

Archive Shuttle 11.4 **Migration Guide**



© 2024 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Metalogix are trademarks and registered trademarks of Quest Software Inc. and its affiliates. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are the property of their respective owners.

Legend



CAUTION: A caution icon indicates potential damage to hardware or loss of data if instructions are not followed.



IMPORTANT, NOTE, TIP, MOBILE OR VIDEO: An information icon indicates supporting information.

Archive Shuttle
Updated August 2024
Version 11.4

Contents

Introduction	5
Architecture	7
Enterprise Vault to Enterprise Vault	9
Migrating an archive from Enterprise Vault to Enterprise Vault	15
Enterprise Vault to Exchange	19
Migrating an archive from Enterprise Vault to Exchange	24
Enterprise Vault to Office 365	27
Migrating an archive from Enterprise Vault to Office 365	32
Mapping retention categories from Enterprise Vault to Office 365 retention tags	35
Restore Enterprise Vault items associated foreign shortcuts within a user's mailbox to Office 365	37
Enterprise Vault to Proofpoint	39
Enterprise Vault Discovery Accelerator to Office 365	44
Enterprise Vault to PST	47
Migrating an archive from Enterprise Vault to PST	52
PST as a source	56
Migrating from PST to a target environment	59
EAS as a source	64
Migrating from EAS to a target environment	67
Metalogix as a source	71
Migrating from Metalogix Archive Manager to a target environment	72
Dell Archive Manager as a source	76
Migrating from DAM to a target environment	80
SourceOne as a source	83
Migrating from SourceOne or EmailXTender to a target environment	87
Migrating a journal archive	91
Migrating a journal archive to PST or Exchange	91
Performing the journal mapping	91
Migrating a journal archive to Enterprise Vault	92
Pre-Enterprise Vault 10.0.0.4 source	92
Enterprise Vault 10.0.0.4 or later source	94
Migrating a journal archive to Office 365	97
Journal Transformation migration	97
Enterprise Vault to Office 365	103
Journal archive folder split settings	105
Journal Splitting	107

Suggested values for splitting journal data	108
Convert journal messages to Office 365 journaling format	109
Migrating leavers to Office 365	111
Configuring the Office 365 environment	112
Performing the Migration	113
Migrating a shared mailbox archive	115
About Us	117
Contacting Quest	118

Introduction

This guide contains information that will help you plan your Archive Shuttle migrations.

Key terms

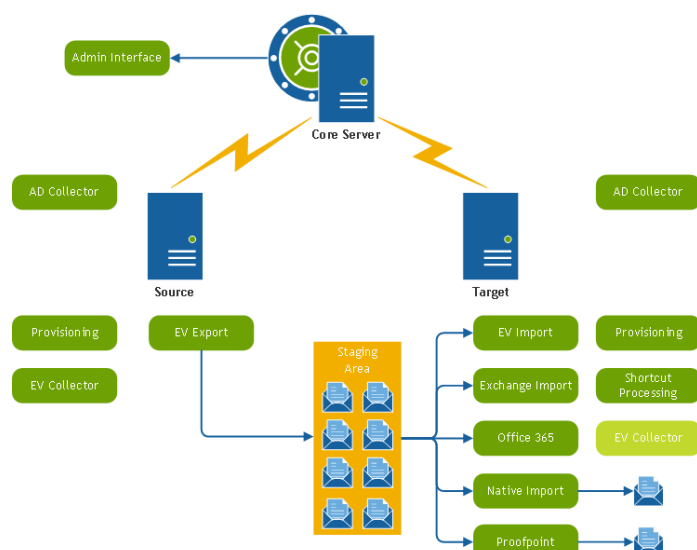
The following table introduces the terminology that is used throughout Archive Shuttle documentation, videos, and the user interface.

TERM	ITEM + DESCRIPTION
Link	Enterprise Vault A Link is a Vault Store Exchange A Link is an Exchange database Office 365 A connection to Office 365 PST A connection to a PST Output Area Proofpoint A connection to a Proofpoint output area EAS A connection to an EAS IIS Server Metalogix A connection to a Metalogix server
Container	Enterprise Vault A Container is a Vault / Archive Exchange A Container is an Exchange mailbox Office 365 A Container is an Office 365 mailbox or Personal Archive

TERM	ITEM + DESCRIPTION
	<p>PST</p> <p>A container is a PST file</p> <p>Proofpoint</p> <p>A container currently has no applicable context</p> <p>EAS</p> <p>A container is an archive relating to a user.</p> <p>Metalogix</p> <p>A container is an archive</p>
Item	<p>Enterprise Vault</p> <p>An Item is an archived Item.</p> <p>Exchange</p> <p>An Item is an item in the mailbox</p> <p>Office 365</p> <p>An Item is an item in the mailbox or personal archive</p> <p>PST</p> <p>An item is an item inside a PST file</p> <p>Proofpoint</p> <p>An item currently has no applicable context</p> <p>EAS</p> <p>An item is a message in the archive</p> <p>Metalogix</p> <p>An item is a message in the archive</p>

Architecture

The following diagram shows the Archive Shuttle components:



Archive Shuttle Core

The Archive Shuttle Core consists of three parts:

- Archive Shuttle User Interface
- Archive Shuttle Web Services
- Archive Shuttle Service

These are explained below.

Archive Shuttle user interface (UI)

The main point of interaction with the administrator is through the web-based user interface. It allows access from anywhere within the enterprise by just using a web browser. It is even possible for archive migrations to be performed remotely from outside of an organization; for example, partner-managed migrations.

Using the wide range of options in the Archive Shuttle User Interface, the administrator can configure and manage archive migrations within an enterprise. The interface also provides at-a-glance progress monitoring, as well as reporting.

Archive Shuttle Web Services

All interaction between Archive Shuttle and Archive Shuttle Modules is done through the Archive Shuttle Web Services. All modules communicate with the Web Services using HTTP(s). If HTTPS is to be used, there is additional configuration that needs to be performed.


Archive Shuttle service

The Archive Shuttle Service is a Windows Service which periodically run Archive Shuttle scheduled tasks. These can be database maintenance tasks or other tasks that need periodic execution in Archive Shuttle itself. It also execute database updates over the directory and item databases.

Enterprise Vault to Enterprise Vault

This section explains what the requirements are for an Enterprise Vault (EV) to Enterprise Vault migration, and then covers the basic steps that should be followed to start such a migration.

Stage 1 – Synch

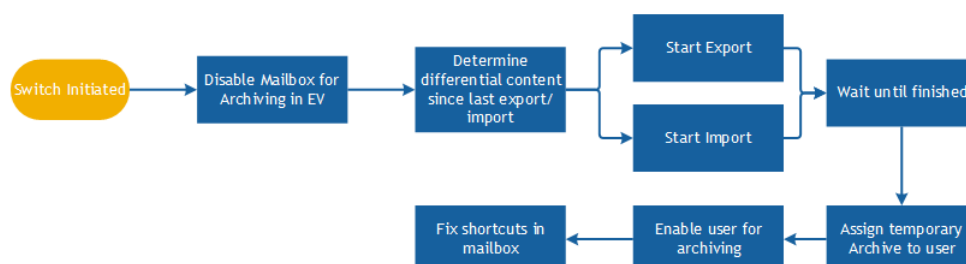


As soon as a user is enabled for migration, the synch process-flow shown above starts.

Archive Shuttle Core sends a command to collect all needed metadata information. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.

The Enterprise Vault Export and Enterprise Vault Ingest process is then started. This runs continuously in the background until the Stage 2 (Switch) phase is initiated for the user, or a specified “external” trigger is fired (e.g the mailbox has been moved from Exchange 2003 to Exchange 2010).

Stage 2 - Switch



After the Stage 2 ('Switch') phase to the target environment is initiated, the configured workflow starts, and by default determines the gap between the source archive and the last imported item. Archive Shuttle then synchronizes the difference one last time to the target mailbox. The archive then gets assigned to the user and enabled. The last step is to cleanup shortcuts in the target mailbox.

Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The EV-type modules should be installed on each appropriate source Enterprise Vault server involved in the migration. The EV import module is not required on these machines.
Modules installed on target	The EV-type modules should be installed on each appropriate target Enterprise Vault server involved in the migration. The EV export module is not required on these machines.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Add Source EV Environment	EV Environment	Add the Enterprise Vault source environment.

Step	Screen	Description
Add Target EV Environment	EV Environment	Add the Enterprise Vault target environment. This is not necessary if both the source and target are the same.
Add Retention Category Mappings	Retention Category Mappings	Map the source and target environment retention categories. If the migration is intrasite, they can be automatically created in the admin interface.
Add Source Links	Links	Configure all appropriate source links. Select the links, and then click on 'Map Modules', create the item database(s) and enable archive gathering. Ensure there is no import module specified on the links as this may impact the process later on.
Add Target Links	Links	Add all appropriate target links. An EV Collector module should be associated with one of the links so that shortcut building information can be gathered from the site settings in Enterprise Vault.
Add Staging Area	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Configure Folder-Less Item Handling	System Configuration	Many versions of Enterprise Vault allow archiving of items in a folder-less area called the Top Of Information Store. On the System Configuration, enter a folder name to be used for placing the items in the target archive.
Stop the EV Provisioning task from auto-enabling users	Enterprise Vault Administration Console	Ensure that the source and target environments do not have any provisioning groups set to 'auto enable' users throughout the duration of the migration. This prevents users being targeted by alternating environments depending on when and which provisioning task runs.
Stop Storage Expiry running	Enterprise Vault Administration Console	Ensure that the source and target environments do not have storage expiry running. If it is enabled then it is possible that items which are collected can not be found when it is time to export them.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration can begin.

Determining Enterprise Vault archive owner(s)

Archive Shuttle must determine the owner of a mailbox archive. This section explains how this takes place..

Archive Shuttle takes the owner information of a mailbox archive from the Auto Security Descriptor field within the Enterprise Vault Directory database. The first entry for an archive will be the owner. Take the following three examples:

Simple 1

- Bill usage to: somedomain\simple1
- Permissions tab: somedomain\simple with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple1'.

Simple 2

- Bill usage to: somedomain\simple2
- Manually grant an additional user read, and delete permissions on the archive
- Permissions tab: somedomain\simple2 with inherited read, write and delete permissions.
somedomain\someotheruser with manually set read, and delete permissions

Archive Shuttle will see the owner of this archive as 'simple2'.

Simple 3

- Bill usage to: somedomain\simple3
- Grant full mailbox access via Exchange Management Console or Exchange Management Shell
- Permissions tab: somedomain\simple3 with inherited read, write and delete permissions.
somedomain\someotheruser also with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple3'.



NOTE: The order which the archives are listed in the permissions tab in Enterprise Vault does not reflect the order that the accounts are described in the Auto Security Descriptor field. The permissions tab shows the archive permissions in alphabetical order.

More detail

1. Archive Shuttle resolves the 'Owner' of an Archive as follows:

In the EnterpriseVaultDirectory database, the following SQL is used:

```
select ADMbxDN, LegacyMbxDN from ExchangeMailboxEntry where DefaultVaultId=@VAULTID
```

2. Archive Shuttle then does a lookup for the ADMbxDN in Active Directory to get the SID:

LDAP Query: (ADMBxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

3. Archive Shuttle then does a lookup for the LegacyMbxDN in Active Directory to get the SID:

LDAP Query: (LegacyMbxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

4. We then compare the SID's we got from ADMbxDN and from LegacyMbxDN, and if both of them match, we have an Owner. If they do not match, they will be marked as 'Ownerless' and move to the next option.
5. If we do not find an entry in ExchangeMailboxEntry, we resolve using the BillingOwner or the AutoSecurityDesc:

select SID from Root inner join trustee on OwningTrusteeIdentity = TrusteeIdentity where VaultEntryId=@ARCHIVEID



NOTE: Regarding OwnerUserSid versus SID fields in the EVArchive Table on in the Archive Shuttle Directory Database:

- SID is always filled out with the SID which is retrieved from the above steps
 - OwnerUserSid has a foreign key constraint on the [User] Table. So Archive Shuttle only fill this out if the SID obtained from EV is actually contained in the Archive Shuttle Directory Database (i.e. has been synced from AD)
5. The UserSidHistory table is stored Sid history for each user when they were migrated from AD1 to AD2. This SidHistory attribute is part of Active Directory, and Archive Shuttle collects this data from Active Directory directly.

Example

User1 was migrated from AD1 to AD2, SidHistory attribute was populated to Active Directory and when the Active Directory sync is performed, table UserSidHistory is populated in parallel. User1 was migrated to another Active Directory, so it is supposed to be the same user and therefore in owner resolution logic is UserSidHistory as the last step. Sometimes, even the migrated user from AD1 to AD2 is not supposed to be the same user as samAccountNames are different in both Active Directories. In this case, owner resolution will take AD2 user as the owner which could be wrong in some special situations.

Workaround

As a workaround, you can:

- remove SidHistory parameter from Active Directory, and allow Archive Shuttle to synchronize Active Directory entries (recommended option), run Enterprise Vault collection, to assign correct users in ContainerToUser table
- use the SQL to assign HistoryUserSid instead of actual Sid, example of such script is below:

```
--update ContainerToUser, take UserHistorySid and assign to Container type
1 (EV containers only)
SET NOCOUNT OFF;
declare @sid varchar(128)
DECLARE ITEM_CURSOR CURSOR FOR
Select UserSid From ContainerToUser a
inner join [dbo].[Container] b on a.ContainerId = b.ContainerId and
```

```

ContainerTypeId = 1
OPEN ITEM_CURSOR;
FETCH NEXT FROM ITEM_CURSOR INTO @sid
WHILE @@FETCH_STATUS = 0
BEGIN
update x set UserSid = (SELECT UserHistorySid FROM [dbo].[UserSidHistory]
where UserSid = @sid)
from [dbo].[ContainerToUser] x
where UserSid = @sid
print 'Updated containerid ' + @sid;
FETCH NEXT FROM ITEM_CURSOR INTO @sid;
END;
CLOSE ITEM_CURSOR;
DEALLOCATE ITEM_CURSOR;

```

- remove entries from UserSidHistory table, disable Active Directory sync to avoid entries are coming back with next Active Directory sync, run Enterprise Vault collection to assign correct users in ContainerToUser table

i | NOTE: Re-using of sAMAccountsName/MbxNTUser or MbxDisplayName in one domain can cause an issue with Enterprise Vault Archive assigning process.

General considerations

The following general considerations should be taken into account for this type of migration:

- Ingest speed is likely to be reduced versus the speed of export since ingesting is also going to place a load on the indexing engine on the target environment. In general do not expect the export and ingest speeds to be comparable.
- When migrating from a pre EV 10.0.3 system extraction of data may be slower than expected if Enterprise Vault Collections have been used. Access to the data inside the CAB file is single threaded. This issue does not exist in Enterprise Vault 10.0.3 and later.
- An issue is currently being investigated by Symantec. Items extracted from EV 9, may not be ingested into EV 11.0.1. This affects only a small number of items. This may be addressed in the future by Symantec.
- EVPM may have difficulty in connecting to, and processing, mailboxes which are homed on Exchange 2013. This knowledge base article may help.
- EVPM may need to have the DS Server registry key set. This knowledge base article may help.

Migrating an archive from Enterprise Vault to Enterprise Vault

For this scenario, the migration takes place between two vault stores in the same Enterprise Vault environment.

Prerequisites

The following steps must be completed before the migration begins:

- Two or more vault stores must exist in the Enterprise Vault environment.
- All modules must be enabled.
- Appropriate Active Directory domains must be enabled for scanning.
- An Enterprise Vault environment has been added for migration.
- A link database has been created for the source vault store/archives.
- Module mappings are configured for the source and the target vault stores. The source vault store needs to be linked to an EV Export module, and the target needs to be linked to an EV import module. In addition, an EV Provisioning Module needs to be linked to the source vault store and to the target vault store.
- During switchover, both the archiving and storage tasks need to run on the same Enterprise Vault server.

Add retention category mappings

The final stage of the setup of the archive migration is to map retention categories between the source and target environments.

1. Go to **Configuration > EV Retention Mappings**.
2. Click **Create Mapping**.
3. Select the values from the drop-down lists, and click **[Add]** to add the mapping to the retention category grid.

Retention mappings are managed from this screen.

There is an option to use “**Add Intrasite Migration Mappings**”, which maps each retention category to itself. This can be used in a situation where archives are to be migrated in the same environment. It simply moves the archives to the new location.

Map containers for the test archive

In order to migrate a test archive from one vault store to another, the source and target containers must be mapped. This can be done as follows:

1. Go to **Manage > Bulk Mapping**.
2. Type the beginning of the archive name in the **Container Name** filter. Click the button in the far right of the filter row in the **Apply Filter** column.

i | **NOTE:** If the test archive is not displayed, go back to the **Configuration > EV Environment** page and click **Sync all AD Users**. Then, select the vault store where the source archive is located and issue a **Run Now** for Archive Gathering.

3. Select the checkbox next to the test archive, and click **Add Mappings**.

i | **NOTE:** A wizard begins gathering information related to the mapping for the select archive. (Multiple archives can also be selected.)

4. When you're prompted to select a target container type, select **Enterprise Vault**, and then click **[Next]**.
5. When you're prompted to select a target user, select **Same User**, and then click **[Next]**.
6. When you're prompted to select a container strategy, select **Create new containers**, and then click **[Next]**.
7. Select the link that corresponds to where the archive is to be migrated, and then click **Next**.

i | **NOTE:** If the target link drop-down list is empty, then it is likely that an EV Import module has not been correctly associated with the link. Review the Links page.

8. On the **Workflow Policy** screen, select **Standard Workflow** – Within same EV Site (without Archive deletion), and then click **[Next]**. You don't need to select the ownerless workflow policy in this situation; it's optional.

i | **NOTE:** If a group of users was selected containing a mixture of ownerless and normal containers, this screen gives an administrator the option to specify which workflow to use for the ownerless containers.

9. For the filter policy, select **Default (No Filter)**, and then click **[Next]**.
10. On the **Container Mapping Settings** screen, select **Enabled** for both **Migration Status** and **Item Gathering Status**, and then click **[Next]**.
11. If desired, set the mapping's priority, and then click **[Next]**.
12. Review the summary screen. If everything looks as expected, click **[Confirm]**.

Review Stage 1 status

A few minutes after the mapping is created, the software tells the appropriate modules to start the actions defined in the mapping. Review the progress for this stage of the migration using these steps:

1. Go to **Manage & Operations > Stage 1 (Sync Data)**.
2. Type the beginning of the archive name in the **Name** filter, and then click the button in the **Apply Filter** column, found at the far right of the filter row.

3. Once the source archive displays, click the **Refresh** button to see the status of the export of the archive and the ingestion of data into the target archive. Continue to click Refresh until both export and import are complete.

i | **NOTE:** If the data progress bars reach 100% for export, but show no progress for import, the Retention Category mappings have likely not been configured.

If more data is added to the source archive, it's synchronized to the target archive every 24 hours using the connection made with the mapping you created in the previous section. In addition, part of the Stage 2 workflow is to perform a final sync. Therefore, it's not necessary to ensure that export and import has reached 100% before moving on to the next steps.

