

Quadrotech Archive Shuttle 11.1 **Administration Guide**



© 2023 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.

Quadrotech Archive Shuttle

Updated July 2023

Version 11.1

Contents

Introduction	6
Configuring Archive Shuttle	8
Initial Archive Shuttle Setup	8
Configuring module schedules	8
Enabling Active Directory domains	9
Adding Enterprise Vault environment(s)	10
Adding link databases	11
Configuring module mappings	12
Validating the setup	13
Adding a new license	13
Using HOTS	14
Using the Archive Shuttle user interface	16
General information	16
Dashboards	20
Progress and Performance	20
System Health	22
System Health events	25
User dashboard	26
Module dashboard	27
Manage	31
Bulk mappings	31
Manual mappings	34
Existing mappings	35
Mapping templates	37
Setting templates	38
Creating a setting template	39
Assigning a setting template	40
Offline media	41
Pre-migration	43
Stage 1 (Sync data)	46
Stage 2 (Switch user)	49
Configuration	51
User groups	51
Tags and Tag Management	52
Creating a .csv file	54
Enterprise Vault environment	55
Enterprise Vault Discovery Accelerator environment	56
Enterprise Vault retention mappings	58
SourceOne/EMX environment	59

Links	61
Advanced configuration	66
Naming policies	66
Considerations when migrating to the recoverable items folder	67
Workflow policies	68
Workflow commands	70
Migration filters	77
Filter policies	78
System configuration	81
Monitoring	105
Day to day administration of migration activity	105
Monitoring a migration	108
Completing a journal migration	113
Logging	115
Logging levels	115
Reporting	117
Link statistics	117
Performance dashboard	118
Scheduled reports	119
Failed items	120
Index fragmentation	121
License information	122
Journal Transformation migration	124
Enterprise Vault to Office 365	129
Selective migration based on recipients collected from EV Index	132
Journal Transformation detail grid	134
Journal Splitting	135
Which Office 365 license to use for leavers or virtual journal	137
Suggested values for splitting journal data	138
Journal envelope processing	138
Journal archive folder split settings	139
EV collector journal collection logs	140
Alternate mailbox support for discovered addresses associated with active users	140
Retry failed remediations	141
Support handling of invalid or changed SMTP addresses	142
Providing accepted domains for a migration	142
Why legal hold needs to be applied for data retention after migration	143
Automatically identify known distribution lists and mark them as such	143
Enable logging	144
Shared mailbox over quota rollover	147
Re-exporting a PST for a certain user	147
Adding Sender-Recipient SMTP addresses	148
Monitoring a migration	150
Completing a migration	154

PowerShell guide	156
System health events with PowerShell	166
Running PowerShell with Archive Shuttle	169
Variables that can be used in PowerShell execution	170
Troubleshooting	171
EV export issues	171
EV import issues	173
EV collector issues	175
Disk space reducing unexpectedly	177
Handling ARCHDVS files whilst migrating	179
Office 365 import issues	180
Exchange import issues	183
Native Format Import issues	185
Stage 2 command issues	187

Introduction

This guide contains detailed information on configuring and using Quadrotech Archive Shuttle.

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

Configuring Archive Shuttle

This section outlines the basic steps required to configure Archive Shuttle. This is applicable for most archive migration projects. The user interface and other components of Archive Shuttle are explained in subsequent sections. Archive Shuttle recommends reading each section before performing the post-installation configuration.

Initial Archive Shuttle Setup

After Archive Shuttle has been installed, the following five tasks need to be performed before you start migrating archives:

- Configure module schedules.
- Enable Active Directory domains.
- Add Enterprise Vault environment(s), if Enterprise Vault is the source or target environment.
- Add link databases.
- Configure module mappings.
- Validate the setup.

Configuring module schedules

During Archive Shuttle installation, the core product, databases, and a number of modules will have been deployed in the environment in order to fulfill migration requirements. During the installation of the modules, each module is enabled. If modules are not enabled, they will not receive any work from the Archive Shuttle Core.

In order to check that all appropriate modules are enabled and operational, follow these steps:

1. From the Archive Shuttle web interface, click **Configuration > AS Modules** in the navigation bar.
2. Review the list of modules:
 - Ensure all required modules are present. For more information see the *Quadrotech Archive Shuttle Installation Guide*, as well as the *Quadrotech Archive Shuttle Planning Guide*, which will help you validate this part of the migration.
 - Ensure that the computer name, domain name, and version are as expected.
3. Verify that all modules are reported appropriately and that none of the modules have a red background. If they do, there is a communication problem between the module and the Archive Shuttle Core.

All of the modules are configured to run continuously by default. If this does not meet your requirements for the migration, modules can be individually scheduled to more suitable times using the **[Set Schedule]** button on the AS Modules page.

Enabling Active Directory domains

A few minutes after the Active Directory Collector Module is enabled, a list of domains where the Active Directory Collector Module is running become visible on the **Configuration > Active Directory** page.

By default, all domains are discovered, but won't be scanned for user accounts. In order to enable one or more domains to be scanned, follow these steps:

1. From the Archive Shuttle web interface, click **Configuration > Active Directory** in the navigation bar.
2. Review the list of domains.
3. Select the check box next to one or more domains.
4. Click **Enable** in the navigation bar.
5. Make sure the **Scan Domain** column shows a check mark.

Active Directory information is collected through the assigned Archive Shuttle Active Directory Collector module.

Each domain in the forest where the Active Directory Collector Module is running will be shown, but by default user scanning for each domain will be disabled.

If there is an Active Directory Collector Module in additional forests, they too will be shown on this screen.

Actions to be performed on the Active Directory page

Sync AD Users

This instructs the appropriate Active Directory Collector modules to synchronize user information from all enabled Active Directory domains.

Sync AD User

This allows you to specify a single Active Directory user that should get synchronized. This is useful if you have issues with a single User and allows you to easily troubleshoot this one user, or simply if you know you had changes for a single user and you need to show up ASAP.

Sync AD Domains

This instructs the Active Directory Collector module to gather a list of all domains in the environment/forest.

Enable

When one or more domains have been selected, clicking the 'Enable' button in the Actions Bar will enable synchronization for those domains.

Disable

When one or more domains have been selected, clicking the 'Disable' button in the Actions Bar will disable synchronization for those domains.

Refresh

Refreshes the current screen

Adding Enterprise Vault environment(s)

i | NOTE: Skip this step if Enterprise Vault is not the source or target of a migration.

The next step to perform in the migration workflow is to add Enterprise Vault environment(s). This is performed from the **Configuration > EV Environment** page. This page also displays environments that are currently configured.

If the migration is between two Enterprise Vault directory databases, then this page is used to enter details for both directory databases. If the migration is between Enterprise Vault and an external system, such as Microsoft Exchange or Office 365, then only one Enterprise Vault Directory database is added. Likewise, if the migration is between two Enterprise Vault sites within the same Enterprise Vault directory, only one entry is listed on this page.

To add an Enterprise Vault environment:

1. From the Archive Shuttle web interface, click **Configuration > EV Environment**.
2. Click **Add** to add an Enterprise Vault directory.
3. Enter the required information (for example, the fully qualified domain name and instance name, if it's not the default instance).
4. Click **Add/Update** to apply the changes to the Archive Shuttle database.

Once you specify required EV Directory databases, the Enterprise Vault Collector Module gathers and displays information about the environment. These are then displayed in the EV Environment page after a few minutes.

i | TIP: Click **Refresh** in the navigation bar to reload the list of vault stores.

To enable archive collection for each vault store that you will use in the migration:

1. From the Archive Shuttle web interface, click **Configuration > EV Environment** in the navigation bar.
2. Select the checkbox next to each vault store.
3. In the navigation bar, click **Enable**.
4. Verify that the Archive Collection Enabled column now contains a check mark.

New Enterprise Vault Directories can be added, and Archive Shuttle collects information about this Enterprise Vault directory through the assigned Archive Shuttle Enterprise Vault Collector module.

For every Enterprise Vault Directory, the table lists the following information:

- Vault Store Name
- Number of Archive in this Vault Store
- Total Item Count in this Vault Store
- Total Item Size in this Vault Store
- Total Shortcut Count in the archives' associated mailbox
- Total Legal Hold Count
- Archive Gathering Enabled
- Backup Mode enabled

i | **TIP:** The data values will only be shown on screen when “Archive Gathering” has been enabled, and an EV Collector module has been linked. (See below)

Adding link databases

Once archive collection is enabled, the next step in the migration is to add link databases for each of the vault stores. A link database is needed for each link in the source environment.

The **Configuration > Links** page may already contain some links that were collected from the Active Directory Collector Module (the Exchange databases) and from the Enterprise Vault Collector Module (the Vault Stores, as a result of adding each Enterprise Vault environment). These links are created automatically. You must manually create other links, for example, links to Office 365 and PST. To create a link database:

1. Open the **Configuration > Links** page.
2. Select the tab of the type of link database you're creating (for example, Office 365), and then click Create Link.
3. Follow the prompts to create the link.

After a link database is created, check the EV Environment page again to review the progress of the archive and item gathering stages.

i | **NOTE:** The number of archives and size of data will not be populated on the EV Environment page until the module linkage is added, but the vault stores will display.

Configuring module mappings

The next stage of the configuration involves mapping modules to the Links in the migration project. To do it, follow these steps:

1. Go to **Configuration > Links**.
2. Select the check box next to a vault store/Exchange database/source link.
3. Click **Map Modules** in the navigation bar.
4. From the pop-up window, select appropriate modules from the drop-down lists, and then click **[Save]**.

Edit module mappings

Module mappings can be modified by following these steps:

1. Go to **Configuration > Links**.
2. Select the check box next to a vault store/Exchange database.
3. Click **Map Modules**.
4. Make desired changes to the current mapping for this link, and then click **[Save]**.

Shortly after module mappings have been added, the appropriate modules perform their tasks and the Number of containers column updates to show the number of containers in the link.

For the modules that are performing export and import functions, it's also important to set the staging area path. The corresponding modules need to be paired up to facilitate the flow of archived data. For example, the EV Export module path should correspond to the Exchange import path if the migration is needed to flow in that direction.

To configure the staging area path:

1. Go to the **Configuration > Links** page.
2. Select the checkbox next to a vault store/Exchange database/source link.
3. Click **Path** and follow the prompts.

Repeat this step for each export and each import module. For the second and subsequent paths you define, the Staging Area Path is pre-populated with the path that was previously used. This makes it possible to quickly configure many modules with the same storage location.

The current export/import path and module mappings can be viewed on the **Links** page. Make sure all links are correctly defined with the appropriate export/import path before proceeding with the migration of any archived data.

It's also possible to configure the **Default Staging Area Path** that all modules and links will use unless the setting is overridden on an individual link.

Validating the setup

Once the stages above have been performed and all necessary setup and configuration steps are complete, the archive migration is ready to begin.

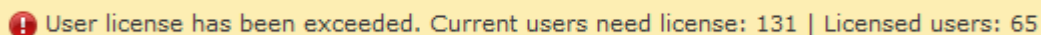
At this point, a collection of links have been defined and mapped to modules that perform tasks on those links. The source and target environments have been defined within Archive Shuttle and all of the Archive Shuttle Modules are enabled and set to an appropriate schedule. Finally, the Enterprise Vault retention categories have been mapped according to the migration needs.

i | **NOTE:** Before proceeding with any migration of archived data, confirm that each stage has been performed.

Depending on the migration, additional steps may be required. These are described later in this guide.

Adding a new license

From time to time it might be necessary to update the license information for Archive Shuttle. This may be because additional modules have been purchased, or it may be because the following type of warning has been displayed, and then additional licenses purchases:

A yellow rectangular warning box with a red exclamation mark icon on the left. The text inside reads: "User license has been exceeded. Current users need license: 131 | Licensed users: 65".

! User license has been exceeded. Current users need license: 131 | Licensed users: 65

In order to update the license information for Archive Shuttle, the following steps need to be performed:


1. Copy the new license.dlv file to the following folder:

 Webservice\bin\

2. Run an IISReset command.

Using HOTS

HOTS is a feature with Archive Shuttle which allows for less bandwidth usage for migrations to Office 365 by combining a highly-optimised transfer system with storage of extracted data taking place in Microsoft Azure. The following diagram presents an overview of HOTS:



Requirements

For migrations from legacy archives to Office 365 using HOTS the following needs to be considered:

- Using HOTS is supported for all currently supported sources, when migrating to Office 365 mailboxes or Personal Archives.
- More CPU usage might be needed on the source environment in order to create the HOTS-format data
- An Azure Storage Account must be configured and used for storing the extracted data
- All export and import modules must have been configured with the connection string to the Azure Storage Account
- A bridgehead server running the ingest and shortcut processing module should be deployed in Azure to facilitate the ingestion of the data from the Azure Storage Account. (For bridgehead server specification, refer to the *Quadrotech Archive Shuttle Installation Guide*).

i | **NOTE:** It is possible to limit the storage used in Azure by configuring a high watermark for the links.

Additional configuration

The following additional configuration needs to be performed in Archive Shuttle in order to utilize HOTS:

System setting

In Archive Shuttle HOTS must be enabled in **System Configuration -> General**

Configuration on each module

Each Export Module must be configured with a connection string in order to be able to send data to the Azure Storage Account. This is configured in the **Credential Editor**. The connection string is obtained from the Azure Portal. The corresponding Import Module needs to also be configured with the credentials.

Configuration on link

The source and target links need to use an appropriate Azure Storage Account connection. This is done from the **Links** page within Archive Shuttle.

High watermark on link

By default Microsoft Azure places a very high maximum allowance for Storage Accounts. Customers may wish to restrict the usage of the Storage Account because it may involve additional cost to Microsoft. To place a limit on the amount of storage to be used the high watermark can be set on each of the source links.

Using the Archive Shuttle user interface

The Archive Shuttle User Interface is fully web-based. This allows archive migrations to be performed from anywhere in an organization from an Internet browser. It also means that archive migrations can be managed externally by partners of the migration organization. Archive Shuttle recommends that HTTPS be enabled in that situation.

General information

All pages share a common look and feel within the Archive Shuttle user interface.

On the left side is the main navigation area. It is used to jump from page to page, where every page controls a different aspect of Archive Shuttle.

Along the top are “Actions”, which are context sensitive to each page.

The following pages are available:

Page	Description
Progress and Performance	Provides an overview of the status of the migration.
System Health	Shows key information about the health of the migration environment.
User Information	Allows an administrator to see detailed information gathered about users.
Bulk Mapping	Provides tasks to create container mappings.
Manual Mapping	Provides tasks to create single container mappings.
Existing Mapping	Provides tasks to manipulate existing mappings.
Offline Media	Allows configuration of offline media (e.g., disk shipping).
Pre-migration	Run workflow commands prior to migration on a per mapping basis.

Page	Description
Stage 1 (Sync Data)	Provides detailed information on the migration status for each Container mapping. This is the “Item Sync” status for the container.
Stage 2 (Switch User)	Provides detailed information on the migration status for each Container mapping. This is the “Workflow” status after switching a container to the target environment.
Link Statistics	Provides detailed information on the migration status for each Link.
Performance	Provides a detailed hourly performance overview of the migration and allows selection of date/time ranges to see this data historically and graphically
Failed Items	Provides information on items that have failed to export or import.
Index Fragmentation	Shows information about SQL Index Fragmentation.
License	Provides information on the licenses available for migration.
Support Bundles	Allows the creation and manipulation of log file bundles that are helpful to QuadrotechSupport Engineers
AS Modules	Allows an administrator to control the configuration of modules deployed in the environment.
Active Directory	Provides a list of Active Directory domains and allows for them to be scanned for users.
User Groups	Provides an administrator the ability to manage Archive Shuttle groups. These are used as aids when selecting containers to migrate.
Tag Management	Add, edit or delete tags assigned to mappings.
EV Environment	Provides an overview of the Enterprise Vault Environment. It’s also used to add new environments and various Enterprise Vault related collection tasks.
EV DA Environment	Provides an overview of the Enterprise Vault Discovery Accelerator environment. It’s also used to add new environments

Page	Description
	and various Enterprise Vault Discovery Accelerator related collection tasks.
EV Retention Mappings	Provides insight into Enterprise Vault Retention Category mappings that are needed for the migration.
EAS Environment	Provides an overview of the Enterprise Vault Discovery Accelerator environment. It's also used to add new environments and various EAS related collection tasks.
SourceOne/EMX Environment	Provides an overview of a SourceOne/EMX environment. It's also used to add new environments and various SourceOne/EMX related collection tasks.
Dell Archive Manager Environment	Provides an overview of a Dell Archive Manager environment. It's also used to add new environments and various Dell Archive Manager related collection tasks.
Metalogix Environment	Provides an overview of a Metalogix environment. It's also used to add new environments and various Metalogix related collection tasks.
Office 365 Environment	Provides an overview of an Office 365 environment. It's also used to add new environments and various Office 365 related collection tasks.
Links	Provides an overview of Links, allows creation of an Archive Shuttle Item databases as well as enabling Map Modules to Links.
Workflow Policies	Allows an administrator to review the commands within a workflow, modify workflows and add new workflows.
Migration Filters	Allow an administrator to construct filters that can filter the data during migration.
System Configuration	Allows an administrator to fine-tune the migration environment.

Grids/Tables

Filtering

Many of the grids / tables can be sorted and filtered as required.

Sorting can be achieved by clicking on the appropriate column header. Ascending versus descending order is achieved by clicking on the column heading a second time.

Filtering is provided on many columns by the means of the textbox below a column header. Data entered into this text box will cause Archive Shuttle to “filter by” on that column. The small “antenna” symbol behind the textbox can be used to change the type of filter, with the following being common options that are available:

- Begins with
- Contains
- Does not contain
- Ends with
- Equals
- Does not equal

It is possible to combine multiple column filters to further enhance the displayed data in the grid / table.

In order to apply the filter, which may contain multiple columns/selections click on the ‘Apply’ button at the right of the filter row.

A second filtering option is provided by the little “antenna” symbol behind the header name of a column. It provides a Microsoft® Excel® like filtering interface which lists the following values for an easy selection:

- (All): Does not apply a filter, displays All Items
- (Blanks): Only displays items that do not have a value
- (Non blanks): Only displays items that have a value, hiding blank values
- Values: Aggregated list of values from that column to choose from

Paging

All grids provide standard paging functionality. It is possible to see the current page, and select to view another page from the pager at the bottom of the grid.

Dashboards

Progress and Performance

The Progress and Performance Dashboard provides an administrator with an overview of the current state of the archive migration project. The following information is displayed:

Item	Description
Status	<p>User finished migration:</p> <p>Percentage of mapped containers that have been through the complete workflow</p> <p>Stage 1 (Synced):</p> <p>Percentage of users that have been synced.</p> <p>Stage 2 (Switched):</p> <p>Percentage of enabled users switched to target</p>
Basic statistics	<p>Data gathered for the migration, including:</p> <p>Users</p> <p>Total number of users</p> <p>Mailboxes</p> <p>Total number of mailboxes</p> <p>Archives</p> <p>Total number of archives</p> <p>Archived Items</p> <p>Total number of archived items</p> <p>Archived Items Size</p> <p>Size Total size of the archived items</p> <p>Mappings</p> <p>Total number of containers that have been mapped</p>

Item	Description
	<p>Enabled for Sync (Stage 1)</p> <p>Total number of containers that are in Stage 1</p> <p>Enabled for Switch (Stage 2)</p> <p>Total number of containers that are in Stage 2</p> <p>Finished Migration</p> <p>Total number of containers that have finished migration</p>
Last 10 Workflow Completed Users	<p>Last ten users that have fully completed their migration workflow</p> <p>Clicking in this area will move you to the Stage 2 screen within the Admin Interface.</p>
Export Speed GB/h Extraction Speed Items/h Ingestion Speed GB/h Ingestion Speed Items/h	Export and import speed per hour, calculated from the last ten minutes values
Overall Performance	<p>Collection, Export and Import performance broken down into a timetable.</p> <ul style="list-style-type: none"> • Last ten minutes • Last hour • Current day • Current week • Current month • Current quarter • Last 6 months • Current year • All
Archive Statistics	<p>For how many archive have archive statistics been received.</p> <p>Archive statistics are:</p>

Item	Description
	<ul style="list-style-type: none"> • Count of items in an archive • Size of items in an archive • Count of shortcuts in mailbox belonging to this archive
Collected Items	How many items need to be collected in total (based on archive statistics), and how many have already been collected.
Current Activity	<ul style="list-style-type: none"> • Running Export Item Count • Running Imports Item Count • Failed Export item Count • Failed Import Item Count • Running Export Item Size • Running Import Item Size • Failed Export Item Size • Failed Import Item Size
Exported Items Exported Size Imported Items Imported Size	Exported and imported count/size as a percentage

The Dashboard can optionally be switched to show all of this information for a specific link order to more closely monitor a particular part of the migration project.

It is also possible to show the extraction and ingestion speed in items per second, rather than items per hour.

Optionally, you can configure the page's widgets to display compressed sizes.

System Health

The dashboard provides an administrator with an overview of storage and other health parameters that may affect the migration being undertaken. If a problem is detected a yellow triangle will be visible in the user interface and a number. The number represents the number of problem types.

The page is divided into the following sections:

Section	Description
Modules	Information relating to the status of modules. Those with issues will be shown with a warning symbol. Hover over those items and a tooltip will give more information. In addition this page shows details relating to any credentials which have been configured to be used by the system to access resources. These credentials are updated upon module restart.
Storage Paths	Shows information about the storage paths used by the system.
Miscellaneous	Any retention categories that are not mapped to a target environment are displayed here.
Events	Status events will be shown on this page. A list of events that may be displayed are described in the topic System Health Events.
Sync	Details relating to environment sync activities will be shown on this page.

Free space on a staging area location is color coded as shown in the following table:

Highlighting	Reason
No highlighting	Free space is above 100 GB
Yellow	Free space is below 100 GB
Red	Free space is below 50 GB

In addition, the System Health page gives an overview of any modules that are deactivated or otherwise not running. This is shown at the top left of the screen in the Admin Interface.


If all available Office 365 licenses are already assigned and there is no available license left, on the Events tab on the System Health page displays a warning message.

The Link information at the bottom of the page also highlights the ‘Used’ space on each link. Used space is defined as the used space from disk. This is also color coded for ease of identifying a potential problem, as follows:


Highlighting	Reason
No highlighting	Used space is within normal parameters
Yellow	Used space is between 75% and 90% of the maximum allowed
Red	Used space is above 90% of the maximum allowed

If the Used Space reaches 100% on a link, exporting of data will be stopped until data has been imported, even if there is still free disk space available on disk. This is further illustrated as follows:

i | **NOTE:** High and low watermarks can be adjusted from the Links page in the Archive Shuttle UI.



Export will function normally



Export for the highlighted link will stop, because the watermark has been reached – even though plenty of free space on the Staging Area.

The link information that is included also gives an indication of the indexing performance on the target (if the target is Enterprise Vault). The Index Incomplete column shows the number of

items that are waiting to be indexed on the target environment. This information is updated once per minute by the EV Import module.

i | **NOTE:** If the Index Incomplete column shows a large value, it might indicate an issue with indexing in the target environment, which should be investigated.

Actions to be performed on the System Health page

- **Cleanup Staging Area:** Issues a command to the Archive Shuttle modules to clean up the staging area of already imported files.
- **Run Now:** Run a scheduled task immediately.

System Health events

The events tab on the system health page in Archive Shuttle reports issue that might impact a migration. This section explains the events that might be generated. The activities are checked when the system health page loads rather than at scheduled intervals. This ensures that the events which are shown are up to date.

Title	Description
Office 365 License Issue	If the number of active licenses is more than the consumed plus suspended licenses a message will be shown.
SourceOne index issue	There is an index which is needed to greatly improve the migration from SourceOne. If SQL index is not found, a message is displayed.
Cargo Bay issue	If SQL Server Integrations Services (SSIS) 2017 is not installed, then Cargo Bay will not work and a message will be displayed.
Folderless item setting	If the system configuration does not have the Folderless Item setting configured, then a message is displayed. To resolve the issue, set the Folderless Item entry in the System Configuration.
License validity	If the current license is not valid, a message is displayed.
License issue	If there is some other problem with the current license, a message is displayed.

Title	Description
Mailbox quota issue	<p>If you have exceeded the limit specified in the setting Mailbox Quota Exceeding Threshold (%), which you can set by going to System Configuration > Office 365 Module > General, a message will be displayed. Clicking on this error will bring you to the Container Mapping Health page, which has additional information.</p> <p>i NOTES:</p> <ul style="list-style-type: none"> • If you do not see the error immediately, click the Run Health Check button to. • Applies only to Office 365 mappings.

User dashboard

The dashboard provides comprehensive information that has been discovered about a specifiable user. The following information is displayed:

Item	Description
Basic User Information	<p>Shows information about the user such as the User SID, login name, OU where the Active Directory account resides.</p> <p>In addition, the user address can be displayed. This might be helpful to confirm the chosen user identify.</p>
Mailboxes	Shows detailed information relating to the user mailbox, and personal archive as discovered by Archive Shuttle.
Archives	Displays detailed information relating to the primary Enterprise Vault archive that has been discovered for the user.

Item	Description
Migration Status	<p>Information relating to the migration can also be viewed, if this user is currently mapped for migration. The information includes the number of items, priority, number of failed items and information relating to the size of the archive.</p> <p>The information is also displayed graphically as a series of pie charts at the bottom of the screen.</p>
Stage 2 Status	This section shows the current state of any Stage 2 commands which have run on this user.

When migration is underway, this page will show additional information related to the migration. This includes statistics and overall progress graphs. It is also possible to see the Failed Items for a selected user/migration by clicking on the hyperlink on this page.

Module dashboard

Valuable information about module-level responsiveness can be gained from the Module Dashboard in Archive Shuttle. Using the dashboard it is possible to see detailed information relating to the operations being performed by particular modules on particular servers involved in the migration. For example, if new mappings have just been created, and yet one or more Enterprise Vault Export modules are showing no items being exported, it may indicate an area that should be investigated further.

The display also shows whether a module is enabled or not and the current logging level for the module and Core.

Modules have to be enabled in order for them to receive work to do.

The modules are monitored every 5 minutes to check if they are still running, or have failed. If they have failed an attempt will be made to restart a module, by issuing a command to the Admin Module on the affected machine. By default Archive Shuttle will try 5 times to restart a module. Every retry to start the module is by default one more minute apart from the previous attempt for a maximum of 10 minutes. The status of the module will be on the System Health page.

Actions to be performed on the Modules page

- **Enable:** Enable selected modules
- **Disable:** Disable selected modules
- **Configure:** Most modules allow specific configuration relating to parallelism and other performance/load related elements to be configured. For example you could configure one EV

Export module to have archive and item parallelism of 10 and 2 respectively, and another EV Export module could have archive and item parallelism of 15 and 1 respectively.

It's also possible to have these configuration changes scheduled to be effective at particular times of day and days of the week.

The data grid showing module information also includes columns which show which configuration is active, and the time left before the next change in configuration.

- **Start:** Start the selected module immediately
- **Stop:** Send a command to the selected module to stop processing
- **Restart:** Send a command to the selected module to restart
- **Delete:** Remove a module. It may be necessary to choose a replacement.
- **Enable performance statistics:** Enables the collection of performance statistics for a module
- **Disable performance statistics:** Disables the collection of performance statistics for a module
- **Set Schedule:** Define a schedule for when the selected module should run
- **Update:** Send a command to the selected module to perform an automatic update to the latest version.

i | **NOTE:** For the "Update" command to work correctly, there are additional steps that should be performed on the Archive Shuttle Core Server. These are covered later in this section.

- **Download:** Allows you to download the module MSI package.
- **Set Log Level:** Allow the logging level of a module (or set of modules) to be changed in real-time. Increasing the logging level may help with providing more detail to Quadrotech Support for troubleshooting.
- **Refresh:** Refreshes the module grid

i | **NOTE:** Disabled modules still transmit backlog data to the Core Web Service. These modules do not get new work to perform.

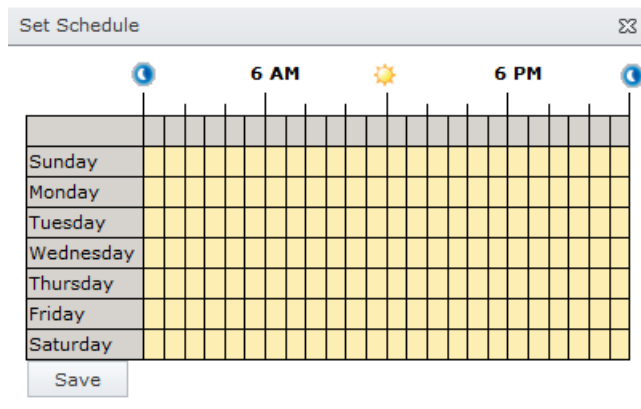
Setting a schedule for a module

Each of the Archive Shuttle modules can have a schedule defined for when the module should execute tasks. To set a schedule, follow these steps:

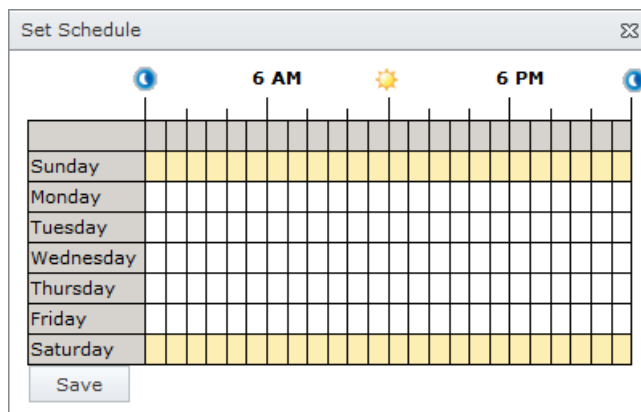
1. Navigate to the "Modules" page.
2. Select a Module by clicking on the checkbox next to the name of the module
3. Click on "Set Schedule" in the navigation bar

i | **NOTE:** When setting a schedule remember to click on "Save" in the "Set Schedule" window to commit the schedule to the database.

The following screenshot shows a module schedule where the module is configured to run 24x7:



The following screenshot shows a module schedule where the module is configured to run just on Saturday and Sunday:



Update installed modules

During a complex migration there may be many modules installed on different servers throughout an environment. Archive Shuttle has a method for simplifying the process of updating the modules when new versions become available. In order to update modules from the Archive Shuttle Web Interface, follow these steps:

1. Download the new MSI file, the file required will have a name formatted in this manner:
ArchiveShuttleModulesInstall-X.x.x.xxxxx
2. Copy the MSI file to the following folder on the server that hosts the Archive Shuttle Web Service:
 - a. Webservice\bin\ModuleInstallers
3. Navigate to the "Modules" page in the Archive Shuttle User Interface.
4. Select one or more modules, and click "Update".

The MSI file will then be transferred to the server that was selected, and then the MSI file will be executed in order to update the installation.

Restart Enterprise Vault Services

From the 'Services' tab it is possible to restart some or all of the Enterprise Vault services on a particular Enterprise Vault server. You can also direct the command to a particular Archive Shuttle Admin Service.

Manage

Bulk mappings

This page is used to map containers to either new or to existing containers based on certain criteria.

The following basic information about each container is visible in the grid view:

- Name
- Indicator field to show whether the container has an owner or not
- Owner Full Name
- Group
- Type of container (e.g., Enterprise Vault)
- Archive Type
- Indicator field to show whether the container has a mailbox or not
- Link Name
- EV Archive Status
- Item Count
- Total Size
- Indicator field to show whether the container is mapped or not

If a Container's owner / user could not be determined in Active Directory, it is considered "ownerless", and is marked as such in the Name column.

Actions to be performed on the Bulk Mapping page

Add Mappings: Add mappings for selected containers. A pop-up wizard allows the mapping to be defined. The information that is required depends on the target:

Enterprise Vault - The target user strategy must then be chosen as follows:

- Same User
- Different User (Choose this option if the migration is to take place to another domain. The target domain can be chosen from a drop-down list, and matching criteria must be specified (e.g., Legacy Exchange DN, SID History, User Name))

Run Item Collection: Collect full or delta items for containers.

Run Shortcut Collection: Run full shortcut collection.

Auto Create Mappings: Assigns mapping templates and automatically creates mapping for selected containers(s) once Office 365 user's mailbox is present in the database. To assign a template to automatically create a mapping:

1. Select one or multiple containers to assign a mapping template to.
2. Click **Auto Create Mappings**.
3. Select the template to assign to the container, then click **Save**.
4. After a moment, the mapping template should be assigned to the container.

You can check which auto-created mapping templates have been assigned under the columns **Autocreate** and **Template For Autocreate** in the grid.

To remove a mapping template for auto-creation:

1. Select the container which has been assigned a mapping template for auto-creation.
2. Click **Auto Create Mappings**.
3. Select **No template**, then click **Save**.



NOTE:

- Only templates with Office 365 as the target can be assigned. Templates for leavers currently cannot be used.
- You currently cannot assign mapping templates to ownerless containers.
- You currently cannot change the template once a mapping has been created.

The container strategy must then be chosen as follows:

Create new containers - Choose this option if Archive Shuttle should create new containers in the target environment.

Using existing containers - Choose this option if Archive Shuttle should use existing containers in the target environment.

Create new if there is no existing - Choose this option if Archive Shuttle should primarily use existing containers in the target environment, and if they do not exist, they will be created.

The target link must then be selected.

Choose the Link where you want to migrate to: - Select an entry from the drop down list that will correspond to a link in the target environment.

Exchange - The mailbox type must then be chosen as follows:

- **Primary mailbox:** Choose this option if Archive Shuttle should ingest the data in to the users primary mailbox.

- **Secondary (Archive) Mailbox:** Choose this option if Archive Shuttle should ingest the data in to the users secondary (archive) mailbox. If this option is selected, then an option of what action should be taken if a secondary mailbox does not exist can be specified. The choice of options are to either skip the ingestion or ingest into the primary mailbox instead.

Office 365 - The mailbox type must then be chosen as follows:

- **Primary mailbox:** Choose this option if Archive Shuttle should ingest the data in to the users primary mailbox.
- **Secondary (Archive) mailbox:** Choose this option if Archive Shuttle should ingest the data in to the users secondary (archive) mailbox. If this option is selected then an option of what action should be taken if a secondary mailbox does not exist can be specified. The choice of options are to either skip the ingestion or ingest into the primary mailbox instead.
- **Enable Secondary (Archive) mailbox:** Select this option to create an archive for Office 365 during provisioning.

The target link must then be selected.

- Choose the Link where you want to migrate to:

Select an entry from the drop down list that will correspond to a link in the target environment.

PST - The format must then be chosen as follows:

- PST: The output format will be PST

The target link must then be selected.

- Choose the Link where you want to migrate to:

Select an entry from the drop down list that will correspond to a link in the target environment.

Later in adding this type of mapping, it is required to choose a PST File name policy. This is selected from a drop down list.

The remaining elements in the pop-up wizard are then the same regardless of the target for the migration:

- **Workflow Policy:** The workflow policy must then be chosen, the list of workflows that are available will be determined by the chosen target. For example, if the migration is to Enterprise Vault, only those workflows applicable to an Enterprise Vault migration will be shown. Additional workflow policies can be defined. This is described in the section Workflow Policies. You are also able to select a pre-migration policy to your mappings.
- **Filter Policy:** The filter policy must then be chosen from the drop-down list. There will be a default filter which performs no filtering. Additional filter policies can be defined. This is described in the section Filter Policies.
- **Choose Container Mapping Settings:** Properties for the new mapping can then be set.
- **Migration Status: Enabled:** Enable immediate migration for the newly added mappings.

- Migration Status: Disabled. Does not immediately start migration for the newly added mappings.
- Item Gathering Status: Enabled. Enables item gathering for the newly added mappings.
- Item Gathering Status: Disabled: Does not start item gathering for the newly added mappings. There will be a summary page displayed at the end of the wizard when creating a new container mapping. This contains all of the information gathered during the wizard and should be reviewed before committing the changes.
- Run Item Gathering: Start Full or Delta item collection for selected containers. The source container will be examined and data relating to each item will be collected and added to the appropriate Archive Shuttle Link database.

Dissociate Archive User

There is an ability to disassociate the user of a source container so the container may be migrated as a leaver. To do this:

1. Select the row(s) with a source container.
2. Click **Dissociate Archive User** from the top ribbon.
3. This container will be marked as ownerless.

Manual mappings

Allows an administrator to create mappings for single containers. The source as well as target container must already exist. Performing a mapping operation in this manner is useful for containers such as Enterprise Vault Journal Archives, or where containers need to be merged to a target container.

The following basic information about each container is visible in the grid view:

- Name
- Owner Full Name
- Group
- Type
- Mailbox Type
- Link Name
- Content Count
- Content Size
- Link name
- Mapping Counts for the source and target

Actions to be performed on the Manual Mapping page

- **Add Mapping:** Adds a mapping for the selected containers. A source and target container must have been selected before this navigation bar item is available.
- **Add to Group:** One or more containers can be added to an existing or new group. This group membership can then be used for filtering and batching of users for migration.
- **Refresh:** Refreshes source and target container tables.
- **Reset:** Resets the source and target container tables back to the default list of columns, and removes any currently defined filters.

Existing mappings

Containers that have already been mapped are listed on the Existing Mappings page. This page gives the ability to Enable/Disable mappings for migration and perform other container mapping related actions.

i | **NOTE:** Additional columns are available if the 'Advanced view' is enabled (under the 'Views' tab) though this makes the user interface slower if there are a large number of mappings.

The table shows the following information for each existing container mapping:

- Container Mapping ID
- Source Container Type
- Source Container User
- Source Group
- Source Container Link Name
- Flag to indicate if the container has an owner
- Flag to indicate if the container has a mailbox
- Target Container Type
- Target Container User
- Target Group
- Target Container Link Name
- Workflow Policy
- Pre-migration
- Filter Policy
- Mapping Template
- Priority

- **Item Gathering Enabled:** Indicates whether the container mapping is enabled for item gathering (container level).
- **Migration Enabled:** Indicates whether the container mapping is enabled for migration. Archive Shuttle does NOT start to export / import when a container mapping is disabled for migration.
- **Pre-Migration Enabled**
- **Pre-Migration Finished**
- **Stage 2:** Indicates whether the container mapping has been switched to the target. If it has, Stage 2 of the Workflow has been started.
- **Stage 2 Finished:** Indicates whether Stage 2 has finished for this mapping
- **Target Switch Date Utc**

Actions to be performed on the Existing Container Mappings page

These actions can be performed on this page:

Item Collection

- **Enable:** enable mappings for item collection.
- **Disable:** disable mappings for item collection.
- **Items:** Run full or delta items for collection.
- **Shortcuts:** Run full shortcut collection.

Migrations

- **Enable:** enable mappings for migration.
- **Disable:** disable mappings for migration.
- **Delete:** Delete selection mappings.

Misc.

- **Set Priority:** Adjust the priority for the selected container mapping(s)
- **Enable Stage 2:** Switch the selected container(s) to the target
- **Set Template:** Assign container mapping template.
- **View Mappings Health:** Navigate to Container Mappings Health page.
- **Comment:** Comment on selected mapping.
- **Switch Target:** Change target of migration from a primary mailbox to a secondary or archive mailbox. Select one or more mappings to switch targets, and click **OK** to confirm.

NOTES:

- You can currently only switch targets for Office 365 migrations.

- Once a target has been switched, the action cannot be reverted.
- Leavers, virtual journals and journal transformations are not currently supported to switch targets.

Workflow

- **Change & Reset:** Change and reset workflow for selected mappings (Stage 2 or pre-migration policies)

i | **NOTE:** You cannot change or set a pre-migration workflow policy for a mapping already enabled in Stage 1.

Filter

- **Set Policy:** Set filter policy for selected workflows.

Failed Items

- **Set Threshold:** Set the failed items threshold.

Staging Area

- **Cleanup per Mapping:** Start cleanup staging area for selected mappings.

User

- **Add to Group:** One or more containers can be added to an existing or new group. This group membership can then be used for filtering and batching of users for migration.
- **Tag Assignment:** You can assign tags to workflows to identify special subset of containers.

Mapping templates

On the Mapping Templates page, you can create new bulk mapping templates that you can then apply on the Bulk Mapping page when creating a new mapping.

Below is a list of actions you can take on the Mapping Templates page:

- **Add:** Create a new mapping template.
- **Copy:** Duplicate a previously created template.
- **Edit:** Edit a template.

i | **NOTE:** Mapping templates cannot be edited if the template is applied to a mapping.

- **Delete:** Delete a template.

i | **NOTE:** If a user removes a workflow policy used by the template, the template will be marked as invalid. This will be shown with an exclamation mark in the IsValid column of the grid.

- **Validate:** This validates the links and policies associated with the template.
- **Refresh:** Refresh the guid.
- **Export:** Export the grid in a .pdf, .xls, .xlsx or .csv format.

Creating a mapping template

1. On the Mapping Templates page, click **Add**.

Enter a name and description for the template, and click **Next**.


2. Select the target container type for the template, and click **Next**.
3. Select the migration type, and click **Next**.
4. Select the target link from the drop down menu, and click **Next**.
5. Select the Workflow Policy and Ownerless Workflow Policy from the drop down menus. Optionally, add a pre-migration Workflow Policy from the drop down menu. Then click **Next**.
6. Select a filter policy from the drop down menu, then click **Next**.
7. Select a settings template from the drop down menu, then click **Next**.
8. Select a file name policy from the drop down menu, then click **Next**.
9. Choose whether to enable or disable item collection and migration to begin immediately, and click **Next**.
10. Enter a numeric value to set priority, and click **Next**.
11. Review the summary of the bulk mapping template, and click **Confirm**.

Applying a mapping template

1. Go to the Bulk Mappings page, and click on the container to apply the template to. Then click **Add Mappings**.
2. Select the template from the drop down menu, then wait for a few moments as the template is applied. Click **Next**.
3. Review the summary, then click **Next**. The template is now applied to the container mapping.

Setting templates

Using the Setting Templates feature, you can create a container mapping template where you can set the total mailbox item count limit or item size limit. Then, when creating a mapping, you can assign a template and override settings for the Office 365 Module in SysConfig.



NOTE: Currently the Mapping Templates feature is available for Virtual Journal only.

Creating a setting template

Create a setting template before creating a mapping by following these steps:

1. Select **Configuration**, then **Templates**, then **Mapping Templates** from the left menu bar.
2. Click **Add**.
3. Enter a name/description and desired limits using the Mapping Template Configuration window (shown below).
4. Save the template.

Mapping Template Configuration

Template Details

Name:

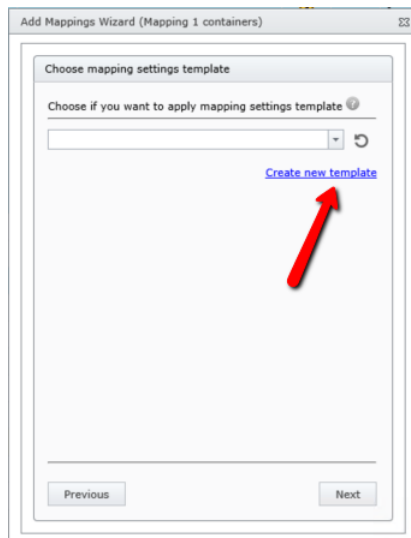
Description:

Settings

Virtual Journal Total Mailbox Item Count Limit

Virtual Journal Total Mailbox Item Size Limit (MB)

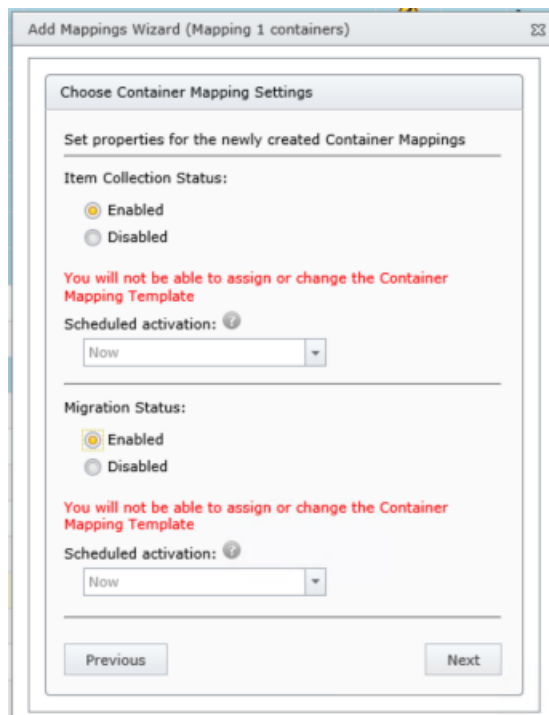
Or, if you don't create a mapping template in advance, you can create one using the **Create new Template** link available on the **Add Mappings Wizard** (shown below).



Assigning a setting template

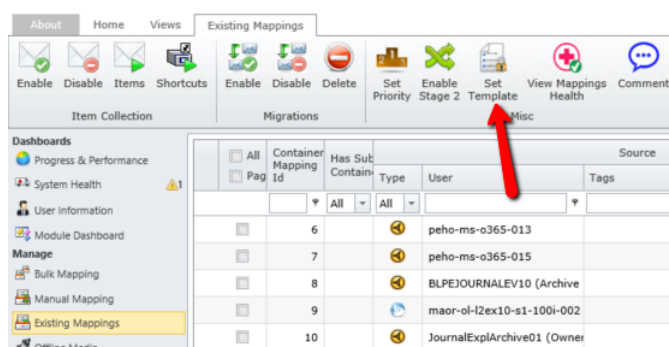
For new mappings, the Add Mappings Wizard prompts you to choose the mapping template you want to assign (see image above).

If you don't assign a template (or an empty one is assigned) during mapping and item collection/migration is enabled, you'll see this warning message (in red) within the wizard:



If you opt to not choose a mapping template, the default settings configured in SysConfig – Office 365 Module are applied.

If you haven't enabled item collection/migration for the mapping, you can change or assign a mapping template on **Existing Mapping** page using the **Set Template** button shown below.



Offline media

The Offline Media configuration helps support the process whereby the source and target staging areas need to be different.

This is currently supported for:

- Enterprise Vault to Enterprise Vault, Office 365 or Exchange
- EAS to EV, Office 365 or Exchange
- Dell Archive Manager to EV, Office 365 or Exchange

Some examples of when this might occur are:

- Source and target environments are in different non-trusted forests
- Slow WAN links mean that the bulk of the data will be migrated via 'disk shipping'

There are a number of steps to follow in order to successfully use the Offline Media feature in Archive Shuttle, and these are described below:

Create mapping set

On the Offline Media a new source and target mapping set can be created by clicking on 'Create' in the Actions bar. This is the first step and in the mapping the source link is specified along with the target link, and name can be given to the set.

Add Export and Import Locations

With the Mapping Set selected multiple export and import paths can then be defined. Export paths can be configured to be one of three states:

Status	Description
Open	Data can be written to this path
Closed	Data will not be written to this path

Status	Description
Ready	Data path is on standby

If an open export path becomes full, then the next ready container will be used.

Configure the links to use offline media

i | NOTE: The import path is restricted to UNC paths only. Local disk paths are not supported.

The source and target links need to have 'Offline Mode' enabled. This is performed by selecting the links on the Links page in the Admin Interface and selected 'Enable' in the 'Offline Mode' of the Actions bar.

When this has been enabled the Staging Area Path for those links will no longer be used. Instead, the paths used are those defined on the 'Offline Media' page.

Creating container mappings

The container mappings are configured in the same way as ordinary mappings.

Export data

Data will then begin to be collected and exported from the source containers and placed on the staging areas defined in the Offline Media set.

Ship the data

Periodically during the export, or at the end of the export, the exported data then needs to be 'shipped' to the target import path. This can be simply a case of copying the data manually or by shipping the disk and performing the necessary actions to make it available on the given name in the target environment.

If the data is copied, ensure that the very top level (below the share name itself) is copied to the target.

Enable scanning of the import location

When the data is ready to be imported the Import Path has to be enabled for scanning on the "Offline Media" page.

When scanning is enabled, the location is scanned by the import module each hour for new data. If found, the data is then made available for the import module to process.

Once the scanning has occurred the data will be made available for the Import Module.

Import data

Once the target Import Path has been scanned and data recorded in the Item Database as being available for import, the import module will process the data in the container mapping as normal.

Pre-migration

The pre-migration page gives an overview on pre-migration workflow policies that have been created. Pre-migration policies that allow you to automate a series of commands prior to beginning a migration on a per mapping basis.

i | **NOTE:** You cannot change or set a pre-migration workflow policy for a mapping already enabled in Stage 1.

Pre-migration workflow commands

There are several pre-migration workflow commands to choose from. These are:

- **CloseDACase** - Collect cases from EV Discovery Accelerator
- **DellArchiveManagerDisableArchive** - Disable Dell Archive Manager archive for archiving
- **DisableMailbox** - Disable mailbox for Enterprise Vault archiving
- **DisableOrphanedArchive** - runs SQL commands to set an Enterprise Vault archive to 'closed' and the archiving state (MbxArchivingState) to 'disabled'
- **EasDisableArchive** - Disable EAS archive for archiving
- **Office365RevertPendingShortcuts** - Revert 'Pending' Enterprise Vault shortcuts back to the origin message class in the Office 365 mailbox
- **PowerShellScriptExecutionRunScript** - Execute custom PowerShell script
- **RestoreSourceOneForeignShortcuts** - Restores SourceOne foreign shortcuts
- **RevertPendingShortcuts** - Revert 'Pending' Enterprise Vault shortcuts back to the origin message class in the Exchange on-premises mailbox
- **SendEmailNotification** - Send an email notification

User interface

This is where your pre-migration workflow policies will be visible.

Source	Container Name	Has Owner	User Name	Tags	Type	Link	Workflow State	Current Status	Command	Status	Error	Date	Next Status	Command	Finished	Apply Filter
<input checked="" type="checkbox"/>	peho-qaex13-wide-2	<input checked="" type="checkbox"/>	peho-qaex13-wide-2		All		QA EV10 01 VSG1 V Enabled	SendEmailNotificati	Failure		SMTP Server is emp	13/08/2021 13:16:5				All

Actions you can take on the pre-migration page are:

- **Add to Group:** One or more containers can be added to an existing or new group. This group membership can then be used for filtering and batching of users for migration.
- **Tag Assignment:** You can assign tags to workflows to identify special subset of containers.
- **Retry:** Retries the selected processing/failed workflow status and reprocesses them.
- **Skip:** Skips the workflow.
- **Pause:** Pauses the workflow for the mapping.
- **Resume:** Resumes the workflow for the selected mapping (needs to be in a Paused state).
- **Change & Reset Policy:** You can change the workflow policy for the mapping and reset the workflow immediately.
- **Refresh:** Refresh the current screen.
- **Auto refresh:** Choose time parameters to refresh the current screen automatically.
- **Columns:** Choose which columns to display in the grid.
- **Reset:** Reset the grid to its default state.
- **Load:** Load grid layout.
- **Save:** Save grid layout. Click the **Shared** box to share the layout to other users.
- **Export:** Export the grid in a .pdf, .xls, .xlsx or .csv format.

The table grid shows the following information for each workflow policy:

- Container name
- Has Owner

- User Name
- Tags
- Link
- Workflow State
 - Manually Switched: state assigned manually in the user interface
 - Auto Switched: state assigned automatically via CheckPremigrationAutoSwitch scheduled task
- Command(s)
- Command status
- Error message

Create a pre-migration workflow policy

Follow the instructions below to create a pre-migration workflow policy. Go to the Workflow Policies page to find out more about workflow policies.

1. Go to the **Advanced Configuration** tab, and select **Workflow Policies**.
2. Select **Add**, then click **Pre-migration**. Then select **Create**.
3. Enter a name and, optionally, a description for your workflow policy.
4. Select the container types to apply the workflow policy to.
5. Click on the commands on the right hand side in the order you would like the workflow to process.
 - a. In **Edit options**, you can set retry parameters if the command fails.
6. Once you have created your policy, click **Save**. This will appear in the list of workflow policies.

Apply a pre-migration workflow to a bulk mapping

1. Go to the **Manage** tab, then select **Bulk Mapping**.
2. Select the container to add the mapping to, then select **Add Mappings**.
3. Create the mapping as usual. Go to the Bulk mappings page for more information.
4. On the **Workflow Policy** section, under **Pre-migration**, select your pre-migration policy from the drop down list.



NOTES:

- By using a pre-migration policy, the option to enable a migration is disabled by default.

- If you would like to apply a pre-migration policy to your mapping, but have not created one yet, in **Container Mapping Settings**, ensure that Item **Collection** and **Migration** are both set to **Disabled**.
- You will need to manually enable stage 1 after pre-migration is complete. If you would like to automate this process, go to **Advanced Configuration > System Configuration**, then check the **Autoenable Stage 1** box.
- Pre-migration workflows can also be enabled via PowerShell. Go to the PowerShell guide for more.

Stage 1 (Sync data)

This page shows an overview of all containers enabled for migration, and their synchronization status. It also allows data to be exported to a variety of formats for reporting on the progress of the migration.



NOTE: Additional columns are available if the **Advanced view** is enabled (under the **Views** tab) though this makes the user interface slower if there are a large number of mappings.

The following basic information about each container is visible in the grid view:

- Full Name
- Tags
- Has Owner
- Routed
- Exported
- Retryable Errors
- Permanent Errors
- Total Errors
- Exported Percentage
- Imported
- Retryable Errors
- Permanent Errors
- Total Errors
- Imported Percentage
- Priority
- Item Gathering Enabled
- Migration Enabled

- Stage 2 Enabled
- Stage 2 Finished
- Failed Item Threshold
- Ignore Failed Item Threshold
- Needs Provisioning

Actions to be performed on the Migration Status page

These actions can be performed on the **Actions** page:

Item Collection

- **Enable:** Enable selected mappings for item collection.
- **Disable:** Disable selected mappings for item collection.
- **Items:** Run full or delta items for collection.
- **View Details:** Show item collection details for selected archive.
- **Shortcuts:** Run full shortcut collection.

Migrations

Enable: Enable selected mappings for migration.

Disable: Disable selected mappings for migration.

Journal Transformation Migrations

- **Enable:** Enable selected Journal Transformation mappings for migration.
- **Disable:** Disable selected Journal Transformation mappings for migration.

These actions can be performed on the **Stage 1 (Sync Data)** page:

- **Set Threshold:** Sets the failed item threshold for the selected mapping or mappings
- **Retry Failed Items:** Retry any failed items

i | **NOTE:** A hanging export or import means that it has been running for more than 1 hour.

Retrying permanently failed items will reset the error count for those items and retry them.

i | **NOTE:** Retrying failed items from the UI (failed items page/Stage 1 page) will work only for mappings where:

- Stage 2 is not enabled or is enabled but not finished

- Stage 2 is enabled and Workflow Policy has command WaitForImportFinished running (stating 'success' or 'skipped')

If Workflow Policy doesn't have this command, then only the first rule applies.

- **Re-export Items:** Items can be re-exported. It is possible to re-export all items, just those that failed to export or just those that failed to import.
- **Enable Stage 2:** Switch the selected container(s) to the target
- **Add to Group:** One or more containers can be added to an existing or new group. This group membership can then be used for filtering and batching of users for migration.
- **Set Priority:** Allow the migration priority to be set
- **View Mappings Health:** Switches to a status page which can be used to check mappings. See the section in this document called 'Mappings Health' for more information.
- **Comment:** Set a comment of up to 256 characters on one or multiple mappings.
- **Switch Target:** Change target of migration from a primary mailbox to a secondary or archive mailbox. Select one or more mappings to switch targets, and click **OK** to confirm.



NOTES:

- You can currently only switch targets for Office 365 migrations.
- Once a target has been switched, the action cannot be reverted.
- Leavers, virtual journals and journal transformations are not currently supported to switch targets.

- **Refresh:** Refreshes the tables data.
- **Columns:** Allows the selection of additional data columns, which might then be used to select groups of containers.
- **Reset:** Resets the grid to the default columns, and removes any filters.
- **Load:** Loads a saved filter/page-layout.
- **Save:** Saves the current filter/page-layout. Click the **Shared** box to share the layout to other users.
- **Export to .pdf, .xls, .xlsx, .csv:** Allow data to be exported to several formats.

Stage 2 (Switch user)

This page shows an overview of all containers that have been switched to target and shows the status in the Stage 2 workflow.

i | **NOTE:** Additional columns are available if the 'Advanced view' is enabled (under the 'Views' tab) though this makes the user interface slower if there are a large number of mappings.

The following basic information about each container is visible in the grid view:

- Container name
- User name (source and target)
- Tags (source and target)
- Type (source and target)
- Workflow State
- Command
- Status
- Error
- Next Command
- Finished

Actions to be performed on the Workflow Status page

These actions can be performed on this page:

Refresh: Refreshes the tables data.

Columns: Allows the selection of additional data columns, which might then be used to select groups of containers.

Reset: Resets the grid to the default columns, and removes any filters.

Load: Loads a saved filter/page-layout.

Save: Saves the current filter/page-layout. Click the **Shared** box to share the layout to other users.

Add to Group: One or more containers can be added to an existing or new group. This group membership can then be used for filtering and batching of users for migration.

Reset Workflow Status: Can be used to retry the current command.

Skip: Skip the current command in the workflow.

Pause: Pause the workflow.

Resume: Continue from a paused workflow.

Change Policy: Allows a new policy to be chosen and the workflow restarted with that policy.

Export to .pdf, .xls, .xlsx, .csv: Allow data to be exported to several formats.

Configuration

User groups

This is an overview page of the groups that have been defined in Archive Shuttle. Groups are used to apply a tag to users so that actions relating to their migration, or progress monitoring can be performed easily in the Admin Interface.

As an example, a group might be defined as “Migration Test Users”. There are several places in the Admin Interface (listed below) where users can then be added to this group. Groups can also be managed from this page in the Admin Interface.

Users/Groups can also be imported from a .csv file, if necessary.

Actions to be performed on the User Groups page

Add: Create a new group.

Edit: The name of a group can be modified.

Delete: A User Group can be deleted.

Unassign Users: Remove users from a group.

Assign Users (.csv Import): Imports users from a .csv file to a specific group. Click here to view instructions on creating a .csv file.

i NOTE: .csv files that can be selected for import should contain either a list of User SIDs to add to a particular group, sAMAccountNames, or Container Names. Go to this page for more on creating a .csv file configurable for Archive Shuttle.

Refresh: Refresh the list of users and groups.

Reset: Reset the grid to the default view.

There are several places where containers/users can be added to groups:

- User Dashboard
- Bulk Mapping
- Manual Mapping
- Existing Mappings

- Stage 1 (Sync Data)
- Stage 2 (Switch User)

Once added a group, all these same pages in the Admin Interface allow the data to be filtered and grouped by the group name.

i | **NOTE:** A container/user can only belong to one group.

Tags and Tag Management

Tags are an organizational mechanism that can be used on any source or target container to identify special subset of the containers within a migration.

Tags can be used to:

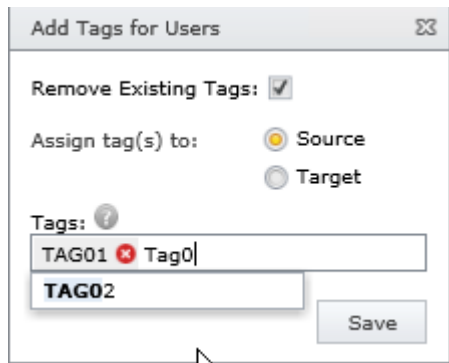
- Flag users under legal consideration
- Group users into migration waves
- Identify problematic containers
- Identify VIPs that require special consideration
- Identify containers that are out of scope

Tags can be created, modified, and assigned. Then, tags can be filtered. Tags can be created and applied to any of these pages:

- User Information
- Bulk Mapping
- Existing Mapping
- