# **Technical Overview**

Exchange Server Archiver - 1.0



# Technical Overview

Red Gate's Exchange Server Archiver enables you to improve the efficiency of your Microsoft<sup>®</sup> Exchange Servers by delivering a transparent archiving solution.

The archived messages and attachments can be retrieved as and when required, allowing you to:

- remove restrictions on mailbox sizes
- spend less time managing users
- free up space on your Exchange server stores

Exchange Server Archiver is simple to install, configure, and maintain.

This document describes the architecture of Exchange Server Archiver, and provides an overview of the archive and retrieval processes. Some further information on security, licensing, and evaluation is also provided.

You are recommended to read this document before you install Exchange Server Archiver. For full details of how to install Exchange Server Archiver, see Installing Exchange Server Archiver (http://www.red-gate.com/esa/installation/v1).

## Contents

Archive modes (page 3) Architecture (page 4) Archive process (page 8) Retrieval process (page 9) Security (page 10) Licensing overview (page 11) Evaluating Exchange Server Archiver (page 12)

## Archive modes

Exchange Server Archiver offers four different archive modes so that you can choose what will happen to messages that are processed during an archive. The selected archive mode is used whenever an archive is run.

If you are evaluating Exchange Server Archiver, you may wish to use one of the 'copy only' modes so that messages and attachments remain on the Exchange server; see Evaluating Exchange Server Archiver (page 12) for more information.

The following archive modes are available:

## Copy attachments only

- copies attachments to the specified Archive Store
- the original attachments remain on the Exchange server

## Copy entire message

- copies the message bodies to the Archive Store
- copies attachments to the Archive Store
- originals of both remain on the Exchange server

## Copy and remove attachments only

- copies attachments to the Archive Store
- removes the attachments from the Exchange server

## Full archive

- copies messages and attachments to the Archive Store
- on the Exchange server, removes each message that has been archived and replaces it with a stub

A stub is a placeholder for an archived message. It contains the first few lines of the message, and links to the full archived message and any attachments.

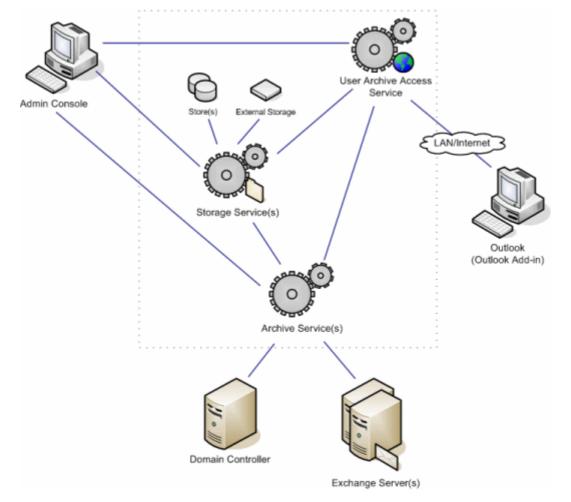
removes the attachments from the Exchange server

If necessary, you can change the archive mode. When you change archive mode, Exchange Server Archiver schedules a tidy up task so that all previously-archived messages and attachments match the new archive mode.

Note that message stubbing occurs only when you choose *Full archive* mode; in the following descriptions of the architecture and processes, it is assumed that you have selected *Full archive* mode.

# Architecture

An overview of the Exchange Server Archiver architecture is illustrated below. The individual components are described in the following sections.



When you install Exchange Server Archiver, no software is installed on your Exchange server (unless you choose to install the optional Outlook<sup>®</sup> Web Access Add-in) and no information is stored in Active Directory<sup>®</sup>.

## **Admin Console**

The Admin Console provides an easy-to-use graphical user interface for you to install and configure the Exchange Server Archiver services, and to manage and schedule archives.

You use the Admin Console to set up rules to:

- specify which mailboxes are archived to which Archive Stores
- specify conditions that define which messages will be archived (such as 'messages over 10MB or messages older than 30 days')
- set up schedules to control when messages will be processed

You can apply a rule to individual mailboxes or to groups of mailboxes using the Exchange server storage hierarchy or Active Directory.

You also use the Admin Console to change the Exchange Server Archiver configuration, for example, to create new Archive Stores as required.