Validate exported data

If there's a large amount of data to export and import, and you want to check the progress, click the **Refresh** button located on the navigation bar on the Stage 1 (Sync Data) page.

In addition, you can use Windows Explorer to browse the staging area data on the disk and view the folder structure.

Also, the Enterprise Vault Admin Console shows that an archive exists in the source Vault Store and that a new archive has been created in the target vault store. It's also possible to grant a service account access to the target archive and perform searches on it to make sure the data matches the source archive.

Enable Stage 2

Before you enable Stage 2, the switch-over for the test archive, check the Stage 1 (Sync Data) page for issues like failed item-export or failed item-import.

To enable Stage 2 for the test archive, perform these steps:

1. Open the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and click **Enable Stage 2** in the navigation bar.
3. When you refresh the Stage 1 status, a check mark should display in the Stage 2 Enabled column.

Stage 2 is the "switch-over" to the target environment. A final synchronization of archived items is performed from the source environment to the target environment, before several additional migration tasks are performed.

Review Stage 2 status

After a few minutes, the progress of the test archive migration displays on the Stage 2 page:

1. Go to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Container Name** filter. Apply the filter by clicking the button in the **Apply Filter** column, located at the far right of the filter row.



NOTE: If the archive isn't displayed, wait 1-2 minutes, and then click **Refresh** in the navigation bar.

3. Once the source archive displays, click **Refresh** to see the progress of the final stages of the archive migration. By default, since the workflow policy that was selected was Standard policy (without archive deletion), these steps are performed:
4. Disable source mailbox from Enterprise Vault archiving.
5. Rename the source archive.
6. Collect any remaining items for migration.
7. Import the remaining items into the target.
8. Zap the mailbox, removing EV settings.
9. Assign the new archive to the user.
10. Enable the mailbox for archiving again.
11. Rename the target archive.

Update all the existing shortcuts to point to the new archive.

Verify data has been migrated


After all Stage 2 operations are complete, the **Stage 2 (Switch User)** page for the test archive shows a check mark in the **Finished** column.

In addition, if Outlook or Outlook Web Access is used to access the test mailbox, then all of the archived items are accessible via the new archive. Shortcuts to archived items will work and open correctly.

Enterprise Vault to Exchange

This section explains what the requirements are for an Enterprise Vault (EV) to Exchange migration, and then covers the basic steps that should be followed to start such a migration.

Stage 1 - Synch

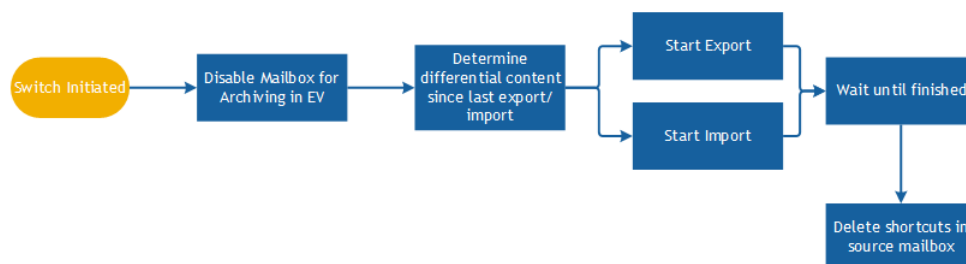


As soon as a user is enabled for migration, the synch process-flow shown above starts.

Archive Shuttle Core sends a command to collect all needed metadata information such as the number of items, size, and Enterprise Vault transaction ID's for the archive of the user. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.

The Enterprise Vault Export and Exchange Ingest process is then started. It runs continuously in the background until the administrator initiates a "switch" (Stage 2) for the user or a specified "external" trigger is fired (e.g., the mailbox has been moved from Exchange 2003 to Exchange 2010).

Stage 2 - Switch



After the Stage 2 ('Switch') phase to the target environment is initiated, the configured workflow starts, and by default determines the gap between the source archive and the last imported item. Archive Shuttle then synchronizes the difference one last time to the target mailbox.

The last step is to cleanup shortcuts in the source mailbox by deleting them.

Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The EV-type modules should be installed on each appropriate source Enterprise Vault server involved in the migration. The EV import module is not required on these machines.
Modules installed on appropriate bridgehead to reach target	Exchange and Post Processing modules are the two which are primarily needed to work with this type of migration.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
Ingest accounts configured for Application Impersonation	5 accounts should be configured with Application Impersonation rights, and those accounts should be added to the credential editor on the machine running the Exchange ingest module.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.

Step	Screen	Description
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Add EV Environment	EV Environment	Add the Enterprise Vault source environment.
Add Source Links	Links	Configure all appropriate source links. Select the links, and then click on 'Map Modules', create the item database(s) and enable archive gathering. Ensure there is no import module specified on the links as this may impact the process later on.
Add Target Links	Links	Add all appropriate target links. An ingest and post-processing module should be associated with each possible target Exchange Database.
Add Staging Area	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Configure Folder-Less Item Handling	System Configuration	Many versions of Enterprise Vault allow archiving of items in a folder-less area called the Top Of Information Store. On the System Configuration, enter a folder name to be used for placing the items in the target archive.
Ensure AIP/EWS is configured correctly	System Configuration	Ensure that the migration providers are listed in the required order and that the AutoDiscoverURL is specified if required. The account which is to be used for ingest needs to have Application Impersonation rights. This article will help with configuring that.
Consider disabling reminders for appointments in the past	System Configuration > Exchange Import Module	Consider enabling the option to 'Disable reminders for appointments in the past' this will prevent pop-ups appearing from Outlook on end-user machines when old appointments are processed.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration can begin.

Determining Enterprise Vault archive owner(s)

Archive Shuttle must determine the owner of a mailbox archive. This section explains how this takes place..

Archive Shuttle takes the owner information of a mailbox archive from the Auto Security Descriptor field within the Enterprise Vault Directory database. The first entry for an archive will be the owner. Take the following three examples:

Simple 1

- Bill usage to: somedomain\simple1
- Permissions tab: somedomain\simple with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple1'.

Simple 2

- Bill usage to: somedomain\simple2
- Manually grant an additional user read, and delete permissions on the archive
- Permissions tab: somedomain\simple2 with inherited read, write and delete permissions.
somedomain\someotheruser with manually set read, and delete permissions

Archive Shuttle will see the owner of this archive as 'simple2'.

Simple 3

- Bill usage to: somedomain\simple3
- Grant full mailbox access via Exchange Management Console or Exchange Management Shell
- Permissions tab: somedomain\simple3 with inherited read, write and delete permissions.
somedomain\someotheruser also with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple3'.

i **NOTE:** The order which the archives are listed in the permissions tab in Enterprise Vault does not reflect the order that the accounts are described in the Auto Security Descriptor field. The permissions tab shows the archive permissions in alphabetical order.

More detail

1. Archive Shuttle resolves the 'Owner' of an Archive as follows:

In the EnterpriseVaultDirectory database, the following SQL is used:

```
select ADMbxDN, LegacyMbxDN from ExchangeMailboxEntry where DefaultVaultId=@VAULTID
```

2. Archive Shuttle then does a lookup for the ADMbxDN in Active Directory to get the SID:

LDAP Query: (ADMBxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

3. Archive Shuttle then does a lookup for the LegacyMbxDN in Active Directory to get the SID:

LDAP Query: (LegacyMbxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

4. We then compare the SID's we got from ADMbxDN and from LegacyMbxDN, and if both of them match, we have an Owner. If they do not match, they will be marked as 'Ownerless' and move to the next option.
5. If we do not find an entry in ExchangeMailboxEntry, we resolve using the BillingOwner or the AutoSecurityDesc:

select SID from Root inner join trustee on OwningTrusteeIdentity = TrusteeIdentity where VaultEntryId=@ARCHIVEID



NOTE: Regarding OwnerUserSid versus SID fields in the EVArchive Table on in the Archive Shuttle Directory Database:

- SID is always filled out with the SID which is retrieved from the above steps
 - OwnerUserSid has a foreign key constraint on the [User] Table. So Archive Shuttle only fill this out if the SID obtained from EV is actually contained in the Archive Shuttle Directory Database (i.e. has been synced from AD)
5. The UserSidHistory table is stored Sid history for each user when they were migrated from AD1 to AD2. This SidHistory attribute is part of Active Directory, and Archive Shuttle collects this data from Active Directory directly.

General considerations

The following general considerations should be taken into account for this type of migration:

- When migrating from a pre EV 10.0.3 system extraction of data may be slower than expected if Enterprise Vault Collections have been used. Access to the data inside the CAB file is single threaded. This issue does not exist in Enterprise Vault 10.0.3 and later.
- EVPM may have difficulty in connecting to, and processing, mailboxes which are homed on Exchange 2013. This knowledge base article may help.
- EVPM may need to have the DS Server registry key set. This knowledge base article may help.

Migrating an archive from Enterprise Vault to Exchange

Complete the steps below to migrate a single test archive within Enterprise Vault to an Exchange 2010/2013 personal archive.

Prerequisites

The steps below must be complete before the migration begins:

- Enable all Modules.
- Enable appropriate Active Directory domains for scanning.
- Add an Enterprise Vault environment for the migration.
- Create a link database for the source archives.
- Configure module mappings for the source archives and the target Exchange environment.
- Run a PowerShell command to allow application impersonation.

In addition, if you'll use an Exchange personal archive, rather than an ordinary mailbox for the test migration, the personal archive needs to be pre-created. The personal archive doesn't need to be created on the same mailbox database as the primary mailbox.

Map containers for the test archive

To migrate a test archive from Enterprise Vault to Exchange 2010/2013, the source and target containers must be mapped. Follow these steps to do it:

1. Go to **Manage > Bulk Mapping**.
2. Type the beginning of the archive name in the **Container Name** filter. Click the button in the **Apply Filter** row, located at the far right of the filter row. Note: If the test archive doesn't display, go back to the **Configuration > EV Environment** page and click **Sync all AD Users**. Then, select the vault store where the source archive is located and click **[Run Now]**.
3. Select the checkbox next to the test archive(s) and click **Add Mappings**. A wizard runs, gathering information related to the mapping(s) for the selected archive(s).
4. When you're prompted to choose a target container type, select **Exchange** and click **[Next]**.
5. When you're prompted to choose the type of migration, select **Normal** and click **[Next]**.
6. When you're prompted to choose the target mailbox type, select **Secondary (Archive)** mailbox. Then, in the second part of the dialog select **Skip**, and then click **[Next]**.
7. When you're prompted to choose the target user, select **Same User**, and click **[Next]**.
8. On the **Choose Workflow Policy** screen, select the **Exchange migration (without archive deletion)** option, leave the second drop-down list blank, and then click **[Next]**.

9. When you're prompted to choose a filter policy, leave the **Default (No Filter)** option selected, and then click **[Next]**.
10. On the **Container Mapping Settings** screen, select **Enabled** for both **Migration Status** and **Item Gathering Status**, and then click **[Next]**.
11. Leave the priority blank and click **[Next]**.
12. Review the summary screen, and then click **Confirm**.

A few minutes after the mapping is created, appropriate modules start the actions defined in the mapping. The progress of this stage of the migration can be reviewed by following these steps:

1. Go to the **Stage 1 (Sync Data)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the **Apply Filter** button at the far right of the filter row.
3. Once the source archive displays, click **Refresh** to see the progress of the export of the archive and the ingestion of data into the Exchange personal archive. Continue to click **Refresh** until both the export and import are complete.

Validate exported data

In addition to checking export/import status (using the steps in the Review Stage 1 status procedure above) you can use Windows Explorer to browse the export/import storage area and data on the disk.

It's also possible to log in to Outlook or Outlook Web Access as the test user. The migrated data will be present in the Personal Archive, and can be freely opened and manipulated.

i **NOTE:** If the test user previously used the Exchange personal archive feature, it might be difficult to locate the migrated data. Because of this, Archive Shuttle recommends that the test user has an empty personal archive before migration.

Enable Stage 2

Before enabling Stage 2 (Switch User) for the test archive, check the **Stage 1 Status** page for issues like failed item-export or item-import.

Complete these steps to enable Stage 2 (Switch User) for the test archive

1. Go to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and click **Enable Stage 2** in the navigation bar.
3. When you refresh the Stage 1 (Sync Data) page, there should be a green and white check mark in the **Stage 2 Enabled** column.

Stage 2 is the switchover to the target environment. A final synchronization of archived items is performed from the source environment to the target environment before several additional migration tasks are performed.

Review Stage 2 Status

After a few minutes, view the progress of the test archive migration on the Stage 2 Status page by following these steps:

1. Go to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Container Name** filter. Click the **Apply** button located at the far right of the row in the Apply Filter column.



NOTE: If the archive is not displayed, wait one to two minutes and click **Refresh** in the navigation bar.

3. Once the source archive displays, click the **Refresh** button to show the progress of the final stages of the archive migration. By default, since the Exchange migration (without archive deletion) workflow policy was selected, the following steps are performed:
 - a) Disable source mailbox from Enterprise Vault archiving.
 - b) Rename the source archive.
 - c) Collect any remaining items for migration.
 - d) Import the remaining items into the target.
 - e) Delete any Enterprise Vault shortcuts in the mailbox

Verify data has been migrated


Once all Stage 2 operations are complete, the Stage 2 (Sync Data) for the test archive will have a check mark in the Finished column.

In addition, if Outlook or Outlook Web Access is used to access the test mailbox, then all archived items will be in the personal archive and no Enterprise Vault shortcuts will remain in the source mailbox.

Enterprise Vault to Office 365

This section explains what the requirements are for an Enterprise Vault (EV) to Office 365 migration, and then covers the basic steps that should be followed to start such a migration.

Stage 1 – Synch

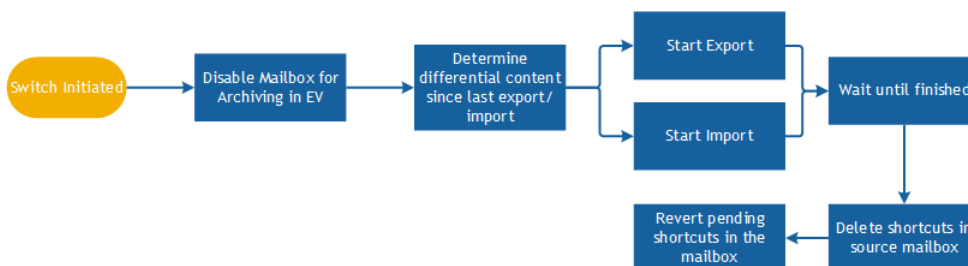


As soon as a user is enabled for migration, the synch process-flow shown above starts.

Archive Shuttle Core sends a command to collect all needed metadata information such as the number of items, size, and Enterprise Vault transaction ID's for the archive of the user. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.

The Enterprise Vault Export and Office 365 module then begin to work. This runs continuously in the background until the administrator initiates a "Switch" (Stage 2).

Stage 2 - Switch



After the Stage 2 ('Switch') phase to the target environment is initiated, the configured workflow starts, and by default determines the gap between the source archive and the last imported item. Archive Shuttle then synchronizes the difference one last time to the target mailbox.

The last step is to 'Zap' the mailbox to remove Enterprise Vault settings, cleanup shortcuts in the source mailbox by deleting them, and change any pending-archive items back to normal items.

Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The EV-type modules should be installed on each appropriate source Enterprise Vault server involved in the migration. The EV import module is not required on these machines.
Modules installed on appropriate bridgehead to reach target	Office 365 and Post Processing modules are the two which are primarily needed to work with this type of migration.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
Ingest accounts configured for Application Impersonation	5 accounts should be configured with Application Impersonation rights, and those accounts should be added to the credential editor on the machine running the Office 365 module. One account should also be configured as a Global Administrator in Office 365, and be marked as that in the credential editor.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.

Step	Screen	Description
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Add EV Environment	EV Environment	Add the Enterprise Vault source environment.
Add Source Links	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Add Staging Area	Links	Select the newly created link, and map a Shortcut Process Module to it.
Configure Folder-Less Item Handling	System Configuration	Many versions of Enterprise Vault allow archiving of items in a folder-less area called the Top Of Information Store. On the System Configuration, enter a folder name to be used for placing the items in the target archive.
Set Office 365 Credentials		OAuth authentication Screen: Office 365 module Run the Archive Shuttle Office 365Module Credential Editor and add OAuth credentials. Then test the connection. Then, from the Office 365 Module, select the Use modern authentication (OAuth) option.
Update Shortcut Processing Module Configuration	System Configuration > Shortcut Processing	Update the settings so that 'Use EWS for Processing' is enabled.
Consider disabling reminders for appointments in the past	System Configuration > Office 365 Module	Consider enabling the option to 'Disable reminders for appointments in the past'.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration started.

Determining Enterprise Vault archive owner(s)

Archive Shuttle must determine the owner of a mailbox archive. This section explains how this takes place..

Archive Shuttle takes the owner information of a mailbox archive from the Auto Security Descriptor field within the Enterprise Vault Directory database. The first entry for an archive will be the owner. Take the following three examples:

Simple 1

- Bill usage to: somedomain\simple1
- Permissions tab: somedomain\simple with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple1'.

Simple 2

- Bill usage to: somedomain\simple2
- Manually grant an additional user read, and delete permissions on the archive
- Permissions tab: somedomain\simple2 with inherited read, write and delete permissions.
somedomain\someotheruser with manually set read, and delete permissions

Archive Shuttle will see the owner of this archive as 'simple2'.

Simple 3

- Bill usage to: somedomain\simple3
- Grant full mailbox access via Exchange Management Console or Exchange Management Shell
- Permissions tab: somedomain\simple3 with inherited read, write and delete permissions.
somedomain\someotheruser also with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple3'.



NOTE: The order which the archives are listed in the permissions tab in Enterprise Vault does not reflect the order that the accounts are described in the Auto Security Descriptor field. The permissions tab shows the archive permissions in alphabetical order.

More detail

1. Archive Shuttle resolves the 'Owner' of an Archive as follows:

In the EnterpriseVaultDirectory database, the following SQL is used:

```
select ADMbxDN, LegacyMbxDN from ExchangeMailboxEntry where DefaultVaultId=@VAULTID
```

2. Archive Shuttle then does a lookup for the ADMbxDN in Active Directory to get the SID:

LDAP Query: (ADMBxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

3. Archive Shuttle then does a lookup for the LegacyMbxDN in Active Directory to get the SID:

LDAP Query: (LegacyMbxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

4. We then compare the SID's we got from ADMbxDN and from LegacyMbxDN, and if both of them match, we have an Owner. If they do not match, they will be marked as 'Ownerless' and move to the next option.
5. If we do not find an entry in ExchangeMailboxEntry, we resolve using the BillingOwner or the AutoSecurityDesc:

select SID from Root inner join trustee on OwningTrusteeIdentity = TrusteeIdentity where VaultEntryId=@ARCHIVEID



NOTE: Regarding OwnerUserSid versus SID fields in the EVArchive Table on in the Archive Shuttle Directory Database:

- SID is always filled out with the SID which is retrieved from the above steps
 - OwnerUserSid has a foreign key constraint on the [User] Table. So Archive Shuttle only fill this out if the SID obtained from EV is actually contained in the Archive Shuttle Directory Database (i.e. has been synced from AD)
5. The UserSidHistory table is stored Sid history for each user when they were migrated from AD1 to AD2. This SidHistory attribute is part of Active Directory, and Archive Shuttle collects this data from Active Directory directly.

