized filter written for this interface, contact your Veritas solutions provider.

Exchange Server and journal mailbox archiving

You can set up an Exchange Server so that a copy of all messages sent and received by the Exchange Server is passed directly to a journal mailbox in addition to the recipient mailbox. This is particularly useful if you want to implement a company email monitoring policy and vital if there is a possibility that you may have to produce email as legal evidence at some later date.

When setting up Enterprise Vault Exchange Journal archiving in the Administration Console, you add Exchange journal targets, policies and tasks. You also create an archive for each target Exchange journal mailbox that Enterprise Vault will archive.

An Exchange Journaling task performs the archiving. One of these tasks can service multiple journaling mailboxes. Exchange Journaling tasks run under the control of the Task Controller Service.

Exchange Server journal mailbox archiving archives all types of messages sent to the journal mailbox; it does not take account of the message classes defined in the properties of the Enterprise Vault Directory.

Items in journal mailboxes are deleted from the mailbox as they are archived, or after the vault store is backed up, and no shortcuts are created. Administrators with access permissions to the journal archives can search for messages. As journaled items may be confidential, it is important to give such access to a few trusted users only.

See [“Types of items to archive with Exchange Server archiving”](#) on page 73.

The Enterprise Vault Exchange Journaling task automatically detects and processes correctly any messages from Exchange Servers with envelope journaling enabled.

See the *Setting up Exchange Server Archiving* guide for more information on how Enterprise Vault supports envelope journaling.

Enterprise Vault Accelerator products can be used on journaled data; Discovery Accelerator enables legal discovery, providing features such as searching, a reviewing system and publishing; Compliance Accelerator provides message monitoring features, such as sampling, searching and a reviewing system.

As journaling generates a large amount of data, Enterprise Vault can be used to control the disk space used on the Exchange Server by continuously archiving the contents of the journal mailbox.

Exchange Server and journal filtering

Selective and group journaling enable you to configure simple filtering of journal mailbox messages. You can set up special filters, called "external filters", to define how specific messages are processed. For example, you may not want to archive system messages or Out of Office messages or you may want to identify messages sent from one group of users to another.

Custom filtering provides more sophisticated journal filtering with a wide variety of filtering rules that you can apply when archiving journal mailboxes.

See ["Exchange archiving filters"](#) on page 71.

Compliance Accelerator and Exchange journaling

Enterprise Vault Compliance Accelerator works with the Exchange Journaling task to take a random sample of messages being archived from the Exchange journal mailbox and add them automatically to the set of messages to be reviewed by compliance officers. The size of the sample required is configurable.

Types of items to archive with Exchange Server archiving

In Exchange Server, items are categorized by message class. Enterprise Vault comes with the following predefined message classes, which identify different types of items.

Table 4-1 Predefined message classes

Type of item	Message class	Archived by default?
Calendar item	IPM.Appointment	No
Contact item	IPM.Contact	No
Document	IPM.Document	Yes
Electronic sticky note	IPM.Stickynote	No
Interpersonal message	IPM.Note	Yes
Journal message	IPM.Activity	No
Message posted to a folder	IPM.Post	Yes
Task	IPM.Task	No

Using the Enterprise Vault Administration Console, you can customize the list of message classes and so choose the types of items that you want to archive. The list is defined on the **Message Classes** tab of the **Directory Properties** dialog box.

With the exception of Exchange Journal archiving, Enterprise Vault archives items with the configured message classes. Exchange Journal archiving stores all items that are sent to the journal mailbox, regardless of the configured classes. If required, you can customize the message classes to archive in Exchange Public Folder Policies or Exchange Mailbox Policies.

Exchange Public Folder archiving

This chapter includes the following topics:

- [Exchange Public Folder tasks, Targets, and Policies](#)
- [How an Exchange Public Folder task archives](#)
- [User access to Exchange Public Folder archives](#)

Exchange Public Folder tasks, Targets, and Policies

You can add one or more Exchange Public Folder tasks to enable archiving from public folders. Exchange Public Folder tasks run under the Task Controller Service.

An Exchange Public Folder task archives a Public Folder target. Each Exchange Public Folder task can process multiple Public Folder targets. A Public Folder target is a single public folder hierarchy, starting from its root path and working down. You can add a Public Folder target with a root path that is higher up a public folder hierarchy than the root path of an existing Public Folder target. You cannot add one with a lower root path.

When you add a Public Folder target, you select the vault store and archive that is to be used for it. The archive must be a Public Folder archive but the vault store may be the same one used for other archive types, such as mailbox archives. If required you can also use the same archive for more than one Public Folder target. To be able to change the archiving settings for a Public Folder target you need owner access to the Public Folder.

An Exchange Public Folder Policy is assigned to each Exchange Public Folder target. This defines how the task will archive the items in the public folders that it services.

The Exchange Public Folder task processes all folders beneath each target's root path, except for folders that are processed by another Exchange Public Folder task and those that have archiving disabled (This can be set in Outlook using the Enterprise Vault Properties for the public folder or in the properties for the Public Folder target).

If required, custom filtering can be configured to select the items to be archived from Public Folders.

New Public Folders under a specified root can be automatically added as targets by setting the public folder target to be an Auto-Enabler. Whenever a new Public Folder is added under the target root, the folder is automatically enabled for archiving and a new archive is created for that folder and any subfolders.

This feature should be used with caution, as it could result in a very large number of archives being created. However, it can be useful in some organizations, such as legal departments, where folders are created for new cases.

Public Folder Auto-Enablers can generate targets for other tasks. On large systems, scanning for new public folders can take a long time, so passing new targets to other tasks means that processing of folder content is not held up by scanning.

How an Exchange Public Folder task archives

When you add a new Exchange Public Folder task, the task runs at the scheduled times, archiving from each of its target folders.


For each folder, the Exchange Public Folder task does the following:

- Matches the permissions on the folder's archive to those on the folder.
- If the folder has been given explicit Enterprise Vault archiving settings, then the Exchange Public Folder task archives according to those settings. Otherwise, the Exchange Public Folder task uses the default settings that you have defined in the Exchange Public Folder Policy.
- If the folder is the root path of a Public Folder target that is processed by a different Exchange Public Folder task, it does not archive the folder, nor any folders beneath it.
- If the folder is not accessible, then the Exchange Public Folder task does not archive the folder, nor any folders beneath it.

As with archiving other types of files, the Storage Service converts the Public Folder items to HTML, if possible, and stores them in the archives. The Indexing Service

Figure 5-1 shows a Public Folder hierarchy that is archived by Exchange Public Folder tasks A and B.

Figure 5-1 Archiving Public Folder hierarchies



When task B finds the root path folder for task A, it archives no further down that branch of the hierarchy.

Note that archiving on root path A was created before archiving on root path B, because you cannot add a Public Folder target with a root path that is already archived by an existing Exchange Public Folder task.

User access to Exchange Public Folder archives

Table 5-1 describes the activities that users with different access levels can perform in Exchange Public Folder archives.

Table 5-1 Permitted activities in Exchange Public Folder archives

Access level	Permitted activities
Read	<ul style="list-style-type: none">■ View the contents of archived items.■ Restore from the Public Folder's archive to their own mailboxes. In Outlook, the users must first copy the shortcuts to their own mailboxes.
Write	As for Read, and also: <ul style="list-style-type: none">■ Modify the folder's Enterprise Vault settings.
Owner	As for Write, and also: <ul style="list-style-type: none">■ Manually archive from a folder.■ Restore items to a folder.

Use the Administration Console if you want to remove a Public Folder target, because this removes the marker that Enterprise Vault places on the root path folder.

File System Archiving

This chapter includes the following topics:

- [About File System Archiving](#)
- [About File archiving policies](#)
- [About shortcut files with File System Archiving](#)
- [About setting up File System Archiving](#)
- [File System Archiving in a clustered environment](#)
- [The process of File System Archiving](#)
- [How File System Archiving handles older versions of archived files](#)
- [How File System Archiving synchronizes permissions](#)
- [File System Archiving reports](#)
- [How to restore files with File System Archiving](#)
- [About FSAUtility](#)
- [How to back up and scan shortcut files with File System Archiving](#)
- [Pass-through recall for placeholder shortcuts with File System Archiving](#)
- [Retention Folders and File System Archiving](#)
- [FSA Reporting](#)

About File System Archiving

You can set up Enterprise Vault File System Archiving (FSA) to archive files from network shares. Users can then access the archived files through facilities such as Enterprise Vault Search, or by using shortcuts in the original locations.

The Enterprise Vault [Compatibility Charts](#) provide a full list of the target platforms, operating systems, and protocols that Enterprise Vault supports for FSA. They also list the supported operating systems for client access of archived items, including opening Internet and Placeholder shortcuts to archived items.

By archiving from the file system, you can gain the following immediate benefits on the volumes that are being archived:

- It is easy to archive files. You may have files that you want to add to your archive system, perhaps because of legal requirements. You can create an archiving policy to archive them all immediately.
- Files that are archived are indexed, so they are searchable.
- Previous versions of archived files are retained. When a user creates a new version of a file that has been archived, that new version will be archived when it is matched by the rules you define. All the earlier archived versions of the file are retained and are searchable.
- There may be an immediate space usage reduction.

The Retention Folder feature enables you to create a hierarchy of folders automatically on file servers, to be managed by Enterprise Vault and archived according to assigned policies. For example, you could create a hierarchy of retention folders in every user's home folder.

FSA Reporting provides summary reports on the active data on your file servers, and on the data that has been archived from them.

A separate guide contains best practices information for implementing File System Archiving with Enterprise Vault. See the following article on the Veritas Support website:

<https://www.veritas.com/docs/100007315>

For more information on migrating and consolidating file servers that have content that has been archived with Enterprise Vault, see the following article on the Veritas Support website:

<https://www.veritas.com/docs/100004422>

About File archiving policies

In the Enterprise Vault Administration Console you define File archiving policies to control which files are archived by FSA. You can apply policies to whole disks (volume policies), or to folders and subfolders (folder policies) as required.

A policy contains one or more archiving rules that you define to select the files you want Enterprise Vault to archive. You can apply the archiving rules in any order, as required. In combination with the other policy settings, the result is that you have a flexible mechanism to archive precisely what is required.

For example, you can create archiving policies that do the following:

- Start archiving when the volume is 80% full and continue until the volume is 60% full.
- Archive all files older than 30 days except Hidden and System files.
- Archive *.zip and *.avi files that are older than three days and larger than 20 MB.
- Delete *.bak files that have not been accessed in the last week, without archiving them.
- Archive *.doc files and do not create a shortcut for each file until one month after it was last modified.

A number of predefined file groups are available to enable you to quickly add the required file types to the policy.

File System Archiving can archive all file types. However, some file types such as executable files and .PST files are not suitable candidates for file archiving. The Default Volume Policy and Default Folder Policy include archiving rules that you can use to exclude unsuitable file types from archiving and shortcut creation. For more details, see *Setting up File System Archiving*.

About shortcut files with File System Archiving

When a file is archived, Enterprise Vault can optionally leave one of the following types of shortcut:

- An internet (URL) shortcut. This is a .url text file that contains a hypertext link to the archived file.
- A placeholder. This is a special file that appears exactly as the original file but, when opened, forces Enterprise Vault to fetch the archived file.
- A folder shortcut. This is a .url file that contains a hypertext link to the archived folder.

Internet shortcuts can be placed on any network share. When a user double-clicks an internet shortcut, the archived file is retrieved and is shown in the appropriate application.

If you open an internet shortcut from within an application, the application opens the contents of the shortcut, not the archived file.

Internet shortcuts have a suffix of `.url`. This suffix is appended to the file's existing suffix. For example, the shortcut for a Word document file named `document1.docx` is named `document1.docx.url`. The inclusion of the original suffix enables you to determine the original file type that the internet shortcut references.

Note: If you choose the Windows Explorer option "Hide known file types", Windows still displays the original file type of an internet shortcut. For example, the internet shortcut `document1.docx.url` appears as `document1.docx`.

Internet shortcuts have the advantage that they can be used on both Windows and non-Windows devices.

Placeholder shortcuts behave exactly as the original files. A placeholder shortcut has the same file extension as the file to which it is a shortcut. When a user opens a placeholder shortcut, the original file is automatically retrieved.

[Table 6-1](#) shows the behavior of placeholder shortcuts when you open, copy, move, or delete them.

Table 6-1 Characteristics of placeholder shortcuts

Action on placeholder	Effect
Open	<p>The file is recalled from the archive.</p> <p>Note: If pass-through recall is in effect, Enterprise Vault recalls the file to disk only if the calling application requires a writeable version.</p> <p>See “Pass-through recall for placeholder shortcuts with File System Archiving” on page 94.</p> <p>A file that is recalled to the file server replaces the placeholder shortcut.</p> <ul style="list-style-type: none"> ■ If the recalled file remains unmodified, then Enterprise Vault converts the file back to a placeholder on the next archiving service run. The only exception is if the archiving policy's shortcut creation rules are based on the last access time. In that case, Enterprise Vault reverts the file only when the shortcut creation rules are met. ■ If the recalled file becomes modified, then Enterprise Vault converts the file back to a placeholder according to the archiving policy's shortcut creation rules.
Copy	<p>The source file is restored and then copied. The destination file is a copy of the restored original file.</p> <p>Note: The copy operation does not restore the source file to disk if pass-through recall is in effect.</p> <p>See “Pass-through recall for placeholder shortcuts with File System Archiving” on page 94.</p> <p>Enterprise Vault converts a restored original file back to a placeholder on the next archiving service run. The only exception is if the archiving policy's shortcut creation rules are based on the last access time. In that case, Enterprise Vault reverts the file only when the shortcut creation rules are met.</p>
Move	<p>If the destination is on the same volume, the placeholder is moved.</p> <p>If the destination is on a different volume, the archived file is restored and then moved to the destination.</p>

Table 6-1 Characteristics of placeholder shortcuts (*continued*)

Action on placeholder	Effect
Delete	<p>You can configure Enterprise Vault to delete archived files when their placeholders are deleted, if you want. You must configure some settings for the file server, and apply an archiving policy with the appropriate settings.</p> <p>For more information, see the <i>Setting up File System Archiving</i> guide.</p>

Placeholder shortcuts are supported on NTFS devices, NetApp Filers, and Dell EMC Celerra/VNX devices. For details of the exact requirements, see the Enterprise Vault *Compatibility Charts*.

Note: On NetApp C-Mode filers, recall of large files (larger than 50 MB) may time out. For information about increasing the timeout value on the Vserver, see the NetApp documentation.

In the archiving policy you can control the time at which Enterprise Vault creates shortcuts. For example, you can create a rule to archive Microsoft Office files. The rule can make Enterprise Vault leave the original files on the disk and create shortcuts to them later. Enterprise Vault can create a shortcut to a file according to any of the following:

- Archived time
- Last accessed time
- Last modified time
- Created time

By creating a rule like this one you can ensure that files are archived for safety but are still available for editing. When a file is no longer being changed frequently, Enterprise Vault creates a shortcut to the archived copy.

A folder shortcut is a .url file that Enterprise Vault creates within a target folder when the File System Archiving task archives the folder. When a user double-clicks the .url file, Enterprise Vault displays the archived files in the original folder structure in Enterprise Vault Search. Enterprise Vault Search displays only the folder and the sub folder in the archive, along with a link to show the complete folder hierarchy.

Folder shortcuts are disabled by default. You can enable folder shortcuts by editing the settings of the associated FSA volume policy. Note that there is no option to enable folder shortcuts at the folder policy level. If you disable folder shortcuts, the existing .url files continue to work as normal.

Folder shortcuts have a .url file extension. The default file name is [View Archived Files].url. You can change the name of the folder link file by editing the Site properties. The File System Archiving task creates a new folder link file with the new name on the subsequent run. If you have customized the folder shortcut name, you may see multiple folder shortcut files because the File System Archiving task does not delete old .url files.

The File System Archiving task records details of folder shortcuts that are created and updated in the report file in `Reports\FSA`.

About setting up File System Archiving

Very briefly, setting up File System Archiving involves the following tasks:

- Preparing the file server as necessary, and then adding it as a target file server. You must install the Enterprise Vault FSA Agent on a Windows file server on which you want to leave placeholder shortcuts, or collect data for FSA Reporting.
- Creating volume policies to define how and what to archive from target volumes. Optionally you can also create folder policies, to override the volume policies for specific target folders.
- Adding the target volumes, and assigning the volume policies.
- Adding target folders, and assigning the parent volume policy or a folder policy. You can define an archive point for each folder that you want to associate with a separate archive. A folder with an archive point forms the top of an archive. Files from the folder and its subfolders are stored in the same archive.
- Configuring other features as required, such as retention folders, and FSA Reporting.
- Configuring the File System Archiving tasks to schedule archiving and associated activities, and to determine the mode in which the archiving is to run.

[Table 6-2](#) shows the properties of the Enterprise Vault Administration Console containers that you can use to control File System Archiving.

Table 6-2 Controlling File System Archiving from the Administration Console

Item	Properties
Target file server (under Targets\File Servers)	<ul style="list-style-type: none">■ Whether to archive the file server.■ Configuration settings for pass-through recall for placeholder shortcuts.■ Configuration settings for deletion of archived files on placeholder deletion.■ Configuration settings for FSA Reporting.

Table 6-2 Controlling File System Archiving from the Administration Console
(continued)

Item	Properties
Target Volume (under Targets\File Servers\<server>)	<ul style="list-style-type: none"> ■ Whether to archive the volume. ■ The File System Archiving task that is to process the volume. ■ The File System Archiving policy to apply when processing the volume. ■ For NTFS volumes, whether to enable pass-through recall for placeholder shortcuts on this volume. <p>A target volume is processed according to the File System Archiving task schedule, but can be processed manually by using the Run Now option.</p>
Target Folder (under Targets\File Servers\<server>\<volume>)	<ul style="list-style-type: none"> ■ Whether to archive the folder. ■ Whether to archive the subfolders of the folder. ■ The FSA policy to apply when processing the folder. ■ The location of archive points, which mark a folder that forms the top of an archive.
Volume policy (under Policies\File)	<p>Each target volume is assigned a volume policy, which defines the following:</p> <ul style="list-style-type: none"> ■ For NTFS volumes, whether to use quotas. ■ The type of shortcut to leave, if the archiving rules specify that a shortcut is to be created. ■ For placeholder shortcuts: <ul style="list-style-type: none"> ■ Whether to delete archived files on placeholder deletion. ■ Whether to delete placeholders for the items that are deleted from archives. ■ The retention category or retention plan to use for archived files. ■ The archiving rules to apply. These rules determine which files to archive, and when to create shortcuts. ■ Whether to archive files that have explicit permissions, and files that are under Dynamic Access Control. These files are subject to a change in permissions when archived.

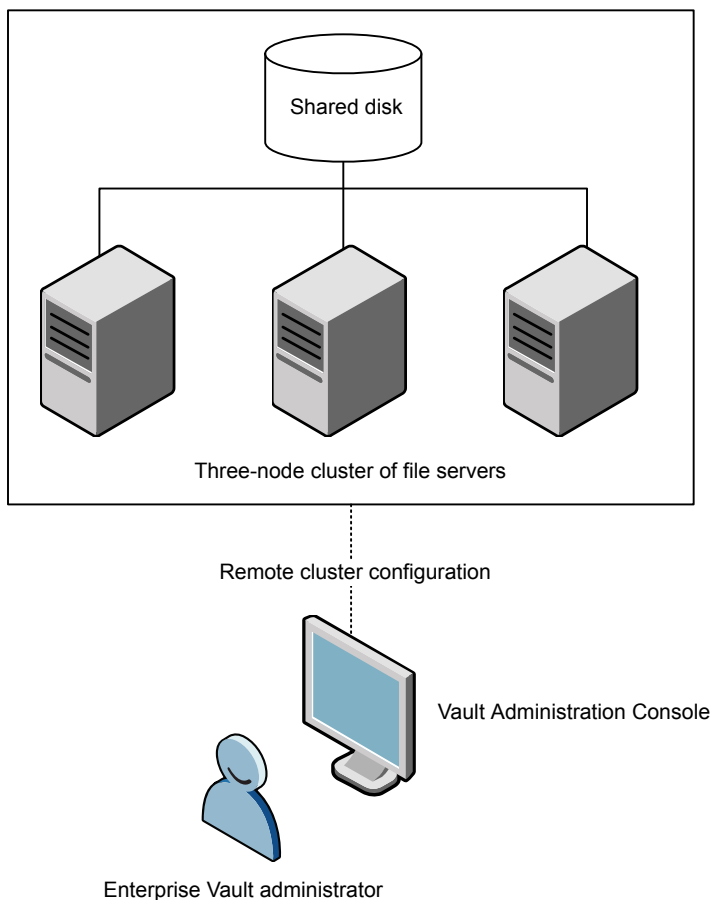
Table 6-2 Controlling File System Archiving from the Administration Console
(continued)

Item	Properties
Folder policy (under Policies\File)	<p>Folder policies are optional. Use them when you want to override the volume policy for specific folders.</p> <p>A folder policy defines the following:</p> <ul style="list-style-type: none">■ The type of shortcut to leave, if the archiving rules specify that a shortcut is to be created.■ For placeholder shortcuts:<ul style="list-style-type: none">■ Whether to delete archived files on placeholder deletion.■ Whether to delete placeholders for the items that are deleted from archives.■ The retention category or retention plan to use for archived files.■ The archiving rules to apply. These rules determine which files to archive, and when to create shortcuts.■ Whether to archive files that have explicit permissions, and files that are under Dynamic Access Control. These files are subject to a change in permissions when archived.
File System Archiving Task (under Enterprise Vault Servers\<server>\Tasks)	<p>Processes target volumes and folders. The task properties define the following:</p> <ul style="list-style-type: none">■ Whether to run in report mode or normal mode.■ Schedule settings, including the option for Run Now.■ Settings to control generation of normal and pruning reports.■ Synchronization schedule.■ Pruning options and schedule.

File System Archiving in a clustered environment

If your Windows file servers are grouped in a cluster, you can make the FSA services that run on them highly available. You must add an FSA resource to the cluster resource group or service group, and configure the FSA resource for high availability. The FSA resource monitors the state of the FSA services on the online node. If a problem occurs with the FSA services on the online node, then the cluster resource group or service group that contains the FSA resource fails over to the next available node.

Figure 6-1 shows an environment in which three file servers are clustered together.

Figure 6-1 Example FSA cluster configuration

Note that you can make the FSA Agent services highly available only when there is a shared disk resource.

The FSA clustering feature works with the following cluster software:

- Windows Server Failover Clustering (formerly known as *Microsoft Cluster Server*, or *MSCS*)
- Veritas Cluster Server (VCS)

See the Enterprise Vault [Compatibility Charts](#) for details of the supported versions of this software, and the supported versions of Windows.

The following cluster types are supported:

- Active/passive cluster. To support high availability, the shared cluster resources are made available on one node of the cluster at a time. If a failure on the active cluster node occurs, the shared resources fail over to the passive node and users may continue to connect to the cluster without interruption.
- Active/active cluster. To support load balancing and high availability, the cluster resources are split among two or more nodes. Each node in the cluster is the preferred owner of different resources. In the event of a failure of either cluster node, the shared resources on that node fail over to the remaining cluster nodes.

Enterprise Vault supports multiple nodes in any combination of active/passive and active/active. We have validated configurations with up to four nodes.

For guidelines on how to make the FSA Agent services highly available in a clustered environment, see the *Setting up File System Archiving* guide.

The process of File System Archiving

Once you have configured File System Archiving, Enterprise Vault archives files according to the schedule for each task. Enterprise Vault selects files that meet the criteria you have defined in the archiving rules.

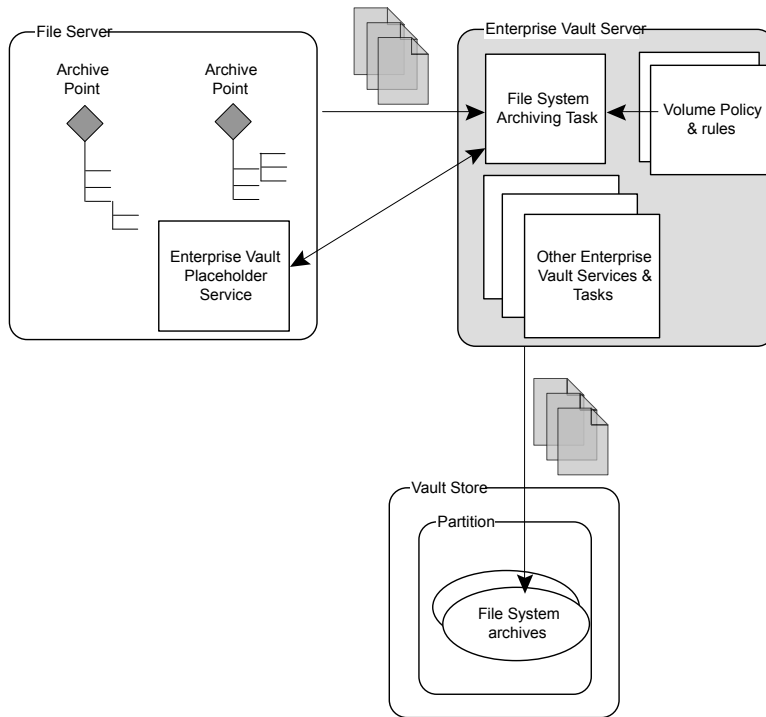
For example, the rules cover the following:

- File type
- File size
- Last access time or modification time
- Creation time
- File attributes

Enterprise Vault automatically creates an archive for each Archive Point and then selects and stores files according to the rules. The files are then replaced with shortcuts, or deleted, according to the policy setting.

Depending on the safety copy settings in the vault store properties, the original files may remain in place until the vault store has been backed up and then be replaced with shortcuts.

[Figure 6-2](#) shows the Enterprise Vault components involved when archiving files from a file server.

Figure 6-2 File System Archiving process example

In this example, placeholder shortcuts are being used on an NTFS file system, so the Placeholder Service (part of the FSA Agent) is installed on the file server. Two Archive Points have been created on the file server to mark the top of the folder structures to be archived. Two archives will be created, one for each Archive Point.

On the Enterprise Vault server, a volume policy defines what files are to be archived. The File System Archiving Task runs under the control of the Task Controller Service according to the task schedule.

When a file is archived, the Enterprise Vault Storage Service converts the file to HTML, if possible, and stores the file in the archive. Using the HTML, the Indexing Service indexes the details and content of the file, so that users can later search for the file.

If many of items are being archived at once, you can speed up the archiving process by disabling indexing. Items can be indexed later, when the Enterprise Vault storage service computer is quiet, by rebuilding the archive indexes.

How File System Archiving handles older versions of archived files

File System Archiving version pruning enables you to control the number of versions of files that are stored in Enterprise Vault archives.

When an archived file is recalled and modified, Enterprise Vault will archive the new version. This means that there are now two versions of the file in the archive. Each time a file is recalled and modified, subsequent archiving means that another version of the file is stored in the archive.

Pruning is the process of deleting the earlier versions of archived files. It is configured in the property settings for the File System Archiving task.

How File System Archiving synchronizes permissions

Enterprise Vault automatically synchronizes archive permissions with folder permissions. When Enterprise Vault archives files, the version stored in the archive is given the same permissions as the folder that contained the original file.

This means the following:

- If there is a file with permissions that are different from those of the containing folder, the archived version has the folder permissions. A shortcut has the same permissions as the file.
- Someone who has access to the original folder can find and access the archived version of the file, even if the file permissions denied access. Such a person cannot use a shortcut to access the file, however.

Synchronization can be configured on the Synchronization tab in the properties of the File System Archiving task. Synchronization runs a maximum of twice each day and always runs until completion.

File System Archiving reports

File System Archiving can create the following types of report:

Archiving report	<p>This report is generated for each normal mode archiving run. The report shows details of what was archived. You can use the following settings to control the level of detail in the report:</p> <ul style="list-style-type: none">■ Brief. Generate summary reports.■ Full. Generate detailed reports that list each file that was matched by a rule.
Pruning report	<p>This report is generated by each pruning run. You can use the following settings to control the level of detail in the report:</p> <ul style="list-style-type: none">■ Brief. Generate summary reports that show the number of items deleted from each archive.■ Medium. Generate reports that list each deleted item.■ Full. Generate reports that list all items, with a status for each, such as "Deleted" or "NeverExpires".

You can configure the reports on the **Reports** tab in the properties of the File System Archiving task. You can also use the options in the list of File System Archiving settings on the **Advanced** tab of the site properties to control how you want Enterprise Vault to store and retain the reports.

If you have installed the Enterprise Vault component, you can configure the FSA Reporting feature to provide detailed reports on the status of your file servers.

See ["FSA Reporting"](#) on page 96.

How to restore files with File System Archiving

Users can restore files in the following ways:

- Clicking a placeholder shortcut recalls the file from the archive.

Note: If pass-through recall is in effect, Enterprise Vault recalls the file to the file server only if the calling application requires a writeable version.

See ["Pass-through recall for placeholder shortcuts with File System Archiving"](#) on page 94.

A file that is recalled to the file server replaces the placeholder shortcut.

- If the recalled file remains unmodified, Enterprise Vault converts the file back to a placeholder on the next archiving service run. The only exception is if the archiving policy's shortcut creation rules are based on the last access time. In that case, Enterprise Vault reverts the file only when the shortcut creation rules are met.

- If the recalled file becomes modified, Enterprise Vault converts the file back to a placeholder according to the archiving policy's shortcut creation rules.
- Clicking an internet shortcut gives the user the option of opening or saving the file. Both of these options involve recalling the file from the archive.
- Using Enterprise Vault Search, the user can select files and then right-click to copy or move them.
Copy places a copy of the archived file in the selected location. Move places a copy of the archived file in the selected location and deletes the file from the archive.

Note: If you attempt to recall a file that is larger than 4 GB from an internet shortcut using Internet Explorer the file may be inaccessible. Enterprise Vault displays a message stating that files larger than 4 GB cannot be opened. This restriction is due to a limitation in Microsoft Internet Explorer. Note that placeholder shortcuts are not affected.

To work around this restriction you can restore the file by using the **Copy to File System** or **Move to File System** menu option in Enterprise Vault Search.

About FSAUtility

FSAUtility is a command-line utility with which you can do the following:

- Recreate archive points on the original path.
- Recreate the placeholders for archived files in their original location.
- Move placeholders from one location to another location and move the archived files to the corresponding destination archive, which is represented by the archive point on the path.
- Migrate placeholders from a source path to a destination path without any movement of the archived data.
- Delete orphaned placeholders for which no corresponding item exists in the archive.
- Restore all archived files, or archived files of the specified file types, to their original location or a new location.
- Recall the archived files that correspond to placeholders that are present in a folder.

The utility works with archive points and placeholders on Windows file servers, NetApp filers, and Dell EMC Celerra/VNX devices.

For details of the utility, see the *Utilities* guide.

For more information on migrating and consolidating file servers that have content that has been archived with Enterprise Vault, see the following article on the Veritas Support website:

<https://www.veritas.com/docs/100004422>

How to back up and scan shortcut files with File System Archiving

Enterprise Vault placeholder shortcuts appear to the operating system as markers for offline files. Some backup and antivirus programs can be configured to ignore offline files, but others cannot.

Note the following:

- If you can configure your antivirus or backup program to ignore offline files, do so before running the application on disks with Enterprise Vault placeholder shortcuts.
- If you cannot configure your antivirus or backup program to ignore offline files, every placeholder that it checks will result in an offline file being recalled. In this case, you can use the Enterprise Vault backup mode program and settings to place the file server into backup mode before a scan or backup is run.

The Enterprise Vault backup mode program and settings enable you to specify the following:

- Programs that are prohibited from recalling archived items. This is most likely to be useful if you use an antivirus or backup program that does not honor the file system offline attribute.
- Specific user groups that backup mode is to be used for. When you restrict backup mode in this way, people who are not in those groups can still recall files as normal.

Pass-through recall for placeholder shortcuts with File System Archiving

For Windows and NetApp file servers, you can configure Enterprise Vault to perform pass-through recall for placeholder shortcuts. Enterprise Vault then passes the data directly through to the calling application on receipt of a read request for a placeholder. Enterprise Vault recalls the file to the file server, subject to permissions, only if the calling application makes a write request. For example, this is the case

if the application requires a writeable file, or the user attempts to save changes to a file.

Pass-through recall can be useful in the following circumstances:

- With placeholders on read-only file systems, such as snapshots. A normal placeholder recall to a read-only file system fails because Enterprise Vault cannot write the recalled file to the file system.
- With Windows file servers where there is limited space on the file server, or when users have strict quotas for space usage. Recalled files normally occupy space on the target file system, and therefore count towards a user's space quota. For Windows file servers, you can enable or disable pass-through recall for each file server volume.

Pass-through recall uses a disk cache to reduce recall times for large files. For Windows file servers, the disk cache is located on the file server. For NetApp file servers, the disk cache is located on the Enterprise Vault server.

Note the following:

- Some applications such as Excel always recall to disk even when pass-through recall is enabled.
- For Dell EMC Celerra/VNX file servers, Enterprise Vault supports the Celerra/VNX pass-through facility.
- For NetApp file servers, the pass-through recall feature works only with read-only file systems. Pass-through recall is ignored for read-write file systems.
- There is a setting "Delete archived file when placeholder is deleted" on the Shortcuts tab of volume policy properties and folder policy properties. That setting is ignored on Windows file server volumes if pass-through recall is enabled on the volume.

Retention Folders and File System Archiving

The Retention Folder feature enables you to create a single folder or a hierarchy of folders automatically on file servers, to be managed by Enterprise Vault and archived according to assigned policies. The folder hierarchy can be added to a specified target folder or to its subfolders. For example, you could create a hierarchy of retention folders in every user's home folder. Items placed in the retention folders are archived by Enterprise Vault according to the particular policy assigned to each folder. You define the archives to use for the retention folders by specifying where archive points are to be created. If a user deletes any folders in the retention folder hierarchy, Enterprise Vault recreates the folders during the next run of the FSA archiving task in normal mode.

You configure retention folders using the Administration Console. The required steps are as follows:

- Create a suitable folder policy to use as the default folder policy for the retention folders.
- Create a Retention Folder policy, to define the hierarchy of folders to be created on the FSA target, and the folder policy to use on each folder
- Add the FSA target on which you want the retention folders created, assign the Retention Folder policy, and specify where archive points are to be created. You can specify that the retention folder hierarchy is added to the root of the FSA target, or to each subfolder.

The folders are created on the file server on the next normal mode archiving run. To test the effect of an assigned retention folder policy you can perform an archiving run in report mode. You can also assign policies to folders using a command line interface.

FSA Reporting

FSA Reporting provides summary reports on the active data on your file servers, and on the data that has been archived from them. FSA Reporting's reports include data on a wide range of items including the following:

- The number of archived files for each file server, and the space used and saved as a result of archiving. You can also view the 10 largest files in a volume.
- Active and archived space usage by different file groups, per server and per archive point.
- Numbers of unaccessed or duplicated files, and the space they are occupying.
- Used and free space on the drives of each file server.

Many of the reports can provide either an overall view for all file servers with FSA Reporting configured, or a detailed view for a named file server.

In order to access FSA Reporting's reports, the Enterprise Vault Reporting component must be installed and configured on a computer with the required prerequisites, including Microsoft SQL Server Reporting Services. You use the SQL Server Reporting Services Report Manager web application to view the reports.

You must also configure FSA Reporting for each file server target for which you want to obtain reports. The Administration Console provides wizards to help you do the following:

- The first time that you configure a file server target for FSA Reporting, a wizard helps you to set up an FSA Reporting database to hold the FSA Reporting scan data.

When you configure another file server target for FSA Reporting, you can assign the file server to an existing FSA Reporting database, or create another database. Multiple FSA Reporting databases can provide scalability if you obtain FSA Reporting data for many file servers.

- For a Windows file server, install the FSA Agent on the file server if the agent is not already present
- For a non-Windows file server, that is NetApp Filers and Dell EMC Celerra/VNX devices, select another server to act as the FSA Reporting proxy server. The FSA Reporting proxy server gathers the FSA Reporting data for one or more non-Windows file servers.

Any of the following can act as an FSA Reporting proxy server, subject to some additional prerequisites:

- An Enterprise Vault server in the Enterprise Vault site.
- A Windows server that is configured as a file server archiving target in the Enterprise Vault site.
- A Windows server on the network.

Note: To enable FSA Reporting on NetApp C-Mode filers, you must have the Enterprise Vault 11.0.1 or later FSA Agent installed.

For more information, see the *Reporting* guide.

Archiving Microsoft SharePoint servers

This chapter includes the following topics:

- [About archiving Microsoft SharePoint servers](#)
- [How to configure SharePoint archiving](#)

About archiving Microsoft SharePoint servers

You can use Enterprise Vault to archive content from servers running any of the following:

- Microsoft SharePoint Server 2010
- Microsoft SharePoint Foundation 2010
- Microsoft SharePoint Server 2013
- Microsoft SharePoint Foundation 2013
- Microsoft SharePoint Server 2016

At scheduled times, Enterprise Vault automatically copies content from the SharePoint server and stores them in Enterprise Vault SharePoint archives. Archived documents can be left on the SharePoint server or deleted, as required, and shortcuts to archived documents can be created on the SharePoint Server.

If versioning is enabled for a document library, you can configure the number of versions of a document that are to be left on the SharePoint server after archiving. An Enterprise Vault version history link enables users to view archived versions of a document from the SharePoint version history page.

An Archive Search web part can be added to SharePoint site pages to enable users to search the SharePoint archives. Documents can be viewed and saved.

The benefits of archiving SharePoint Servers can be summarized as follows:

- Control of SharePoint storage growth by moving older content to online archives. Archiving is automatic and policy driven.
- Enable the use of cost effective storage for long term retention of older information. Archives are storage independent, enabling the use of cost-effective storage solutions such as disk, optical, tape, SAN, NAS, and SSP.
- Retain the intellectual property represented by SharePoint stored content.
- Archive information to meet legal retention requirements.
- Central information retention for distributed departmental systems.
- Ease system consolidation.

SharePoint holds the newest, changing information, while Enterprise Vault provides managed, long term, high volume, online storage of older information.

SharePoint archiving requires a separate Enterprise Vault SharePoint license.

How to configure SharePoint archiving

To enable SharePoint archiving, the Enterprise Vault Microsoft SharePoint component needs to be installed and configured on the SharePoint server computer. SharePoint must be installed and configured on the target server to be archived before Enterprise Vault components are installed. The Enterprise Vault Admin Service is installed automatically with the Enterprise Vault Microsoft SharePoint component.

Before configuring the Enterprise Vault Server for archiving, the SharePoint site collections for archiving must exist on the SharePoint server. In Enterprise Vault Administration Console, you can then create objects for SharePoint archiving tasks, Policies and Targets.

SharePoint archiving tasks

The SharePoint task is the process that archives the SharePoint site collections. Task properties include the archiving schedule, the account the process is to use and the archiving reports to generate.

A single task can support several SharePoint targets, which may be web applications or site collections. Alternatively, you can create multiple SharePoint tasks and assign targets to each task, as required. The SharePoint task runs under the control of the Task Controller Service.

To archive regularly all the target sites associated with a task, you can use an Enterprise Vault site schedule or you can set a separate schedule for the SharePoint

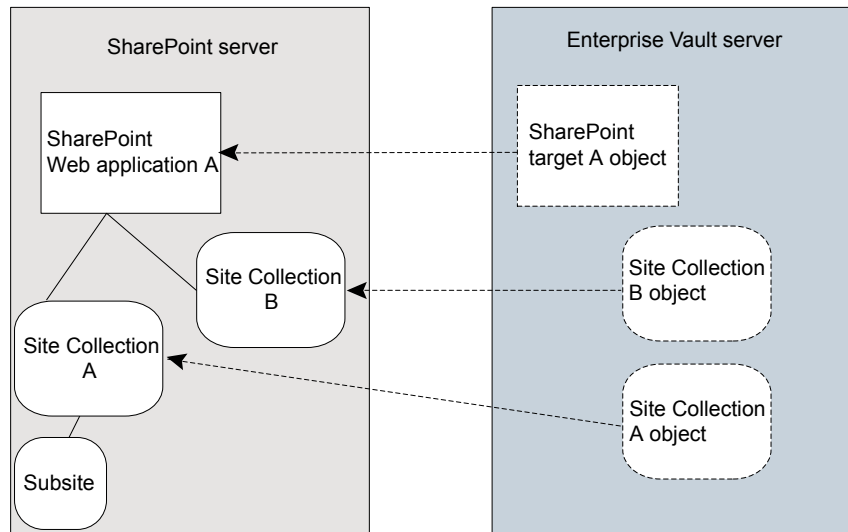
task. Alternatively, an immediate archive run can be started for all sites serviced by an archiving task, or for a particular site collection.

SharePoint archiving targets

The SharePoint Servers and site collections to be archived by the archiving tasks are defined using Targets in the Administration Console. These reflect the SharePoint structure: site collection targets point to the top-level sites on the SharePoint server, and site targets point to subsites.

Before you can configure the archiving targets in Enterprise Vault, the site collections must exist in SharePoint and the Enterprise Vault components must have been installed and configured on the SharePoint server

Figure 7-1 SharePoint archiving targets



When you add a SharePoint target URL, you can auto-enable site collection archiving and assign default values for the archiving task, vault store, policy and Retention Category or Retention Plan to be used.

Subsites will be archived automatically using the default settings for the target URL. If new subsites are created, these will be automatically included in the archiving. Content archived from a site collection are stored in the same archive.

If automatic site collection archiving is not enabled, or you want to override the default settings for a site collection, you can create site collection targets and subsite targets manually under the target URL. For example, you may want to use a different archiving policy for a particular site collection.

For a site collection, you can limit the scope of archiving to the top-level website only, the subsites only, or both.

When the archiving task runs, a SharePoint archive is created automatically for each SharePoint site collection. In the Administration Console tree, you can see the archives under **Archives > SharePoint**. Content in the top level site and all subsites of that site collection are stored in the same archive.

[Figure 7-1](#) illustrates the relationship between a SharePoint 3.0 web application and site collections on the SharePoint server (on the left) and associated target objects in the Enterprise Vault Administration Console (on the right).

SharePoint archiving reports

Reports can be generated for each run of the archiving task. You can select the amount of detail you want included in reports and the number of reports for a task that you want kept in the `Reports` folder (for example `C:\Program Files (x86)\Enterprise Vault\Reports`).

SharePoint archiving policies

A policy defines which documents are to be archived and how they are to be archived.

In a SharePoint policy you can configure the following archiving actions:

- Leave the document in SharePoint. This means that the document will not be deleted from SharePoint once it is archived; users will be able to access all versions of the document both on the server and in the archive.
- Delete document from SharePoint once archived, and leave a shortcut to the archived document.
- Delete document from SharePoint once archived, without leaving a shortcut. This means that an archived document is deleted from SharePoint and is only available in the archive.
- Prune to a specified number of versions of the document. If versioning is enabled for the document library, you can set the number of versions of an archived document that you want left in SharePoint after archiving. Earlier versions will be available in the archive only.
- You can specify how to archive items with unique permissions and whether to archive drafts of items.
- You create rules to select the documents that you want to archive with the policy.

You can make Enterprise Vault delete shortcuts automatically, according to the age of the shortcuts. For example, you can choose to delete those shortcuts that are more than one year old.

How to access archived SharePoint documents

By installing the optional Archive Search web part on the SharePoint Server, and adding the Archive Search web part to site pages, you can enable SharePoint users to search for documents stored in the Enterprise Vault SharePoint archive. The search is very similar to the SharePoint Portal Server search.

Using Archive Search, documents can be viewed, saved, restored to the SharePoint Server and, if allowed, deleted.

When users view the version history of a document, the versions of the document on SharePoint are displayed on the versions history page. After the archiving task has run for the first time, a new link is displayed under the SharePoint versions that gives users access to archived versions of the document.

About Enterprise Vault shortcuts in SharePoint

Shortcuts behave exactly like SharePoint documents. They use the same icons as the corresponding archived documents. Shortcuts are not created for social content.

Domino mailbox archiving

This chapter includes the following topics:

- [About Domino mailbox archiving and Enterprise Vault](#)
- [Domino provisioning groups](#)
- [Domino mailbox archiving tasks](#)
- [Domino mailbox archiving policies](#)
- [Domino mailbox archiving retention folders](#)
- [Domino mailbox archiving desktop policies](#)

About Domino mailbox archiving and Enterprise Vault

In the context of Enterprise Vault, the term "Domino mailbox archiving" refers to archiving items from user mail files on Domino mail servers.

Domino mail files hold many types of information, for example, messages, documents, spreadsheets, and graphics. You can specify the default types of items that Enterprise Vault archives (according to the Domino forms they use) in the properties of the Directory (Domino Forms tab). The list can be tailored in individual Domino Mailbox Policies.

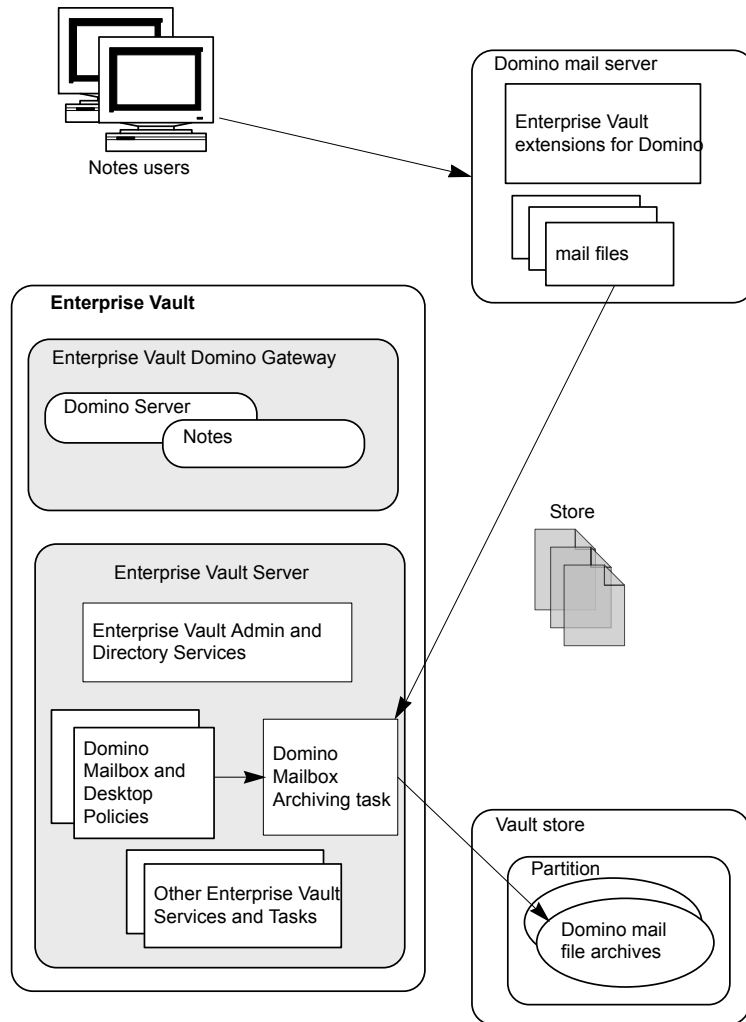
Domino mailbox archiving does not automatically archive content held in NSF files stored on users' computers. However, you can use NSF migrator to archive the content from NSF files.

See ["How to archive NSF file contents"](#) on page 57.

Enterprise Vault for Domino mailbox archiving comprises an Enterprise Vault Domino Gateway and an Enterprise Vault server. These can be on the same computer, but typically will be on separate computers to support larger Domino environments.

[Figure 8-1](#) provides an overview of Domino mailbox archiving.

Figure 8-1 Overview of Domino mailbox archiving



Domino mailbox archiving requires at least one Enterprise Vault Domino Gateway for each Domino domain. This is a computer with Domino Server, Notes and Enterprise Vault installed. If this is on a separate computer from the Enterprise

Vault server, then only Enterprise Vault Admin Service and Directory Service are required. The Domino Server installed on the Enterprise Vault Domino Gateway computer is used by Enterprise Vault to access the Domino Directory and target Domino mail servers.

Enterprise Vault servers that are archiving Domino mail servers must have Notes installed.

On the Enterprise Vault server, using the Administration Console, you create a Domino Mailbox Archiving task to archive items from user mail files. You define as archiving targets the target Domino mail servers and user mail files that the task is to archive. Enterprise Vault automatically creates an archive for each user mail file to be archived.

Enterprise Vault can archive from mail-in databases. This feature enables you to archive mail items that several users share. Vault Cache is not available for the mail that is archived from mail-in databases.

Domino mailbox policies define how mail files are to be archived. Domino desktop policies define the features available in Notes and control how the client functions.

If required, there can be several Enterprise Vault servers to one Enterprise Vault Domino Gateway. Each Enterprise Vault server can have a maximum of one Domino Mailbox Archiving task that archives to a local vault store.

If you use secondary Domino servers to keep replicas of users' mail files, it is possible to archive from those secondary servers instead of from the mail servers. When there are many mail servers and only a few secondary servers this approach can simplify the configuration.

Enterprise Vault supports clustered Enterprise Vault Domino Gateway servers. These servers are clustered using IBM Domino server clustering.

Domino provisioning groups

You use provisioning groups to group the user mail files that are to be archived using a particular mailbox policy and desktop policy. You then run the Domino Provisioning task to apply the settings in the provisioning group.

You can select any of the following target types to associate with a provisioning group:

- Directory Group
- Mailbox (mail file)
- Mail-in database
- Organizational Unit

- Corporate Hierarchy

A mail file must be part of a provisioning group before it can be enabled for archiving.

If a mail file appears in more than one provisioning group, it belongs to the first provisioning group in which it appears. You can change the order of provisioning groups by editing the properties of the "Provisioning Groups" container (under Targets > Domino > *domain*) in the Administration Console.

In the properties of a provisioning group you can choose to enable mail files for archiving and you can specify the indexing level to use when archiving. It is possible to override the indexing level for an individual archive by editing its properties.

Domino mailbox archiving tasks

In the Administration Console, on the required Enterprise Vault server under Enterprise Vault Servers, you create a Domino Mailbox Archiving task. This Domino Mailbox Archiving task can process mail files on more than one Domino Server and more than one Domino domain.

If necessary, it is possible for multiple Enterprise Vault servers to process the same Domino server by provisioning different groups of users and archiving each group to a different vault store.

These tasks are controlled by the Task Controller Service.

The Domino Mailbox Archiving task is responsible for the following:

- Accessing each mail file and archiving items according to the policy set for the mail file. The task works in cooperation with the Indexing Service, which converts and indexes the items that are being archived.

To obtain an estimate of the number of items that will be archived, without actually archiving anything, you can run the task in Report Mode.

By default, Domino Mailbox Archiving tasks run automatically according to the schedule defined for the Enterprise Vault site. You can override this schedule for individual tasks.

Each mail file is processed by the Domino Mailbox Archiving task that runs on the same Enterprise Vault server as the Storage Service that hosts the associated vault store; in other words, the archive is created locally to the archiving task that processes the Domino mail file.

By default, the vault store that is used is the one defined in the Domino server properties, but you can override this setting as required by editing the properties of individual provisioning groups.

Domino mailbox archiving policies

A Domino mailbox policy supplies information for the archiving task to use when processing the target mail files, including the following:

- The archiving strategy. This can be any of the following:
 - All items over a certain age are archived.
 - Archiving actions, such as deleting the original item or creating shortcuts after archiving an item.
 - Whether shortcuts to archived items are created and what the shortcuts contain.
- Which items to archive according to the Domino form types that they use.
- Whether to archive unread items.

The mailbox policy also defines whether Enterprise Vault deletes old shortcuts. You can configure the behavior as follows:

- Deletion based on shortcut age.
- Deletion of orphaned shortcuts. These are those shortcuts that no longer have corresponding archived items. Typically the archived items have been deleted by users or storage expiry.
- Deletion of shortcuts when the retention period has elapsed. The corresponding archived items may be removed by storage expiry. You can delete the shortcuts without deleting the archived items.

You create Domino mailbox archiving policies in the Enterprise Vault Administration Console under Policies > Domino > Mailbox.

Domino mailbox archiving retention folders

The Retention Folder feature enables you to create a single folder or a hierarchy of folders automatically in users' mail files. Enterprise Vault archives these folders according to policies that you assign. If a user deletes any folders in the retention folder hierarchy, Enterprise Vault automatically recreates them.

You specify the retention folders and their retention categories in retention plans. You can create as many retention plans as you require.

Caution: These retention plans differ from those that you can create with the Enterprise Vault retention plan feature, which was introduced in Enterprise Vault 12. With that feature you can set up retention plans that associate a retention category with a number of other settings, such as a classification policy, and apply them all to one or more archives. This is not true of the retention plans that are described in this section.

For information on the Enterprise Vault retention plan feature, see the *Administrator's Guide*.

You use Enterprise Vault provisioning groups to apply retention plans to mail files. Thus, different users can have different retention folders with the appropriate retention categories. You can also define a default retention plan that Enterprise Vault applies to all users for whom a specific plan is not defined.

If a user moves a retention folder, the folder does not retain the retention plan settings. Items that are archived in the future will be archived according to the policy that applies to the folder in its new location. Items that have already been archived from the folder are unaffected and retain the original retention category.

If a user creates a subfolder beneath a retention folder, that subfolder inherits the retention folder settings. For example, if you create a 'Projects' folder users could then create a subfolder for each project. The subfolders would automatically use the retention folder settings from the parent 'Projects' folder.

You create an XML file in which you define the retention plans. You then use the `EVDominoRetentionPlans` command-line tool to upload the XML file to Enterprise Vault.

For details of how to create Domino retention plans, see the "Domino Retention Plan Tool" chapter of the *Utilities* guide.

Domino mailbox archiving desktop policies

A Domino desktop policy defines the end user's experience when using the Enterprise Vault Notes client. Its settings determine the Enterprise Vault features and functionality that the client provides.

The desktop policy settings include following options:

- Show or hide Enterprise Vault menu options, such as Search, Store, Restore, and Delete.
- Control the availability of Vault Cache and its maximum size.
- Control advanced settings for Vault Cache.

You create Domino desktop policies in the Enterprise Vault Administration Console under Policies > Domino > Desktop. When you create a provisioning group you assign a desktop policy to it. You can create multiple desktop policies if you want different provisioning groups to use different policy settings.

Domino Journal archiving

This chapter includes the following topics:

- [About Domino Journal archiving](#)
- [Domino Journal archiving policies](#)
- [Domino journal archiving database considerations](#)
- [How to set up Domino Journal Archiving](#)
- [Support for clustered Domino Journal databases](#)

About Domino Journal archiving

This section describes how Enterprise Vault archives the contents of Domino Journal databases.

You can set up Domino so that a copy of every message sent or received is saved in a Journal database. This is particularly useful if you want to implement a company email monitoring policy and vital if there is a possibility that you may have to produce email as legal evidence at some later date.

You can set up Enterprise Vault to archive all items from Domino Journal databases.

Domino Journal archiving policies

In Enterprise Vault Administration Console you create archiving policies under Policies > Domino Journaling.

A Domino Journaling Policy supplies information for the archiving task to use when processing databases in the target location; currently the only option is whether or not to expand distribution lists when archiving.

Domino journal archiving database considerations

Enterprise Vault archives from all databases in a specified folder in the server's data directory.

The normal Enterprise Vault configuration is to retain the original item until the vault store that contains the archived item has been backed up. Enterprise Vault then deletes the original item. The Domino Database Management method must not interfere with this Enterprise Vault process, which means that the "Purge and Compact" method (specified in the Journaling section of the server configuration document) is unsuitable, because there is the potential to lose items that have, for some reason, not been archived.

Thus, the Domino Journal database must have its Database Management method set to one of the following in the Journaling section of the server configuration document:

- Periodic Rollover or Size Rollover — Domino automatically places the old database in the server's data folder. In order to ensure that all items are archived from the old database, and to remove items when the archive has been backed up, you must move the old databases back to the Journal database folder.
- None — This method has the side-effect that the database will keep growing and will require manual maintenance.

You must configure your Domino Journal database appropriately so that it can be archived by Enterprise Vault.

How to set up Domino Journal Archiving

To configure Domino Journal Archiving, you need to set up the following in the Enterprise Vault Administration Console:

- A Domino Journal archive for each Domino Journal database location that Enterprise Vault will archive.
- A Domino Journaling task to perform the archiving. If required, multiple Domino Journaling tasks can be created on one computer. Domino Journaling tasks run under the control of the Task Controller Service.
- A Domino Journaling Policy. You assign this to the target location and it enables you to tailor how the Domino Journaling task archives items from the database. For example, you may want to stop the Domino Journaling task from expanding distribution lists.
- The target location to archive. In order to do this you must add the Domino domain and server and then the location of the Domino Journal database. You configure these in the Targets > Domino section of the Administration Console.

Depending on the Enterprise Vault configuration, items in Journal databases are deleted either at the time they are archived or after the vault store is backed up.

Users with access permissions to a journal archive can search for messages within the archive. Because journaled items may be confidential, it is important to give such access to a few trusted users only.

Support for clustered Domino Journal databases

Enterprise Vault can archive from Domino Journal databases on Domino Servers that are clustered using Domino application clustering.

To support clustered Journal databases, the following requirements must be satisfied:

- Each Domino Server in the cluster should be independently journaling to a local database.
- Mail journaling databases should not be configured to replicate to other Domino servers in the cluster. This includes both cluster replication and scheduled replication.
- Enterprise Vault should be configured to archive from the Domino Journal databases on each server in the cluster.

SMTP Archiving

This chapter includes the following topics:

- [About Enterprise Vault SMTP Archiving](#)
- [SMTP Archiving components](#)
- [Security and authentication](#)
- [Scalability and disaster recovery](#)

About Enterprise Vault SMTP Archiving

Enterprise Vault SMTP Archiving lets Enterprise Vault archive data that is sent to the Enterprise Vault server using SMTP protocol. Any application that can send information to an SMTP server can send data to Enterprise Vault.