The settings you define using the Admin Console are replicated in the Archive Services. You can therefore view and manage the Exchange Server Archiver settings from Admin Consoles that have been installed on different computers.

To use the Admin Console, you must be a *domain administrator*.

## Archive Service

The Archive Service coordinates the scheduling and archiving of messages.

Using the rules set up using the Admin Console, the Archive Service accesses the mailboxes, and determines which messages to extract, archive, and stub. It then compresses the messages to be archived, and passes them to the appropriate Storage Service.

The Archive Service also handles administration tasks. It:

- stores the system configuration
- performs tidy-up tasks, for example if you choose to change archive mode

The Archive Service uses MAPI (Messaging Application Programming Interface) to communicate with Exchange Server.

The Archive Service requires a user with permissions to access all messages when it communicates with Exchange Server. For more information, see Creating the Archive Service user (http://www.red-gate.com/ESA/1.0?i=9190).

## User Archive Access (UAA) Service

The UAA Service enables your users to access archived messages. It is a Web service and Web site hosted on IIS.

When a user opens an archived message, the UAA Service communicates with the Archive Service to confirm that the user is authorized to view the message. It then locates and retrieves the message from the Archive Store, via the Storage Service.

When a user searches for a message, the UAA Service delegates the search request to the appropriate Storage Services, and then combines the results retrieved from the Storage Services so that they can be displayed for the user.

If you choose not to install the Outlook Add-in or the Outlook Web Access Add-in, users can use the Web pages provided by the UAA Service to access their messages.

Administrators with permissions to read all messages on the Exchange servers can set up the UAA Service so that they can search all archived messages.

The UAA Service uses MAPI if you choose to export messages to EML or MSG format.

For more information, see Using the UAA Service (http://www.red-gate.com/ESA/1.0?i=9430).

#### **Storage Service**

A Storage Service is a Microsoft<sup>®</sup> Windows<sup>®</sup> service that provides fast, reliable, space efficient storage and retrieval of archived messages.

You can set up any number of Storage Services, and for each Storage Service you can set up any number of Archive Stores.

You are recommended to install the Storage Service on the same server as the Archive Stores that it hosts. However, there may be some exceptions; for example, if you create an Archive Store on a NAS (network attached storage) unit, the Storage Service is installed on a different computer.

The Storage Service uses the Microsoft<sup>®</sup> Indexing Service to index and search the Archive Stores.

#### **Archive Stores**

An Archive Store is a folder in which messages are archived.

Each message body and attachment is stored as a file in the Archive Store. Message bodies are compressed. Once written, an archive file cannot be modified. Because the archived messages and attachments are stored as individual files, they are easy to back up.

Single-instance storage (SIS) ensures that duplicate messages and attachments are stored only once in each store. Messages or attachments are considered to be duplicates if their checksums are identical. The header is stored for each instance of the duplicated message.

You create and manage Archive Stores using the Admin Console.

You can change the Archive Store to which a mailbox is archived. For example, you may want to do this when the store becomes large: you specify that you want to archive messages to a new Archive Store in the archive rules, and you can then make the old store read-only, and perform a final backup on it. Messages in the old store can still be accessed.

Because Archive Stores and archived messages are uniquely identified using GUIDs (Globally Unique Identifiers), you can move an Archive Store to a different Storage Service if necessary.

## Outlook

When you install the Outlook Add-in, Exchange Server Archiver provides a transparent interface for your users to access archived messages using Microsoft<sup>®</sup> Outlook 2007 or Outlook 2003.

Messages that have been archived are identified with  $\square$ .

A search feature is provided so that users can search archived and unarchived messages within a single search when working online.

If you do not install the Outlook Add-in, messages that have been archived are identified with A in Outlook, but only a stub is displayed when an archived message is selected. Users can access the full message from within Microsoft Outlook by using a hyperlink to the UAA Service which is provided in the stub. Searches from within Outlook will search only the message stubs.

#### **Outlook Web Access**

When you install the Outlook Web Access Add-in, Exchange Server Archiver provides a transparent interface for your users to access the archived messages using Exchange Server 2007 Outlook Web Access or Outlook Web Access 2003.

Messages that have been archived are identified with  $\triangle$ .

Searches from within Outlook Web Access will search only the message stubs. Users can use the UAA Service Web page for a full-text search of archived messages.

If you do not install the Outlook Web Access Add-in, archived messages are not identified with an icon. When a user selects an archived message, a stub is displayed and users can access the full message by using a hyperlink to the UAA Service.

#### **Outlook Anywhere**

You can use Outlook Anywhere with Exchange Server Archiver if the server on which you install the UAA Service is visible to users accessing their messages remotely.

## The archive process

For a full archive:

1. When a mailbox is scheduled to be archived, the Archive Service accesses the mailbox on Exchange Server, and determines which messages to process.

Note that messages in the Drafts or Outbox folders, (or any subfolders within these folders) are not archived.

- 2. For each message to be processed, the Archive Service checks whether the message is already in the Archive Store:
  - If the message does not exist already in the store, the Archive Service extracts the message body, header, and any attachments, and passes them to the corresponding Storage Service. The message body and header are compressed.
  - If the message already exists in the store (because it is a duplicate of a message that has already been archived), the Archive Service passes the message header to the appropriate Storage Service, and uses single-instance storage (SIS) to link it to the existing archived message body (and attachments, if it has any).
- 3. The Storage Service stores the messages.
- 4. When the messages have been safely stored, the Archive Service replaces the messages in Exchange Server with stubs.

Note that the original message is removed from the Exchange server only when the message and attachments have been written to the Archive Store, ensuring no data is lost if there is an interruption to the process.

You are advised to defragment your Exchange Server databases after the first archive to take advantage of the reduction in the size of Exchange Server stores; if you do not defragment the databases, you will not see any reduction in size. You should then continue to defragment the databases according to your usual maintenance schedule.

Using Outlook, the user chooses to view a message that has been archived.

The User Archive Access (UAA) Service confirms, via the Archive Service, that the user is authorized to view the message, and then retrieves the message from the Archive Store and returns it to the user.

Exchange Server Archiver uses the permissions that have been set up in Exchange Server and Active Directory. The UAA Service and Archive Service prevent unauthorized access to messages.

HTTPS can be used to protect messages in transit between the UAA Service and the Outlook client.

For more detailed information about security, see Data security in Exchange Server Archiver (http://www.red-gate.com/ESA/1.0?i=9450).

Exchange Server Archiver is licensed per mailbox.

When the number of mailboxes approaches your licensing limit, the Admin Console displays a warning, so that you can upgrade your license if required.

If a mailbox is no longer scheduled to be archived (for example, because the user has left the company), the license is automatically made available so that you can reuse it with a different mailbox. The old mailbox archives remain accessible.

For more detailed information about licensing, see Licensing (http://www.red-gate.com/ESA/1.0?i=9400).

When you install Exchange Server Archiver, you can choose an archive mode that leaves the original messages and attachments on your Exchange server and only copies them to the Archive Server (*Copy entire message* or *Copy attachments only*). Therefore, you can evaluate it with no risk to your data.

Because the archived messages and attachments are not removed from the Exchange Server, you will not be able to see the reduction in size of your Exchange Server stores. However, you will be able to evaluate the performance of Exchange Server Archiver and assess the storage requirements for the archives, and you will have the reassurance that if you choose to uninstall Exchange Server Archiver, you do not need to copy back messages and attachments to the Exchange Server store.

If you prefer, you can evaluate Exchange Server Archiver in *Full archive* mode. Messages will be archived to the Archive Stores according to the rules you set up, and replaced with stubs in Exchange Server, and attachments will be removed from the Exchange server. You will therefore be able to see the resulting reduction in the size of Exchange Server stores when you defragment your Exchange Server databases. If you later choose to uninstall Exchange Server Archiver, you can first copy back archived messages and attachments to the Exchange Server stores; however, you will lose any single-instance storage that was used on those messages.

Whichever archive mode you choose, you can specify which mailboxes and messages will be archived by setting up rules. For example, you may want to process only a few test mailboxes initially, and then roll out the evaluation to further mailboxes later.

The installation is a simple two-stage process:

- 1. Install the Admin Console.
- 2. Install the services and set up the initial archive.

For full details, see Installing Exchange Server Archiver (http://www.red-gate.com/esa/installation/v1).

For details of how to install and configure a trial environment for evaluating Exchange Server Archiver, see Setting up a trial environment (http://www.red-gate.com/products/exchange/Trial\_Environment\_Configuration\_Guide.pdf).