General considerations

The following general considerations should be taken into account for this type of migration:

- A call should be logged with Microsoft to get the throttling limits raised for the duration of the migration, otherwise ingest performance will be reduced.
- When migrating from a pre EV10.0.3 system extraction of data may be slower than expected if Enterprise Vault Collections have been used. Access to the data inside the CAB file is single threaded. This issue does not exist in Enterprise Vault 10.0.3 and later.
- Different flavors of Office 365 license have different size limits, as described here. It is advised to run SQL Queries on the source environment in order to determine if any archives are likely to have.
- Many large items- Many archives have more data in them than is allowed by the Office license.
- It is advised to ensure a high level of Mailbox Parallelism and a low level of item/batch parallelism. This gives least chance to get to throttling limits.

Migrating an archive from Enterprise Vault to Office 365

This section outlines steps of the migration of a single test archive within Enterprise Vault to an Office 365 mailbox.

Prerequisites

Complete these steps before beginning the migration:

- Enable all modules.
- Enable appropriate Active Directory Domains for scanning.
- Add an Enterprise Vault Environment for the migration.
- Create a Link Database for the source archives.
- Configure Module Mappings for the source archives.



NOTE: It is assumed that the mailbox has been moved to Office 365 already and that it is the EV Archive that is to be moved.

Office 365 Global Administrator

In addition, it is necessary to configure the Archive Shuttle to Office 365 connection to use credentials that have sufficient privileges.

User Principal Name (UPN)

When mailbox information is gathered from Office 365, it is matched against local Active Directory users by using the Primary SMTP Address.

This means that the local Active Directory users must have the same UPN as the Office 365 users, or have the same Primary SMTP Address.



NOTE: Archive Shuttle does not support automatic matching of mailboxes to more than one set of Active Directory users (for example, matching to both on-prem and cloud users) when the same UPNs are present in both Active Directory instances.

Configure the Office 365 Module

Details can be found in the *Archive Shuttle Installation Guide*.

Create Office 365 link

In order to migrate an archive to Office 365, an Archive Shuttle Office 365 Link is required. Follow these steps to configure it:

1. Go to the **Links** page.
2. Click **Office 365** in the Actions Bar.
3. Click **Create Link**.
4. Give the link a name, for example, Office 365.
5. Click **Create**.

The new link needs to have Archive Shuttle Modules associated with it, and the Staging Area Path needs to match the EV Export Module.

Sync mailbox data

Before you can map a container in the source environment to an Office 365 target, select the **Office 365** link on the **Links administration** page, and click **Sync Mailboxes** in the Actions Bar.

Map containers - for the test archive

In order to migrate a test archive from Enterprise Vault to Office 365, the source and target containers must be mapped. This can be done as follows:

1. Go to the **Bulk Mapping** page.
2. Type the beginning of the archive name in the **Name** filter. Click **Apply**, located at the right of the filter row.

i **NOTE:** If the test archive is not displayed, go back to the EV Environment page and click **Sync all AD Users**, then select the Vault Store where the source archive is located and issue a **Run Now for Archive Gathering**.

3. Select the checkbox next to the test archive, and click **Add Mappings**.

i **NOTE:** A short wizard runs, gathering information related to the mapping for the select archive. (Multiple archives can also be selected.)

4. On the **Target Container Type** screen, select **Office 365**, and click **[Next]**.
5. Choose whether the migration should be performed to an **Office 365 Mailbox** or **Personal Archive**, and click **[Next]**.
6. Select the link to use for this operation, and click **[Next]**.
7. Choose the **Workflow Policy** for this migration, and click **[Next]**.
8. Choose the **Filter Policy** for this migration, and click **[Next]**.
9. On the **Container Mapping Settings** screen, select **Enabled** for both **Migration Status** and **Item Gathering Status**, and click **[Next]**.
10. Review the summary screen before clicking **[Confirm]**.

Review Stage 1 Status

A few minutes after the mapping is created, Archive Shuttle instructs the appropriate modules to start the actions defined in the mapping. The progress of this stage of the migration can be reviewed as follows:

1. Go to the **Stage 1 (Sync Data)** page.
2. Type the beginning of the archive name in the **Name** filter. Click **Apply** at the right of the filter row.
3. Once the source archive displays, click **Refresh** from time to time to show the progress of the export of the archive and the ingesting of the data into Office 365. Continue to click **Refresh** until both export and import reach 100%.

Validate exported data

If there is a large amount of data to export and import, then it will be possible to see the progress bars for the Stage 1 Status move by clicking **Refresh** from the navigation bar on the Stage 1 (Sync Data) page of the user interface.

In addition, you can use Windows Explorer to view the export/import storage area data on the disk.

Also, at this time it will be possible to login to Outlook or Outlook Web Access as the test user. The migrated data will be present in the Office 365 mailbox, and can be freely opened and manipulated.

Enable Stage 2

Before enabling Stage 2, the switch, for the test archive, the **Stage 1 (Sync Data)** page should be checked for issues such as failed item-export or item-import.

To enable Stage 2 for the test archive, perform the following steps:

1. Go to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and click **Enable Stage 2** in the navigation bar.
3. When the Stage 1 Status page is refreshed, there should be a green and white check mark in the Stage 2 Enabled column.

Stage 2 is the switchover to the target environment. A final synchronization is performed of archived items from the source environment to the target environment, before several additional migration tasks are performed.

Review Stage 2 Status

After a few minutes, the progress of the test archive migration will be visible in the Stage 2 Status page, as follows:

1. Go to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Name** filter. Click Apply at the right of the filter row.

i | **NOTE:** If the archive is not displayed, wait one to two minutes and click **Refresh** in the navigation bar.

3. Once the source archive is displayed, you can click **Refresh** from time to time to show the progress of the final stages of the archive migration. By default, since the workflow policy which was selected was 'Office 365 (without archive deletion)', the following steps will be performed by Archive Shuttle:
 - a. Rename the source archive
 - b. Collect any remaining items for migration
 - c. Import the remaining items in to the target
 - d. Delete any Enterprise Vault shortcuts in the Office 365 mailbox
 - e. Change any archive-pending items back to normal items in the Office 365 mailbox

Verify data has been migrated


Once all of the Stage 2 operations have been completed, the 'Stage 2 (Switch User)' for the test archive will have a green and white check mark in the 'Finished' column.

In addition Outlook or Outlook Web Access can be used to access the Office 365 mailbox to verify that all data from the Enterprise Vault archive is present and accessible.

Mapping retention categories from Enterprise Vault to Office 365 retention tags

You can collect retention tags from Office 365 and map/apply them in a migration where Enterprise Vault is used as the source, and Office 365 is used as a target.

By default the collection of retention tags is disabled. If you want to make use of this new feature, then the option in System Settings needs to be enabled:



When this is enabled, retention tags will be collected by the Office 365 module.

Retention category mapping can then be performed. The user interface has been extended to allow the mapping between Enterprise Vault retention category and Office 365 retention tag.

During a migration, items will be ingested into Office and two additional attributes will be set/populated:

- PR_POLICY_TAG (GUID of the retention tag)
- PR_RETENTION_PERIOD (Age limit for retention)

PowerShell commands have also been extended to allow for this mapping to process to be performed.

i | NOTE: If an item is exported, and there is no retention mapping for it, the item will not ingest.

There is also the ability to bind Enterprise Vault retention categories with the Office 365 Compliance Retention labels. When the feature is active, **PR_POLICY_TAG** and **PR_RETENTION_PERIOD** properties are put into messages properties. This works only with AIP as an ingest protocol. The validation on the existing retention mapping is processed; if this doesn't exist, the ingest is not processed (items are exported only). In the case when Enterprise Vault archive contains items with more retention categories (in case when PST file was imported into archive under another RC as is origin of the EV archive) and only one 1 RC mapping exists, then these items will ingested and stamped with the Office 365 Compliance Retention label.

This is reflected in the UI below:

EV Retention Mappings screen:

Create Retention Mapping

Source Site: QA

Source Retention Category: MIRY Retention Category

Target Category Type: Office365

Target Link: O365 (QA)

Target Retention Category:

Name	Retention Type
1 Week Delete	Personal (EAC)
1 Month Delete	Personal (EAC)
Personal (EAC) archive	Personal (EAC)
6 Month Delete	Personal (EAC)
1 Year Delete	Personal (EAC)
5 Year Delete	Personal (EAC)
MIRY LABEL 3 (Forever)	Retention label (O365 Compliance)
MIRY 6	Retention label (O365 Compliance)
MIRY LABEL 2 (Complete Desc)	Retention label (O365 Compliance)
MIRY LABEL 1 (delete after 1 day)	Retention label (O365 Compliance)

Page 2 of 3 (22 items)

Source Site	Source Retention Category	Source Type	Target Name	Target Type	Last Check Date	Details
QA	MIRY Retention Category	999 999 Years	N/A	O365 (QA)	MIRY LABEL 3 (Forever)	
QA	ASQATA Retention Category	999 999 Years	N/A	O365 (QA)	MIRY LABEL 4 (Trigger a	

Unique combination of EV retention category and O365 retention label is validated

Create Retention Mapping

Source Site: QA

Source Retention Category: MIRY Retention Category

Target Category Type: Office365

Target Link: O365 (QA)

Target Retention Category: MIRY LABEL 5 (test month)

Already existing mapping!

Already existing mapping!

Add

Mapping health page has been updated to handle with this missing retention mapping for Office 365.

Mapping Id	Source Name	Source Type	Target Name	Target Type	Last Check Date	Details
Issue: Missing Retention Category Mappings (Count: 2)						
2	miry-qaex10ev10-Stage2RC		miry-0365-05		06-05-2020 18:38:46	More Info...
3	miry-qaex07-1001-005		miry-0365-03		06-05-2020 18:38:46	More Info...

Archive Shuttle collects **ALL** retention labels, published or unpublished.

- Published labels are visible via **Outlook > Assign Policy**.
- Unpublished are hidden for users via Outlook.

Restore Enterprise Vault items associated foreign shortcuts within a user's mailbox to Office 365

You to restore original items from foreign shortcuts in a users' mailbox, and remove that shortcut during stage 2.

A foreign shortcut is one which is not from this users archive, but it's from another archive which the user has access to. With Enterprise Vault, it will retrieve that item when needed by the user.

Archive Shuttle also has the ability to retrieve these items, and migrate the full item to the users chosen target, and remove the shortcut during stage 2.

To enable this functionality, click the 'Convert EV foreign shortcuts to regular items' under the Shortcut Process Configuration options.

i | NOTES:

- Archive Shuttle has new workflows and stage 2 commands, eg “Enterprise Vault to Office 365 (with restore foreign shortcuts)”.
- Archive Shuttle does not do retention tagging on the restored items, because information about retention is not known for these items.
- The shortcut process module performs the retrieval action from Enterprise Vault using the Enterprise Vault API.
- It is not required to have an item database associated with the vault store where the item resides.
- Emails, appointments, tasks and sticky notes are supported. Contacts are not supported.
- Ensure to assign a staging area path to the import link.

Enterprise Vault to Proofpoint

Introduction

This section explains what the requirements are for an Enterprise Vault (EV) to Proofpoint migration, and then covers the basic steps that should be followed to start such a migration. When doing a migration to Proofpoint, Archive Shuttle make use of PST files as intermediate step in getting the data into Proofpoint.

Requirements


The following is a list of requirements.

Requirement	Description
Modules installed on source	The EV-type modules should be installed on each appropriate source Enterprise Vault server involved in the migration. The EV import module is not required on these machines.
Modules installed on appropriate bridgehead	The Native Format Import module, and if needed, the post processing module should be installed on a bridgehead server.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
PST Output Area created	One or more PST Output areas should be setup, of an appropriate size and should be excluded from virus scanners. It is this location where 'finalized' PSTs will be place.
PST Temporary Path created	A PST Temporary Path should be setup, of an appropriate size and should be excluded from virus scanners. It is this location where "temporary" PSTs will be placed.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.

Requirement	Description
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Add EV Environment	EV Environment	Add the Enterprise Vault source environment.
Create PST Link	Links	Create a PST link, and ensure that the Target modules are associated with it. Configure the PST Output Path and the PST Temporary Path.
Set Watermarks for Output Path & Temporary Path	Links	With the PST link selected, go to the PST tab, and set the watermarks for the PST Output location.
Update System Configuration	Native Format Import, System Configuration	

Step	Screen	Description
		Ensure that the option 'Stamps a header to imported messages for Proofpoint to identify message source' is selected in the System Configuration for the Native Format Import module.
Add Source Links	Links	Configure all appropriate source links. Select the links, and then click on 'Map Modules', create the item database(s) and enable archive gathering. Ensure there is no import module specified on the links (this will only confuse things later).
Add Staging Area	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration can begin. It is often easiest to use a filter on the bulk mapping screen to show only DAM as a source.

Determining Enterprise Vault archive owner(s)

Archive Shuttle must determine the owner of a mailbox archive. This section explains how this takes place..

Archive Shuttle takes the owner information of a mailbox archive from the Auto Security Descriptor field within the Enterprise Vault Directory database. The first entry for an archive will be the owner. Take the following three examples:

Simple 1

- Bill usage to: somedomain\simple1
- Permissions tab: somedomain\simple with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple1'.

Simple 2

- Bill usage to: somedomain\simple2
- Manually grant an additional user read, and delete permissions on the archive
- Permissions tab: somedomain\simple2 with inherited read, write and delete permissions.
somedomain\someotheruser with manually set read, and delete permissions

Archive Shuttle will see the owner of this archive as 'simple2'.

Simple 3

- Bill usage to: somedomain\simple3
- Grant full mailbox access via Exchange Management Console or Exchange Management Shell
- Permissions tab: somedomain\simple3 with inherited read, write and delete permissions.
somedomain\someotheruser also with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple3'.

i **NOTE:** The order which the archives are listed in the permissions tab in Enterprise Vault does not reflect the order that the accounts are described in the Auto Security Descriptor field. The permissions tab shows the archive permissions in alphabetical order.

More detail

1. Archive Shuttle resolves the 'Owner' of an Archive as follows:

In the EnterpriseVaultDirectory database, the following SQL is used:

```
select ADMbxDN, LegacyMbxDN from ExchangeMailboxEntry where DefaultVaultId=@VAULTID
```

2. Archive Shuttle then does a lookup for the ADMbxDN in Active Directory to get the SID:

LDAP Query: (ADMBxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

3. Archive Shuttle then does a lookup for the LegacyMbxDN in Active Directory to get the SID:

LDAP Query: (LegacyMbxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

4. We then compare the SID's we got from ADMbxDN and from LegacyMbxDN, and if both of them match, we have an Owner. If they do not match, they will be marked as 'Ownerless' and move to the next option.

5. If we do not find an entry in ExchangeMailboxEntry, we resolve using the BillingOwner or the AutoSecurityDesc:

select SID from Root inner join trustee on OwningTrusteeIdentity = TrusteeIdentity where VaultEntryId=@ARCHIVEID

i | **NOTE:** Regarding OwnerUserSid versus SID fields in the EVArchive Table on in the Archive Shuttle Directory Database:

- SID is always filled out with the SID which is retrieved from the above steps
 - OwnerUserSid has a foreign key constraint on the [User] Table. So Archive Shuttle only fill this out if the SID obtained from EV is actually contained in the Archive Shuttle Directory Database (i.e. has been synced from AD)
5. The UserSidHistory table is stored Sid history for each user when they were migrated from AD1 to AD2. This SidHistory attribute is part of Active Directory, and Archive Shuttle collects this data from Active Directory directly.

General considerations

The following general considerations should be taken into account for this type of migration:

- When migrating from a pre EV10.0.3 system extraction of data may be slower than expected if Enterprise Vault Collections have been used. Access to the data inside the CAB file is single threaded. This issue does not exist in Enterprise Vault 10.0.3 and later.
- When choosing a PST filename policy for Mailbox Archives ensure that SAMAccountName is used in the definition.
- Set the PST Rollover Threshold to no more than 4 Gb.
- When migrating Journal Archives, it is recommended to use data-based filters to allow greater control over the mappings.

Enterprise Vault Discovery Accelerator to Office 365

This section explains what the requirements are for an Enterprise Vault Discovery Accelerator (EV DA) to Office 365 migration, and then covers the basic steps that should be followed to start such a migration.

Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The EV-type modules should be installed on each appropriate source Enterprise Vault server involved in the migration. The EV import module is not required on these machines.
Modules installed on appropriate bridgehead to reach target	Office 365 and Post Processing modules are the two which are primarily needed to work with this type of migration.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
OAuth configuration	For modern authentication (OAuth) only: If OAuth is being used for authentication, click here for more information.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.


Step	Screen	Description
Check Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Add Enterprise Vault Discovery Accelerator Environment	EV DA Environment	Add the Enterprise Vault Discovery Accelerator source environment.
Add Source Links	Links	Ensure that the default staging area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Add Staging Area	Links	Select the newly created link, and map a Shortcut Process Module to it.
Configure Folder-less Item Handling	System Configuration > Folder > Folder name for folder-less items	All migrations allow archiving of items in a folder-less area called the Top Of Information Store. On the System Configuration page, enter a folder name to be used for placing the items in the target archive.
Set Office 365 Credentials		OAuth authentication Screen: Office 365 module Run the Archive Shuttle Office 365 Module Credential Editor and add OAuth credentials. Then test the connection. Then, from the Office 365 Module, select the Use modern authentication (OAuth) option.

Step	Screen	Description
Update Shortcut Processing Module Configuration	System Configuration > Shortcut Process Module	Update the settings so that 'Use EWS for Processing' is enabled.
Consider disabling reminders for appointments in the past	System Configuration > Office 365 Module	Consider enabling the option to 'Disable reminders for appointments in the past'.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration started.

Enterprise Vault to PST

This section explains what the requirements are for an Enterprise Vault (EV) to PST migration, and then covers the basic steps that should be followed to start such a migration.

Stage 1 - Synch




As soon as a user is enabled for migration, the synch process-flow shown above starts.

Archive Shuttle Core sends a command to collect all needed metadata information such as the number of items, size, and Enterprise Vault transaction ID's for the archive of the user. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.

The Enterprise Vault Export and Native Format Import module then begin to work. This runs continuously in the background until the administrator initiates a "Switch" (Stage 2).

i | **NOTE:** PST files will be created by default of about 2 Gb in size. A lower size can be set in the UI if required.

Stage 2 - Switch



After the Stage 2 ('Switch') phase to the target environment is initiated, the configured workflow starts, and by default determines the gap between the source archive and the last imported item. Archive Shuttle then synchronizes the difference one last time to the target mailbox.

Archive Shuttle performs a 'Zap' of the mailbox to remove Enterprise Vault settings before closing the PST files and moving them to the PST Output Area. Finally, shortcuts are cleaned up in the source mailbox by deleting them, and any pending-archive items are changed back to normal items.

Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The EV-type modules should be installed on each appropriate source Enterprise Vault server involved in the migration. The EV import module is not required on these machines.
Modules installed on appropriate bridgehead	The Native Format Import module, and if needed, the post processing module should be installed on a bridgehead server.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
PST Output Area created	One or more PST Output areas should be setup, of an appropriate size and should be excluded from virus scanners. It is this location where ‘finalized’ PSTs will be place.
PST Temporary Path created	A PST Temporary Path should be setup, of an appropriate size and should be excluded from virus scanners. It is this location where ‘temporary’ PSTs will be placed.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important

Step	Screen	Description
		to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Add EV Environment	EV Environment	Add the Enterprise Vault source environment.
Create PST Link	Links	Create a PST link, and ensure that the Target modules are associated with it. Configure the PST Output Path, and PST Temporary Path.
Set Watermarks for Output Path & Temporary Path	Links	With the PST link selected, go to the PST tab, and set the watermarks for the PST Output location.
Add Source Links	Links	Configure all appropriate source links. Select the links, and then click on 'Map Modules', create the item database(s) and enable archive gathering. Ensure there is no import module specified on the links (this will only confuse things later).
Add Staging Area	Links	<p>Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.</p> <p>Note: Additional space may be required on the staging area as temporary PST files are created on there, before being placed on the PST Output Path. It is recommended to use double the normal size of staging area (or use a PST Temporary Path).</p>
Configuring High and Low Watermarks	Links	It is recommended to set the high and low watermarks to around 25% of the available disk space. For example if there is 200 Gb of disk allocated for the staging area, and one link, then the low watermark should be 50 Gb and the high watermark should be 150 Gb. Having this additional free space allows for some room for

Step	Screen	Description
		operating with the temporary PST files that reside on the staging area.
Configure Folder-Less Item Handling	System Configuration	All migrations allow archiving of items in a folder-less area called the Top Of Information Store. On the System Configuration screen, enter a folder name to be used for placing the items into the target archive.
Configure File Name Policy	File Name Policy	Review and configure any required PST file name policy. This is the naming standard which will be used when placing the PST file on to the PST Output Path.
Change PST Roll-Over Threshold	Links	The default PST Rollover threshold is 2 Gb. If it is required to adjust this, it can be done on the Links page when the PST link has been selected.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration can begin. It is often easiest to use a filter on the bulk mapping screen to show only DAM as a source.

Determining Enterprise Vault archive owner(s)

Archive Shuttle must determine the owner of a mailbox archive. This section explains how this takes place..

Archive Shuttle takes the owner information of a mailbox archive from the Auto Security Descriptor field within the Enterprise Vault Directory database. The first entry for an archive will be the owner. Take the following three examples:

Simple 1

- Bill usage to: somedomain\simple1
- Permissions tab: somedomain\simple with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple1'.

Simple 2

- Bill usage to: somedomain\simple2
- Manually grant an additional user read, and delete permissions on the archive
- Permissions tab: somedomain\simple2 with inherited read, write and delete permissions.
somedomain\someotheruser with manually set read, and delete permissions

Archive Shuttle will see the owner of this archive as 'simple2'.

Simple 3

- Bill usage to: somedomain\simple3
- Grant full mailbox access via Exchange Management Console or Exchange Management Shell
- Permissions tab: somedomain\simple3 with inherited read, write and delete permissions.
somedomain\someotheruser also with inherited read, write and delete permissions

Archive Shuttle will see the owner of this archive as 'simple3'.

i | NOTE: The order which the archives are listed in the permissions tab in Enterprise Vault does not reflect the order that the accounts are described in the Auto Security Descriptor field. The permissions tab shows the archive permissions in alphabetical order.

More detail

1. Archive Shuttle resolves the 'Owner' of an Archive as follows:

In the EnterpriseVaultDirectory database, the following SQL is used:

```
select ADMbxDN, LegacyMbxDN from ExchangeMailboxEntry where DefaultVaultId=@VAULTID
```

2. Archive Shuttle then does a lookup for the ADMbxDN in Active Directory to get the SID:

LDAP Query: (ADMBxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

3. Archive Shuttle then does a lookup for the LegacyMbxDN in Active Directory to get the SID:

LDAP Query: (LegacyMbxDN=)

If Archive Shuttle does not find the AD object, the archive will be marked as Ownerless and move to the next option.

4. We then compare the SID's we got from ADMbxDN and from LegacyMbxDN, and if both of them match, we have an Owner. If they do not match, they will be marked as 'Ownerless' and move to the next option.
5. If we do not find an entry in ExchangeMailboxEntry, we resolve using the BillingOwner or the AutoSecurityDesc:

```
select SID from Root inner join trustee on OwningTrusteeIdentity = TrusteeIdentity where VaultEntryId=@ARCHIVEID
```

i | NOTE: Regarding OwnerUserSid versus SID fields in the EVArchive Table on in the Archive Shuttle Directory Database:

- SID is always filled out with the SID which is retrieved from the above steps
 - OwnerUserSid has a foreign key constraint on the [User] Table. So Archive Shuttle only fill this out if the SID obtained from EV is actually contained in the Archive Shuttle Directory Database (i.e. has been synced from AD)
5. The UserSidHistory table is stored Sid history for each user when they were migrated from AD1 to AD2. This SidHistory attribute is part of Active Directory, and Archive Shuttle collects this data from Active Directory directly.

Migrating an archive from Enterprise Vault to PST

This section outlines the steps that are necessary in order to complete the migration of a single test archive within Enterprise Vault to a PST file.

Prerequisites

It's essential that the following steps are complete before the migration begins:

- All Modules are enabled.
- Appropriate Active Directory Domains are enabled for scanning.
- An Enterprise Vault Environment has been added for migration.
- A Link Database has been created for the source archives.
- Module Mappings is configured for the source archives.

In addition, a PST Link needs to be created.

Creating a new PST naming policy

Finalized PST files are named according to a chosen PST naming policy. Policies are managed on the **Configuration > Naming Policies** page.

The Naming Policies page shows existing policies. It's used to view, edit, and delete existing policies, and create new policies.

Token	Description
username	Username of the owning user (sAMAccountName)
firstname	First name of the owning user
lastname	Last name of the owning user
fullname	Full name of the owning user
email	E-mail address of the owning user
upn	User principal name of the owning user
pstid	ID of the PST file; continuous integer over all PST files

Token	Description
pstnumber	Number of PST files; continuous integer per user/mapping
archivename	Name of the archive
archiveID	The Enterprise Vault archive ID associated with the archive.

Map containers - for the test archive

To migrate a test archive from Enterprise Vault to a PST file, the source and target containers must be mapped. To do this, follow these steps:

1. Go to **Manage & Operations > Bulk Mapping**.
2. Type the beginning of the archive name in the **Container Name** filter. Then, click the **Apply** button located at the far right of the filter row.

i **NOTE:** If the test archive doesn't display, go back to the **Configuration > EV Environment** page and click Sync all AD Users. Then select the vault store where the source archive is located and click Run Now.
3. Select the checkbox next to the test archive, and click **Add Mappings**. A brief wizard runs, gathering information relating to the mapping for the selected archive (multiple archives can also be selected).
4. On the **Choose Target Link** page, select the name of the PST Link from the drop-down list, and click **[Next]**.
5. On the **Workflow Policy** page, select **Native Import (PST) (without archive deletion)**, and then click **[Next]**. You don't need to select the ownerless workflow policy in this situation; it's optional.
6. On the **Filter Policy** page, select **Default (No Filter)**, and then click **[Next]**.
7. On the **PST Filename Policy** page, select the name of a policy from the drop-down list, and click **[Next]**.
8. On the Container Mapping Settings page, select Enabled for both Item Gathering Status and Migration Status, and click Next.
9. If desired, set the mapping's priority, and then click **[Next]**.
10. Review the summary page, and then click **[Confirm]**.

Review Stage 1 Status

A few minutes after the mapping is created, modules begin the actions defined in the mapping. Check this stage's progress by following these steps:

1. Go to **Manage & Operations > Stage 1 (Sync Data)**.
2. Type the beginning of the archive name in the **Container Name** filter. Click the button located at the far right of the filter row to apply the filter.
3. Once the source archive displays, click **Refresh** to see the progress of the export of the archive and the ingestion of the data into a temporary PST file. Continue to click **Refresh** until both export and import are complete.

Validate exported data

If there is a large amount of data to export and import, then you can see the progress bars for the Stage 1 Status move by clicking **Refresh** from the navigation bar on the **Stage 1 (Sync Data)** page.

In addition, you can use Windows Explorer to browse the export/import storage area data on the disk.

In the case of PST Migration, in the folder on disk containing the exported data, there will be a further sub-folder containing a temporary PST file. You can copy it to a new location and attach to Outlook in order to review its contents.

Enable Stage 2

Before enabling Stage 2, the switch, for the test archive, check the **Stage 1 (Sync Data)** page for issues like failed item-export or item-import.

Then, to enable Stage 2 for the test archive, perform the following steps:

1. Go to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and then click **Enable Stage 2** in the navigation bar.
3. When the Stage 1 (Sync Data) page is refreshed, there should be a green and white check mark in the Stage 2 Enabled column.

Stage 2 is the switch-over to the target environment. A final synchronization is performed of archived items from the source environment to the temporary PST file, before several additional migration tasks are performed.

Review Stage 2 Status

After a few minutes the progress of the test archive migration will be visible in the Stage 2 Status page, as follows:

1. Go to the **Stage 2 (Switch User)** page.

2. Type the beginning of the archive name in the **Name** filter. Click **Apply** at the right of the filter row.

i | **NOTE:** If the archive is not displayed, wait 1-2 minutes and then click **Refresh**.

3. Once the source archive displays, click **Refresh** from time to time to see the progress of the final stages of the archive migration. By default, since the **Native Import (PST) (without archive deletion)** workflow policy was selected, the following steps will be performed:
 1. Disable source mailbox from Enterprise Vault archiving.
 2. Rename the source archive.
 3. Collect any remaining items for migration.
 4. Import the remaining items in to the temporary PST file.
 5. Finalize the data in the temporary PST file.
 6. Move the PST file to the PST Output Path, as defined on the PST Link, using the naming standard defined in the filename policy.
 7. Delete any Enterprise Vault shortcuts in the mailbox.
 8. Change any pending-archive items in the source mailbox back to normal items.
 9. Zap the permissions on the source archive.

Verify data has been migrated


Once all of the Stage 2 operations have been completed, the Stage 2 (Sync Data) for the test archive will have a green and white check mark in the Finished column.

In addition, the PST file can be opened with Outlook from the PST Output Path.

PST as a source

This section explains what the requirements are for migrating from PST files to a target environment, and then covers the basic steps that should be followed to start such a migration.

Stage 1 - Synch



The first task which has to be performed is the discovery of the PSTs, followed by the assignment of those PSTs to target users.


Then when a PST is enabled for migration, the synch process-flow, shown above, starts.

Archive Shuttle Core sends a command to collect all needed metadata information from each PST file. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.


The PST Export and chosen target environment import module then start to process the items from the PSTs. This runs continuously in the background until the administrator initiates Stage 2 ('Switch') for the archive.

Stage 2 - Switch

PST To Exchange or Office 365



PST To Enterprise Vault



After the administrator has initiated a “Switch” (Stage 2) to the target, Archive Shuttle then synchronizes the difference one last time to the target environment. Depending on the target environment there then may be some additional steps (see the diagrams above).

Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The PST Export module should be installed on an appropriate server.
Modules installed on target	The appropriate modules should be installed on each appropriate target server involved in the migration.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
PSTs made available	Ensure that the PST files that are the source for the migration are available in a high speed manner. In other words if they’re on network locations ensure that they are close to the server which is running the PST Export module. If they’re on multiple USB drives ensure as many of those are accessible, or at least the first is accessible ready to start the migration process.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (indicating that Archive Shuttle Core has not

Step	Screen	Description
		had contact with the module) and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Create PST Link	Links	Create a PST link, and ensure that 'Source' modules are associated with it. The PST Output Path in this case is not used, so any UNC Path can be given.
Create Link Database	Links	Create a link database for the newly created PST Link.
Add PST Source Folders	Links	Add one or more source folder locations where the PST files that are to be migrated can be found. It is not necessary to run the 'scan' command. Newly added paths are automatically scanned.
Add Target Links	Links	Ensure that the link for the migration target is added, and modules associated with it.
Add Staging Area	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Configure Folder-Less Item Handling	System Configuration	Many versions of Enterprise Vault allow archiving of items in a folder-less area called the Top Of Information Store. On the System Configuration screen, enter a folder name to be used for placing the items into the target archive.
Associate owner to PST	Bulk Mapping	<p>Note: It is recommended to associate an owner to each discovered PST. If this is not actioned, then the Bulk Mapping page can not be used to migrate the PST, instead the Manual Mapping screen must be used.</p> <p>It is easiest to use a filter on the bulk mapping screen to show only PST as a source archive type. If the PST file names match a users samAccountName, then the ownership of that PST will be shown, otherwise many PST files may show as ownerless.</p>

Step	Screen	Description
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration can begin.

General considerations

The following general considerations should be taken into account for this type of migration:

- Scanning PST files themselves for items can take considerable time .Exports will not begin for that PST until the scanning is complete.

Migrating from PST to a target environment

This section outlines the steps to migrate a single test PST to any supported target environment.

i **NOTE:** The steps are specific to migrating to an Exchange 2010/2013 Personal Archive, but the steps are largely the same when migrating from source PST files to any supported target environment.

Prerequisites

It is essential that the following steps have been completed prior to beginning the migration:

- All Modules are enabled
- Appropriate Active Directory Domains have been enabled for scanning
- The Exchange target links should be configured with a specific staging area, and with modules associated to the link.

Create a PST link

To use PST as a source for migration, a PST link needs to be created and configured:

1. Go to **Configuration > Links**.
2. On the **PST** tab click **Create Link**.
3. Give the link a meaningful name (for example, PST Source).
4. Specify a **UNC path**, in this case it can be any path, because it will not be used in the case of using PST as a source.
5. With the link still selected, click on **Map Modules** on the **Link** tab.
6. A Native Export Module should be selected underneath the Source option.
7. With the link still selected, click **Create Database** to create a link database.

Scan PST location(s)

The PST link needs to be configured with one or more source UNC paths to scan for PSTs, so that they can be migrated from. This can be achieved as follows:

1. Go to **Configuration > Links**.
2. Select the PST link that was created in the previous step.
3. On the **PST tab** select **PST Source**.
4. Click **New**, which is located in the center of the dialog, and enter a UNC path to scan for PSTs.
5. If required, multiple paths can be entered.

i | **NOTE:** Clicking **Scan** after adding a new UNC path for the module is not necessary, it will scan new paths automatically.

Auto-assign user/owner to PST files

PST files are not assigned to users automatically after the scanning runs on the PST Export UNC path.

To automatically assign PST files to users when using the PST Export module, you must use the user's primary SMTP address or email address or SAMAccount Name in the file name (in the format below), and then run the Scan on the PST Source link from the Links screen.

Here are the supported PST file name formats for auto-assignment:

- PrimarySMTPAddress_id.pst
- SAMAccountName_id.pst

For example:

- firstname.lastname@somedomain.com_1.pst
- firstinitiallastname_1.pst

i | **NOTE:** In the case when a SAM Account Name exists in multiple domains that are enabled in Archive Shuttle, PST is not assigned to an owner and you will see entry below in the PST export module log file:

```
WARN|GetUserBySAMAccountName(PstExportScanDirectoryResultProcessor.cs:97)|More users with SAMAccountName ["samaccountname"], skipping auto assignment of PST owner
```

Map containers - for the test archive

To migrate a test archive from one vault store to another, the source and target containers must be mapped:

1. Go to the **Bulk Mapping** page.

2. Change the filter for **Type**, so it is set to **PST**. This will show all the discovered PST files.

i NOTE: The PSTs will show as ownerless, and will display the file path and file name. The number of items in the PST will show as zero at this stage

3. Locate the PST you wish to use to test the migration, and select it.
4. Click on **Assign Archive User** and pick an Active Directory user whose mailbox will be the target of the migration.
5. With the container still selected, click **Add Mappings** from the navigation bar.

i NOTE: A short wizard will start, which gathers information relating to the mapping for the select archive. (Multiple archives can also be selected)

6. On the **Target Container Type** screen, select **Exchange**, and click **[Next]**.
7. On the **Choose target mailbox type** screen, select **Secondary (Archive) mailbox**. In the second part of the dialog select **Map to primary mailbox** if the secondary mailbox does not exist, and then click **[Next]**.
8. On the **Target User** screen, select **Same User** and click **[Next]**.
9. On the **Workflow Policy** screen, select the **PST to Exchange/Office 365(without archive deletion)** option, and click **[Next]**.
10. On the **Filter Policy** screen, select the **Default (No Filter)** option, and click **[Next]**.
11. On the **Container Mapping Settings** screen, select **Enabled** for both **Migration Status** and **Item Gathering Status**, and click **[Next]**.
12. Review the summary screen before clicking **[Confirm]**.

i NOTE: If the PST is not assigned an owner, it is still possible to use the Manual Mapping screen to perform a migration from PST to a target.

Review Stage 1 Status

A few minutes after the mapping has been created, the program will instruct the appropriate modules to start the actions defined in the mapping. The progress for this stage of the migration can be reviewed by doing the following:

1. Go to the **Stage 1 (Sync Data)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the **Apply** button at the right of the filter row.
3. Once the source archive is displayed, click **Refresh** from time to time to show the progress of the export of the archive and the ingesting of the data into the target Archive. Continue to click **Refresh** until both export and import are 100%.

If further data is added to the source archive, it will be synchronized every 24 hours to the target archive via the connection made with the mapping performed in the previous section. In

addition, part of the Stage 2 Workflow is to perform a 'final synch'. Therefore, it's not necessary to ensure that export and import has reached 100% before moving on to the next steps.

Validate exported data

If there is a large amount of data to export and import, you can view the progress bars for the Stage 1 (Sync Data) move by clicking **Refresh** from the navigation bar on the **Stage 1 (Sync Data)** page of the user interface.

In addition, you can use Windows Explorer to browse the export/import storage area data on the disk.

Also at this time it will be possible to login to Outlook or Outlook Web Access as the test user. The migrated data will be present in the Personal Archive, and can be freely opened and manipulated.

i | **NOTE:** If the test user previously used the Personal Archive feature, it might be difficult to locate the migrated data. Because of this, Archive Shuttle recommends that the test user has an empty personal archive before migration.

Enable Stage 2 Status

Before enabling Stage 2, the switch, for the test archive, the **Stage 1 (Sync Data)** page should be checked for issues such as failed item-export or item-import.

To enable Stage 2 for the test archive, perform the following steps:

1. Navigate to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and then click **Enable Stage 2** in the navigation bar.
3. When the **Stage 1 (Sync Data)** page is refreshed, there should be a green and white check mark in the Stage 2 Enabled column.

Stage 2 is the switchover to the target environment. A final synchronization is performed of archived items from the source environment to the temporary PST file, before several additional migration tasks are performed.

Review Stage 2 Status

After a few minutes, the progress of the test archive migration will be visible in the Stage 2 Status page, as follows:

1. Navigate to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the **Apply** button at the right of the filter row.

i | **NOTE:** If the archive is not displayed, wait one to two minutes and click Refresh in the navigation bar

Once the source archive is displayed, you can click **Refresh** from time to time to show the progress of the final stages of the archive migration. By default, since the workflow policy that was selected was **PST to Exchange/Office 365 (without archive deletion)**, the following steps will be performed:

- a) Export any remaining items from the PST file
- b) Import the remaining items in to the target

Verify data has been migrated


Once all of the Stage 2 operations have been completed the Stage 2 (Sync Data) for the test archive will have a green and white check mark in the **Finished** column.

In addition, if Outlook or Outlook Web Access is used to access the test mailbox then all of the archived items will be in the Personal Archive, and no Enterprise Vault shortcuts will remain in the source mailbox.

EAS as a source

This section explains what the requirements are for migrating away from EAS, and then covers the basic steps that should be followed to start such a migration.

Stage 1 - Synch

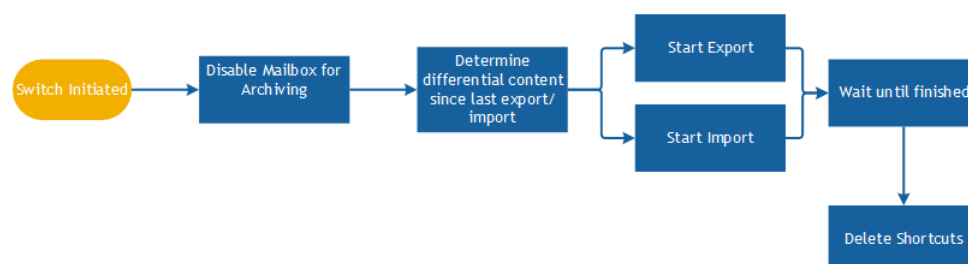


The first task which has to be performed is the discovery of archives in EAS. One or more of these archives are mapped and enabled for migration, Archive Shuttle Core sends a command to collect all needed metadata information from each archive. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.

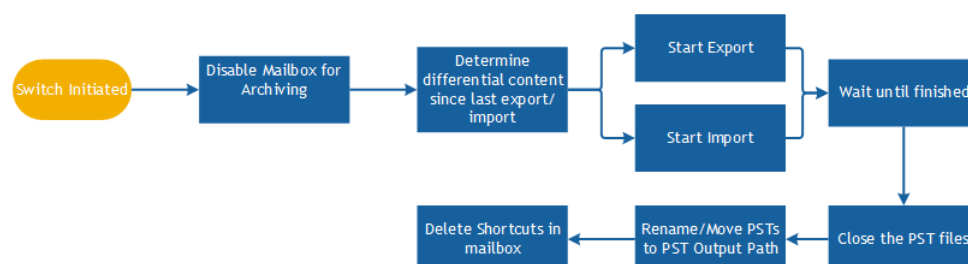
The EAS module and chosen target environment import module then start to process the items from the archives. This runs continuously in the background until the administrator initiates Stage 2 ('Switch') for the archive.

Stage 2 - Switch

EAS to Exchange or Office 365



EAS to PST



Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The EAS module should be installed on each the source server.
Modules installed on target	The appropriate modules should be installed on each appropriate target server involved in the migration.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Configure EAS Environment	EAS Environment	Add a new EAS environment, specifying the following:

Step	Screen	Description
		<ul style="list-style-type: none"> • Module to associate with this link • Display name to use for this link • EAS SQL Server Name • EAS SQL Database Name • SQL Table schema
Assign Post Processing module to EAS Link	Links	Select the newly created link, and map a Shortcut Process Module to it.
Map Modules to Target Link	Links	Ensure that the link for the migration target is added, and modules associated with it.
Add Staging Area	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration can begin. It is often easiest to use a filter on the bulk mapping screen to show only EAS as a source.

Determining EAS archive owner(s)

Archive Shuttle must determine the owner of a mailbox archive. This section explains how this takes place.

1. We use the below query to match user to EAS archive by comparing ObjDistNam result to exchangeLegacyDN column.

```
SELECT t5.ROW_NUMBER, t5.USERID AS UserId, t5.USERNAME AS UserName,
t5.OBJDISTNAME AS ObjDistName,
t5.ItemsCount AS ItemsCount, t5.ItemsSize
FROM (
SELECT ROW_NUMBER() OVER(ORDER BY t4.USERID) AS ROW_NUMBER, t4.USERID,
t4.USERNAME, t4.OBJDISTNAME, t4.ItemsCount, t4.ItemsSize
FROM(SELECT COUNT(t3.MSGID) AS ItemsCount, SUM(t3.MSGSIZE) as ItemsSize,
t0.USERID, t0.USERNAME, t0.OBJDISTNAME
FROM USERS AS t0
INNER JOIN REFER AS t1 ON(t0.USERID) = t1.USERID
INNER JOIN FOLDER AS t2 ON t1.FOLDERID = t2.FOLDERID
INNER JOIN PROFILE AS t3 ON t1.MSGID = t3.MSGID
WHERE (t1.MSGID >= 0) AND (t1.FOLDERID >= 0)
```

```
GROUP BY t0.USERID, t0.USERNAME, t0.OBJDISTNAME ) AS t4) AS t5
ORDER BY t5.ROW_NUMBER
```

2. If the user record is not found. The matching is executed by comparing address column in UserProxyAdress table againsts ObjDistName.

General considerations

The following general considerations should be taken into account for this type of migration:

- Ensure Shortcut Processing is configured to use EWS.
- The database server hosting EAS must be SQL Server.

Migrating from EAS to a target environment

This section outlines the steps to migrate a single test EAS to any supported target environment.

i | **NOTE:** The steps are specific to migrating to Enterprise Vault, but the steps are largely the same when migrating from source EAS to any supported target environment.

Prerequisites

It is essential that the following steps have been completed prior to beginning the migration:

- All Modules are enabled
- Appropriate Active Directory Domains have been enabled for scanning
- The Enterprise Vault target links should be configured with a specific staging area, and with modules associated to the link.

Add an EAS environment

To use EAS as a source for migration, an EAS Environment needs to be added, and configured as follows:

1. Go to the **EAS Environment** page (**Configuration -> EAS Environment**)
2. Click **Add**.
3. Enter the details required to connect to the EAS Environment.
4. Go to the System Configuration page for EAS, and enter the URL to be used to connect to the EAS WebServer.

Map modules to link

The newly added EAS Environment will be added as a Link to the Links page, but modules may still need to be associated with it. This can be achieved as follows:

1. Go to **Configuration > Links**.

2. Select the EAS Link which was created previously.
3. Click **Map Modules** and check and assign the modules on the link.

Map containers - for the test archive

In order to migrate a test EAS archive to an Enterprise Vault environment, the source and target containers must be mapped. This can be done as follows:

1. Navigate to the **System Configuration** page
2. On the **EV Import Module** tab, enter the **retention category ID** to use for the items which will be ingested into Enterprise Vault.
3. Navigate to the **Bulk Mapping** page.
4. Change the filter for **Type**, so it is set to **EAS**. This will show all the discovered EAS archives.
5. Locate the archive you wish to use to test the migration, and select it.
6. With the container still selected, click **Add Mappings** from the navigation bar.



NOTE: A short wizard will start, which gathers information relating to the mapping for the select archive. (Multiple archives can also be selected)

7. On the **Target Container Type** screen, select **Enterprise Vault**, and click **[Next]**.
8. On the **Choose target user type** screen, select **Same User**, and click **[Next]**.
9. On the **New or Existing Container** screen, select **Create new containers** and click **[Next]**.
10. On the **Choose Target Link** screen, choose the appropriate Enterprise Vault target link.
11. On the **Workflow Policy** screen, select the **EAS to Enterprise Vault (without archive deletion)** option, and click **[Next]**.
12. On the **Filter Policy** screen, select the **Default (No Filter)** option, and click **[Next]**.
13. On the **Container Mapping Settings** screen, select **Enabled** for both Migration Status and Item Gathering Status, and click **[Next]**.
14. Review the summary screen before clicking **[Confirm]**.

Review Stage 1 Status

A few minutes after the mapping has been created, the program will instruct the appropriate modules to start the actions defined in the mapping. The progress for this stage of the migration can be reviewed by doing the following:

1. Navigate to the **Stage 1 (Sync Data)** page.
2. Type the beginning of the archive name in the **Name** filter. Click on the **Apply** button at the right of the filter row.

3. Once the source archive is displayed, click **Refresh** from time to time to show the progress of the export of the archive and the ingesting of the data into the target Archive. Continue to click **Refresh** until both export and import are 100%.

Validate exported data

If there is a large amount of data to export and import, you can view the progress bars for the Stage 1 (Sync Data) move by clicking Refresh from the navigation bar on the **Stage 1 (Sync Data)** page of the user interface.

In addition, you can use Windows Explorer to browse the export/import storage area data on the disk.

Enable Stage 2

Before enabling Stage 2, the switch, for the test archive, the **Stage 1 (Sync Data)** page should be checked for issues such as failed item-export or item-import.

To enable Stage 2 for the test archive, perform the following steps:

1. Navigate to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and then click **Enable Stage 2** in the navigation bar.
3. When the **Stage 1 (Sync Data)** page is refreshed, there should be a green and white check mark in the Stage 2 Enabled column.

Stage 2 is the switchover to the target environment. A final synchronization is performed of archived items from the source environment to the temporary PST file, before several additional migration tasks are performed.

Review Stage 2 Status

After a few minutes, the progress of the test archive migration will be visible in the Stage 2 Status page, as follows:

1. Navigate to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Name** filter. Click on the Apply button at the right of the filter row.

i **NOTE:** If the archive is not displayed, wait one to two minutes and click **Refresh** in the navigation bar

3. Once the source archive is displayed, you can click Refresh from time to time to show the progress of the final stages of the archive migration. By default, since the workflow policy that was selected was **EST to Enterprise Vault (without archive deletion)**, the following steps will be performed:

- a) Disable archiving in EAS

- b) Import the remaining items in to the target
- c) Assign the EV archive to the end user
- d) Enable the mailbox for archiving in Enterprise Vault
- e) Rename the target archive
- f) Update all existing EAS shortcuts in the mailbox and convert them to Enterprise Vault shortcuts.


Verify data has migrated

Once all of the Stage 2 operations have been completed the Stage 2 (Sync Data) for the test archive will have a green and white check mark in the **Finished** column.

In addition, if Outlook or Outlook Web Access is used to access the test mailbox then all of the archived items will be in the Personal Archive, and no Enterprise Vault shortcuts will remain in the source mailbox.

Metalogix as a source

Stage 1 - Synch

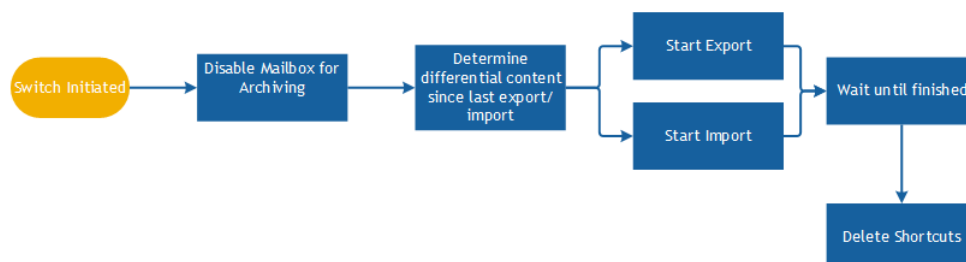


The first task which has to be performed is the discovery of archives in Metalogix. One or more of these archives are mapped and enabled for migration, Archive Shuttle Core sends a command to collect all needed metadata information from each archive. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.

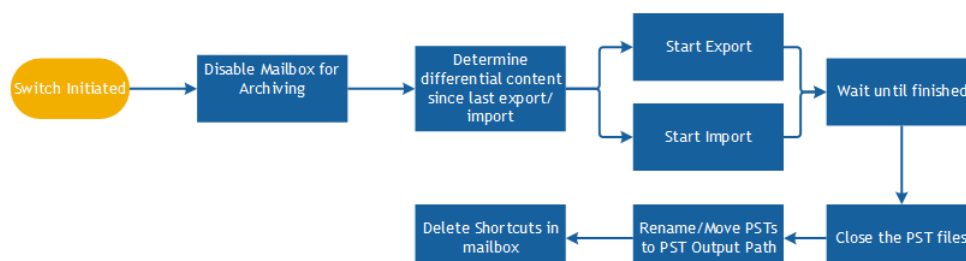
The Metalogix module and chosen target environment import module then start to process the items from the archives. This runs continuously in the background until the administrator initiates Stage 2 ('Switch') for the archive.

Stage 2 - Switch

Metalogix to Exchange or Office 365



Metalogix to PST



Migrating from Metalogix Archive Manager to a target environment

This section outlines the steps to migrate a single test Metalogix Archive Manager to any supported target environment.

i | **NOTE:** The steps are specific to migrating to Enterprise Vault, but the steps are largely the same when migrating from Metalogix to any supported target environment.

Prerequisites

It is essential that the following steps have been completed prior to beginning the migration:

- All Modules are enabled.
- Appropriate Active Directory Domains have been enabled for scanning.
- The Metalogix target links should be configured with a specific staging area, and with modules associated to the link.

Add a Metalogix environment

In order to use Metalogix as a source for migration, a Metalogix Environment needs to be added, and configured as follows:

1. Go to the **Metalogix Environment** page (**Configuration -> Metalogix Environment**).
2. Click **Add**.
3. Enter the details required to connect to the Metalogix Environment

Map modules to link

The newly added Metalogix Environment will be added as a Link to the **Links** page, but modules may still need to be associated with it. This can be achieved as follows:

1. Go to **Configuration > Links**.
2. Select the Metalogix Link which was created previously.
3. Click on **Map Modules** and check and assign the modules on the link.

i | **NOTE:** If you are using the HasShortcut migration filter, ensure to assign the Shortcut Processing Module to the Source Link.

Map containers - for the test archive

In order to migrate a test Metalogix archive to an Enterprise Vault environment, the source and target containers must be mapped. This can be done as follows:

1. Navigate to the **System Configuration** page.
2. On the **EV Import Module** tab, enter the retention category ID to use for the terms which will be ingested into Enterprise Vault.
3. Open the **Bulk Mapping** page.
4. Change the filter for **Type** so it is set to **Metalogix**. This will show all the discovered Metalogix archives.
5. Locate the archive you want to use to test the migration, and select it.
6. With the container still selected, click **Add Mappings** from the navigation bar.

i | **NOTE:** A short wizard will start, which gathers information relating to the mapping for the select archive. (Multiple archives can also be selected)

7. On the **Target Container Type** screen, select **Enterprise Vault**, and click **[Next]**.
8. On the **Choose target user type** screen, select **Same User**, and then click **[Next]**.
9. On the **New or Existing Container** screen, select **Create new containers** and click **[Next]**.
10. On the **Choose Target Link** screen, choose the appropriate Enterprise Vault target link.
11. On the **Workflow Policy** screen, select the **Metalogix to Enterprise Vault (without archive deletion)** option, and click **[Next]**.
12. On the **Filter Policy** screen, select the **Default (No Filter)** option, and click **[Next]**.
13. On the **Container Mapping Settings** screen, select **Enabled** for both **Migration Status** and **Item Gathering Status**, and click **[Next]**.
14. Review the summary screen before clicking **[Confirm]**.

Review Stage 1 Status

A few minutes after the mapping has been created, the program will instruct the appropriate modules to start the actions defined in the mapping. The progress for this stage of the migration can be reviewed by doing the following:

1. Go to the **Stage 1 (Sync Data)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the Apply button at the right of the filter row.
3. Once the source archive is displayed, click Refresh from time to time to show the progress of the export of the archive and the ingesting of the data into the target Archive. Continue to click Refresh until both export and import are 100%.

Validate exported data

If there is a large amount of data to export and import, you can view the progress bars for the Stage 1 (Sync Data) move by clicking **Refresh** from the navigation bar on the **Stage 1 (Sync Data)** page of the user interface.

In addition, you can use Windows Explorer to browse the export/import storage area data on the disk.

Enable Stage 2

Before enabling Stage 2, the switch, for the test archive, the **Stage 1 (Sync Data)** page should be checked for issues such as failed item-export or item-import.

To enable Stage 2 for the test archive, perform the following steps:

1. Navigate to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and then click **Enable Stage 2** in the navigation bar.
3. When the **Stage 1 (Sync Data)** page is refreshed, there should be a green and white check mark in the **Stage 2 Enabled** column.

Stage 2 is the switchover to the target environment. A final synchronization is performed of archived items from the source environment to the temporary PST file, before several additional migration tasks are performed.

After a few minutes, the progress of the test archive migration will be visible in the Stage 2 Status page, as follows:

1. Navigate to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the Apply button at the right of the filter row.



NOTE: If the archive is not displayed, wait one to two minutes and click **Refresh** in the navigation bar.

3. Once the source archive is displayed, you can click **Refresh** from time to time to show the progress of the final stages of the archive migration. By default, since the workflow policy that was selected was **EAST to Enterprise Vault (without archive deletion)**, the following steps will be performed:

Disable archiving in Metalogix

- a) Import the remaining items in to the target.
- b) Assign the EV archive to the end user.
- c) Rename the target archive.
- d) Update all existing Metalogix shortcuts in the mailbox and convert them to Enterprise Vault shortcuts.

Verify data has migrated

Once all of the Stage 2 operations have been completed the Stage 2 (Sync Data) for the test archive will have a green and white check mark in the **Finished** column.

In addition, if Outlook or Outlook Web Access is used to access the test mailbox then all of the archived items will be in the Personal Archive, and no Enterprise Vault shortcuts will remain in the source mailbox.


Additional information

- Archive Shuttle uses LegacyExchangeDN to pair a user with a Metalogix archive. The values from the LegacyExchangeDN column from the ExchangeMailbox table in Archive Shuttle's database must match the values from the MAILBOXDN column from the EXCHANGEPAMMAILBOXES table in the Metalogix database. If there are no results in the ExchangeMailbox table, then the Address column in the UserProxyAddress table will be used.

Dell Archive Manager as a source

This section explains what the requirements are for migrating from Dell Archive Manager (DAM)/Quest Archive Manager (QAM) to a target environment, and then covers the basic steps that should be followed to start such a migration.

Stage 1 - Synch

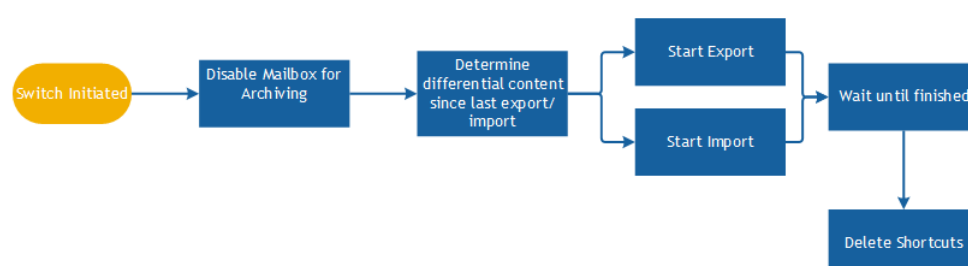


The first task which has to be performed is the discovery of archives in DAM. One or more of these archives are mapped and enabled for migration, Archive Shuttle Core sends a command to collect all needed metadata information from each archive. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.