Enterprise Vault SMTP Archiving enables you to do the following:

- Archive Exchange journal email directly into an archive using SMTP. You could use this functionality instead of Exchange journaling with Enterprise Vault Exchange journal archiving.
- Archive email from other messaging platforms that can send to an SMTP server, for example Office 365, Google Apps, Zimbra, and Sun Mail System.
- Archive data from any application that is capable of sending email to an SMTP server.
- Archive data from physical devices such as printers, fax machines, or scanners. This provides opportunities for electronic mailroom automation and records management workflow.
- Populate user mailbox archives from the SMTP journal feed. All email that is sent from or received by the user can be stored in the user's mailbox archive. There is no folder synchronization with SMTP Archiving.

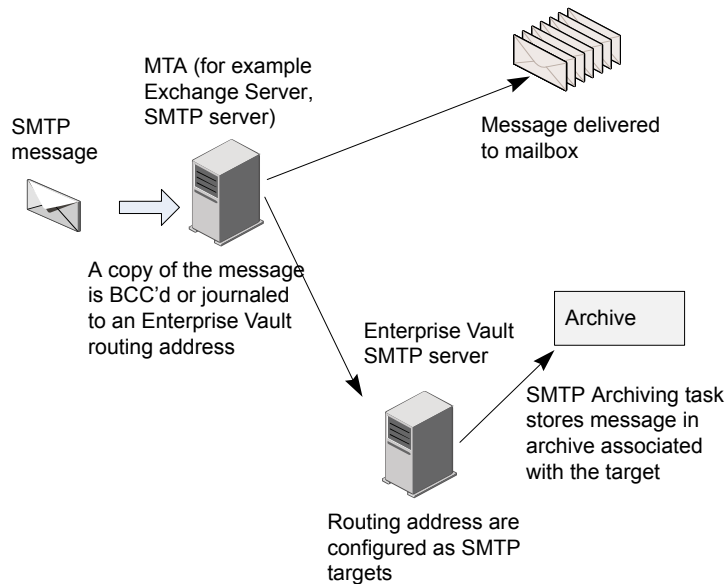
Figure 10-1 SMTP Archiving overview

Figure 10-1 shows an example of a simple SMTP Archiving environment:

- A Message Transfer Agent (MTA) receives an SMTP message from an application such as an Exchange Server.
- The MTA sends the message to the destination mailbox, and also copies or journals the message to the SMTP routing address for the Enterprise Vault SMTP server.
- The Enterprise Vault SMTP server receives the message, and adds the routing address to the message as an X-RCPT-TO header. The SMTP server then places the message as an email (.eml) file in the SMTP holding folder.
- The SMTP Archiving task processes the message file in the holding folder, and archives it in the archive that is specified for the target address. During processing, the task applies the retention category that is specified in the target properties. It ensures that Enterprise Vault indexes any X-Headers that are listed in the policy.

Enterprise Vault SMTP Archiving provides the following benefits:

- Simplified architecture and reduced infrastructure costs for journaling, with no dependency on MAPI or the need to manage large journal mailboxes. The journal feed can be sent directly to Enterprise Vault instead of maintaining multiple journal mailboxes on dedicated Exchange servers.

- Single Instance Storage. Enterprise Vault uses single instance storage to deduplicate messages.
- Provisioning groups. You can provision multiple users for SMTP Archiving at the same time. Target users that you add to the groups can be Active Directory users or SMTP addresses that are not associated with an Active Directory account. Depending on the type of provisioning group, the members' messages are all stored in the same archives, or each member is provisioned with their own archive.
- Enterprise Vault can capture and index metadata, such as BCC addresses, and journal report information.
- You can add X-Headers to messages that are sent to Enterprise Vault to override the default retention category or destination archive. You can also use X-Headers to add information to the Enterprise Vault index.
- Messages can be stored in existing archives, or in dedicated SMTP archives.
- Scalable architecture. Enterprise Vault can support multiple SMTP servers.
- You can enable message tracking to log the details of messages that each SMTP server receives.

SMTP Archiving components

Table 10-1 provides an overview of the main components of SMTP Archiving. You can configure SMTP Archiving using the Enterprise Vault Administration Console, or Enterprise Vault PowerShell cmdlets.

Table 10-1 Overview of SMTP Archiving components

Component	Description
Enterprise Vault SMTP server	<p>The SMTP server is implemented as the Windows service, Enterprise Vault SMTP service. This service is displayed in the Windows Services Console, but not in the Enterprise Vault Administration Console.</p> <p>The SMTP server manages SMTP connections. It receives messages that are sent to the Enterprise Vault SMTP server by relay Message Transfer Agents (MTAs), such as Exchange Server, or SMTP servers. The Enterprise Vault SMTP server stores the messages as .eml files in the SMTP holding folder.</p>

Table 10-1 Overview of SMTP Archiving components (*continued*)

Component	Description
SMTP Archiving task	<p>The SMTP Archiving task processes the email files in the holding folder as follows:</p> <ul style="list-style-type: none">■ Checks if the routing address in the message is an SMTP target that is enabled for archiving. If SMTP Group or Mailbox Journaling are configured, it also searches for SMTP target addresses in the To, From, CC, BCC, and Sender fields.■ Applies a policy, stores the messages in the archive, and applies a retention category.■ When an email has been archived, deletes the file from the holding folder. <p>In SMTP Group Journaling and SMTP Mailbox Journaling, you can change the default behavior for certain messages.</p>
SMTP Provisioning task	<p>The SMTP Provisioning task processes the SMTP provisioning groups. The task assigns the group policy and archives to the target users in each group. The task runs once or twice each day. It processes the highest priority groups first.</p>
SMTP holding folder	<p>The SMTP holding folder is a local folder that is assigned to the SMTP Archiving task. The folder location is in the SMTP Archiving task properties. The Enterprise Vault SMTP server places messages in the folder for the archiving task to process.</p> <p>Messages that the archiving task fails to archive are not deleted automatically from the holding folder. The messages are placed in a Failed subfolder.</p>
SMTP policies	<p>An SMTP policy is assigned to an SMTP target address. The policy defines how the SMTP Archiving task manages journal reports and X-Headers, when archiving messages that contain the target address.</p> <p>The target addresses and provisioning groups to which the policy applies are also displayed in the policy properties.</p> <p>The SMTP policies are displayed in the Administration Console, under Policies > SMTP.</p>

Table 10-1 Overview of SMTP Archiving components (*continued*)

Component	Description
SMTP target addresses	<p>The SMTP target addresses are the SMTP addresses that the Enterprise Vault SMTP server and SMTP Archiving task look for in the messages that are sent to the Enterprise Vault SMTP server.</p> <p>You can add SMTP target addresses as manual targets, or using provisioning groups.</p> <p>SMTP targets are displayed in the Administration Console, under Targets > SMTP.</p>
SMTP provisioning groups	<p>The SMTP provisioning groups let you apply policies and other settings to multiple SMTP targets at the same time. SMTP targets in the groups can be Active Directory users or SMTP addresses that are not associated with an Active Directory account.</p> <p>The SMTP provisioning groups are displayed in the Administration Console, under Targets > SMTP > Provisioning Groups.</p>
SMTP archives	<p>You can create SMTP journal archives. These are displayed in the Administration Console, under Archives > SMTP.</p> <p>If you configure SMTP Journaling or SMTP Group Journaling, the SMTP Archiving task can store SMTP messages in the following types of archives: SMTP, Shared, Exchange Journal, and Domino Journal. For SMTP Mailbox Journaling, only Internet Mail and Exchange Mailbox archives are supported.</p>

Security and authentication

The Enterprise Vault SMTP server provides the following configuration options for enhanced security and authentication for connections:

- SMTP authentication; use SASL authentication for communication with the Enterprise Vault SMTP server.
- Connection security; use TLS to communicate with the Enterprise Vault SMTP server.
- Connection Control; specify the hosts or IP addresses that can communicate with the Enterprise Vault SMTP server. SMTP connections from any other source will not be accepted.

Scalability and disaster recovery

To handle incoming journal data in large environments, you can configure multiple Enterprise Vault SMTP servers with the same SMTP routing address, for example `journal@ev.example.com`. Alternatively, you can configure multiple SMTP routing addresses for the Enterprise Vault SMTP servers, for example, `journal1@ev.example.com`, `journal2@ev.example.com`, `journal3@ev.example.com`, and so on. Load balancing can be handled using load balancing technologies, such as DNS MX records, or network load balancers.

If you assign multiple archives to a routing address or an SMTP Group Journaling provisioning group, Enterprise Vault automatically spreads the archiving load over several archives across servers.

Deploying multiple Enterprise Vault SMTP servers provides for disaster recovery. If one of the Enterprise Vault SMTP servers fails, the load balancer can route the connection request to the next available server.

You can also use clustering services to make the Enterprise Vault SMTP server highly available.

Skype for Business Archiving

This chapter includes the following topics:

- [About Skype for Business Archiving](#)
- [About Skype for Business Archiving licensing](#)
- [Skype for Business Archiving overview](#)
- [Skype for Business Archiving components](#)

About Skype for Business Archiving

Enterprise Vault enables you to archive Skype for Business Server 2015 and Lync Server 2013 instant messaging (IM) and conferencing (meeting) communications. This in turn helps support the compliance requirements that are specified in industry and government regulations.

Enterprise Vault archives the following:

- Peer-to-peer instant messages.
- Conferences (meetings), which are multiparty instant messages.
- Conference content, including uploaded content (for example, handouts) and event-related content (for example, joining, leaving, uploading, sharing, and changes in visibility).
- Whiteboards, polls, and Q&A sessions that are shared during a conference.

Note: Data for some conference features, including whiteboards and polls, is stored in a proprietary Microsoft XML format. Enterprise Vault cannot index this XML, so you cannot search for text in these items. In addition, Enterprise Vault cannot display them in a visual format.

Enterprise Vault cannot archive the following:

- Persistent Chat conversations.
- Peer-to-peer file transfers.
- Audio/video for peer-to-peer instant messages and conferences.
- Desktop and application sharing for peer-to-peer instant messages and conferences.
- Items that are collected via Microsoft Exchange Server storage.

About Skype for Business Archiving licensing

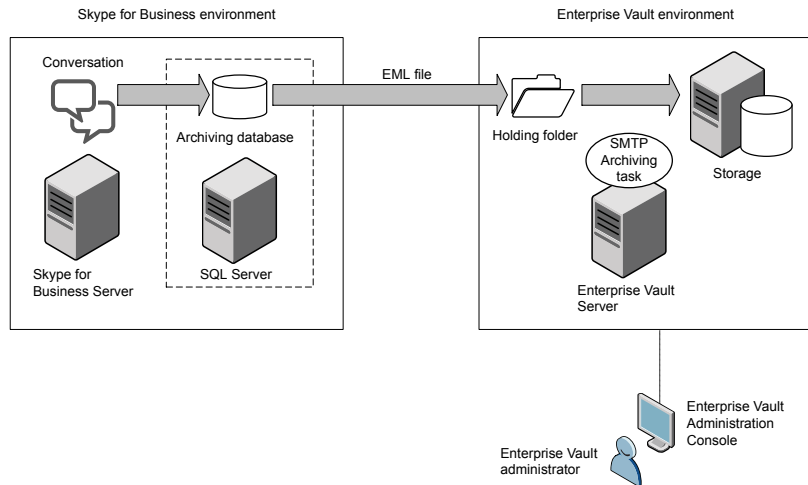
You require a license to archive conversations from Skype for Business. See *Installing and Configuring* for information about licenses.

Skype for Business Archiving overview

Skype for Business archiving with Enterprise Vault links the environments of Skype for Business Server and Enterprise Vault together. It is assumed that archiving in Skype for Business has been enabled for groups of users or individuals, and that the appropriate user-level archiving policies have been applied. Enterprise Vault requires Skype for Business to be configured to actively archive conversations to a SQL Server store. These conversations are then transferred to the Enterprise Vault environment as EML files for archiving.

[Figure 11-1](#) illustrates the process when archiving a conversation from Skype for Business.

Figure 11-1 Skype for Business archiving overview



The archiving process is as follows:

- 1 A conversation takes place between users of Skype for Business.
- 2 Skype for Business archives the conversation in a database, according to the Skype for Business archiving policy.
- 3 Enterprise Vault exports the conversation from the database into the holding folder as an email (.eml) file. One .eml file is created per conversation.
- 4 Enterprise Vault marks the conversation as purged in the database.
- 5 The SMTP Archiving task processes the email (.eml) file in the holding folder, and stores it in the archive that is specified for the Skype for Business target. During processing, the task applies the retention category or retention plan that is specified in the target properties.

Skype for Business Archiving components

[Table 11-1](#) provides an overview of the main components of Skype for Business archiving. You can configure Skype for Business archiving using the Enterprise Vault Administration Console or Enterprise Vault PowerShell cmdlets.

Table 11-1 Overview of Skype for Business Archiving components

Component	Description
SMTP Archiving task	<p>The SMTP Archiving task processes the .eml conversation files in the holding folder as follows:</p> <ul style="list-style-type: none"> ■ Applies the policy that is associated with the Skype for Business target. ■ Stores the conversations in the archive that is associated with the target. ■ Applies the target retention category or retention plan. ■ By default, deletes the conversation files from the holding folder when archiving is completed successfully. If errors occur, the task does not delete the files.
SMTP holding folder	<p>The SMTP holding folder is a local folder that is assigned to the SMTP Archiving task. The folder location is in the SMTP Archiving task properties. Conversation files are placed in the folder for the archiving task to process.</p> <p>Conversations that the archiving task fails to archive are not deleted automatically from the holding folder. The conversations are placed in a Failed subfolder.</p>
SMTP policies	Skype for Business Archiving does not require an SMTP policy.
Retention categories and retention plans	<p>Retention categories are assigned to conversations when they are archived. This categorization makes it easier to retrieve conversations because it is possible to search by retention category.</p> <p>Users can select retention categories and assign them to Skype for Business conversations. When Enterprise Vault archives a conversation, it is stored with the appropriate retention category.</p> <p>With a retention plan, you can associate a retention category with a number of other settings, such as a classification policy and the criteria for discarding expired conversations. All these settings can be applied to one or more archives.</p> <p>For more information, see "Working with retention categories and retention plans" in the <i>Administrator's</i> guide.</p>

Table 11-1 Overview of Skype for Business Archiving components
(continued)

Component	Description
Skype for Business target	<p>The target Skype for Business server from which conversations are archived. The target properties contain the following settings:</p> <ul style="list-style-type: none"> ■ A Skype for Business computer. ■ An Enterprise Vault server. ■ A retention category or retention plan. ■ An archive. ■ A vault store. <p>The Skype for Business targets are displayed in the Administration Console, under Targets > Skype for Business. They can also be listed out using the <code>Get-EVSkypeForBusinessTarget</code> PowerShell cmdlet.</p> <p>Only one Skype for Business target per Enterprise Vault server is supported.</p>
Skype for Business conversation archives	<p>Skype for Business conversations can be stored in an existing archive, or you can create new archives to hold the conversations. Although you can use any archive type, the archive types that are typically used for Skype for Business Journaling are SMTP, Shared, Exchange Journal, or Domino Journal archives. These archive types support sampling. We recommend that you do not use Exchange Mailbox or Internet Mail archives for Skype for Business Journaling.</p>

Enterprise Vault Accelerators

This chapter includes the following topics:

- [About the Enterprise Vault Accelerators](#)
- [Differences between the Enterprise Vault Accelerators](#)
- [About Compliance Accelerator](#)
- [About Discovery Accelerator](#)

About the Enterprise Vault Accelerators

The Accelerator products are specialized add-on applications to Enterprise Vault.

Compliance Accelerator enables sampling and monitoring of an organization's electronic messages. Features include monitored employee management, message sampling and item reviewing and exporting.

Discovery Accelerator is designed for data-mining activities, such as finding messages and documents to present as evidence in legal cases or that relate to an internal inquiry. Features include case management, advanced, multi-archive searching and item reviewing and publishing.

Both Compliance Accelerator and Discovery Accelerator require separate licenses.

Running the Compliance Accelerator and Discovery Accelerator server software concurrently on the same computer is not currently supported.

Differences between the Enterprise Vault Accelerators

It is important to understand that Compliance Accelerator and Discovery Accelerator were created for different purposes.

Discovery Accelerator is an electronic discovery and review system that integrates with Enterprise Vault services and archives. Discovery Accelerator lets authorized users search for, retrieve and preserve, analyze, review, mark, and export or produce emails, documents, and other electronic items for lead counsel examination or court-ready production—rapidly and in a cost-effective manner.

Using attorneys and external counsel to review large numbers of items is costly. With Discovery Accelerator, you can create a hierarchy of reviewers for a discovery action or case, with different levels of reviewers able to assign certain review marks. In this way, paralegal staff and non-legal staff can perform an initial review of search and collection results and leave only the privileged, relevant, or questionable items for counsel. Optionally, you can then produce the relevant items with an appropriate "Bates" number or else simply export them from Discovery Accelerator in various formats, such as PST, Domino NSF database, HTML, MSG, and ZIP.

Compliance Accelerator enables companies to implement an ongoing electronic message monitoring policy, in order to meet requirements set by the company or by an industry regulatory body, such as SEC. As a company would set up monitoring for departments, work in Compliance Accelerator is set up in terms of departments and monitored employees within the departments. Compliance Accelerator deals exclusively with email, social media messages, instant messages, Bloomberg messages, and faxes; you cannot search for or view documents held in file system archives.

Compliance Accelerator samples messages as they are archived and adds the messages automatically to the review set for the department. Although reviewers have a similar role to Discovery Accelerator reviewers, the marking scheme in Compliance Accelerator is simpler.

Both Accelerator products have an automatic search facility. In Compliance, this can be used for ongoing monitoring of company mail for particular behavior, such as unacceptable language or sending confidential information outside the company.

About Compliance Accelerator

The use of Compliance Accelerator does not in itself make an organization compliant with regulatory requirements, such as NASD 3010 and 3110. However, it provides a tool for implementing a compliance strategy.

A company's compliance strategy may typically require the following:

- A certain percentage of employees' electronic messages to be captured and checked by compliance officers on a regular basis. Electronic messages may include email, instant messages, social media messages, faxes, and, particularly in the financial sector, Bloomberg messages. Depending on the compliance strategy, monitoring may be required of internal messages (messages sent between employees in certain departments), or external messages (messages sent between employees in certain departments and people outside the company).
- Regular searches to be run on electronic messages to capture any instances of unacceptable language or illegal business practice, such as insider dealing.
- Messages to be stored securely for several years and retrieved quickly at any time.
- Detailed audit information showing the review history of a message.

Enterprise Vault can be used to provide secure archiving of all company messages. Compliance Accelerator then builds on Enterprise Vault to provide the following additional features to assist an organization in implementing their compliance strategy:

- The option to collect a configurable, random sample of items from a variety of archives (Exchange or Domino journal mailbox, SMTP, or shared), and add them to the set of messages for review.
- A system for defining the employees that are to be monitored and grouping them in an organizational structure that reflects the departments within the company. The messages of certain employees (called "exception employees"), such as senior managers, can be kept separate and reviewed by specially assigned reviewers.
- A client application that enables compliance administrators to configure Compliance Accelerator to fit the requirements of the company's compliance strategy. For example, departments of employees can be managed, granular access permissions can be assigned to designated compliance officers and administrators, automatic searches of messages can be scheduled, and sets of words for searches can be added and managed.
- The client application also enables designated compliance officers or reviewers to read and mark the messages that are captured.
- History information about all employees monitored, messages captured and the review process applied to each message is kept securely in a SQL database.

