

Symantec Technology Network

Enterprise Vault Best Practice

Managing Mailbox Moves within Enterprise Vault Archiving Environments

- by -

Glenn Martin

Symantec Regional Product Manager, Enterprise Vault

October 31, 2007



Table of Contents

Table of Contents	2
Document Overview	3
Introduction	3
Scope of Document	3
Target Audience	3
Acknowledgements	3
Enterprise Vault Component Review	4
Exchange Mailbox Moves - Impact on Enterprise Vault	5
EV Component Overview	5
Simple Single EV Server Model	6
Distributed Server Model – Single Vault Store	7
Distributed Server Model – Multiple Vault Stores	8
One to One Model	9
Mailbox Moves – Operational Considerations for EV	10
Moving Mailbox and User Archive	11
The administrator at the original location	11
The administrator at the new location	11
User Archive Move Process	12
Moving Across Regions or Sites	13
Option 1: Move mailbox but leave user archive	13
Option 2: Move archive to EV server in the new location	14
Option 3: Restore Shortcuts before Mailbox is Moved	14
Option 4: Plain export to PST	14
About the Author	15
About Symantec Technology Network (STN)	15

Document Overview

Introduction

This best practice document discusses the operational issues when managing an EV environment where mailboxes are moved between Exchange servers. There are various ways to deal with this situation and considerations that need to be taken into account before selecting each method.

Scope of Document

This document is concerned with the mailbox moves within a Microsoft Exchange environment (version independent) and is not relevant to a Domino environment.

Target Audience

This document is aimed at customers, consultants and support staff and it is assumed the reader has a good understanding about the architecture and operational aspects of an Enterprise Vault server, Microsoft Exchange as well as Active Directory terminology.

Acknowledgements

I would like to acknowledge the contribution that other individuals made towards making this a successful and informative document. I am most grateful to a number of customers who provided their feedback and comments about the real world usage of EV. Also other contributions and feedback came from the following teams; technical product management, technical field enablement, engineering, development and the performance team. Thanks to them all.

Enterprise Vault Component Review

Before getting into detail about how to manage mailbox moves within an EV environment it's worth spending a little bit of time reviewing the various components and configuration types that occur within typical EV environments.

EV Directory - All configuration information for Enterprise Vault is stored within a SQL Server database with a default name of EnterpriseVaultDirectory. The Directory Service is used to access this database and all information that it contains.

EV 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 (Microsoft Exchange Servers). A Site also contains a collection of Vault Stores, archiving policies that define how and when items are to be archived, and Retention categories that define how long items are to be stored before being deleted. An Enterprise Vault Site is located in an Enterprise Vault Directory on one or more Enterprise Vault servers.

EV Server - An EV server is contained within an EV Site which is in turn part of the overall Enterprise vault directory

Exchange Archive Task - The mailbox 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 deletes the original, if applicable, and creates a shortcut for it, if specified. Each enabled Exchange Server will have one Archiving Task assigned to it. When Enterprise Vault archives an item, the EV Directory database is accessed to find out where the required archive is.

Storage Service - 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.

Index - On instruction from the Storage Service, the Indexing Service indexes items as they are archived. It uses the AltaVista® Search technology to create and store indexing data about each item. There is one index for each archive, when users search the archives the index files are used to find the relevant items but the items are then retrieved from the Vault Store.

Vault Store - A Vault Store consists of an SQL database and either a single or multiple partitions which define the physical location of the storage where the archived data is stored. When an item is archived, metadata is written to the database and this identifies who has access to the archived item and where the item is stored within the Vault Store.