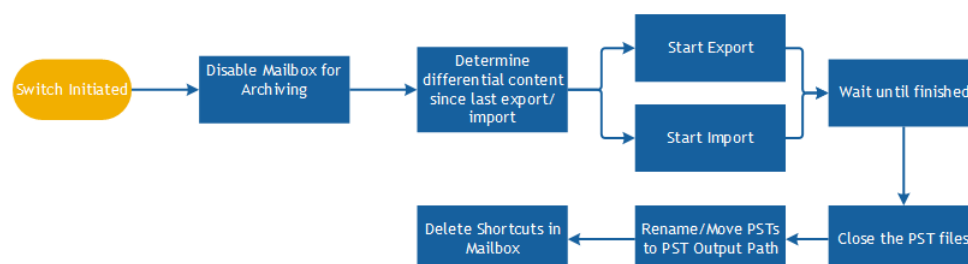
The DAM module and chosen target environment import module then start to process the items from the archives. This runs continuously in the background until the administrator initiates Stage 2 ('Switch') for the archive.

Stage 2 - Switch

DAM to Exchange or Office 365



DAM to PST



Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The DAM/QAM module should be installed on an appropriate server.
Modules installed on target	The appropriate modules should be installed on each appropriate target server involved in the migration.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and optionally a schedule for them defined. It is important to verify that none of the modules have a red background (this indicates that Archive Shuttle Core has not had contact with the module) and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization.
Configure Dell Archive Manager	Dell Archive Manager	Add a new Dell Archive Manager environment, specifying the following:

Step	Screen	Description
Environment	Environment	<ul style="list-style-type: none"> • Module to associate with this link • Display name to use for this link • SQL Server Name • SQL Database Name <p>Finally indicate that you want to 'Create Item Database', an 'Enable Archive Collection'. The 'Link SQL Server' should be set to the server where the Archive ShuttleItem databases are hosted.</p>
Assign Post Processing module to DAM Link	Links	Select the newly created link, and map a Shortcut Process Module to it.
Map Modules to Target Link	Links	Ensure that the link for the migration target is added, and modules associated with it.
Add Staging Area	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of vault stores being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Map one or more containers	Manual Mapping/Bulk Mapping	One or more containers can now be mapped and migration can begin. It is often easiest to use a filter on the bulk mapping screen to show only DAM as a source.

Determining DAM archive owner(s)

When the Dell Archive Manager (DAM) module sends results of DellArchiveManager collection to the core, the user is assigned to an archive based on LegacyMailboxDN.

First, Archive Shuttle looks into the [ExchangeMailbox] table in the Archive Shuttle Directory DB for the particular ExchangeLegacyDn. If no record is found, then it looks in the [UserProxyAddress] table.

The steps are:

Step 1

```
SELECT * FROM [User]
INNER JOIN ExchangeMailbox ON [User].UserSid = ExchangeMailbox.OwnerUserSid
WHERE ExchangeMailbox.LegacyExchangeDN = '{SourceDN}' -- SourceDN ->
ExchangeLegacyDN from source DB (DAM DB)
```

If a user is found, they're assigned to the archive. If no user is found, Step 2 below runs.

Step 2

```
SELECT * FROM [User]
INNER JOIN UserProxyAddress ON [User].UserSid = UserProxyAddress.UserSid
WHERE UserProxyAddress.[Address] = '{SourceDN}' -- SourceDN ->
ExchangeLegacyDN from source DB (DAM DB)
```

If a user is found, the user is assigned to the archive.

If the user record is not found, then the archive is considered ownerless. Archive Shuttle is not able to pair the DAM archive with a user, which exists in the Archive Shuttle Directory database.

If you need to troubleshoot why a user is not assigned to the DAM archive, it might be worth to know which columns/values are fetched when DellArchiveManager collection is completed by Archive Shuttle.

Archive Shuttle fetches the following columns/values from the [MailBox] table from Dell Archive Manager DB:

1. MailboxID
2. Name
3. MailBoxType
4. ActiveDirectoryExchangeLegacyDN

Data in the [ExchangeMailbox] table and the [UserProxyAddress] table in Archive Shuttle Directory DB is collected by the Active Directory Collector module.

First, Archive Shuttle looks into the [ExchangeMailbox] table in Archive Shuttle Directory DB for particular ExchangeLegacyDn and if no record is found, then it looks also in [UserProxyAddress] table.

If no record/ExchangeLegacyDn is found in the [ExchangeMailbox] table and the [UserProxyAddress] table Archive Shuttle is not able to pair the DAM archive with a user that exists in Archive Shuttle Directory DB.

Value 'ActiveDirectoryExchangeLegacyDN' in the [MailBox] table in the Dell Archive Manager DB should also be available in the [ExchangeMailbox] table or the [UserProxyAddress] table in Archive Shuttle Directory DB to be able to pair the DAM archive with some user that exists in Archive Shuttle Directory DB.

General considerations

The following general considerations should be taken into account for this type of migration:

- Ensure Shortcut Processing is configured to use EWS.

Migrating from DAM to a target environment

This section outlines the the steps to migrate a single test DAM to any supported target environment.

i | **NOTE:** The steps are specific to migrating to PST, but the steps are largely the same when migrating from DAM to any supported target environment.

Prerequisites

It is essential that the following steps have been completed prior to beginning the migration:

- All Modules are enabled.
- Appropriate Active Directory Domains have been enabled for scanning.
- The PST target links should be configured with a specific staging area, and with modules associated to the link.
- The PST Rollover Threshold should be configured to an appropriate value.
- The PST Filename policy should be defined.

Add a DAM environment

To use DAM as a source for migration, a DAM Environment needs to be added, and configured as follows:

1. Go to the **DAM Environment** page (**Configuration -> Dell Archive Manager Environment**).
2. Click **Add**.
3. Enter the details required to connect to the DAM Environment

Map modules to link

The newly added DAM Environment will be added as a Link to the **Links** page, but modules may still need to be associated with it. This can be achieved as follows:

1. Go to **Configuration > Links**.
2. Select the DAM Link which was created previously.
3. Click on **Map Modules** and check and assign the modules on the link.

Map containers - for the test archive

In order to migrate a test DAM archive to PST file, the source and target containers must be mapped. This can be done as follows:

1. Navigate to the **Bulk Mapping** page.

2. Change the filter for **Type** so it is set to **DAM**. This will show all the discovered DAM archives.
3. Locate the archive you wish to use to test the migration, and select it.
4. With the container still selected, click **Add Mappings** from the navigation bar.



NOTE: A short wizard will start, which gathers information relating to the mapping for the select archive. (Multiple archives can also be selected)

5. On the **Target Container Type** screen, select **PST**, and click **[Next]**.
6. On the **Choose a format** screen, select **PST**, and then click **Next**.
7. On the **Choose Target Link** screen, choose the appropriate PST link.
8. On the **Workflow Policy** screen, select the **Dell Archive Manager to PST** option, and click **[Next]**.
9. On the **Filter Policy** screen, select the **Default (No Filter)** option, and click **[Next]**.
10. On the **Container Mapping Settings** screen, select **Enabled** for both **Migration Status** and **Item Gathering Status**, and click **[Next]**.
11. Review the summary screen before clicking **[Confirm]**.

Review Stage 1 Status

A few minutes after the mapping has been created, the program will instruct the appropriate modules to start the actions defined in the mapping. The progress for this stage of the migration can be reviewed by doing the following:

1. Go to the **Stage 1 (Sync Data)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the Apply button at the right of the filter row.
3. Once the source archive is displayed, click Refresh from time to time to show the progress of the export of the archive and the ingesting of the data into the target Archive. Continue to click Refresh until both export and import are 100%.

Validate exported data

If there is a large amount of data to export and import, you can view the progress bars for the Stage 1 (Sync Data) move by clicking **Refresh** from the navigation bar on the **Stage 1 (Sync Data)** page of the user interface.

In addition, you can use Windows Explorer to browse the export/import storage area data on the disk.

Enable Stage 2

Before enabling Stage 2, the switch, for the test archive, the **Stage 1 (Sync Data)** page should be checked for issues such as failed item-export or item-import.

To enable Stage 2 for the test archive, perform the following steps:

1. Navigate to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and then click **Enable Stage 2** in the navigation bar.
3. When the **Stage 1 (Sync Data)** page is refreshed, there should be a green and white check mark in the **Stage 2 Enabled** column.

Stage 2 is the switchover to the target environment. A final synchronization is performed of archived items from the source environment to the temporary PST file, before several additional migration tasks are performed.

After a few minutes, the progress of the test archive migration will be visible in the Stage 2 Status page, as follows:

1. Navigate to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the Apply button at the right of the filter row.



NOTE: If the archive is not displayed, wait one to two minutes and click **Refresh** in the navigation bar.

3. Once the source archive is displayed, you can click **Refresh** from time to time to show the progress of the final stages of the archive migration. By default, since the workflow policy that was selected was **Dell Archive Manager to PST**, the following steps will be performed:

Disable archiving in DAM

- a) Import the remaining items in to the target.
- b) Close the PST file.
- c) Rename the PST file, according to the file name policy.
- a) Delete the shortcuts from the mailbox.

Verify data has migrated


Once all of the Stage 2 operations have been completed the Stage 2 (Sync Data) for the test archive will have a green and white check mark in the **Finished** column.

The PST file that was created can then be attached to Outlook and reviewed to ensure that the data was successfully migrated.

SourceOne as a source

This section explains the requirements for migrating from SourceOne or EmailXtender to another environment, and then covers the basic steps that should be followed to start such a migration.

Stage 1 - Synch




The first task which has to be performed is the discovery of archives in SourceOne. One or more of these archives are mapped and enabled for migration, Archive Shuttle Core sends a command to collect all needed metadata information from each archive. The results are reported back to Archive Shuttle Core to allow item level tracking and auditing.


The SourceOne module and chosen target environment import module then start to process the items from the archives. This runs continuously in the background until the administrator initiates Stage 2 ('Switch') for the archive.

Stage 2 - Switch

SourceOne to Exchange or Office 365



SourceOne to PST



Requirements

The following is a list of requirements.

Requirement	Description
Modules installed on source	The SourceOne module should be installed on an appropriate source server.
Modules installed on target	The appropriate modules should be installed on each appropriate target server involved in the migration.
AD Collector module installed	An AD Collector module should be installed so that it can collect user-level information for the migration project.
Staging Area created	One or more staging areas should be setup, of an appropriate size and should be excluded from virus scanners.
Workflows reviewed	The Stage 2 workflows should be reviewed, and if necessary, customized to meet the needs of the project.
Apply appropriate failed item threshold	It is best practice to find an appropriate failed item threshold for the project/customer and apply that at the links level. This way all mappings which are created will inherit this value.

Basic steps

The following are the basic steps to start this type of migration.

Step	Screen	Description
Enable Modules	Modules	All Archive Shuttle modules must be enabled, and a schedule for them defined if necessary. It is important to verify that none of the modules have a red background (indicating that the Archive Shuttle Core has not had contact with the module), and that the module versions are as expected.
Enable Domains	Active Directory	Select and enable one or more domains for synchronization. Note: Configuration of EmailXtender is done in the same way as SourceOne.

Step	Screen	Description
Configure SourceOne Environment	SourceOne Environment	<p>Add a new SourceOne environment, specifying the following:</p> <ul style="list-style-type: none"> • Module to associate with this link • Display name to use for this link • SourceOne Basepaths (These are the share paths for each archive folder. Multiple paths should be semicolon separated.) • SourceOne/EMX SQL server <p>A Link Database will be created for this environment. It is possible to override the SQL Server and instance, if required.</p>
Enable Organization	SourceOne Environment	The module will scan for organizations and present them in a list. The required organizations should be enabled so that the module can scan for archives.
Map Modules to Target Link	Links	Ensure that the link for the migration target is added, and modules associated with it.
Map Modules to Source Link	Links	Ensure that the source links have a Shortcut Processing module assigned to them. This module will scan mailboxes to find shortcuts. This is required before the import can begin.
Add Staging Area	Links	Ensure that the Default Staging Area is configured correctly. Depending on the migration, the free space, and the number of archive folders being used as the source for the migration, the 'high water mark' can be adjusted upwards.
Configure Folder-less folder name	System Configuration	<p>SourceOne does not maintain folder information for items. Therefore, a folder will be created to migrate items into. The name of this folder can be specified under 'Folder name for folder-less items' on the General System Settings Page.</p> <p>The module scans for shortcuts in a mailbox and any items which have a shortcut will be migrated to the current folder where that shortcut resides.</p>
Map one or more containers	Manual Mapping/Bulk	One or more containers can now be mapped and migration can begin. It is often easiest to use a filter on

Step	Screen	Description
	Mapping	the bulk mapping screen to show only SourceOne as a source.

Determining SourceOne archive owner(s)

Archive Shuttle follows steps below to assign an owner to SourceOne archive:

1. If there is still an Exchange on-premises mailbox, then Archive Shuttle matches it with EmailAddress in dbo.SourceOneArchive table and LegacyExchangeDN in dbo.ExchangeMailbox table and UserSID in dbo.User
2. If there is only an Office 365 mailbox, then Archive Shuttle tries to match owner with SourceOne archive using Address in dbo.UserProxyAddress and EmailAddress in dbo.SourceOneArchive
 - dbo.ExchangeMailbox and dbo.UserProxyAddress table are populated during the Azure Directory sync.
 - dbo.SourceOneArchive archive is populated during the archive collection.

General considerations

The following general considerations should be taken into account for this type of migration.

- If you need to migrate the primary S1 archive into a primary mailbox, and a secondary S1 archive into a personal archive (in Exchange or Office 365), then two mappings will be required.
- It is recommended to add the following index to the SQL databases relating to S1. It will improve the item collection process: `CREATE NONCLUSTERED INDEX [QT_IX_Route_EmailId] ON [dbo].[Route] ([EmailId] ASC) INCLUDE ([MessageID], [RouteMask], [PartitionId])`
- It is recommended to add the following index to the SQL databases relating to EmailXtender. It will improve the item collection process: `CREATE NONCLUSTERED INDEX [QT_IX_Route_EmailId] ON [dbo].[Route] ([EmailId] ASC) INCLUDE ([MD5HashKey], [TimeStamp])`
- Disabling source archives is currently not available as part of a Workflow Policy and will need to be done manually post migration.
- You may find that items exported from these sources are not placed in the folder it originated from. For more information on this, [click here](#).

The following filters are not supported:

- ArchivedDate
- RetentionCategory
- HasLegalHold

Migrating from SourceOne or EmailXTender to a target environment

This section outlines the steps to migrate a single test SourceOne archive to PST. The same steps can be used to migrate from EmailXTender.

i | **NOTE:** The steps are largely the same when migrating from SourceOne to any supported target environment.

Prerequisites

It is essential that the following steps have been completed prior to beginning the migration:

- All Modules are enabled.
- Appropriate Active Directory Domains have been enabled for scanning.
- The PST target links are configured with a specific staging area, and with modules associated to the link.
- The PST Rollover Threshold is configured to an appropriate value.
- The PST Filename policy is defined.
- The Shortcut Process Module must be installed, regardless if shortcuts are being used.

Add a SourceOne Environment

To use SourceOne as a source for migration, a SourceOne Environment needs to be added, and configured as follows:

1. Go to the SourceOne Environment page (**Configuration -> SourceOne Environment**)
2. Click **Add**.
3. Enter the details required to connect to the SourceOne Environment

Map Modules to Links

The newly added SourceOne Environment will be added as a Link to the Links page, but modules may still need to be associated with it. This can be achieved as follows:

1. Go to the Links page (**Configuration -> Links**).
2. Select the **SourceOne Link** which was created previously.
3. Click **Map Modules** and check and assign the modules on the link.

Map containers - for the test archive

To migrate a test SourceOne archive to PST file, the source and target containers must be mapped. This can be done as follows:

1. Navigate to the **Bulk Mapping** page.
2. Change the filter for **Type**, so it is set to **SourceOne**. This will show all the discovered SourceOne archives.
3. Locate the archive you wish to use to test the migration, and select it.
4. With the container still selected, click **Add Mappings** from the navigation bar.



NOTE: A short wizard will start, which gathers information relating to the mapping for the select archive. (Multiple archives can also be selected)

5. On the **Target Container Type** screen, select **PST**, and click **[Next]**.
6. On the **Choose a format** screen, select **PST**, and then click **[Next]**.
7. On the **Choose Target Link** screen, choose the appropriate PST link.
8. On the **Workflow Policy** screen, select the **SourceOne to PST** option, and click **[Next]**.
9. On the **Filter Policy** screen, select the **Default (No Filter)** option, and click 'Next'.
10. On the **Container Mapping Settings** screen, select **Enabled** for both **Migration Status** and **Item Gathering Status**, and click on 'Next'.
11. Review the summary screen before clicking 'Confirm'.

Review Stage 1 Status

A few minutes after the mapping has been created, the program will instruct the appropriate modules to start the actions defined in the mapping. The progress for this stage of the migration can be reviewed by doing the following:

1. Go to the **Stage 1 (Sync Data)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the **Apply** button at the right of the filter row.
3. Once the source archive is displayed, click **Refresh** from time to time to show the progress of the export of the archive and the ingesting of the data into the target Archive. Continue to click **Refresh** until both export and import are 100%.

If further data is added to the source archive, it will be synchronized every 24 hours to the target archive via the connection made with the mapping performed in the previous section. In addition, part of the Stage 2 Workflow is to perform a 'final synch'. Therefore, it's not necessary to ensure that export and import has reached 100% before moving on to the next steps.

Validate exported data

If there is a large amount of data to export and import, you can view the progress bars for the Stage 1 (Sync Data) move by clicking **Refresh** from the navigation bar on the **Stage 1 (Sync Data)** page of the user interface.

In addition, you can use Windows Explorer to browse the export/import storage area data on the disk.

Also at this time it will be possible to login to Outlook or Outlook Web Access as the test user. The migrated data will be present in the Personal Archive, and can be freely opened and manipulated.

i | **NOTE:** If the test user previously used the Personal Archive feature, it might be difficult to locate the migrated data. Because of this, Archive Shuttle recommends that the test user has an empty personal archive before migration.

Enable Stage 2 Status

Before enabling Stage 2, the switch, for the test archive, the **Stage 1 (Sync Data)** page should be checked for issues such as failed item-export or item-import.

To enable Stage 2 for the test archive, perform the following steps:

1. Navigate to the **Stage 1 (Sync Data)** page.
2. Select the checkbox next to the test archive, and then click **Enable Stage 2** in the navigation bar.
3. When the **Stage 1 (Sync Data)** page is refreshed, there should be a green and white check mark in the Stage 2 Enabled column.

Stage 2 is the switchover to the target environment. A final synchronization is performed of archived items from the source environment to the temporary PST file, before several additional migration tasks are performed.

Review Stage 2 Status

After a few minutes, the progress of the test archive migration will be visible in the Stage 2 Status page, as follows:

1. Navigate to the **Stage 2 (Switch User)** page.
2. Type the beginning of the archive name in the **Name** filter. Click the **Apply** button at the right of the filter row.

i | **NOTE:** If the archive is not displayed, wait one to two minutes and click Refresh in the navigation bar

Once the source archive is displayed, you can click **Refresh** from time to time to show the progress of the final stages of the archive migration. By default, since the workflow policy

that was selected was **PST to Exchange/Office 365 (without archive deletion)**, the following steps will be performed:

- a) Export any remaining items from the PST file
- b) Import the remaining items in to the target

Verify data has been migrated

Once all of the Stage 2 operations have been completed the Stage 2 (Sync Data) for the test archive will have a green and white check mark in the **Finished** column.