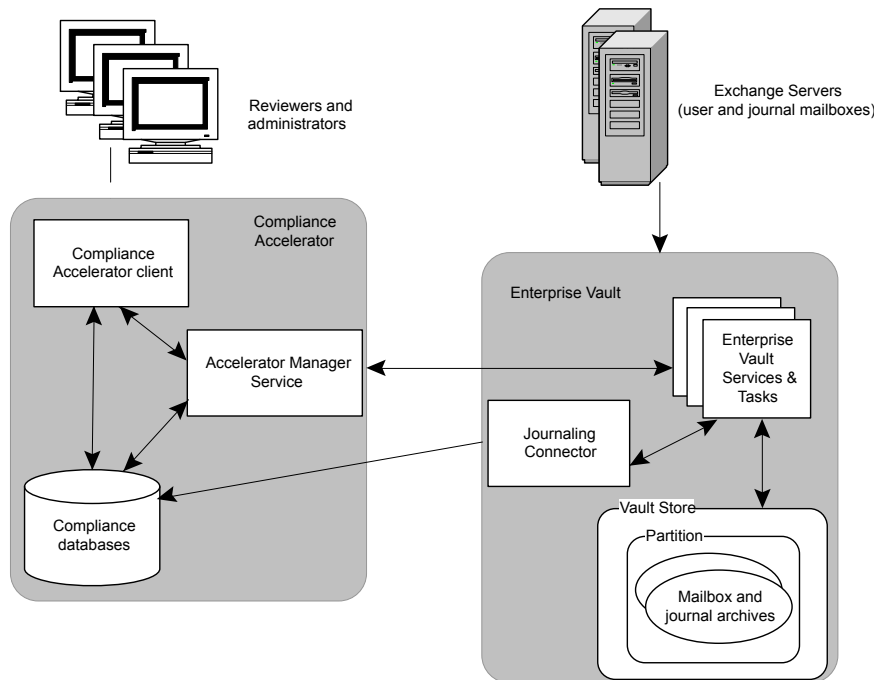
Compliance Accelerator components

Table 12-1 describes the Compliance Accelerator components.

Table 12-1 Compliance Accelerator components

Component	Notes
Compliance Accelerator client	The client is used by Compliance Accelerator administrators to set up and manage the system and by reviewers to access the items that they are to mark.
Accelerator Manager website	This website lets you set up and manage multiple Compliance Accelerator databases in which to store your data. For example, this facility lets you split your data by date range or organizational unit.
Enterprise Vault Accelerator Manager service	This service handles the requests from the Compliance Accelerator client and works with the Enterprise Vault components to access archives, perform searches, and so on.
Customer database	The customer database is a SQL database in which Compliance Accelerator stores details of departments, user roles, search results, and more. You can set up multiple customer databases.
Configuration database	The configuration database is a SQL database that specifies the location of the customer databases and stores details of the SQL Server, database files, and log files to use.

Figure 12-1 shows how Compliance Accelerator components work with Enterprise Vault to access archived data. In this diagram, data is archived from Exchange servers. Alternatively, data could be archived from Domino mail servers.

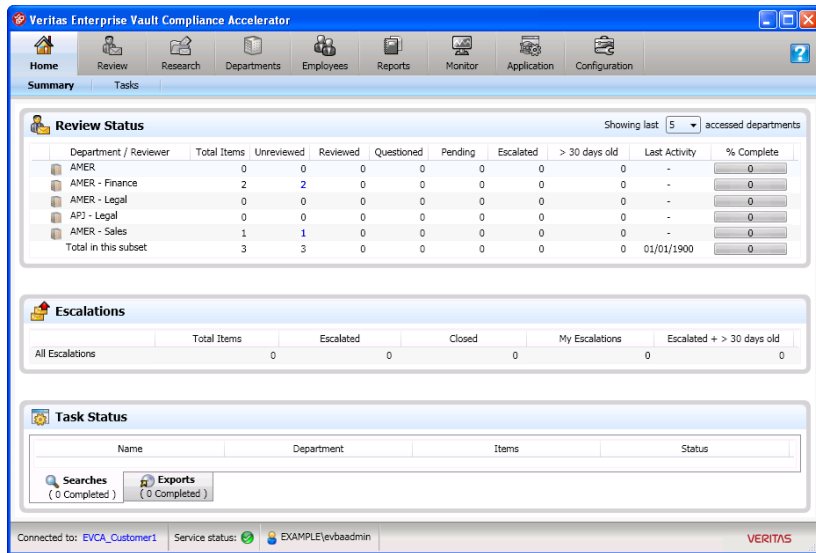
Figure 12-1 How Compliance Accelerator works with Enterprise Vault

The Accelerator Manager service finds out the location of archives and Enterprise Vault services and tasks from the Enterprise Vault Directory. It uses the Enterprise Vault Indexing service to run searches and the Enterprise Vault Storage service to view messages.

In small installations, the Compliance Accelerator components can all reside on the same computer as Enterprise Vault. In larger installations, a more distributed setup is recommended, with Compliance Accelerator components on a separate computer from Enterprise Vault.

The Compliance Accelerator client application

The client application provides administrators with application and department management facilities and reviewers with facilities for reviewing and marking messages. The role that is assigned to the logged-in user determines which facilities are available.

Figure 12-2 The Compliance Accelerator client

Department administration with Compliance Accelerator

Within the Compliance Accelerator system, monitored employees are organized into departments, which would normally reflect the company organization. Tasks such as adding employees to departments and assigning reviewers can be done using drag and drop or right click options.

To perform any tasks within the Compliance Accelerator system, users must be assigned an appropriate role. Roles contain a configurable group of permissions. New roles to suit your organization can be created and the permissions assigned to most roles can be changed.

Administrators can be assigned system roles that allow them to perform application-level tasks. Other administrators can be assigned roles for managing a particular department only. Reviewer roles can enable a compliance officer to mark messages and add comments or enable them only to view messages. The Compliance Supervisor role enables senior compliance officers to check (or appraise) the work of more junior officers, assign exception status to certain employees, and ensure that they are reviewed by an appropriate reviewer.

A reviewer can also be made a delegate for another reviewer.

Compliance Accelerator searches

In addition to taking a random sample of messages, you may want to monitor employees' messages for unacceptable language or business conduct. For this requirement, you can use Compliance Accelerator to run searches on the messages that Enterprise Vault has archived.

Compliance Accelerator offers a wide range of search criteria: words and phrases to look for, date ranges, message size, type and direction, author and recipient details and attachment details.

Note that for phrase searching in the content of a message or attachment, the indexing on the Enterprise Vault archives must be set to full.

Searches can be run automatically, on a regular basis using a schedule for search runs, or manually for specific tasks. An application administrator can set up a search to run in several departments concurrently, while searches run by a department administrator are limited to the department that the administrator is responsible for. Messages returned by the searches can then be added to the review set for the department reviewers to read. Accepting search results can be automatic or manual.

By default, Compliance Accelerator identifies and remove duplicate items from the results of a search, and so prevents them from appearing in the review set. To determine whether one item is a duplicate of another, Compliance Accelerator compares the metadata properties of the items, such as their author display names, subjects, and number of attachments.

When a search is created in Compliance Accelerator, the Accelerator Manager service contacts the Enterprise Vault Indexing service to run the search. If the search results are accepted, details of the search and results are stored permanently in the customer database.

Schedules for searches use the SQL Server Agent service.

Reviewing messages with Compliance Accelerator

Reviewers use the Compliance Accelerator client to read each message captured and allocate a review status mark to indicate whether the message is acceptable or requires further investigation. The status marks can be customized to suit your review system.

There is an audit trail of the review process. This makes it possible to see when a message was reviewed, who reviewed it, and what action they advocated.

Compliance Accelerator reporting

A reporting system enables you to generate a variety of reports, such as the roles assigned to users, the percentage of messages captured for each employee in a department, and the progress of department message reviewing.

Exporting messages with Compliance Accelerator

Messages can be exported from the Compliance Accelerator system together with any reviewing marks and comments that have been assigned. This allows users who do not have access to the Compliance Accelerator system to view messages in a department review set.

Compliance Accelerator configuration data

Configuration data for the Compliance Accelerator system, such as departments, employees, roles and monitoring policies, will typically be added using the Compliance Accelerator client. This lets you add employees to be monitored by synchronizing with Windows user accounts and groups. Alternatively, you can bulkload the data using XML files.

About Discovery Accelerator

A company can use Discovery Accelerator to search across their Enterprise Vault archives and quickly find documents and messages that meet criteria for inclusion in a particular inquiry or legal case. All types of archives can be searched: user and journal mailbox archives, file system archives, SharePoint archives, and public folder archives.

Although Discovery Accelerator uses the search facility available in Enterprise Vault, it adds the necessary security that is vital in legal discovery. To ensure that search criteria and results are secure, Discovery Accelerator stores details of all the searches performed, the criteria used and the items found. These details can be viewed but cannot be changed or deleted from the system.

Discovery Accelerator's review system provides an orderly and efficient process for checking all the items found by searches. Using this system, permitted users view items found by the searches and assign review marks depending on the item's relevance to the case. As reviewers can see the marks applied by other reviewers, they can quickly select only the items that they need to work on, avoiding duplication of effort. The case administrator can track the progress of all reviewers for a case.

Using lawyers to review large numbers of items can be very costly. With Discovery Accelerator, a hierarchy of reviewers can be created for a case, with different levels of reviewers able to assign certain review marks. In this way, less expensive,

non-legal staff can perform an initial review of search results, leaving only the relevant or questionable items for lawyers.

The relevant items can then be assigned an appropriate Bates number and published, typically in a PST file, for presentation as evidence in court. Once an item has been published as evidence in a particular case, it is secured, together with its review history, in the Discovery Accelerator system. No further marks can be added by reviewers, and the item cannot be republished in that case. If required by the court, a report can be produced that shows the review process applied to a particular item.

By default, Discovery Accelerator automatically identifies and removes duplicate items from the review set and from the items in an export run. To determine whether one item is a duplicate of another, Discovery Accelerator compares the metadata properties of the items, such as their author display names, subjects, and number of attachments. In addition, for items in analytics-enabled cases only, Discovery Accelerator compares the content of the items.

To ensure that items in a case are not deleted, an administrator can assign legal hold status to the items associated with a case. This means that the items cannot be deleted manually or automatically (by Enterprise Vault expiry deletion).

The analytics facility in Discovery Accelerator

Optionally, you can choose to enable *analytics* on a Discovery Accelerator case. This facility provides additional analyses of the metadata and content of items that are collected in the case. Among the extra benefits that analytics provides are the options to do the following:

- Set up rules by which Discovery Accelerator automatically marks or categorizes the items that it adds to the case.
- Classifying large numbers of items without much human intervention ultimately results in better and smaller review sets for manual review.
- Examine and review entire conversation threads in one view.
- Conduct quick or advanced searches within the items in a case.

These facilities deliver a new review experience that is known as *Guided Review*.

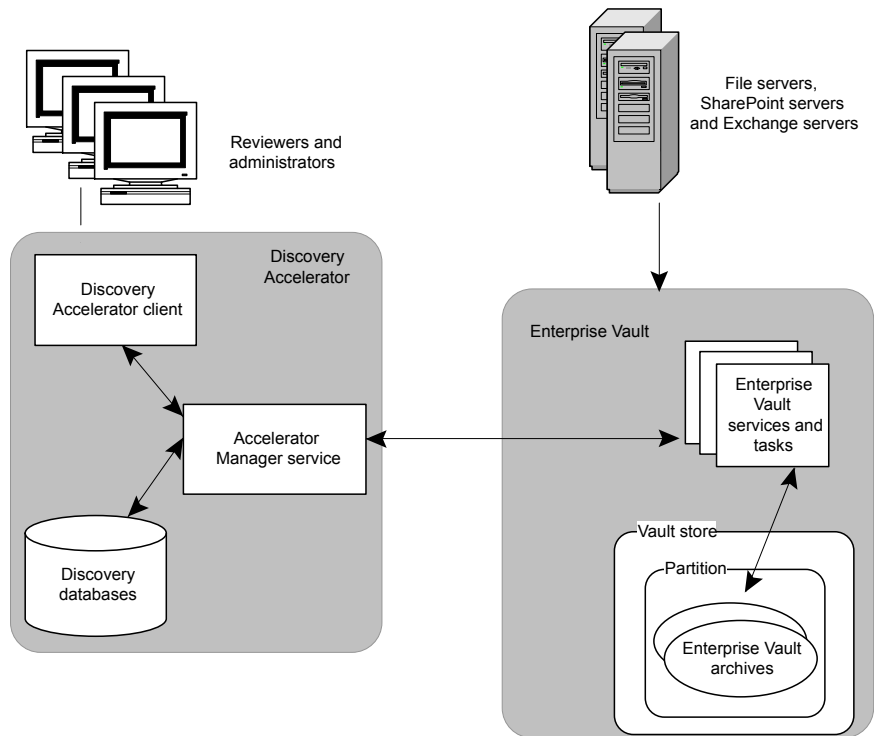
Discovery Accelerator components

[Table 12-2](#) describes the Discovery Accelerator components.

Table 12-2 Discovery Accelerator components

Component	Notes
Discovery Accelerator client	The client is used by Discovery Accelerator administrators to set up and manage the system and by reviewers to access the items that they are to mark.
Accelerator Manager website	This website lets you set up and manage multiple Discovery Accelerator databases in which to store your data. For example, this facility lets you split your data by date range or organizational unit.
Enterprise Vault Accelerator Manager service	This service handles the requests from the Discovery Accelerator client and works with the Enterprise Vault components to access archives, perform searches, and so on.
Customer database	<p>The customer database is a SQL database in which Discovery Accelerator stores details of cases, user roles, search results, review marks and tags, and more.</p> <p>You can set up multiple customer databases.</p>
Configuration database	The configuration database is a SQL database that specifies the location of the customer databases and stores details of the SQL Server, database files, and log files to use.
Custodian Manager website (optional)	This website lets you store the details of the <i>custodians</i> (individual employees) and custodian groups for which you want to search with Discovery Accelerator. A custodian group is any collection of employees, such as Windows or Domino groups and distribution lists, Active Directory or Domino LDAP searches, and Active Directory containers.
Discovery Accelerator API website (optional)	<p>This website lets you use the Discovery Accelerator API to integrate third-party tools with the software, and thereby retrieve data from or export it to a Discovery Accelerator customer database.</p> <p>For more information on the Discovery Accelerator API, contact Veritas Support.</p>

Figure 12-3 shows an overview of the Discovery Accelerator components and how the application integrates with Enterprise Vault.

Figure 12-3 How Discovery Accelerator works with Enterprise Vault

Note the following:

- All communication between Discovery Accelerator and Enterprise Vault is through the Enterprise Vault Accelerator Manager service.
- The Enterprise Vault Accelerator Manager service uses the Enterprise Vault Directory service to find archives, databases and services. Enterprise Vault Storage services are used to preview items and fetch original items. Enterprise Vault Indexing services are used to search for items in the archives.

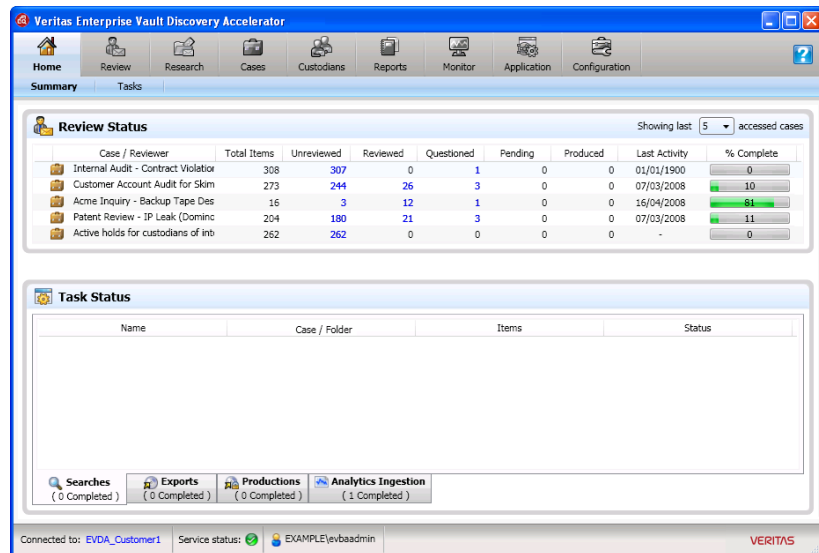
Before installing Discovery Accelerator, you need to consider your present Enterprise Vault and SQL server installation and plan how you want to integrate Discovery Accelerator. In small installations, Discovery Accelerator can be installed on the same computer as Enterprise Vault and the SQL Server, but it is often installed on a different computer.

Large companies may want to distribute the Enterprise Vault and Discovery Accelerator installation even more in order to spread the load over several servers.

Discovery Accelerator client application

The client application provides administrators with application and case management facilities and reviewers with facilities for reviewing and marking the items in cases. The role that is assigned to the logged-in user determine which facilities are available.

Figure 12-4 The Discovery Accelerator client



Case administration with Discovery Accelerator

Within the Discovery Accelerator system, information is organized in cases.

A case includes the following:

- Case administration and reviewer roles
- Case marking schemes to use when reviewing items
- Searches and search results
- Target addresses to use when searching messages
- Items that are *produced* (published) or exported from Discovery Accelerator

To perform any tasks within the Discovery Accelerator system, users must be assigned an appropriate role. Roles contain a configurable group of permissions. New roles to suit your organization can be created and the permissions assigned to most roles can be changed.

Administrators can be assigned system roles that allow them to perform application-level tasks, such as creating cases and running searches in multiple cases. Other administrators can be assigned roles for managing a particular case only. Reviewer roles can enable a user to review items and apply particular marks and add comments.

Discovery Accelerator searches

The criteria for a Discovery Accelerator search can include items in a particular date range, messages to or from particular addresses or words or phrases in the file content (including attachments to messages). If an organization indexes custom attributes in Enterprise Vault, these custom attributes can also be used in Discovery Accelerator searches.

Note that for phrase searching in the content of an item, the indexing on the archives must be set to full.

Search results can be accepted or rejected. Accepted results are added to the review set for the case reviewers to read and can be assigned to particular reviewers. Accepting search results can be automatic or manual.

When a search is created in Discovery Accelerator, the Accelerator Manager service contacts the Enterprise Vault Indexing service to run the search. If the search results are accepted, details of the search and results are stored permanently in the customer database.

Reviewing items with Discovery Accelerator

Reviewers use the Discovery Accelerator client to check each item and allocate a mark from the case marking scheme to indicate whether the item is relevant to the case or requires further investigation. The marks that a reviewer can use depends on the role that the user has been assigned.

The marking scheme can be customized to suit your cases. For example, you may want a hierarchy of reviewers, so that only the minimum number of items need to be reviewed by lawyers. A set of marks can be assigned to each level of reviewer.

There is an audit trail of the review process. This makes it possible to see when an item was reviewed, who reviewed it, and what action they advocated.

Although reviewers can view stored items and add marks, they cannot change the actual items in any way.

Producing and exporting items with Discovery Accelerator

When all reviewers have finished checking the items in the case, they can be produced in Discovery Accelerator in a suitable format to present as evidence in

court. When items have been produced, their status is changed to Produced and they are locked so that no further changes can be made to their status. Produced items are also assigned a formal Bates number that you define for documents associated with each case.