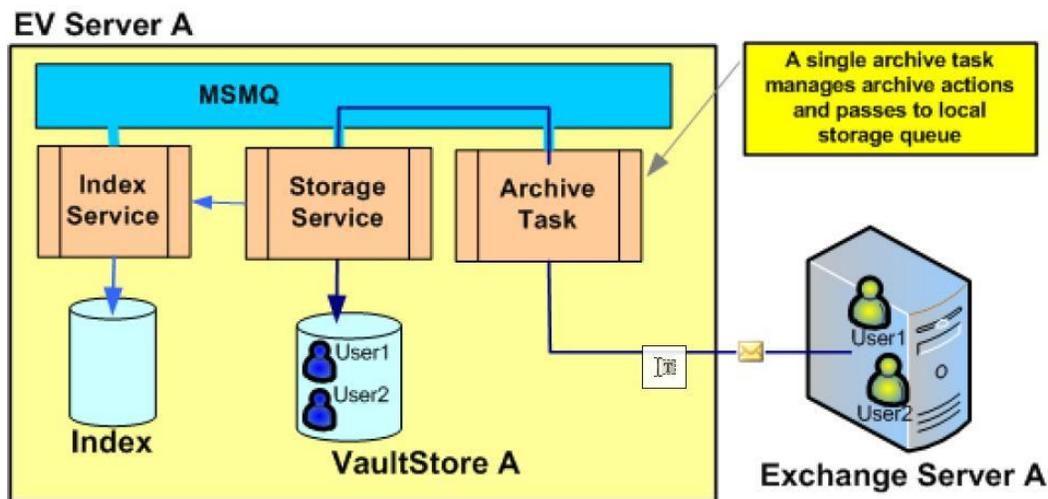
User Archive - A user archive is associated to an Exchange mailbox and is created when a mailbox is enabled for archiving, the user archive has permissions assigned and determines the level of access to the users' archived data (typically this is synchronised from the mailbox). In earlier releases of EV this object has also been called a User Vault ; some people still use this expression to describe the user archive.

Exchange Mailbox Moves - Impact on Enterprise Vault

Before going into more detail about the impact of mailbox moves it is worth spending some time reviewing the basics of archiving and the relationship between EV and Exchange.

When a mailbox is enabled for archiving it will be allocated a user archive in a particular Vault Store on an EV server. The mailbox is then archived by the specific exchange archive task assigned to the Exchange server on which the mailbox resides. Archived items are stored into the Vault Store that the user archive has been provisioned against; this is done by the Storage Service. So in simple terms, the archive task collects the items which need archiving and the Storage Service stores them into the Vault Store (see diagram below)

EV Component Overview



From this overview you can see that a relationship is formed between a mailbox, the EV server that archives it and the EV server that stores its archived items.

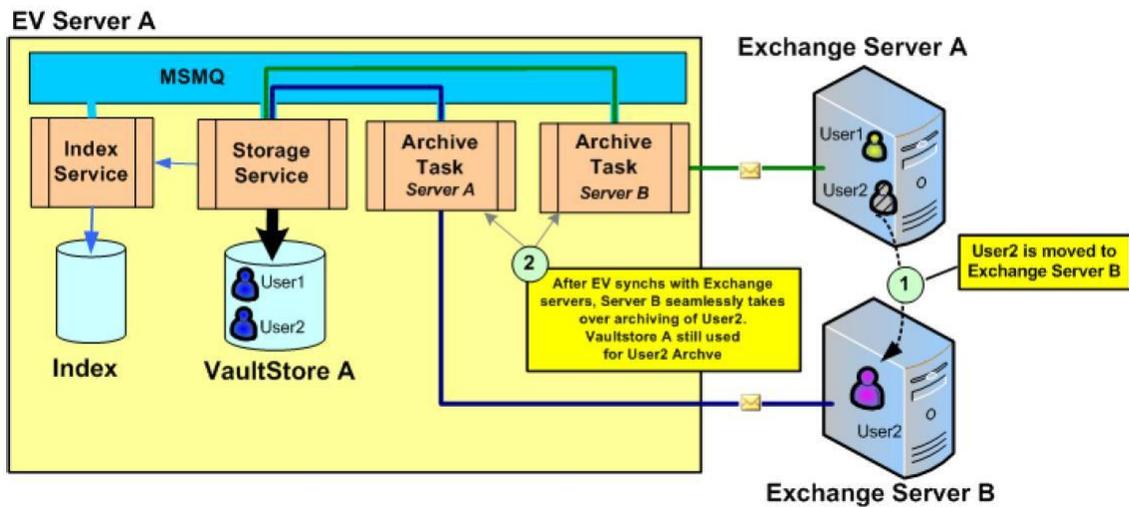
When a mailbox is moved to a different Exchange server it will be archived by a different archive task. If the archive task is on a different server the archived items still need to end up in the Vault Store that was originally provisioned for the user archive. This means that the archive task must now communicate with Storage Service on the server where the user's original Vaultstore is configured. This is not a problem for EV because it's been designed to work in this manner but you will need to evaluate the operational issues that this presents in your own environment.