In addition, if Outlook or Outlook Web Access is used to access the test mailbox then all of the archived items will be in the Personal Archive, and no Enterprise Vault shortcuts will remain in the source mailbox.

Migrating a journal archive

Migrating a journal archive can take a considerable length of time due to the size of the archive itself in the source environment. It is possible to migrate journal archives to any of the supported targets, as described in the following sections.

Migrating a journal archive to PST or Exchange

When migrating a journal archive out of Enterprise Vault to a non-Enterprise Vault target environment, it is recommended to make use of a new feature relating to folder splitting. Many target environments have limits on the number of items in a single folder, or perform poorly when there are a large number of items in a single folder, as is the case with PST.

i | **NOTE:** For information relating to migrating journal archives to Office 365, using Journal Splitting, see Office 365 Configuration Changes.

Item	Description
Journal Archive Split to Folders – Base Folder Name	This is the name of the top-level folder to ingest items into in the target environment.
Journal Archive Split to Folders – Max Items Per Folder	This is the number of items that should be placed in each folder before a new folder is created.

If the migration is to take place to PST or Exchange, navigate to the System Configuration page in the Admin Interface, and configure the following settings.

i | **NOTE:** The splitting of items is performed after the Enterprise Vault Collector process has gathered metadata about the items to migrate.

Performing the journal mapping

Once the configuration changes have been made, then the journal archive can be mapped in the normal way except that you should choose the option in the wizard to process the mapping as journal splitting.

Migrating a journal archive to Enterprise Vault

This section outlines the steps that are required in order to complete the migration of a single journal archive within Enterprise Vault to a new Vault Store.

The first section explains how to perform the migration using a single container mapping. It is necessary to perform these steps if the migration is from a version of Enterprise Vault prior to 10.0.4. In Enterprise Vault 10.0.4, changes were made to the Enterprise Vault API meaning that a new way can be used to extract and ingest data. The second part of this topic explains how to perform a journal archive migration if Enterprise Vault 10.0.4 or later is the source.

Pre-Enterprise Vault 10.0.0.4 source

Prerequisites

It is essential that the following steps have been completed prior to beginning the migration:

- All Modules are enabled
- Appropriate Active Directory Domains are enabled for scanning
- An Enterprise Vault Environment has been added for migration
- A Link Database has been created for the source archive
- Module Mappings are configured for the source Vault Store and also the target Vault Store

The target for the journal archive migration must also:

- Be a 'user' and be setup with a mailbox
- Have a journal archive created for it, in an appropriate Vault Store
- Be listed as a Journal Target in the Vault Administration Console

i | **NOTE:** The Enterprise Vault journal task is used to ingest the data in to the target journal archive container.

Map containers - For the journal archive

The first thing to do when migrating a journal archive to a new vault store is to do a container mapping. This is slightly different to the process followed earlier in the document since Archive Shuttle must do a 'single container mapping'. To setup the mapping, follow these steps:

1. Click **Manual Mapping**.

You will now see two lists of containers.

2. Locate the source journal archive in the left hand list, and select the radio button next to the archive.
3. Locate the target journal archive in the right hand list, and select the radio button next to the archive.

i | **NOTE:** All the normal filtering and searching is available on the Single Container Mapping screen, which can help identify the correct archives when there are many journal archives displayed on screen.

4. Click **Add Mapping**.

Set policies and enable migration

Once the container mapping has been selected, it is necessary to add two more components to the link, before the migration can be enabled. To setup the migration, follow these steps:

1. Click **Existing Mappings**.

i | **NOTE:** All the normal filtering and searching is available on the **Existing Mapping** screen, which can help identify the mapping.

2. Locate the mapping that was created in the previous step.
3. Select the checkbox next to the mapping.
4. From the toolbar, select **Set Workflow Policy** and choose the **Journal Workflow**.
5. From the toolbar, select **Set Filter Policy** and choose a filter if one is required.
6. Finally, from the toolbar select **Enable for Migration**, and **Enable Item Gathering**.

i | **NOTE:** The container mapping will now look much like other mappings that may have been created. Ensure that the **Item Gathering** and Migration **Enabled** flags are displayed before moving on to monitoring the migration.

Monitoring the migration

Monitoring the Stage 1 synch process for a journal archive is the same as with other types of archive migration. A few minutes after setting the policies and enabling migration (see previous section), the container will appear on the **Stage 1 (Sync Data)** screen in the administration console.

Performing Stage 2 and extra steps

Migrating a journal archive is slightly different than migrating an ordinary Enterprise Vault archive. There are a few additional steps, as outlined below:

1. Stop the Enterprise Vault Journal Task, and set the startup of the task to **Disabled**.

i | **NOTE:** The journal mailbox that the task was targeting will now start to grow in size because Enterprise Vault is no longer processing items from it.

2. Go to the **Stage 1 (Sync Data)** screen.
3. Locate the container mapping for the journal archive, and select the check box next to it.
4. From the toolbar, select **Enable Stage 2**.
5. After a few minutes, observe the **Stage 2 (Switch User)** screen and ensure that after a few more minutes the migration is marked as complete.

The Journal Archive has now been moved to a new Vault Store, and a new archive. The final step is to change the journaling process in Enterprise Vault. Usually this is a matter of changing the journal target to be the new archive, and restarting the journal task after setting the task startup back to Automatic.

i | **NOTE:** Depending on the migration being performed, it may also be necessary to reconfigure the Exchange journaling configuration to journal to the new mailbox.

Enterprise Vault 10.0.0.4 or later source

Prerequisites

It is essential that the following steps have been completed prior to beginning the migration:

- All Modules are enabled
- Appropriate Active Directory Domains is enabled for scanning
- An Enterprise Vault Environment has been added for migration
- A Link Database has been created for the source archive
- Module Mappings are configured for the source Vault Store and also the target Vault Store

Map containers - For the Journal Archive

The first thing to do when migrating a journal archive to a new vault store is to do a container mapping. This is slightly different than the process followed earlier in the document since Archive Shuttle must do a 'single container mapping'.

To setup the mapping, follow these steps:

1. Navigate to the **Bulk Mapping** or **Manual Mapping** page.
2. Type the beginning of the archive name in the **Name** filter. Click the Apply button at the right of the filter row.
3. Select the checkbox next to the archive, and click **Add Mappings** from the navigation bar.

i | **NOTE:** A short wizard will start, which gathers information relating to the mapping for the archive.

4. The destination of the migration can then be chosen from the available list and then subsequent pages of the wizard will prompt for further information.

i | **NOTE:** It is recommended that a workflow policy is created that suits the requirements, or the **Journal Archive** workflow is chosen.

Set Policies and Enable Migration

Once the container mapping has been selected, it is necessary to add two more components to the link, before the migration can be enabled.

To setup the migration, follow these steps:

1. Click **Existing Mappings**.

i | **NOTE:** All the normal filtering and searching is available on the Existing Mapping screen, which can help identify the mapping.

2. Locate the mapping that was created in the previous step.
3. Select the checkbox next to the mapping.
4. From the toolbar, select **Set Workflow Policy** and choose the **Journal Workflow**.
5. From the toolbar, select **Set Filter Policy** and choose a filter if one is required.
6. Finally, from the toolbar select **Enable for Migration** and **Enable Item Gathering**.

i | **NOTE:** The container mapping will now look much like other mappings that may have been created. Ensure that the 'Item Gathering' and 'Migration Enabled' flags are displayed before moving on to monitoring the migration.

Monitoring the Migration

Monitoring the 'Stage 1' synch process for a journal archive is the same as with other types of archive migration. A few minutes after setting the policies and enabling migration (see previous section), the container will appear on the 'Stage 1 (Sync Data)' screen in the administration console.

Performing Stage 2 and Extra Steps

Migrating a journal archive is slightly different to migrating an ordinary Enterprise Vault archive. There are a few additional steps, as outlined below:

1. Stop the **Enterprise Vault Journal Task**, and set the startup of the task to **Disabled**.

i | **NOTE:** The journal mailbox that the task was targeting will now start to grow in size because Enterprise Vault is no longer processing items from it.

2. Go to the **Stage 1 (Sync Data)** screen.
3. Locate the container mapping for the journal archive, and select the check box next to it.
4. From the toolbar, select **Enable Stage 2**.
5. After a few minutes, observe the **Stage 2 (Switch User)** screen and ensure that after a few more minutes the migration is marked as complete.

The Journal Archive has now been moved to a new Vault Store, and a new archive. The final step is to change the journaling process in Enterprise Vault. Usually this is a matter of changing the journal target to be the new archive, and restarting the journal task after setting the task startup back to Automatic.

i | **NOTE:** Depending on the migration being performed, it may also be necessary to reconfigure the Exchange journaling configuration to journal to the new 'mailbox'.

Migrating a journal archive to Office 365

This section outlines actions you can take to perform a migration of a single journal archive to Office 365.

Journal Transformation migration

Journal Transformation migrates a journal archive to a target environment in a special way that fans-out the recipient / custodian list. If an original message was addressed to three different people, then the item is exported, and ingested into up to four target containers.

Here are some reasons you might want to use Journal Transformation:

- It causes deliver/ingest messages to be sent to user mailboxes in the same way the Exchange or Office 365 modules do in regular migrations, rather than a separate mailbox.
- If a user's data becomes corrupt or lost, journal archives can be used to recreate the user's data.
- It can be used for leavers who lack archives.
- It is easy to distinguish external email addresses.
- It can help fulfill legal/compliance/regulatory requirements.

Requirements for Journal Transformation

All journaled Enterprise Vault content originating from Exchange or an SMTP stream is a supported source data and Archive Shuttle supports exploding data into Office 365.

Sizing information depends on several project details and should be reviewed within the scope of a project.

At a minimum, it is generally recommended to leverage a x64 EV Export module, a dedicated SQL Enterprise instance, and a minimum of 1 TB space for staging the data between source and target.

Procedure for a journal archive migration

There are a number of steps that need to be followed to configure and perform a journal archive migration. They are covered below.

Perform base configuration for Journal Transformation

Set up the base configuration of Journal Transformation (see how to do it [here](#)).

- Enable Modules
- Enable Domains
- Add EV Environment
- Add Source Links
- Add Staging Area
- Set Office 365 Credentials

Configure journal archive migration link

On the Configuration > Links page, select the Journal Transformation tab and create a new Journal Transformation link.

The link can be given a name, and the target type must be selected from the available list.

Once the link has been created, an ingest module must be mapped to the link.

Office 365 environment configuration

If you will be migrating leavers, the Journal Transformation link must be configured in the Configuration > Office 365 Environment. This will include configuring the Office 365 Rolling license settings, the naming scheme for DL, Groups and E-Mail domain, and leaver mailbox settings; such as enabling Litigation Hold.

Confirm the general folder location for ingesting items

The Journal Transformation default folder location for active users, is the Recoverable Items Folder in the target container and for a leaver user, the default folder location is Journal Items in the target container

These folder locations can be overridden on an individual user mapping basis by creating a Folder Name Policy. On the **Configuration > Naming Policies** page, select the Folder Name Policy tab and create the required policy.

Journal Transformation system configuration

User Domain Whitelist

All exploded SMTP addresses are stored in Journal Transformation. However, only the user domains added to the Accepted Domains List will be enabled for further processing.

To configure the domain whitelist, go to **Configuration > System Configuration > Journal Transformation Settings** and add the allowed user domains that are required to the Accepted Domains List.

The Journal Transformation System settings at **Configuration > System Configuration > Journal Transformation Settings** will require configuration and are dependent on the migration project dependencies and requirements.

Ensure you plan and consider the relevant settings during the design phase of a project before embarking on a journal migration.

Create mapping

The source environment, associated staging areas, and module mappings then need to be configured. Once those are done, a mapping can be created for the source container.

When creating the mapping, there is an option to use Journal Transformation in the Bulk Mapping wizard, rather than normal migration.

i | **NOTE:** If you have a large source journal container, you may want to consider using date range filters to break down the source journal into smaller/multiple mappings. This should be carefully planned before creating any mappings.

Sender-Recipient Collection Starts

After mapping created and sender-recipient collection begins, the P1 information is reviewed for SMTP addresses.

These SMTP addresses can then be mapped and enabled, as described in the next section.

Journal Transformation user management

As items in the source container are analyzed for recipient information, a user list is built up that needs to be reviewed, mapped, and enabled for migration. This is done on the **Journal Transformation User Management** page in the user interface.

You can enable/add users for import manually or automatically.

To perform the process manually, you first need to add the user manually or by .csv import. Click [here](#) to view instructions on creating a .csv file.

Add an active user SMTP manually

1. Go to **Journal Transformation** and select **User Management**.
2. Click **Add**.
3. Select the required Journal Transformation container mapping/s from the filter menu.
4. In the Enable Users for Journal Transformation window, select the required active user SMTP and click Save & Close.
5. The selected user SMTP is now enabled for use in Journal Transformation and is added to the User Management screen.

i | **NOTE:** If you need to manually change the Mapped SMTP Address or assigned Folder Name Policy on the mapping, select the relevant mapping and use the appropriate toolbar button.

Add active user SMTP using a .csv file

A .csv file can be used to add users for Journal Transformation.

The csv file requires the format to be: OriginalSMTPAddress,MappedSMTPAddress

1. Go to Archive Shuttle > Configuration and select Journal Transformation User Management.

2. Click Import from CSV.
3. Select the required Journal Transformation container mapping/s from the drop-down.
4. Click Browse and select your .csv file.
5. Ensure the option, Update existing user mappings is NOT selected.
6. Click Import.
7. The selected user SMTP is now enabled for use in Journal Transformation and is added to the User Management screen.

i | **NOTE:** If you need to manually change the **Mapped SMTP Address** or assigned **Folder Name Policy** on the mapping, select the relevant mapping and use the appropriate toolbar button.

Enable the users for import manually

Once the user mappings are made either from a .csv file or manually, and a Folder Name policy is applied (if required), one or more of the user mappings can then be enabled for migration.

Manually enable for import an active user SMTP

To manually enable a user SMTP:

1. On the **User Management** page, select the required user SMTP.
2. Ensure the selected Sender-Recipient SMTP status is:
 - Has Mailbox = Yes
 - Is Ready for Import = Ready
 - Enabled for Import = No
 - Is Deleted = No
3. Click the Enable for Import button.
4. Ensure the Enabled for Import status changes to Yes (Enabled) for the selected Sender-Recipient SMTP.

Automatically Add and Enable Users

Rather than enabling users for import using .csv file mappings, you can set configuration rules and then auto-enable certain users and domains for import based on those rules.

Use wild card operators to create general rules that result in a larger number of Journal Explosion user mappings.

Create user auto creation rules

1. On the Auto-Creation Rules Management page, click Add.
2. Use the Add User Rule window to create rules. The rule editor works with these wild cards operators: “*”, “%”
3. Click Add.
4. The configured rules display on the Email Address Rules page.

The mappings that were created to existing users/target containers will progress as normal. For example, if you have mapped a source SMTP address of abc@somedomain.com to an existing mailbox with an address of abc@somedomain.com, then the import process will begin for that mapping.

Leavers

For leaver mappings, the configuration defined on the Office 365 Environment page will be used. Mailboxes will be created and licensed, as required, using a rolling license count. This operates in much the same way as the process for migrating ownerless archives to Office 365. The difference is that the Office 365 license will be added, and removed, as required, during Stage 1. The addition and removal of the license may happen multiple times.

i | **NOTE:** Bulk actions cannot be performed when you create manual mappings to external SMTP addresses.

During a journal archive migration, leavers are handled on a global level. So, for example, you can ingest the same user from different journals or mappings into the same mailbox. You can manage leavers, including adding/removing/merging them, from the Journal Explosion Leavers Management page.

Leavers Finalization

1. Click the Manage Leaver(s) Finalization button on the Journal Explosion Leavers Management page.
2. The status page shows only leavers that are enabled for import. Each row represents an SMTP address (target mailbox).
3. After you click Finalize, mailboxes progress through these statuses:
 - Finalization Started
 - User Mailbox Licensed
 - Placed On Hold
 - Finalized

i | **NOTE:** When finalization is complete, the target leaver mailbox is removed (made into an InActive Mailbox), and no other actions can be performed against the leaver.

Viewing pre-migration statistics for Journal Transformation leavers

On the Journal Transformation leavers management page, it is possible to see statistics relating to overall item counts and overall item sizes. Having this information helps operators determine what actions might be possible with the leavers in terms of migration.

Here is how it looks in the user interface:

report

Preview Smtip Address 1

Grouping from these 2 columns (only 1 level of grouping)

New added columns on JTL M page. Total row added for doing summary for particular mapping (base on grouping)

Container Mapping	Original Smtip Address	Mapped Smtip Address	Is Valid Preview Address	Group Name	Folder Name Policy	Has Owner	Is Distributor List	User Mapping Association Status	Is Existing Preview User	Is Staged/Finished	Item Count	Total Size	Apply Filter
▼ Preview Smtip Address: miryje_asqata-leaver-01@miryquadrotech.onmicrosoft.com (Count=1)													
<input type="checkbox"/>	1 asqata-leaver-01@hotmail.com	miryje_asqata-leaver-01@miryquadrotech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			1	91 KiB	
Total:											1	91 KiB	
▼ Preview Smtip Address: miryje_asqata-leaver02@miryquadrotech.onmicrosoft.com (Count=4)													
<input type="checkbox"/>	1 asqataleaver02@gmail.com	miryje_asqataleaver02@miryquadrotech.onmicrosoft.ci	✓		Leavers (Default)	✓					1	90 KiB	
<input type="checkbox"/>	1 asqataleaver02@outlook.com	miryje_asqataleaver02@miryquadrotech.onmicrosoft.ci	✓		Leavers (Default)	✓		✓			1	91 KiB	
<input type="checkbox"/>	1 asqata-leaver-01@outlook.sk	miryje_asqataleaver02@miryquadrotech.onmicrosoft.ci	✓		Leavers (Default)	✓		✓			1	91 KiB	
<input type="checkbox"/>	1 asqataleaver02@hotmail.com	miryje_asqataleaver02@miryquadrotech.onmicrosoft.ci	✓		Leavers (Default)	✓		✓			2	186 KiB	
Total:											5	458 KiB	
▼ Preview Smtip Address: miryje_asqataleaver03@miryquadrotech.onmicrosoft.com (Count=1)													
<input type="checkbox"/>	1 asqataleaver03@outlook.com	miryje_asqataleaver03@miryquadrotech.onmicrosoft.ci	✓		Leavers (Default)	✓		✓			1	92 KiB	
Total:											1	92 KiB	
▼ Preview Smtip Address: miryje_asqata-qex10JE-recycle@miryquadrotech.onmicrosoft.com (Count=1)													
<input type="checkbox"/>	1 asqata-qex10JE-recycle@qa.quadrotech-it.com	miryje_asqata-qex10JE-recycle@miryquadrotech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			6	559 KiB	
Total:											6	559 KiB	

Management and Monitoring

Tracking Progress

A Journal Transformation mapping can be viewed on the Stage 1 page in the user interface. It is also possible on this page to expand the mapping and see the progress on a per user basis. This can also be exported to supported formats (ex. .xls, .pdf.)

Additional Notes

There are some additional notes to be aware of when using Journal Transformation:

- Items may be skipped on ingest. This happens when an item to be ingested into a user mapping is older than the creation date of the user. These can be reviewed in the log files and database. If assistance is needed, contact support.
- User mappings can be imported from .csv file to aid with the handling of a large list of recipients.
- In JE to Office 365 migrations, you can create user-defined mailbox naming tokens per container mapping using the Bulk Mappings Wizard.
- When an item has been imported into all applicable user mappings it will be removed from the staging area. Later, if a new user mapping is created, requiring that same item, it will be re-exported. This will not affect the license usage.
- When using Journal Transformation with Enterprise Vault as a source, you are able to process messages without P1 headers by enabling the option to **Process messages without P1 header** on the **System Configuration > EV Export Module** page.
- Enterprise Vault 9 and below is excluded from Journal Transformation selective migrations.

Enterprise Vault to Office 365

The following is a general recommendation for large journal migration projects using the journal transformation process based on the current understanding of the project. Additional hardware may be needed depending on expected performance of the migration.

High level requirements

- Core and SQL servers.
- Bridgehead servers: 2 per EV server to run EV Export and Office 365 Import modules.
- Staging areas: 1 per EV server – sized to the largest individual journal on the server (1.2-1.5x).
- Office 365 service accounts (five per Bridgehead server).
- Available Exchange Online licenses (around 100).
- Remote access for managed migration.

Machine specifications and additional information

Core server

- Windows Server 2012R2, 2016, 2019 and 2022.
- Minimum 4 CPU (recommended 8 CPU).
- Minimum 8GB RAM (recommended 16GB RAM).
- IIS 7.5 or later is required.

SQL server

- Windows Server 2016, 2019
- SQL server 2016, 2017.
 - Enterprise Edition highly recommended
 - Use the latest versions (Windows 2016/2019, SQL 2016/2017).
 - Minimum 16 CPU, recommended 32 CPU (cores or vCPUs).
 - If using less than 32 cores, please ensure that CPU clock speed is equal to or more than 2.8Ghz.
 - Minimum 512GB RAM (recommended 1TB RAM)
- SQL storage

- Disks should be configured as per Microsoft SQL server best practices with TempDB, databases and logs separate disks.
- Assume approximately 5K per item (as shown in the EV usage report) when sizing the storage for database. This is dependent of the number of recipients per message and is only an initial estimate.
- Use fast tier 1 disk arrays, SSD is recommended.
- The SQL server must be treated as you would a production server.
 - Maintenance plans including full backups are the responsibility of the customer. [Click here](#) for more information.

Remote access for managed migrations

Customer must provide access to assigned staff with sufficient privileges to administer any system involved in the project. Specifically, the following:

- **Required**
 - Access to Archive Shuttle core server
- **Optional, but highly recommended**
 - Access to module servers (bridgehead servers)
 - Access to SQL server or MSSQL Management Studio
 - Access to staging areas
- Remote access to the on-premises server(s) for the Delivery and Operations teams is required, either through Citrix VDI, VPN, or other remote access technology.

i | **NOTE:** WebEx is not suitable nor is anything that requires the customer to initiate the session.

Bridgehead servers (two per EV server)

- Windows Server 2012 R2, 2016 or 2019.
- Minimum 4 CPU (recommended 8 CPU).
- Minimum 16GB RAM (recommended 32GB RAM).
- .net 3.5.1.
- .net 4.6.2 or later.
- Azure AD PowerShell module

Staging areas (one per bridgehead server)

- Volume sized to the largest individual journal on the EV server (1.2-1.5x).
 - Shared via SMB; can be attached to the bridgehead server, recommended on a datastore with a fast disk, or on fast NAS storage.

Available Office 365 (EXO) licenses

- The process of finalizing and managing leaver mailboxes in Office 365 requires Archive Shuttle to have licenses to work with during the migration.
- It is recommended to have around 100 licenses to ensure that Archive Shuttle can migrate to as many targets at once (parallelism), and finalize leaver mailboxes.
- These licenses would only be needed during the migration (finalization and management of leavers will be removed and returned as soon as the migration is complete).
- The number of licenses Archive Shuttle has access to heavily affects the overall performance.

Office 365 throttling limits

In order to achieve the best migration performance possible, it is recommended to contact Microsoft and request that throttling be lifted. [Click here for more information.](#)

Globally increase message size limit from 25 MB (optional)

[Click here for more information on size limits when ingesting into Office 365.](#)

Anti-virus exclusions

Generally, Archive Shuttle requires anti-virus exclusions on any machine running the Archive Shuttle modules. This would include EV servers in addition to the bridgehead server (ingest server):

- Exclude staging area volumes
- Disable Real-time scanning
- Disable In memory scanning (if available)

[Click here for more information.](#)

Journal archive folder split settings

When migrating a large source journal archive into a target, you might want to implement and use Folder Splitting functionality in Archive Shuttle, as well as archive splitting. This article explains why.

When migrating from a source archiving system, such as Enterprise Vault, to a new target archiving system, such as Enterprise Vault, there is usually no need to split the source archive in terms of the archive itself, or into new folders (and subfolders).

However, if the migration is to Exchange or Office 365, then it is often desirable for several reasons:

- Office 365 and Exchange have limits on the number of items that can be stored in individual folders.
- Performance in relation to viewing folders that contain millions of items is poor when using Outlook-type clients (MAPI based).
- Mailbox or Personal Archive size limits on the target may prevent a large journal archive being ingestable into a single target container.

For these reasons, splitting at the folder level (which is a configuration of the Enterprise Vault Collector module) is desirable and easy to do. Splitting at the archive level by querying Enterprise Vault and counting the number and size of items in the journal archive per day, per week, per two weeks, or per month. Once a suitable splitting range has been defined, filters based on a date can be implemented in Archive Shuttle, and the source archive can be mapped multiple times to different target archives with different filters.


User Interface

When migrating a journal archive, using the Archive Shuttle folder splitting feature is recommended since many target environments have limits on the number of items that can be in a single folder.

These settings can be found on the **Configuration > General System Settings** page.

The settings are:

- **Journal Split Base Folder Name:** This is the name of the top-level folder that will be split when items are ingested into the target environment.
- **Journal Split Threshold:** This is the number of items that should be placed in each folder before a new folder is created.



The screenshot shows a configuration window titled "Folder" with several settings. The settings are as follows:

Setting	Value
Folder name for folder-less items	Folderless
Treat all messages as folderless	<input type="checkbox"/>
Split Folder Parallelism	2
Split threshold large folders (except journal / folderless)	900000
Split threshold folderless items	100000
Journal Split Base Folder Name	
Journal Split Threshold	0

The "Journal Split Base Folder Name" and "Journal Split Threshold" fields are highlighted with a red rectangle.

The settings above can be applied only for Enterprise Vault and SourceOne sources.

Other sources (like EAS, DAM/QAM, and PST) are not supported. This is because they are unable to recognize the journal archive as a specialized container and are treated like a normal archive.

For EAS, DAM/QAM, and PST sources, the large folder setting named Split threshold large folders (except journal / folderless) can be applied on their journal archive.

Journal Splitting

Archive Shuttle Journal Splitting setup

The Office 365 Virtual Journal feature brings with it the possibility to migrate large journal archives to Office 365, using multiple target mailbox/personal-archives and specifying criteria such as number of items and size of target. This article explains more detail about this splitting process.

There are two ways that data can be split:

Number of Items

This is the simplest form. The splitting to a new mailbox/personal-archive occurs when the specified number of items are routed to the target mailbox or personal archive.

Size of Items

The splitting here is determined by the size of the items from the source environment. In the case where migration is from Enterprise Vault, the size used is the larger of the original size versus compressed size. When the amount of data routed to a target container reaches the specified value, a new target container will be created.

Journal archives in source environments are typically large. Their multiple terabyte nature makes it currently quite hard to migrate the data to Office 365 as there are limits imposed by Microsoft on both the mailbox and personal archive sizes. To be able to migrate a large journal archive to Office 365, the product introduced a concept of a 'Journal Splitting'.

There are two processes in migrating a journal archive to Office 365. This can be via **Journal Splitting** or a **Journal Transformation**. Below are the steps for Journal Splitting.

The general process for doing this is:

1. Define the naming scheme for the target. It is suggested to prefix or postfix names, eg AL- or -Departed. This makes the data easier to find in the target after the migration has completed.
2. Define the allowed number of rolling licenses.
3. Decide on the type of 'hold' to place on the migrated data.
4. Configure the maximum number of items allowed per child container, and maximum size allowed.
5. Perform the mapping.

In the background what will happen is:

1. Provisioning

- a. A user account will be created according to the naming scheme.

- b. A Personal Archive will be created if it was required in the mapping.
- c. The mailbox/personal archive will be placed on the selected type of hold.
- d. A license will be assigned from the pool.
- e. The data from Office 365 about the 'user' will be synchronized.

2. Migration

- a. Data will be exported after the provisioning process is done.
- b. Data will be imported soon after it is exported, when a particular child container is full, as determined by the system settings, a new one will be created.

3. Stage 2

- a. The license which was assigned in order to be able to ingest data in to the mailbox (or Personal Archive) will be removed (usually as the last step) and returned to the pool so that another mapping can take that license and complete the provisioning step.

Transformation

Journal Transformation information can be found [here](#).

Suggested values for splitting journal data

This article gives an indication of some suggestions relating to splitting of data when migrating with Archive Shuttle.

Virtual Journal Total Mailbox Item Count Limit

900,000 items

Virtual Journal Total Mailbox Item Size Limit (MB)

80 Gb, this way it is lower than the normal size allowed for a personal archive

Journal Split Base Folder Name

This is a customer/migration dependent, or can be left at the default

Journal Split Threshold

100,000 items, this way if it is ever necessary to open a mailbox it won't cause too many problems if using Outlook.

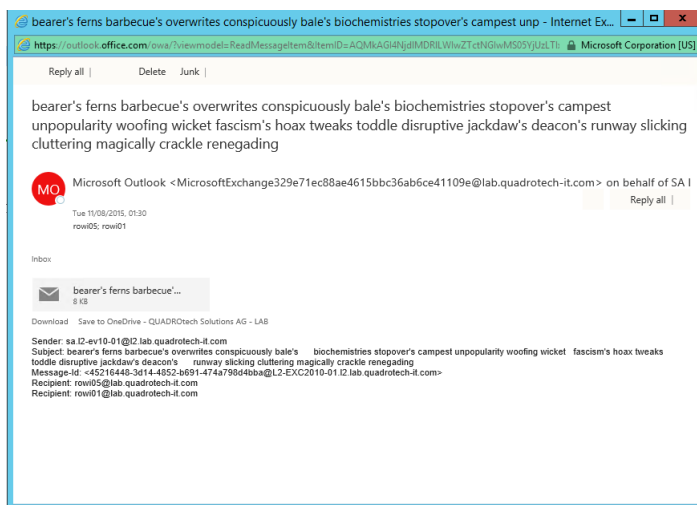
Convert journal messages to Office 365 journaling format

Occasionally, before migrating Journal messages from Enterprise Vault to Office 365, it is important to consider if want to have migrated the entire envelope, or just the P2.

Archive Shuttle has a **Convert Journal messages to Office 365 Journaling Format** setting on Office 365 module tab in System Configuration, and this setting works as follows:

a) If unchecked (default), it will migrate the whole envelope (P1 and P2) to Office 365. In the Compliance Center, then you can do a search, because envelopes were implemented in the search functionality.

Journal message migrated with P1 and P2 item to Office 365:



b) If checked, it will migrate only the P2 message and will remove the P1 header. The metadata of the P1 message (DL Expansion, BCC etc.) is then put into special properties of the P2. Because of this, Office 365 Compliance Search can find them.

Journal message migrated with P2 item to Office 365:



In the background, once Archive Shuttle collected information from the P1 message, it will create new properties and store the P1 info into them.

i | NOTE: Archive Shuttle does not currently support only ingesting P1 and stripping P2.

Considerations when enabling the option

- If you wish for the items to be migrated much more as-is, i.e. matching the source system, then do not enable the option.
- If you wish for the items to be more inline with the way that Office 365 does this natively, then the option should be enabled.
- If you are doing a Journal Transformation migration, then the option is enabled silently, and can not be disabled.
- When looking at emails from a compliance point of view, messages stored with the GERP attribute contain information about who got the original email, but not the why they got the email (that information is stored in the P1 message, for example, the message was sent to a distribution that contained a particular recipient)
- Office 365 works with either formatted message.
- Migrating data with the P1 & P2 data means that the P1 data becomes a free-text search, whereas enabling the option to convert the journal message (and adding the information to the GERP attribute) means it is a fully indexed property of the message. Confusion can arise if you search for 'mails to X', if P1 & P2 data is migrated, this type of query may return fewer results. Instead, you should perform a free-text search instead, but then it is not possible to search for 'mails o X' in the same way.

Migrating leavers to Office 365

Archive Shuttle can manage a pool of licenses to provision and migrate data to Office 365 mailboxes or archives whose source archives have been retained but owners have left the organization. It can even treat an archive as though it is ownerless, and migrate it using this process, even if an owner is shown in the user interface.

The general process for doing this is:

1. Define the naming scheme for the target. It is suggested to prefix or postfix names, ex. AL-
<archivename> or <archivename>-Departed. This makes the data easier to find in the target after the migration has completed.
2. Define the allowed number of rolling licenses.
3. Perform the mappings.

In the background what will happen is:

Provisioning

1. A user account is created according to the naming scheme.
2. A Personal Archive is created if it was required in the mapping.
3. A license is assigned from the pool.
4. The mailbox/personal archive is placed on litigation hold.
5. A license is assigned from the pool.
6. The data from Office 365 about the user is synchronized into Archive Shuttle.

Migration

1. Data is exported after the provisioning process is done.
2. Data is imported soon after it is exported.

Stage 2

1. The familiar parts of the workflow still occur, such as renaming the source archive and doing a final delta.
2. The user associated with the mailbox targeted is removed; resulting Office 365 processes are to reclaim the license and treat the mailbox in an inactive status. The mailbox is removed in Stage2 with the RemoveMailbox command.

Items which belonged in a mailbox can be searched via eDiscovery. Click here for more information.

How to set this up is described here.

Requirements

The normal Office 365 migration requirements are necessary (see the earlier section). In addition, note that Azure management tools are required. These can be downloaded from: <https://msdn.microsoft.com/en-us/library/azure/jj151815.aspx>


If these components are not installed, the normal Office 365 migrations will still be successful, but processing of leavers will not be successful. They can be added at any time during the migration; it is not necessary to reinstall or modify the Office 365 module following their installation.

Configuring the Office 365 environment

Naming scheme

This determines how mailboxes and accounts will be created, using a token system. The following are tokens which can be used:

- Archive Name
- Archive ID
- Container Mapping ID
- PST ID
- PST Number

 **NOTE:** PST ID and PST Number are applicable if PSTs are the source environment.

In addition to those, alphanumeric characters can be appended or prepended to the name of the mailbox in Office 365.

Usage location

This is the location relating to license allocation and is required to be chosen from the drop down list.

E-Mail domain

This is the location relating to license allocation and is required to be chosen from the drop down list.

Hide from GAL

If enabled the mailbox/user will be hidden from the Global Address List.

Litigation hold

Litigation hold can be configured, if required, using the Legal Hold option.

Set rolling license count

A specific number of licenses can be used to process the journal splitting mappings. That number should be configured by clicking **Set Rolling License Count**. Licenses will be consumed up to that limit, and will be freed up when Stage 2 completes on those mappings.

i | **NOTE:** There is a **System Configuration** option which can be enabled to migrate the data into Office 365 journal format, if required.

Office 365 configuration changes

Journal Splitting item count limit

This determines the maximum number of items which will be in a particular mapping before rolling over to a new mapping.

Journal Splitting item size limit

This determines the maximum size of a mapping before rolling over to a new mapping.

Performing the Migration

The migration of the archive/container can then be performed. The process of mapping the users is the same as with a normal Office 365 migration, however, there is an additional option in the Mapping Wizard where you can select to use Leavers, and you should select an Office 365 Leavers Workflow.

A container can be processed as a leaver if one or more of the following is satisfied:

- It is ownerless in the user interface
- The Active Directory account is disabled
- The Active Directory account is deleted and exists in the Active Directory Dumpster. Note: The account running the AD Collector module must have permissions to read the contents of the dumpster.

i | **NOTE:** Archives will be exported, but only the configured number of rolling licensed users will ingest data into the target.

Stage 2

The Stage 2 workflow (EnterpriseVault to Office 365 Leavers) functions in a similar manner to the normal Office 365 workflows. There is an additional step in the workflow (Office365RemoveUser) that handles the removal of the user, resulting in an “inactive mailbox” appropriate for a naturally occurring leaver. Once the license has been reclaimed from the removed user, additional users that require migration via the leavers process will begin.

Handling users who leave during a migration

During a migration, if Archive Shuttle detects that a mailbox is no longer present, the user is marked **IsDeleted**.

Then, you can select the user and select the **Enable to Recreate** button on the **Bulk Mapping** page. This enables the leaver mailbox to be created. The user’s data then gets ingested into the leaver mailbox, as described above.

Or, you can automate the process by selecting the option to **Enable auto-recreate deleted Office 365 Users as Leavers** located on the **Office 365 Module** page within System Configuration.

Migrating a shared mailbox archive

This section outlines the steps that are necessary to complete the migration of an archive shared amongst a group of users. The migration will be from one Enterprise Vault environment to another.

In Enterprise Vault, this may simply be a regular mailbox archive so that it maintains the structure from the mailbox where the data comes from. The mailbox itself will ultimately have one true owner in Active Directory, but a number of other users and/or groups may also have access to both the mailbox and the archive.

Capture Permissions on the Source

There are a number of ways in which the permissions can be shown for an archive, including the following:

- Open the properties of the archive in the **Vault Administration Console**, and check the entries on the permissions tab
- Use a script similar to the one provided on the Symantec Connect web site:

<https://vox.veritas.com/t5/Articles/Script-to-know-AD-permission-assigned-on-Archives-Folders/ta-p/806130>

Build an EVPM File

From the list of permissions on the source archive an EVPM file should be built. There are details about the way to use EVPM in the Enterprise Vault Utilities Guide. At a high level the steps are:

- Directory section should contain information about the Enterprise Vault Directory computer and Enterprise Vault site name.
- A section should be added called **ArchivePermissions**
- The **Archive Name** property should be added to this section with the name of the archive.
- The **Grant Access** property should be added to this section with a list of people to grant access to the archive. This should be listed one per line.

Perform the Migration

The migration of the archive/container can then be performed. Particular care should be taken in relation to migrating these types of archives/containers because the migration has an impact on a number of people.

Apply the EVPM File

Following the successful migration of the Shared Mailbox / Archive, the EVPM file that was created previously should be run to grant permissions on to the target/new archive.

About Us

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. For more information, visit www.quest.com.

Contacting Quest

For sales or other inquiries, visit www.quest.com/contact.

Technical Support Resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>. The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to videos
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product