The export option is less formal than production. You can use this to allow another person, who does not have access to Discovery Accelerator, to view items in the review set. Exported items are not locked and their status is not changed; you can continue to work on them after they have been exported. Exported items are given an export ID, but this is different from a Bates number.

Items can be produced or exported as MSG, HTML, or PST files. With HTML, you can include item comments and review history information. With PST format, you have the option to set a password and size limit for each PST.

Discovery Accelerator configuration data

Configuration data for the Discovery Accelerator system, such as cases, target addresses and roles, is typically added using the Discovery Accelerator client. Alternatively, you can bulkload the data using XML files.

Building in resilience

This chapter includes the following topics:

- [About Enterprise Vault and VCS](#)
- [About Enterprise Vault and Windows Server Failover Clustering](#)
- [About Enterprise Vault building blocks](#)

About Enterprise Vault and VCS

You can integrate Enterprise Vault with Veritas Cluster Server (VCS) to provide a highly-available solution for Enterprise Vault.

Supported VCS configurations and software

Note: This documentation uses the terms *VCS* and *Storage Foundation HA for Windows (SFW HA)* throughout. However, in version 7.0 of the cluster software, these terms became *Veritas InfoScale Availability* and *Veritas InfoScale Enterprise* respectively.

Both active/passive and N+1 configurations are supported, but active/active configurations are not.

In an active/passive configuration, a dedicated spare server is available for each Enterprise Vault server, ready and waiting for the primary server to go down. In an N+1 configuration, there is a computer for each Enterprise Vault server and then one or more spare servers waiting for any of the active servers to fail over.

The following software must be installed:

- A supported version of VCS
- Enterprise Vault

- A supported version of Windows Server

For supported versions of software, see the Enterprise Vault [Compatibility Charts](#).

Neither Compliance Accelerator nor Discovery Accelerator must be installed on any server in the planned cluster. These products are not supported within a cluster. However, an unclustered Compliance Accelerator or Discovery Accelerator can reference a clustered Enterprise Vault virtual server.

About Enterprise Vault and the VCS GenericService agent

The VCS GenericService agent brings online the following Enterprise Vault services, monitors their status, and takes them offline:

- Admin service
- Directory service
- Indexing service
- Shopping service
- Storage service
- Task Controller service
- SMTP service (only if the Enterprise Vault SMTP Archiving components are installed and configured on the Enterprise Vault server)


See the *Cluster Server Bundled Agents Reference Guide* for detailed information on the GenericService agent, including the resource type definitions, attribute definitions, and sample configurations.

The GenericService agent detects an application failure if a configured service is not running. When this happens, the Enterprise Vault service group is failed over to the next available system in the service group's system list, and the services are started on the new system. This ensures continuous availability for the data that Enterprise Vault is managing and archiving.

Typical Enterprise Vault configuration in a VCS cluster

[Figure 13-1](#) illustrates a typical configuration.

Figure 13-1 Active/passive failover configuration



Here, the volumes for the Enterprise Vault services data are configured in a cluster disk group on shared storage. The Enterprise Vault virtual server is configured on the active node (System 1). If System 1 fails, System 2 becomes the active node, and the Enterprise Vault virtual server comes online on System 2.

About Enterprise Vault and Windows Server Failover Clustering

You can cluster Enterprise Vault in a failover cluster on a supported version of Windows Server, to provide a high availability solution for Enterprise Vault.

For details of supported versions, see the Enterprise Vault [Compatibility Charts](#).

High availability is provided by creating an Enterprise Vault cluster server that can fail over between physical nodes in the cluster. When Enterprise Vault services are running on a cluster server they operate with virtual IP addresses, a virtual computer name, virtual Microsoft Message Queues, and highly available shared disks. When a failure occurs, the cluster software can move the server's resources to a different physical node in the cluster.

Supported Windows Server Failover Clustering configurations

An Enterprise Vault cluster consists of:


- One or more primary nodes, each normally hosting an Enterprise Vault cluster server.
- One or more failover nodes: standbys that can take over the job of hosting an Enterprise Vault cluster server if a primary node fails.

Enterprise Vault does not permit "active/active" cluster configurations. That is, only one Enterprise Vault cluster server can run on a clustered node at any one time. You can configure Enterprise Vault in any operation mode that adheres to this restriction, such as:

- An active/passive failover pair: a primary node with a dedicated failover node.
- N+1 (hot standby server): two or more primary nodes share a single failover node. Only one node failure can be accommodated at any one time.
- N+M: an extension of the hot standby concept with N primary nodes and M failover nodes. Only M node failures can be accommodated at one time.
- N+M *any-to-any*: identical to N+M, except that there is no need to fail back to the original node after a failover. When the original node becomes available again, it can operate as a failover node.

Typical Enterprise Vault configuration in a Windows Server failover cluster

Figure 13-2 illustrates a typical configuration.

Figure 13-2 Enterprise Vault in an active/passive failover pair configuration

In this example:

- NODEA and NODEB are the two Enterprise Vault nodes in the failover cluster. NODEA is the primary node. NODEB is the failover node.
- The SQL server and Microsoft Exchange may also be configured in the cluster: this does not affect Enterprise Vault.
- The volumes for the Enterprise Vault services data are configured on shared storage.
- The Enterprise Vault cluster server is configured on the primary node, NODEA. If NODEA fails, the cluster server's resources fail over to NODEB, and the cluster server comes online on NODEB.

About Enterprise Vault building blocks

An alternative to clustering is to implement Enterprise Vault building blocks. This is part of a straightforward methodology for deploying Enterprise Vault in a way that is both scalable and reliable. You can easily extend a solution made from building blocks to increase capacity. In addition, you can configure building blocks for several different failover schemes such as active/passive and active/active.


A building block is a repeatable unit of Enterprise Vault functionality hosted on an application server. By adding sufficient building blocks, Enterprise Vault can be extended to scale to the largest of workloads, as shown in [Figure 13-3](#).

Each building block comprises a set of Enterprise Vault services that deliver the same functionality.

The combined services form a unit known as a building block, which is identified by a building block name.

Each building block is hosted on a particular server, but in the event of a disaster on that server Enterprise Vault can failover a building block to another server (without the need for clustering technology). Failover does not require any reconfiguration on the desktop and users are unaffected.

Figure 13-3 Enterprise Vault building blocks



Building blocks and high availability

The building block approach can be used in an N+1 architecture in which an idle server stands by to replace any one of N servers that may fail. The "+1" server could also be used to run as an extra, hot standby server running IIS to take some of the workload from the other servers.

In an active/active architecture an existing server can take on the additional workload of a failed server. An active/active solution can be implemented if servers have spare capacity to handle extra work. After failover the net performance may drop, depending on the amount of spare capacity. Thought must be given to this during the design stage.

Planning component installation

This chapter includes the following topics:

- [About planning component installation](#)
- [About valid computer names for Enterprise Vault servers](#)
- [Prerequisites for Enterprise Vault components when planning installation](#)
- [Factors to consider when planning deployment of Enterprise Vault components](#)
- [Enterprise Vault Directory Service installation planning](#)
- [Where to set up the Enterprise Vault Services and Tasks](#)
- [Enterprise Vault databases and planning their installation](#)
- [Vault store groups and vault stores installation planning](#)
- [Administration Console installation](#)
- [Installation planning for client components](#)

About planning component installation

This section helps you to plan which components of Enterprise Vault you need to install on the computers in your Enterprise Vault site.

About valid computer names for Enterprise Vault servers

An Enterprise Vault server that has Unicode characters in its computer name may not operate properly. We strongly recommend that the computer names of your Enterprise Vault servers contain ASCII characters only.

Prerequisites for Enterprise Vault components when planning installation

There is a standard list of required software and system settings that is needed for the majority of the Enterprise Vault components. When you install some of the Enterprise Vault components, other components may also be installed automatically. For this reason you need to check that the prerequisites are correctly set up on each computer in your Enterprise Vault site. Details of all required software and settings are in the *Installing and Configuring* guide.

Factors to consider when planning deployment of Enterprise Vault components

There are many different ways in which you can deploy the Enterprise Vault components.

Factors that influence how you deploy Enterprise Vault include the following:

- The amount of data you anticipate storing and how long you intend to retain the data.
- The amount of computing resources you have at your disposal, that is, disk space as well as memory.
- The network connections you have between computers.
- Whether you intend to use Enterprise Vault single instance storage.
- Whether you intend to use Index Server groups.
- Whether you are using offline storage, in which case the storage space you require for the vault stores will be reduced.

Enterprise Vault Directory Service installation planning

The Enterprise Vault Directory Service must be able to connect to the Directory database, but there is no requirement for the Directory Service to be on the same computer as the database.

If necessary, Enterprise Vault sites can share a Directory Service and its associated Directory database. The Directory Service uses different DNS site aliases to distinguish between the different sites. With the exception of the types of files that can be archived (Exchange messages classes or Domino forms), there is no sharing of configuration information across sites.

The Directory Service, with its associated database, must be available at all times to all the computers in the Enterprise Vault sites that it serves, including any servers running Compliance Accelerator or Discovery Accelerator.

The Directory database is managed using SQL management tools.

See [“Enterprise Vault databases and planning their installation”](#) on page 158.

Enterprise Vault initially allocates the following amount of permanent disk space for the Directory database:

- 10 MB for the data files
- 25 MB for the transaction log files

After the database has been populated, the amount of data in the Directory database will not change very much over time.

Where to set up the Enterprise Vault Services and Tasks

The core Enterprise Vault Server component includes the following main services:

- Admin Service
- Directory Service
- Indexing Service
- Shopping Service
- Storage Service
- Task Controller Service

Archiving and PST Migration tasks are created as required after Enterprise Vault is installed and configured. As they run under the Task Controller Service, this must be installed on any computer that is to host a task.

After you have installed the Enterprise Vault Server, the configuration wizard enables you to specify which particular services or items to run on that computer. The following sections give details of where to set up individual services.

You need to install the Enterprise Vault Server component on any computer on which you want to run any Enterprise Vault Service or Task. When you install the Enterprise Vault Server component, the Enterprise Vault Administration Console is also automatically installed.

To configure Enterprise Vault SMTP Archiving on one or more Enterprise Vault servers, you must install the Enterprise Vault SMTP Archiving components on each server on which you want the feature to run.

If you want to install on the Exchange Server computer, perhaps as part of a pilot, then you can do so, but this is not recommended for production use.

[Table 14-1](#) provides summary information that will help you decide where to set up the Enterprise Vault Tasks and Services.

Table 14-1 Where to set up the Enterprise Vault Tasks and Services

Task or Service	Number in Enterprise Vault site	Notes
Client Access Provisioning Task	One per Active Directory or Domino domain.	A single Provisioning Task can process multiple Provisioning Groups.
Domino Journaling Task	As many as required to archive your archiving target Domino journaling location.	<p>You will need to ensure that Enterprise Vault has the correct access to the Domino server, domain and journaling location.</p> <p>Domino Journaling database must have its Database Management method set to Periodic Rollover, Size Rollover or None.</p>

Table 14-1 Where to set up the Enterprise Vault Tasks and Services
(continued)

Task or Service	Number in Enterprise Vault site	Notes
Domino Mailbox Archiving Task	One per Enterprise Vault server.	You cannot install more than one Domino Mailbox Archiving Task on an Enterprise Vault computer. Each task in a single site can archive multiple Domino mail servers in more than one Domino domain. Multiple Domino Mailbox Archiving Tasks can archive the same Domino mail server.
Domino Provisioning Task	One per Domino domain.	You can install only one Domino Provisioning Task on an Enterprise Vault server. One Provisioning Task will typically process multiple Provisioning Groups.
Exchange Journaling Task	As many as resources permit.	An Exchange Journaling Task can process multiple journal mailboxes. You can install more than one Exchange Journaling Task on an Enterprise Vault computer. The Vault Service account must have access in Active Directory to the Exchange domains associated with the Exchange Servers being archived.
Exchange Mailbox Archiving Task	As many as the number of Exchange Servers to be archived; one per Microsoft Exchange Server.	You can install more than one Exchange Mailbox Task on an Enterprise Vault computer. Exchange Mailbox Tasks in a single site can archive multiple Exchange Servers in more than one Exchange domain. The Vault Service account must have access in Active Directory to the domains associated with the Exchange Servers being archived.

Table 14-1 Where to set up the Enterprise Vault Tasks and Services
(continued)

Task or Service	Number in Enterprise Vault site	Notes
Exchange Provisioning Task	One per Exchange domain.	You can install more than one Exchange Provisioning Task on an Enterprise Vault server. One Provisioning Task will typically process multiple Provisioning Groups.
Exchange Public Folder Task	As many as required to process your public folder hierarchy.	A single Exchange Public Folder Task can process one or more branches of the hierarchy. The Vault Service account must have access in Active Directory to the domains associated with the Exchange Servers being archived.
File System Archiving Task	As many as required to process the Archive Points defined.	The Task should run on the same computer as the Storage Service that you select when you create a File Server Archiving Target. The Vault Service account must have access to the file system being archived.
Move Archive Task	One per Enterprise Vault storage server.	Enterprise Vault creates this task automatically when you run the Move Archive wizard. See "How to plan installing the Move Archive task" on page 155.
SharePoint Task	As many as required to archive your SharePoint Archiving Targets.	The SharePoint Task will run on the same computer as the Storage Service that you select when creating the SharePoint Task. You can install more than one SharePoint Task on an Enterprise Vault computer.

Table 14-1 Where to set up the Enterprise Vault Tasks and Services
(continued)

Task or Service	Number in Enterprise Vault site	Notes
SMTP Archiving task	One per Enterprise Vault SMTP server.	<p>A site can have multiple Enterprise Vault SMTP servers. Each SMTP server has an Enterprise Vault SMTP service and an SMTP Archiving task. A message that contains a configured SMTP target address can be processed by any of the SMTP servers in the site.</p> <p>The SMTP Archiving task is assigned a local SMTP holding folder. As the task processes .eml message files in the holding folder, the logon account for the task must have full access to the holding folder.</p>
Accelerator Service	One.	Compliance Accelerator and Discovery Accelerator should not be run on same computer.
Admin Service	One per computer with other Enterprise Vault Services.	Mandatory. The Admin Service is installed automatically when you install the Enterprise Vault Server component.
Indexing Service	At least one.	<p>You cannot install more than one Indexing Service on an Enterprise Vault computer but you can have more than one Indexing Service in a site.</p> <p>There must be sufficient storage space to hold the indexes.</p> <p>For the Indexing Service to function correctly, you must install it on a computer whose year format is equivalent to the Gregorian year (currently 2018).</p>

Table 14-1 Where to set up the Enterprise Vault Tasks and Services
(continued)

Task or Service	Number in Enterprise Vault site	Notes
Shopping Service	One per Enterprise Vault server.	Requires IIS. There must be sufficient disk space on the computer to hold the shopping baskets until users delete them.
Storage Service	At least one.	You cannot install more than one Storage Service on an Enterprise Vault computer. There must be sufficient storage space to hold the vault stores.
SMTP service	One per Enterprise Vault SMTP server	The Enterprise Vault SMTP service is created when you install the Enterprise Vault SMTP Archiving components on the Enterprise Vault server. There can be several Enterprise Vault SMTP servers in a site. The SMTP server settings that you configure for the first Enterprise Vault SMTP server are shared by any Enterprise Vault SMTP servers that you subsequently add to the site.

How to plan installing Exchange Mailbox Archiving Tasks

To archive user mailboxes, there must be one Exchange Mailbox Archiving Task for each Microsoft Exchange Server computer that is being served in the Enterprise Vault site. A Microsoft Exchange Server can be served by only one Exchange Mailbox Archiving Task. However, there can be more than one Exchange Mailbox Archiving Task running on the same computer in an Enterprise Vault site, servicing different Microsoft Exchange Servers.

The Exchange Mailbox Archiving Task interfaces most directly with Microsoft Exchange Server (using MAPI), the Storage Service and the Indexing Service. Therefore, it makes most sense to install the task on the same computer as the

Storage Service and Indexing Service with which it interfaces. This reduces the amount of data that has to cross network links.

How to plan installing Exchange Journaling Tasks

An Exchange Journaling Task can archive from multiple Microsoft Exchange Server journal mailboxes. However, a journal mailbox can be served by only one Exchange Journaling Task. There can be more than one Exchange Journaling Task running on the same computer in an Enterprise Vault site, processing one or more journal mailboxes on one or more Exchange Servers.

Like the Exchange Mailbox Archiving Task and Exchange Public Folder Tasks, the Exchange Journaling Task interfaces most directly with Microsoft Exchange Server, the Storage Service and the Indexing Service. For this reason, you are recommended to install the Exchange Journaling Task on the Enterprise Vault computer that hosts the associated Storage and Indexing Services.

How to plan installing Exchange Public Folder Tasks

There can be many Exchange Public Folder Tasks. Each Exchange Public Folder Task process can archive one or more branches of the Microsoft Exchange Server public folder hierarchy.

When an Exchange Public Folder Task archives from a public folder it uses the archiving settings you specify in the associated Exchange Public Folder Policy in the Administration Console. You can, however, modify the archiving settings for individual public folders, by changing their properties from within Outlook, in the mailbox that owns the public folder. The Enterprise Vault Outlook Add-In must be installed.

The Exchange Public Folder Task requirements are similar to those for the Exchange Mailbox Task; as it interfaces most directly with Microsoft Exchange Server and the Storage and Indexing Services, it makes most sense to install the task on the same computer as the associated Storage and Indexing Services. This reduces the amount of data that has to cross network links.

How to plan installing Domino Journaling and Mailbox Archiving Tasks

As the planning requirements for these tasks are complex, it is important that you read the information in the *Installing and Configuring* guide during the planning stage of your Enterprise Vault environment.

How to plan installing the Move Archive task

The Move Archive task is installed automatically on storage servers when you run the Move Archive wizard. The first time you move an archive associated with a particular storage server, Enterprise Vault creates a Move Archive task on that server. The task is configured to start automatically each time the task controller service starts.

How to plan installing the Storage Service

There must be at least one Storage Service in an Enterprise Vault site. There can be only one Storage Service on a computer. If you want more than one Storage Service in an Enterprise Vault site, they must run on separate computers. It is usually best to put the Storage Service on the same computer as the Indexing Service.

When associated Storage and Indexing Services are on different computers, the computers must be able to communicate over a fast connection.

If you anticipate that your archiving needs are great and require much storage, consider having more than one Storage Service in the site. This would require more than one computer running Enterprise Vault in the site.

How to plan installing the Indexing Service

There must be at least one Indexing Service in an Enterprise Vault site. However, there can be only one Indexing Service on a computer. If you have more than one Indexing Service in an Enterprise Vault site they must run on separate computers. It is usually best to install the Indexing Service on the same computer as the Storage Service.

Note: For the Indexing Service to function correctly, you must install it on a computer whose year format is equivalent to the Gregorian year (currently 2018). For example, setting the computer's regional format to Thai causes the Indexing Service to fail, as the current year in the Thai solar calendar falls outside the date range that the service supports.

In larger or distributed deployments, consider using Index Server groups to spread the indexing load. Ensure that associated Storage and Indexing Services are either collocated, or can communicate over fast connections.

An Indexing Service can manage simultaneously the indexes for many archives, which may be stored in different vault stores on different Storage Service computers.

The index locations assigned to the Indexing Service must have sufficient disk space to store the indexing data.

The indexes are organized as follows:

- There is a separate index for each archive.
- Each index consists of a set of related files.
Always back up and restore these files as a complete set; never restore only some of the files.
The number of files will both increase and decrease over time.
- Files are held in folders; for most archives, there is one folder (index volume) for the index files. When an index volume becomes full, Enterprise Vault automatically creates a new one. This may occur for FSA archives, Exchange or Domino Journal archives, and Exchange Public Folder archives, but is unlikely to occur for normal user mailbox archives.