In order to help understand these concepts, the following pages show some of the different types of EV configurations that are found in the field and considers the impact that mailbox moves will have in each case.

Simple Single EV Server Model

This model comprises a single EV server which archives multiple Exchange Servers and also runs the Storage & Indexing services. This is typical of many small or medium environments that have perhaps 1000-2000 users and two or three Exchange servers.

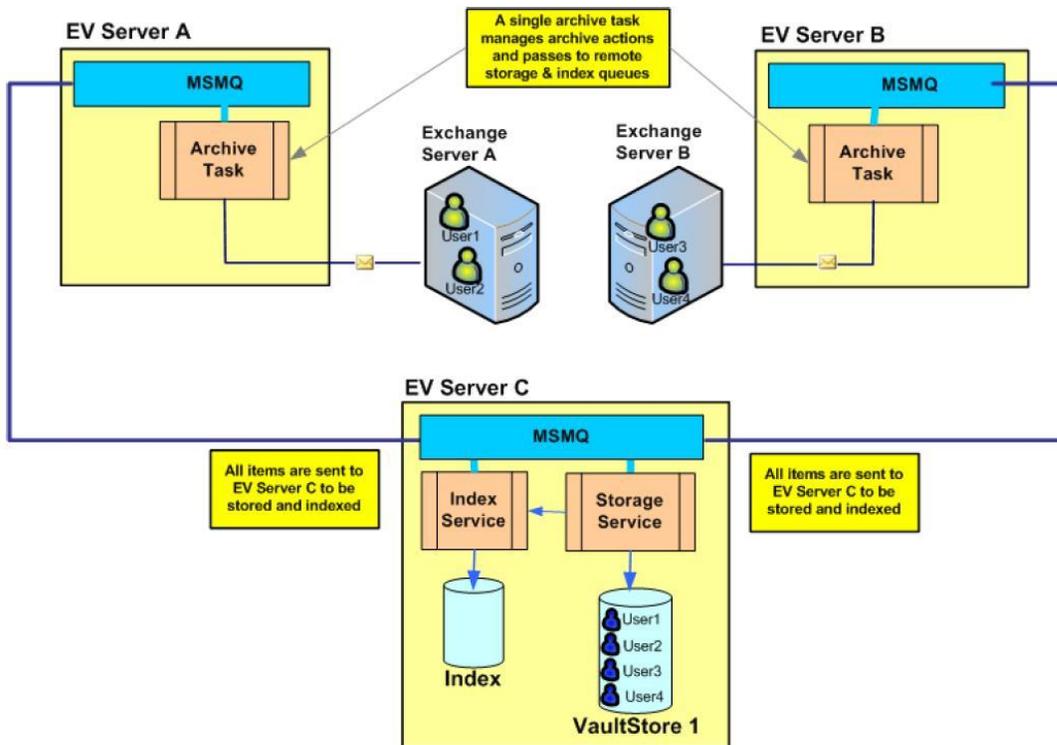
Mailbox Move Impact – A mailbox move in this environment will have little or no impact. Once a mailbox has been moved from one Exchange server to another, it's a simple case of re-synchronizing the mailbox with EV and it will then be automatically archived by a different archive task, The synchronization can be initiated manually otherwise it would occur automatically at the time specified on the archive task.



Distributed Server Model – Single Vault Store

The distributed model is where multiple EV Servers are used to provide an overall archiving environment, some EV servers will only have archive tasks and others may also have the indexing or Storage Services. This type of configuration is more typical of earlier installations of EV and separates the Indexing & Storage Service overhead from the archive task overhead.

Mailbox Move Impact – Moving mailboxes in this type of environment does not present much of a change since the archive tasks are already detached from the Storage Services and will need to send data across the network anyway. All of the user archives are stored on *EV Server C* so if a mailbox is moved between *Exchange Servers A* and *B*, the appropriate archive task will synch up with Exchange and start to archive it.

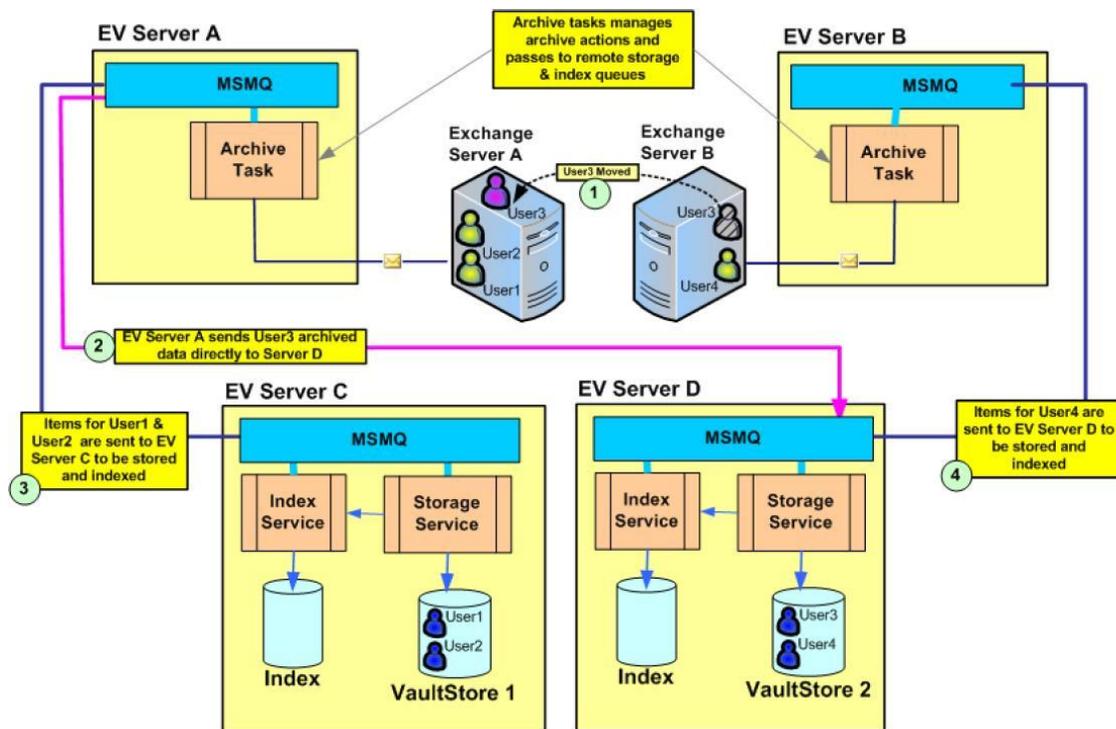


Distributed Server Model – Multiple Vault Stores

This is a variation on the previous example, but here we have an additional EV server that is balancing the indexing and storage load. Therefore an additional Vault Store has been introduced to the environment which increases the complexity of the configuration. This example shows that when User3 is moved to *Exchange Server A* it will be archived by *EV Server A* but the user archive remains on *EV Server D*. This means that *EV Server A* will have an additional overhead in communicating with *EV Server D* as well as *EV Server C*.

Mailbox Move Impact – Other than the additional network overhead, EV will continue to function normally but this introduces an operational dependency that needs consideration when managing EV servers and archive tasks. The EV architecture may have been designed with the fact that *EV Server A* would be aligned with the Storage Service of *EV Server C*. However, after User3 is moved, it also needs access to *EV Server D* so operational staff would need to realise that stopping the Storage Service on *EV Server D* would impact users on *EV Server A*.

Another point to consider is the amount of mailboxes that are moved from *Exchange Server A* to *Exchange Server B*. This is because you could find that the Archive Task on EV Server A could become overloaded whilst EV Server B is idle.