You specify during configuration the index locations where index volumes are to be created. If you want multiple index locations, you are recommended to spread them over different physical devices. (These locations are sometimes referred to as "index root paths" in error and diagnostic messages.)

If you specify more than one index location, indexes for new archives and new index volumes are spread over the locations. If you create Index Server groups, then indexing the associated vault stores is shared by all the Index Servers in the group. The index volumes are spread over the locations assigned to the Index Servers in the group.

You can choose how much information is indexed for items in an archive using the indexing level; this can be brief, or full indexing. If you want to be able to search item content for phrases, for example using Compliance Accelerator or Discovery Accelerator, then you need to use full indexing.

The more information that is indexed about an item, the easier it is to search for it. However, the more information that is indexed about an item, the more disk space is required for the index. The size of index data for an item varies with the indexing level.

Table 14-2 shows the estimated size of an index as a percentage of the unarchived size of the item for the different indexing levels.

Table 14-2 Estimated size of indexing data

Indexing level	Estimated size
Brief	4%
Full	12%

So, if you have allowed in the region of hundreds of gigabytes (or terabytes) of space for your vault stores, you are likely to require in the region of gigabytes (or tens of gigabytes) for your indexes.

How to plan installing the Shopping Service

There is one Shopping Service on each Enterprise Vault server.

When users select items to restore, the items are held in shopping baskets, which are held in Shopping storage locations on the Shopping Service computer.

Make sure you have sufficient disk space on the computer to store the shopping data.

How to plan installing File System Archiving

For Windows file servers you need to install the Enterprise Vault FSA Agent on each file server on which you want to leave placeholder shortcuts, or obtain data for FSA Reporting. The FSA Agent provides the services required by these features.

File System Archiving from NetApp Filers, and Dell EMC Celerra/VNX devices is configured differently: The FSA Agent services or their equivalent do not run on the file server; but on the Enterprise Vault server or another Windows server.

How to plan installing SharePoint Archiving

A single SharePoint Task in the Enterprise Vault site can archive several virtual SharePoint Servers or SharePoint site collections, but if required you can configure multiple SharePoint Tasks on an Enterprise Vault Server. In addition to the SharePoint Task on the Enterprise Vault Server, you need to install Enterprise Vault SharePoint components on each SharePoint Server computer that is to be archived.

For users to be able to use the Archive Search Web Part, this also needs to be installed on each SharePoint Server computer.

How to plan installing SMTP Archiving

You need to install the Enterprise Vault SMTP Archiving components on each Enterprise Vault server that is to archive SMTP messages using the SMTP Archiving feature. On each Enterprise Vault SMTP server there must be adequate space for the local SMTP holding folder.

When you configure SMTP Archiving, the Enterprise Vault SMTP server settings and target configuration information are shared with all the Enterprise Vault SMTP servers in the site. This means that any Enterprise Vault SMTP server in the site can archive messages sent to any SMTP target in the site.

If there is a large volume of SMTP messages for journaling, it is advisable to spread the archiving load over several Enterprise Vault storage servers. You can do this by creating several journal archives in different vault stores. A different Enterprise Vault Storage service should manage each of the vault stores.

If messages containing target SMTP addresses are liable to be stored in several archives, then you may want to ensure that single instance storage is implemented. Check which vault store is used for each of the target archives. You also need to check the vault store sharing that is set for each of the vault stores.

How to plan installing Accelerator Services

The Accelerator Manager Service can be installed on the Enterprise Vault computer, but would typically be installed on a separate computer. You can only have one Accelerator Manager Service in a site.

Currently, running Compliance Accelerator and Discovery Accelerator at the same time on one computer is not supported.

The Accelerator web applications require IIS and are typically installed on the Accelerator Manager Service computer, but they can be installed on a different IIS computer, if required.

In large installations, where frequent large searches are performed using Compliance Accelerator or Discovery Accelerator, the Indexing and Storage Services are likely to be heavily used. Performance may be improved by using computers with more memory and CPU for the Indexing and Storage Services.

Enterprise Vault databases and planning their installation

Enterprise Vault has the following core databases:

- The Enterprise Vault Directory database. There is just one of these and it may be shared by more than one Enterprise Vault site.
- The vault store databases. There is one of these for each vault store in a site.
- The fingerprint databases. There is one of these for each vault store group. One possible exception is the Default Upgrade Group, which Enterprise Vault creates if you previously upgraded to Enterprise Vault 8.0. Enterprise Vault does not create a fingerprint database for the Default Upgrade Group until you configure sharing for it.
- The Enterprise Vault Monitoring database. There is one of these for each Enterprise Vault Directory database. If multiple Enterprise Vault sites share a

Directory database, then they must also share a Monitoring database. The Monitoring database holds the status information that the Monitoring agents gather about the Enterprise Vault servers.

If you configure FSA Reporting, Enterprise Vault also uses one or more FSA Reporting databases to hold the FSA Reporting data that it gathers from the file servers.

SQL Server must be installed and set up before configuring Enterprise Vault. Note that the sort order/collation setting for the SQL Server installation must be case-insensitive; case-sensitive installations are not supported.

Microsoft SQL Server does not need to be on the same computer as the Directory Service computer, nor does it have to be on the vault store computers. On each computer, run Microsoft SQL Enterprise Manager to register the SQL Servers you have installed.

It is possible, perhaps as part of network reconfiguration, to change the instance of SQL Server that manages the Directory database. The process is described in the Administration Console help.

When configuring Enterprise Vault, you are asked to supply the following:

- The Vault Service account details. This enables Enterprise Vault to create the Directory database and vault store databases
- The SQL server location and the data and log file locations for the Enterprise Vault Directory database
- The SQL server location and the data and log file locations for the Enterprise Vault Monitoring database

Enterprise Vault creates databases with the following names:

- EnterpriseVaultDirectory – for the Directory database
- EVvaultstore – for the vault store databases
- EnterpriseVaultMonitoring – for the Monitoring database

Enterprise Vault also creates the following database locations:

- VaultDev – for the Directory database data
- VaultLog – for the Directory database transaction log
- EVvaultstore – for the vault store database data
- EVvaultstore – for the vault store database transaction logs
- A location for the Monitoring database data
- A location for the Monitoring database transaction log

Each vault store database contains an entry for each item that is archived in the associated vault store, so the vault store databases will grow over time. Only when an item is deleted from the archive will references to it be deleted from the relevant vault store database.

When you create a vault store group, Enterprise Vault creates a fingerprint database for the group. The New Vault Store Group wizard provides the following options for configuring the database filegroups:

- A basic configuration, in which Enterprise Vault locates the primary filegroup and all the non-primary filegroups on one device.
- An option to configure additional locations for the 32 non-primary filegroups.

The non-primary filegroups can grow rapidly in size when you use single instance storage. For best performance, spread the non-primary filegroups across multiple locations.

On the first occasion that you configure a file server for FSA Reporting, a wizard helps you to create an FSA Reporting database. When you configure additional file servers for FSA Reporting, you can choose to use the existing FSA Reporting database or to configure an additional one.

For database storage requirements, see "Storage requirements" in *Installing and Configuring*.

If Compliance Accelerator or Discovery Accelerator is installed, there will also be separate databases created for each of these. The Accelerator databases can be managed by the same SQL Server as the Enterprise Vault databases or a different one, as required.

In Compliance Accelerator, data about all departments, captured items, searches and items reviewed is kept permanently in the databases. Similarly, in Discovery Accelerator, data about cases, searches, search results and items reviewed is kept permanently in the databases. Ensure that there is enough storage space for the growing database files. The databases can be managed using standard SQL administration tools.

Vault store groups and vault stores installation planning

Vault stores are contained in vault store groups. A vault store group defines the outer boundary for sharing items in Enterprise Vault single instance storage.

See ["About single instance storage"](#) on page 26.

A vault store can comprise one or more partitions. A vault store partition resides on a storage device. When you set up Enterprise Vault, you specify on which device each partition is created.

If a partition is created on an NTFS volume set, it must be contained within a single NTFS volume set, and cannot span volume sets. Bear in mind that a volume set can contain multiple physical disks, so your partition can be spread over more than one disk. There can be more than one partition on an NTFS volume set. If you need to create or extend NTFS volume sets to accommodate the partition, use the Windows Administrative Tools.

The amount of space you will require for your vault stores depends on the following:

- The average size of items that you anticipate archiving.
- The number of items you anticipate archiving in a day to that vault store.
- The reduction in archive storage space due to single instance storage.
- The reduction in storage space due to compression of files.
- How often you anticipate archiving items from the Archiving Target.
- How long you anticipate archived items remaining in the vault store.

Note: For assistance with storage sizing, contact your Veritas supplier.

The first two items in this list give you some idea of the amount of space that you would need in the vault store for one day, before any savings due to single instance storage or compression.

Single instance storage and compression can provide significant savings in storage space.

If you use expiry to delete archived items automatically at the end of their retention period, then the amount of space required will not increase as much over longer periods.

If you set up an aggressive archiving policy that archives items from the archiving targets after about two weeks, users are likely to want to access those items frequently. Keep the archived items online in the vault store for a period of time so users can easily restore them. You will need adequate online space to accommodate the data in the vault stores.

If, on the other hand, you set up archiving so that items are not archived from the archiving targets for about six months or a year, users are much less likely to want to access the archived items. You could then almost immediately move the archived items from the vault stores to secondary storage.

You can select a default vault store for archives when creating archiving targets.

For Exchange mailbox archiving, you can select a different vault store at provisioning group, Exchange Server or Enterprise Vault server level. For public folder archiving, you can select a different vault store at Exchange Server or Enterprise Vault server level. For Domino mailbox archiving, you can select a different vault store at provisioning group or Domino Server level.

The vault store selected will be used for automatically-enabled archives associated with the target. Similarly, for SharePoint archiving, the vault store selected for a target will be used for any auto-enabled site collection archives associated with the target.

The vault store databases do not need to be on the same computer as the vault stores. The databases are managed using SQL Enterprise Manager.

See [“Enterprise Vault databases and planning their installation”](#) on page 158.

Administration Console installation

When you install the Enterprise Vault Server component, the Administration Console is also automatically installed.

You can install more than one Administration Console in an Enterprise Vault site.

The following are installed with the Administration Console:

- PST Migrator
- NSF Migrator
- Export archive
- Microsoft Exchange Forms
- Enterprise Vault documentation

Installation planning for client components

Users can access archived items from within their email clients by using the following facilities:

- Enterprise Vault Outlook Add-In. Outlook 2013 and later users can also use the facilities in the Enterprise Vault Office Mail App, if you choose to deploy it.
- Enterprise Vault Client for Mac OS X.
- Enterprise Vault extensions for OWA.
- Enterprise Vault extensions for Notes and Domino Web Access clients.

In addition, users can access archived items from a stand-alone browser by using facilities such as the Enterprise Vault Search facilities.

See [“How to access items in archives”](#) on page 16.

Installation planning for Outlook Web App (OWA) and RPC over HTTP components

As the planning requirements for these components are complex, it is important to read the information in the *Installing and Configuring* guide when you plan your Enterprise Vault environment.

Planning your archiving strategy

This chapter includes the following topics:

- [About archiving strategies](#)
- [Where to define default settings for the Enterprise Vault Site](#)
- [How to allow users flexibility](#)
- [How to plan the types of items to archive](#)
- [How to define your archiving policy for user mailboxes](#)
- [How to plan the archiving policy for journal mailboxes](#)
- [How to plan the archiving strategy for Exchange public folders](#)
- [How to plan an archiving strategy for FSA](#)
- [How to plan a strategy for SharePoint archiving](#)
- [How to plan settings for retention categories](#)
- [How to plan the automatic deletion of archived items](#)
- [How to plan PST migration](#)
- [How to plan NSF migration](#)
- [How to plan shared archives](#)
- [How to plan vault stores and partitions](#)
- [How to plan single instance storage](#)

- [About Enterprise Vault reports](#)

About archiving strategies

This section considers the main points that you need to consider when defining how archiving is to be implemented in your organization and describes configurations settings that will enable you to implement your strategy. Your archiving strategy will determine how you set up Enterprise Vault.

The considerations that you must weigh include the following:

- What types of data you want to archive: user mailbox items, journaled items, public folders, files on file servers or SharePoint servers.
- Which mailboxes, journal mailboxes, public folders, file servers and SharePoint servers to target for archiving.
- How often to archive the archiving targets.
- The vault stores to use for particular archives.
- Required indexing levels.
- The policies needed for each archiving target.
- The Retention Categories and Retention Plans to define.
- Whether to use the classification feature to classify the items that Enterprise Vault archives.
- How and when shortcuts and items are to be deleted.
- The type of shortcuts to use in mailboxes and on file servers.
- Whether to enable archives automatically for mailboxes, public folders and SharePoint site collections.
- Whether to keep safety copies, and where to keep them.
- Which settings to apply across the site.
- How to configure single instance storage for your archives.
- Location of PST and NSF files and how they are to be archived.
- How much flexibility you give users.
- How much Enterprise Vault functionality to show on the users' desktops.

These options, and others, are considered in the following sections.

Where to define default settings for the Enterprise Vault Site

Table 15-1 lists default settings that are to apply across the Enterprise Vault Site.

Table 15-1 Site properties

Tab	Settings
General	<ul style="list-style-type: none"> ■ The Vault site alias and description. ■ The protocol and port to use for the Web Access application. ■ A system message for users of the Web Access application, if required. ■ The following site properties settings apply only to Exchange Server archiving: PST holding area details. ■ A note for administrators, if required.
Archive Settings	<ul style="list-style-type: none"> ■ The default retention category. ■ If users perform actions that could potentially update the retention categories of their archived items, whether to allow these updates to take place. ■ Whether users can delete items from their archives. ■ Whether the items that users have deleted can be recovered. ■ The length of time for which the deleted items remain available for recovery. ■ The length of time for which to retain the transaction history for archives.
Storage Expiry	<ul style="list-style-type: none"> ■ The schedule to run storage expiry to delete from archives any items that are older than the retention period assigned. ■ Whether expiry is based on an item's modified date or its archived date.
Site Schedule	<ul style="list-style-type: none"> ■ The schedule to run automatic, background archiving.
Archive Usage Limit	<ul style="list-style-type: none"> ■ If required, you can set limits on the size of archives.
Indexing	<ul style="list-style-type: none"> ■ Indexing level: brief or full. ■ Email content that should not be indexed, such as disclaimers. ■ How long indexing subtasks are retained before they are deleted.

Table 15-1 Site properties (*continued*)

Tab	Settings
Advanced	<ul style="list-style-type: none">Advanced settings that you can use to tune Enterprise Vault indexing within the Enterprise Vault site. <p>Note: Do not change the Indexing settings unless your technical support provider advises you to do so.</p>
Monitoring	<ul style="list-style-type: none">Performance counters for monitoring Enterprise Vault.

Some of these can be overridden for individual archiving targets. For example, the default setting in Archive Settings for Exchange Server archiving can also be defined at the provisioning group, Exchange Server and Enterprise Vault server levels. Consider the planned structure of your Site and decide where these defaults need to be defined.

Site properties enable you to have overall control of some aspects of archiving in the site. These include the schedule for archiving runs and whether users can delete items in archives. Archiving target and policy settings enable you to tailor archiving for particular archiving targets or groups of users.

Site properties settings are described in the Administration Console help. Refer to the help for more details about these settings.

How to allow users flexibility

For Exchange Server mailbox and public folder archiving, you can allow users the ability to change some of the archiving settings for their own profile. You do this by clearing the appropriate Lock check boxes in policy properties. If you select the check box, users are not allowed to change the setting.

Decide which settings you want users to be allowed to change and leave these unlocked.

How to plan the types of items to archive

In Exchange Server archiving, the types of message files to be archived are defined by message classes. In Domino Mailbox Archiving, the types of files to be archived are defined by Domino Forms. The default list of Message Classes or Domino Forms to be archived by Enterprise Vault are defined in Directory properties. These apply across all sites using the Directory.

In archiving policies you can select the types of messages to be archived from the targets.

In your email system there may be other types of messages that you want to archive. These messages may come from a third party, or may be ones that you define yourself. Decide which types of message files that you want Enterprise Vault to archive.

How to define your archiving policy for user mailboxes

You must decide on the different groups of users that you want to archive (Provisioning Groups) and the policies to use for each group for automatic mailbox archiving.

Archiving from mailboxes can be based on one of the following strategies:

- Age: Enterprise Vault archives items automatically as soon as they reach the specified age.
- Quota: archiving keeps a percentage of each user's Exchange mailbox storage limit free. (Archiving based on quota does not apply to Exchange public folders.)
- Age and quota: Enterprise Vault performs age-based archiving first. If age-based archiving does not make the required percentage of mailbox storage limit free, quota-based archiving continues until the required percentage is reached.

Consider archiving based on age and quota when archiving by age only or quota only do not give the results you want. When archiving is based only on age, it may not archive enough items to keep some mailboxes within their quota. When archiving is based only on quota, some mailboxes may not come close to the Exchange mailbox storage limit. In this case, Enterprise Vault does not even archive older items.

You can set a minimum age limit so that recent items are not archived to meet the archiving criteria.

Additionally, it is possible to archive the largest items first, so you get the greatest benefit from archiving relatively few items. This option is particularly helpful if you use archiving based on quota, or on age and quota.

For example, you can use the following settings. With these settings, Enterprise Vault first archives items that are larger than 3 MB and older than 30 days. It then archives all items older than 60 days:

- Never archive items younger than 30 days
- Start with items larger than 3 MB
- Archive items when they are older than 60 days

As another example, you can use the following settings. With these settings, Enterprise Vault first archives items larger than 1 MB. It then archives items until each mailbox has 10% of its storage limit free:

- Never archive items younger than 30 days
- Start with items larger than 1 MB
- Archive items until mailbox available storage reaches 10%

Enterprise Vault obeys the minimum age limit even if the result is that the available storage percentage is not achieved.

For more detailed information about archiving based on quota or age and quota, see the *Administrator's Guide*.

How to plan enabling mailboxes

You have the option of automatically enabling new mailboxes for archiving. If you choose this option, all mailboxes that are new to Enterprise Vault are treated in the same way, using the defaults set for the provisioning group. The archives that are automatically created for the newly enabled mailboxes are stored in the default vault store and items in the mailboxes are assigned the default Retention Category. Items in the mailboxes are indexed according to the default Indexing settings.

If you think that some users will want to change the archiving defaults, you can initially suspend archiving for the new mailboxes that have been automatically enabled. Archiving does not start until the user of the new mailbox enables archiving. This gives users the opportunity to change the archiving defaults.

When you first start using Enterprise Vault, you should gradually introduce users to Enterprise Vault. You are recommended to initially choose not to automatically enable new mailboxes. This way you can control the number of users who use Enterprise Vault and can assess your requirements and the performance of Enterprise Vault. Use the Enable Mailbox wizard in the Administration Console to enable mailboxes manually.

As you get more familiar with your usage of Enterprise Vault and can predict the resources that are required, then you can choose to automatically enable new mailboxes for archiving.

How to plan controlling the appearance of the desktop

You can use settings in the desktop policy to hide Enterprise Vault menu options, buttons, and property sheets in the client interface for the following users:

- Outlook users with the Outlook Add-In installed
- OWA users

- Notes users
- Mac OS X users with the Enterprise Vault Client for Mac OS X installed

Additionally, you can control the appearance of the client for those Outlook users who work with Vault Cache and Virtual Vault.

For users who use other email clients, you can configure customized shortcuts on the mailbox policy to enable users to access archived items from shortcuts.

How to plan the archiving policy for journal mailboxes

Implementing journal archiving is important if your archiving strategy is to satisfy regulatory compliance requirements. If you are installing Compliance Accelerator or Discovery Accelerator in your site, you are strongly recommended to use journal archiving with full indexing.

Enterprise Vault supports Domino Server journal archiving and Exchange Server journal archiving.

In the Journaling Policy you can specify a number of settings, including whether the Journaling Task is to expand distribution lists.

If you want to process certain messages in a particular way, you may want to consider using custom filters. This is described in detail in the *Setting up Exchange Server Archiving* and *Setting up Domino Server Archiving*.

As journal archives contain sensitive information, consider carefully which users should have search access on journal archives.

How to plan the archiving strategy for Exchange public folders

Since items in public folders are usually there for many people to read, some of the options that are acceptable for individual users are not appropriate.

The main control you have for public folders is the age at which you archive items. For example, you could decide to archive items that have been unmodified for 60 days, rather than 90 days. If the items are likely to be in frequent use, though, archiving them too soon means that users will frequently be retrieving them.

You can also control archiving by size and there is also a minimum age limit, below which items are not archived.

For example, you could use the following combination of settings, which would have the effect of archiving all items larger than 10 MB:

- Never archive items younger than 0 days
- Start with items larger than 10 MB
- Archive remaining items, taking oldest items first and stopping when all items older than 99 years are archived

As an alternative, you could use the following settings, which would have the effect of archiving all items older than 90 days, but would first remove items larger than 5 MB and older than 60 days:

- Never archive items younger than 60 days
- Start with items larger than 5 MB
- Archive remaining items, taking oldest items first and stopping when all items older than 90 days are archived

You can also, however, choose to archive larger items sooner than smaller items. The setting is in the Exchange Public Folder Policy.

Also, if storage space is important, you could decide not to keep safety copies of archived items. An alternative is to choose to store safety copies on the storage queue. By moving the safety copies to the storage queue there is an immediate space gain in Exchange Server.

The types of items to archive are defined in the Directory.

See [“How to plan the types of items to archive”](#) on page 167.

How to plan enabling public folders

When you set up archiving for public folders, you add an Exchange Public Folder Task and assign to it one or more public folder root paths in the public folder hierarchy. You also specify a vault store for the task to use.

New public folders can be auto-enabled for archiving, with a new archive created automatically for each new public folder created under the specified root path.

This feature should be used with caution and only where the creation of public folders is strictly controlled.

How to plan an archiving strategy for FSA

Why you are implementing file system archiving will to a great extent determine your archiving strategy. For example, you may need to store legal and financial documents for a set period of time in order to comply with industry regulations. You

need to consider which files need to be stored and where they are in your file system. Obviously, the more organized your file system is, the easier it will be to define the files to archive.

When you have decided which file servers need to be included in archiving, you need to decide where the Archive Points should be created. These mark the top of each folder structure that is to be stored in a single archive. To ensure that the archive does not fill up too quickly, you need to consider the size of the folder structure below each Archive Point.

How files below an Archive Point are archived is determined by assigning an archiving policy to the target Archive Point. Archiving policies include rules to filter what files are archived and settings such as the Retention Category to be assigned, how permissions on files are to be treated and the type of shortcuts to be created. You can use Volume Policies for the whole folder structure and override settings for certain folders using Folder Policies. For example, in most folders under a particular Archive Point you may want to archive documents and leave shortcuts. However, in one folder you just want a copy of the files archived and the original files left in the folder. You can see that archiving policies could potentially become very complicated if your file system is not organized. By keeping policies simple, it is easier to manage your file system archiving.

If files are going to be removed from the file server when they are archived, you need to consider whether to leave shortcuts or just let users search the archives for stored data. Your choice will to some extent depend on the experience that you want users to have when accessing archived files; how transparent do you want the operation to be.

If you want to use shortcuts, you need to decide what kind to use—placeholder or internet shortcuts. Placeholder shortcuts are more transparent than internet shortcuts. However, your choice will also depend on the type of storage devices you are using.

Check that your backup and virus scanning applications honor the file system offline attribute if you use placeholder shortcuts. If they do not, they may try to recall each file during a run or scan. A backup mode program is included with Enterprise Vault to enable you to switch the file server into backup mode before running your application. This prevents files from being recalled.

Using shortcuts on frequently modified files may mean that a large number of versions of the file are stored. You may want to consider running the Enterprise Vault pruning process at intervals, to delete earlier version of files in the file system archives.

If you want to set up special folders on target file systems that are created and managed by Enterprise Vault, you can use the Managed Folders feature.

How to plan a strategy for SharePoint archiving

As with file system archiving, why you are implementing SharePoint archiving will determine the archiving policies that you implement. This may be to meet compliance regulations or to control disk space usage on the server, or both. If compliance is a requirement, then this will obviously take priority.

When you have decided which SharePoint servers are to be archived, you also need to work out which site collections to include as archiving targets and the archiving policy that you want applied to each one.

Auto-enabling site collections means that new sites created under the target top level site are automatically enabled for archiving and the policy for the site collection applied. This reduces the need for manual administration intervention but requires that your sites are well defined and organized. If, for example, on a SharePoint Server you create a site collection but decide that some of the subsites are not to be archived, you should not set auto-enable on and must manually create site collection objects for each new subsite created.

To enable users to search for archived items from SharePoint site pages, you will need to install the Archive Search web part on each SharePoint server being archived. You can then place the Archive Search web part on site pages, as required. The Archive Search web part behaves in a very similar way to the SharePoint Portal Server search, which users may already be familiar with.

If versioning is enabled, the version history link also needs to be added to allow users to see earlier versions of a document that have been archived.

If you are archiving drafts, and some users have access to approved documents but not drafts, then you need to consider carefully the SharePoint draft options that you select in the SharePoint archiving policy.

How to plan settings for retention categories

You can categorize archived items using retention categories. Settings for a retention category include the retention period, which defines how long items are stored for.

A retention category also includes settings that let you do the following:

- Prevent automatic deletion of items at the expiry of the retention period.
- Prevent user deletion of items.

Enterprise Vault comes with a single, predefined retention category called Default Retention Category. The retention period is forever.

You can create other retention categories as required. For example:

- If you have documents that you must retain for ten years for legal requirements, you could create a new retention category, named Legal, with a retention period of ten years.
- You might want to keep the minutes of your regular meetings for three years. In this case, create another new retention category, named Minutes, and give it a retention period of three years.

You select one of the retention categories, predefined or otherwise, to be the default for all archiving in the site.

If permitted, Outlook users can change the retention category for their Exchange Server mailbox or for selected folders. Because you may want to prevent Outlook users from archiving with particular retention categories, you can hide a retention category from the list of available retention categories. Users are still able to search for items that have been archived with retention categories that are now hidden.

Retention categories cannot be deleted, because they may still be assigned to items in archives.

Note that some Enterprise Vault features, such as the retention folders and classification features, can update the retention categories of archived items and prevent users from changing the categories. For more information on retention, see the *Administrator's Guide*.

About retention plans

With a retention plan, you can associate a retention category with a number of other settings and apply them all to one or more archives. The extra settings that you can apply with a retention plan include the following:

- A classification policy
- One or more retention folders
- The criteria for discarding expired items

Applying a retention plan to an archive gives you greater control over the retention periods of the items in the archive. In particular, a retention plan lets you dispose of already-archived items by giving them a different retention period than the one that Enterprise Vault first gave them when it archived the items. For example, you can configure a retention plan so that Enterprise Vault expires the affected items according to the retention category that you have associated with the retention plan, and not the retention categories that Enterprise Vault originally assigned to them.

About classification policies

If you choose to set a classification policy with a retention plan then, for the archives to which you assign the retention plan, the classification policy determines the following:

- Whether to classify items at the same time that Enterprise Vault indexes and archives them. After Enterprise Vault has applied the classification tags, users of applications like Compliance Accelerator and Discovery Accelerator can use them to filter items when they conduct searches and reviews.
- Whether to classify items when users manually delete them or Enterprise Vault automatically expires them.

For more information on the classification feature, see the *Classification* guides.

About retention folders

Note: The retention folders that are described here differ from the Domino and File System Archiving retention folders that are described elsewhere in the Enterprise Vault documentation. You create Domino and File System Archiving retention folders on the sources from which Enterprise Vault archives items, but you create the retention folders that are described here in the archives themselves.

In this release, you can create this second type of retention folder in Exchange archives and Internet Mail archives only.

The retention folder feature lets you control the retention and expiry of archived items at the folder level within your users' archives. Use this feature to create a single retention folder or a hierarchy of folders in these archives. The attributes that you set for each retention folder determine the retention and expiry settings that Enterprise Vault applies to the items in the folder. For example, you can create a folder that applies a retention category with a one-year retention period to the items, overriding the retention categories that Enterprise Vault has previously applied to them. You can further choose whether the subfolders of the retention folder inherit their retention and expiry settings from it or have their own settings.

The retention and expiry settings that you define for a retention folder override those that you define elsewhere in Enterprise Vault, such as in the associated retention plan or at the site level.

Through facilities such as Virtual Vault, Enterprise Vault Search, and IMAP, users can access the retention folders and move items into or out of them.

How to plan the automatic deletion of archived items

Enterprise Vault can automatically delete items from archives when they are no longer needed. Settings for storage expiry are in site properties.

The deletion is based on the retention settings that you have defined in an archived item's retention category. For example, you can specify a retention period of five years from the date on which Enterprise Vault archived the item. Alternatively, you can specify that the retention period ends on a particular date, such as December 31 2021, or even specify that the item is retained forever.

How to plan PST migration

How you implement the archiving of PST files will depend on the number of files to be archived, your company policy on the use of PST files and whether users have infrequent access to a fast network connection.

There are several ways to import PST files into Enterprise Vault:

- Wizard-assisted migration can be used for a small number of PST files.
- Scripted migration using Enterprise Vault Policy Manager is ideal for performing bulk migrations of PST files.
- Locate and Migrate uses Enterprise Vault Tasks to locate PST files on users' computers, copy them to a central location, and then import them. You can configure specific paths to include or exclude during PST searching. Locate and Migrate is designed to minimize the difficulties of collecting PST file from users' computers and is likely to require least effort on your part.
- Client-driven migration is similar to Locate and Migrate but the locating of PST files and sending them to a collection location is done automatically by the user's computer and not by Enterprise Vault Server Tasks. This can be useful if, for example, there are users with laptop computers who are in the office only one or two days a week, thus making it difficult to obtain their PST files by other methods. You can also give users' control over migrating their PST files. Client-driven migration needs to be enabled in the Administration Console.

To aid PST migration you can configure desktop clients so that, when a user starts Outlook, the client writes a marker into each PST file that is listed in the mail profile. When a marked PST file is subsequently imported, the marker indicates the owning mailbox.

PST migration is described in the *PST Migration* guide.

When planning PST migration, take the following points into consideration:

- Do not enable PST migration for all users at the same time. Migrate the PST files for a small group of users and then move onto the next group.
- A typical rate of PST migration is 2 GB/hour.
- An administrator has to allow each PST file before it can be imported.
- It is important that the language setting in the properties for each PST is correct before the PST is imported.

How to plan NSF migration

Enterprise Vault provides two tools that let you migrate items from NSF files:

- The NSF migration wizard lets you migrate items from NSF files into users' archives. The NSF migration wizard is described in the *NSF Migration* guide.
- Enterprise Vault Policy Manager lets you script the migration of items from NSF files. This approach is suitable for large-scale migrations. Policy Manager is described in the *Utilities* guide.

How to plan shared archives

When a user mailbox or folders on a file server are archived, the access permissions set on the original folders are also set on the associated folders in the archive. So users that had access to the original folders will also have access to the folders in the archive. This includes delegate access permissions in Outlook. For folder permissions to be set up on new archives, you need to synchronize folders and permissions.

Alternatively, you can share access to an archive either by modifying the archive's properties in the Administration Console, or using the New Archive wizard in the Administration Console to create an archive specifically for sharing. Archives created using the New Archive wizard do not contain a folder structure.

How to plan vault stores and partitions

Consider where to create vault stores and partitions and the most appropriate type of storage devices to use. For example, some devices provide a large amount of secure storage (WORM) at reasonable cost, which may suit a compliance archiving strategy.

For standard vault store partitions, Enterprise Vault's partition rollover feature provides automatic rollover from one partition to another to support continuous archiving. For example, when the physical disk that hosts your open partition reaches

capacity, Enterprise Vault can automatically close the partition and open another. This rollover capability is not available for smart partitions. On the other hand, multiple smart partitions can be open for archiving at the same time. This is not true of standard partitions, which are limited to one open partition for each vault store.

Where vault store partitions are held on non-WORM devices other than Dell EMC Centera, the speed and efficiency of vault store backups can be improved by using the Enterprise Vault "collection" feature. This feature collects multiple small files into CAB files. Collection is not recommended on devices that perform deduplication, as it causes loss of deduplication.

Collections are handled differently on Dell EMC Centera devices, as follows:

- Centera collection files are used instead of CAB files.
- Files are collected as soon as they are archived (not according to a schedule).

How to plan handling safety copies

Safety copies are copies of the archived items. In a mailbox, they are identified by the pending archive icon. When you set up a vault store, you must specify when Enterprise Vault deletes safety copies.

The type of environment in which you are running Enterprise Vault influences how you configure the process for safety copies.

If you are setting up a pilot system in a test environment, where there is no need to keep users' data, set up the vault store so that Enterprise Vault deletes the safety copies immediately after archiving.

If you are setting up a pilot system in a production environment, it may be necessary for users to get their data back again after testing Enterprise Vault. In this situation, edit the vault store properties so that Enterprise Vault never deletes the original items.

In the properties of each vault store you can choose to store safety copies in the original locations or on the storage queue. If you keep safety copies on the storage queue Enterprise Vault deletes the original items immediately after archiving so space in the original location is recovered soon. For example, when items are archived from an Exchange Server mailbox, mailbox icons change from "pending archive" to "archived" within a minute or two. The storage queue location must have sufficient space to retain the safety copies until after the next backup or replication of the vault store partition.

When you have completed your testing, you can set Enterprise Vault to delete the safety copies at the selected time.

How to choose a location for safety copies

By default Enterprise Vault retains safety copies of items that have been archived. Enterprise Vault keeps the safety copies until the vault store partition in which they are archived has been backed up or replicated.

The default is for Enterprise Vault to use the original items as safety copies. This default means that Enterprise Vault must wait until the vault store partition has been backed up or replicated before it can delete the original items. The original items must stay in the original locations until the backup or replication is complete.

You do not need safety copies if all changes on the device that hosts the vault store partition are immediately replicated to a duplicate device.

There is an alternative location for safety copies: you can configure vault stores so that Enterprise Vault keeps the safety copies on the storage queue. This mechanism means that Enterprise Vault can remove the original items as soon as they have been archived. The storage space in the original location is recovered quickly.

When an item is archived using IMAP access Enterprise Vault always keeps the safety copy on the storage queue, regardless of whether you choose to keep safety copies in the original location or on the storage queue.

The storage queue location should be on a fault-tolerant device (RAID 1 or higher). The device must have sufficient space for all items on the queue. If you intend to use the storage queue for safety copies then there must be enough space to keep them until after the next backup.

You can change the storage queue location on the Storage Queue tab of the Storage service properties.

In the properties of each vault store you can configure specific safety copy options for mail journal items. For example, you could choose to delete journal items as soon as they have been archived but to retain safety copies of items that are archived from other sources.

Enterprise Vault checks for safety copies to remove when the Storage service is started and also when backup mode is cleared from the vault store.

How to plan single instance storage

Enterprise Vault can use single instance storage to optimize archive storage space. Enterprise Vault's single instance storage mechanism archives a single instance of parts of an item that are suitable for sharing, such as large message attachments. Enterprise Vault can share the single instance storage parts (SIS parts) within a vault store, or across two or more vault stores within a vault store group. A vault store group forms an outer boundary for sharing.

If you use Enterprise Vault's single instance storage mechanism, you need to create a sharing regime that meets your organization's data sharing requirements and which is appropriate for your network connection speeds. Consider what sort of sharing regime you require before you start archiving. Enterprise Vault cannot share items retrospectively. Also there are limits to what you can change after Enterprise Vault has started sharing items.

Note the following:

- **Partitions on Dell EMC Centera devices.** Enterprise Vault single instance storage is not performed when items are stored to partitions that are hosted on Dell EMC Centera devices. You can configure a partition for a Dell EMC Centera device to use the Centera's device-level sharing mechanism, if required. Enterprise Vault then stores the shareable parts of a saveset as separate data blobs, so that the Centera device is able to share them.
- **Smart partitions.** Enterprise Vault shares the SIS parts between items in the same smart partition, but it does not share the SIS parts between a smart partition and other partitions.
For example, suppose that two employees receive the same email, which has an attachment. For compliance reasons, Enterprise Vault archives one employee's emails to a smart partition; it archives the other employee's emails to the standard vault store partition. If Enterprise Vault first archives the email and its attachment to the standard partition then, when it subsequently archives the email to the smart partition, it should not normally archive the attachment again. This would mean that the data on the smart partition is not fully compliant, however, so in this case Enterprise Vault archives both the email and the attachment again.

For more information about setting up single instance storage, see "Setting up Storage" in *Installing and Configuring*.

About Enterprise Vault reports

You can run an Enterprise Vault task in Report Mode. When you do this, nothing is processed, but you get a report for each archiving target that shows what would be processed if you ran the task normally.

Reports created by mailbox and Exchange Public Folder Tasks include details about the number and total size of items that would be archived when the task is run in Normal Mode and the number of expired shortcuts ready to be deleted.

In mailbox archiving, reports are generated for enabled, disabled or new mailbox archives, so that you can determine how much space you would save by enabling more mailboxes for archiving.

The Provisioning tasks can be configured to generate brief or full reports. Full reports list the mailboxes that will be processed, the Provisioning Group assigned, and the policies that will be used. Reports are also produced by this task on a normal run.

File system and SharePoint archiving tasks can be run in Report Mode but can also be configured to generate reports for normal archiving runs. You can also specify the level of reporting required.

There are several levels of reporting for file system archiving reports:

- Brief reporting includes a summary for each Volume and Archive Point of the number and total size of items that will be archived. Each policy and rule is listed together with the number and total size of items that meet the criteria.
- Full reporting includes details of the items that satisfy each rule.

Similarly, in SharePoint archiving, there are several levels of reporting:

- Brief reporting includes a summary for each site collection of the number and total size of documents that will be archived. Each policy and rule is listed together with the number and total size of documents that meet the criteria. The summary also contains information about version pruning and the space saved as a result of pruning.
- Full reporting includes details of the documents that satisfy each rule and document versions pruned.

For all types of archiving, reports go into a text file in the `Reports` subfolder of the Enterprise Vault installation folder. The text file has tab-separated fields, so you can easily import it into a spreadsheet program such as Microsoft Excel. You can use these reports to gauge the usage of the archiving task under different setup conditions.

Enterprise Vault Reporting feature

The Reporting feature provides enterprise-level reporting for Enterprise Vault. It uses Microsoft SQL Server Reporting Services as the reporting mechanism. Administrators manage report content and view reports using the Reporting Services Report Manager web application.

For more information on using Enterprise Vault Reporting, see the *Reporting* guide.