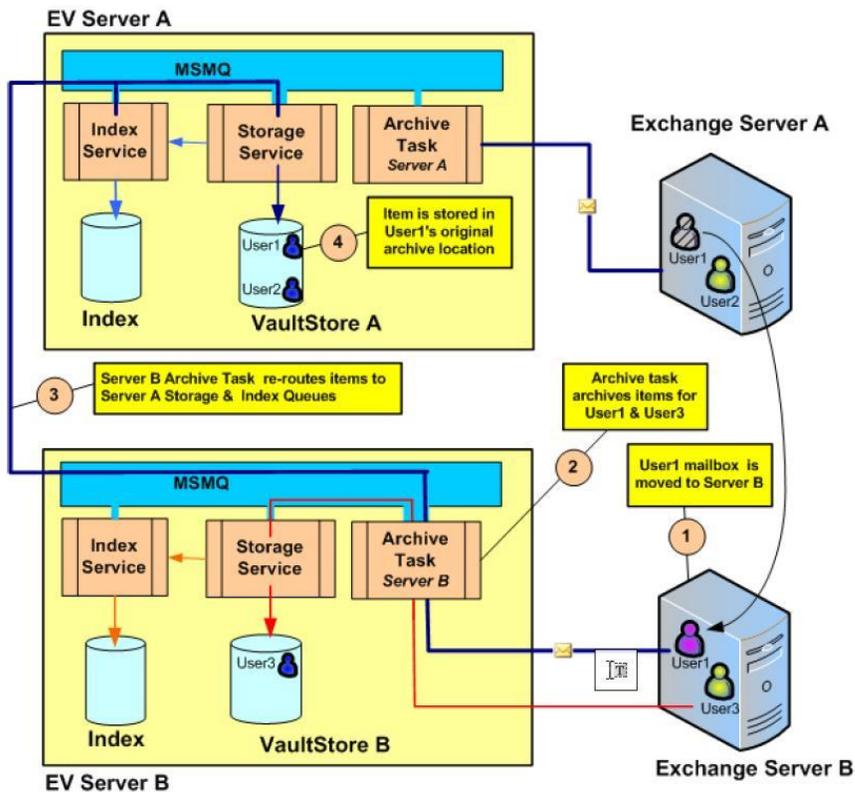


One to One Model

The One to One model is actually a scaled up version of the single server model described earlier. However, this type of configuration is built to operationally maintain a one to one relationship so that a given EV server will be assigned solely to archive a specific Exchange server. All the EV servers are configured and managed as part of the EV site but are operationally self contained so that one EV server outage will not impact the archiving process on another server.

This type of configuration is typically used where there are multiple Exchange servers with a high density of users (eg 2000-3000 users). This arrangement provides a dedicated EV server for each Exchange server with its own core services (Index, Archive and Storage and Vault Store). This configuration will also typically leverage modern NAS and SAN technologies so that whilst the EV services are all server centric, the storage will actually be distributed so the DR requirements can be focused at an EV server level.

Mailbox Move Impact – When mailboxes are moved, while EV will continue to function normally, the one to one relationship will be broken and cross server communication between Archive Task & Storage Service will occur. This will present issues when planning maintenance activities, as you can no longer be sure that taking an EV server down will only cause a service outage for the users of the Exchange server that it archives. This is because you will now have users on another EV server which needs access to the Vault Store on the EV server you are taking down. There is also an increased amount of network traffic because of the cross server MSMQ traffic for the storage and index services.



Mailbox Moves – Operational Considerations for EV

The previous pages have looked at some of the types of EV configurations that are found in the field and has highlighted some of the issues that need to be dealt with when mailboxes are moved between Exchange servers. This section goes into a little more detail about the operational impact on EV when mailboxes are moved between Exchange servers.

Cross Server Dependencies

Your configuration may have been designed so that each EV server is tasked to specifically archive a particular Exchange server (*as described in the complex distributed & one to one models*). In both these examples, when mailboxes are moved, the EV archive process becomes dependent upon multiple EV servers working together. When the archive and storage tasks become separated across servers, there is of course an additional network overhead to consider. In an environment containing multiple EV servers, after many mailboxes get moved round between Exchange servers it is easy to lose sight of exactly what each EV server is doing. It is of course important to understand the interaction between your EV and Exchange servers so that performance bottle necks don't develop.

Network Impact

This is related to the previous topic. When multiple EV servers have to interact as part of the archive process (*the Archive Task and Storage Service*) there could be a change to network traffic. This could be because the EV servers which need to communicate are on separate LAN segments or possibly dependent upon WAN access. In some environments there might even be a firewall between the two servers.

Action: Make sure that you understand the network topology that your EV servers use to communicate and that any WAN latency or firewall usage has been considered.

Archive Task Workload

As an Archive Task processes more mailboxes, the length of time it takes to process them will increase. Also, as the user count increases, the Archive Task workload can increase because of the number of message retrievals that will occur when users restore or recall archived email.

Action: Monitor the number of mailboxes that an Archive Task is processing, this is best achieved by running an archive in report mode for each Archive Task. This will show the number of mailboxes which are EV enabled and the amount of data outstanding for archiving. If the Archive Task has performance problems, this might be indicated in the report because all or some of the mailboxes might have a high volume of data ready for archiving.

Vault Store Balance

Moving users between Exchange servers will not directly effect the location of the user archive within EV, since this remains in the location where it was originally created, but it can become remote from the Archive Task that processes the respective mailbox.

If you to decide that you want to keep user archives on the same server as the Archive Task & Vault Store (*to maintain One to One Model*) you will need to move the user archive from one Vaultstore to another. However, you do need to consider that this can be an involved process, especially if the user archive is in the multi-gigabyte size range. There is currently no feature within the VAC to move a user archive using a single "move" command (*although this is expected in a future release*) so a number of steps are required to move a user between Vault Stores.

Actions: Follow the "move archive" procedure (described in the next section) to move a user archives between Vault Stores but consider the impact of doing this.

Moving Mailbox and User Archive

As explained in the previous section, there isn't currently a "move archive" option in the VAC so a combination of steps are required in order to move a user's archive from one Vault Store to another (see *below and overleaf*). This mechanism is suitable for moving a small number of users between Vault Stores or for ad hoc purposes such as when a user moves between geographic regions. However, if you intend to move many users between Vault Stores you need to consider the amount of time and effort that this will take; it can be considerable when you think how large some user archives can grow to (often multiple GBs). Also – refer to the current EV performance guide to get an idea about the speed of export/import for your environment.

Another point of reference for moving user archives is the EV Admin Guide; this has a section called "Exporting Archives" and covers the basics of using the VAC to export and import user archives. The instructions below are extracted from the admin guide and on the following page you will see a diagram which explains the steps involved to export/import a user archive.

Steps to move a user archive

The administrator at the original location

1. Disable mailbox archiving for the mailbox/user archive being moved
2. Export the user archive to one or more PSTs, making sure to split the PST files by retention category (*this is relevant for the import process*). *When you export, you can filter the output by date and/or retention category. For example, you could export items less than a year old archived with a retention category of Business.*
3. Sends the exported PST files, with their configuration files, to the administrator at the new location.
4. Move the Exchange mailbox to the new server

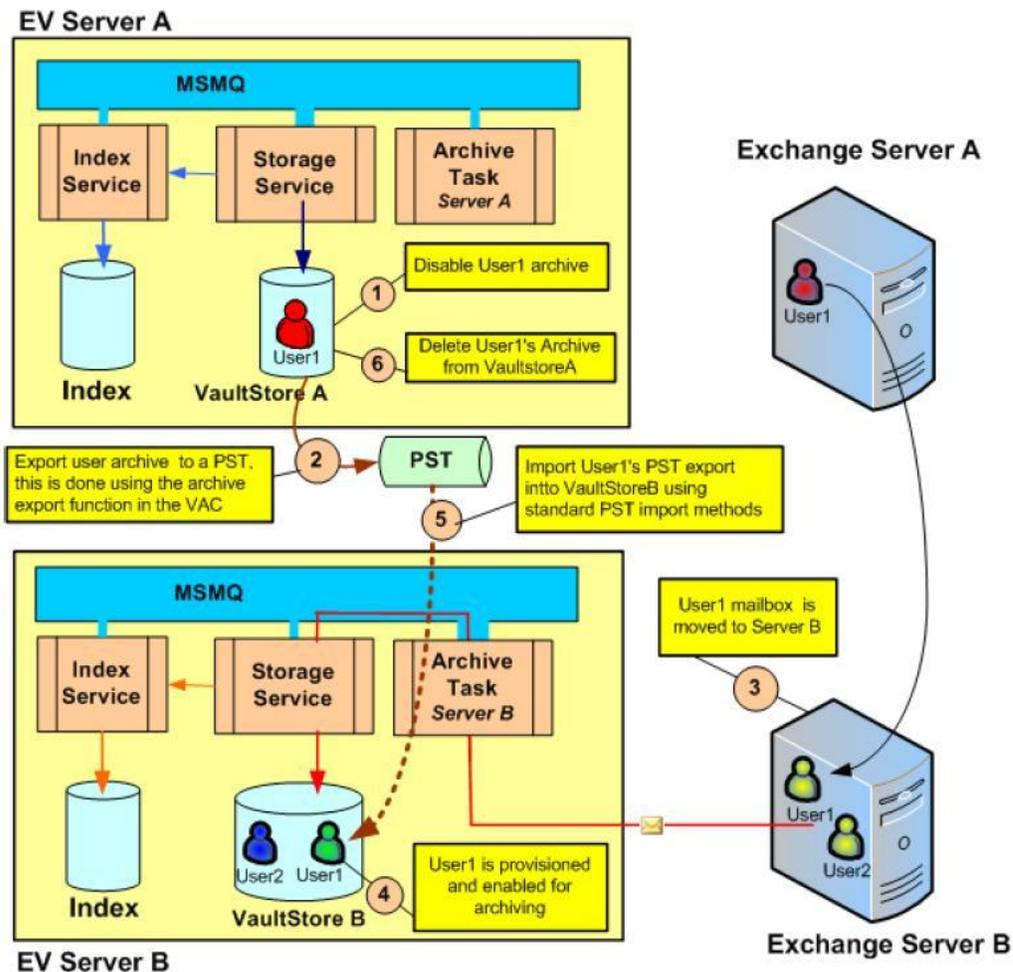
The administrator at the new location

1. Make sure the Exchange mailbox move has completed (*it will still contain shortcuts to the original vault*). *The mailbox can then be enabled for archiving at the new location.*
2. Check the configuration files that were supplied with the exported PST files to determine the retention categories to use when importing.
3. Make sure that each PST file is in the same folder as its corresponding configuration file.
4. Use PST Migrator (*EVPM or wizard*) to import all the items in the PST into the new archive.

NB. *During the import, shortcuts in the mailbox are automatically modified to point to the new location of the archived items. This is relevant in the case where the mailbox has been moved and still contained EV shortcuts because once the user archive has been moved the original shortcuts are no longer valid and need to be updated (using this method) or deleted.*

Note: *A fix was made in EV2007 SP1 to resolve an issue when importing multiple PST files with associated configuration files, shortcuts in the first mailbox were modified to point to the new archive location but shortcuts in the other mailboxes were not modified. The work around for earlier releases is to only import one PST per archive for each PST import iteration.*

User Archive Move Process



1. Disable the mailbox – to avoid further shortcuts being created in mailbox
2. Export the user archive to a PST
3. Move Exchange mailbox to another server
4. Enable mailbox in the new EV environment
5. Import the PST from the old EV environment
6. Delete the original archive.

NB. If the user archive has been exported to another EV environment, you may want to leave the archive on the original EV server for compliance purposes, this will depend upon what your local policy is for retaining content in your region/department.

Moving Across Regions or Sites

In today’s world of global organisations it is quite common for users to move between offices; quite often this is also between regions such as Europe and North America. Many organizations operate a globally managed AD environment with regionalised Exchange/EV server management. This will quite often mean that when a user moves offices, their mailbox will be moved between Exchange servers to ensure they are not subject to WAN latency and that they can use localised remote access (such as Blackberry or OWA). The mailbox move will probably break the relationship between the user’s mailbox and its EV user archive because the user’s new Exchange server will be archived by a separate EV site or perhaps part of a completely separate instance of EV. This situation will lead to a number of different actions depending upon your configuration and the policy of your organisation.

Option 1: Move mailbox but leave user archive

This is useful if you have users that frequently change offices. This might be preferable to moving the user archive around each time the user moves (leading to lots of export/import process). This method assumes the office the user is moving to is managed by a separate instance of EV. So you will move the Exchange mailbox to the new location and it will then start to be archived by the other EV instance but the old shortcuts can still be opened. This will also mean that the user will have two archives (*the original one and the one at the new location*) but only the new one will be active.

Considerations for implementing Option 1

Access Rights – if the user is using the same AD credentials to access their new mailbox then no action is required on the original user archive. If the user has a different user account then their original user archive will need to be modified to allow access to their new account.

Site Alias – Since the user’s mailbox contains EV shortcuts that point to a non-local EV instance it will not know how to locate the original EV server. When a user clicks on one of their old shortcuts it will attempt to access the local EV server and of course fail to recall the archived item. This is because the site alias stored in the shortcut won’t be resolved correctly as it’s no longer valid for the new location. This can be rectified by modifying a setting within the Exchange Policy for the desktop settings (now called Outlook in EV2007). There is a parameter which allows multiple site aliases to be set, this is called WebAppURL and is explained below.

<i>Legacy name:</i> WebAppURL	
Value	Description
One or more URLs	The URLs of one or more Web Access applications in different vault sites. Enter the addresses in this format: [site1.dns.alias]=http://site1.server.com/EnterpriseVault; [site2... If you specify multiple addresses, separate them with semicolons (;). The overall length of the string cannot exceed 255 characters.

Searching Original Vault – Providing the user still has access to their original vault, you can optionally provide them with the search URL for their original EV server. This will allow them to search for old emails but they will not be able to restore archived items because of course their original EV server has no relationship with the Exchange server where their mailbox has been moved to. You might also want to provide them with the ArchiveExplorer URL on the original EV server, this is obviously another method which might make it easier for them to find old email.

Option 2: Move archive to EV server in the new location

This is effectively the “move archive” process which was described earlier. This method involves creating an Export of the archive in one or more PST files; these are then imported into the new user archive at the new location. This method allows you to move the archived content from one location to another; this might be because there of network or security limitations that would prevent Option 1 from being used. This method does mean that the user keeps all their archived content in a single archive but obviously generates more work for the admin to export and re-import the content as well as the overhead in copying large PST files between locations.

Some organisations may not permit archived content to be exported and taken to another location or subsidiary within the company. If this is the case, then a possible workaround is to use Option1 , this still allows access to the archived data but doesn't require that it is moved between EV installation.

Option 3: Restore Shortcuts before Mailbox is Moved

To avoid having any dependency on the user archive when moving mailboxes, one option is to restore all shortcuts before the mailbox is moved. This could be done by the user via the EV Outlook add-in (*also known as User Extensions*) or by the administrator (using the Export to Mailbox feature in the VAC). However, clearly this may not be appropriate in all instances because of the amount of data this would add to the mailbox, possibly causing it to go over quota. This might also add a significant overhead when moving the mailbox between Exchange servers. However, it does mean that once the user has been moved to the new location they would have all their current and archived email in the mailbox.

Option 4: Export to PST

This would be used if the location which the user is moving to does not have EV installed. You simply export the user archive to a PST and they take it with them to the new location. This might not be permitted in some organisations for reasons of security or compliance.

About the Author

Glenn Martin is a Symantec Regional Product Manager for Enterprise Vault operating out of Symantec's offices in the United Kingdom. He is a 25 year veteran of the IT business and has worked in various roles (from hardware vendors to oil companies) but has always been close to email and infrastructure management. Over the last 10 years Glenn has specialized in Exchange, Internet email management and in 2002 was an early adopter of Enterprise Vault whilst working within investment banking in London. Now at Symantec he brings his wealth of experience to the product management group to help shape the future of Enterprise Vault. A key part of his role now is engaging with existing EV customers to understand the challenges they face in the ever evolving world of email archiving.

About Symantec Technology Network (STN)

Symantec Technology Network (STN) is Symantec Corporation's technical information generation and dissemination organization. It distributes a free monthly technical newsletter that discusses timely technology events to 120,000 email subscribers across the globe. STN also publishes technical data storage and security articles each month for large enterprise and Small and Medium Business (SMB) readers, as well as hosts a variety of blogs and product discussion forums discussing Symantec product tips and insights. To subscribe to STN's Technical Newsletter and review other STN materials, please visit STN at:

<http://www.symantec.com/enterprise/stn/index.jsp>

W. David Schwaderer is STN's Editor-In-Chief. He has authored ten technical books, and six commercial software programs. His eleventh book titled *Innovation Survival* is nearing completion. He has a Masters Degree in Applied Mathematics from the California Institute of Technology and an MBA from the University of Southern California. He lectures at Stanford in Silicon Valley on the subject of innovation.

Russ G. Gregg is STN's Video Content Producer. He is presently focusing on producing state-of-the-art digital video for the Web and small form-factor, portable video devices such as Apple Inc.'s iPhone™ and Video iPod®. Russ has a Bachelor of Arts in English Literature, Texas from



A&M University.

About Symantec

Symantec is the world leader in providing solutions to help individuals and enterprises assure the security, availability, and integrity of their information. Headquartered in Cupertino, Calif., Symantec has operations in more than 40 countries. More information is available at www.symantec.com.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
408 517 8000
800 721 3934

Symantec and the Symantec logo are U.S. registered trademarks of Symantec Corporation. Other brands and products are trademarks of their respective holder/s. Any technical information that is made available by Symantec Corporation is the copyrighted work of Symantec Corporation and is owned by Symantec Corporation. NO WARRANTY. The technical information is being delivered to you as-is and Symantec Corporation makes no warranty as to its accuracy or use. Any use of the technical documentation or the information contained herein is at the risk of the user.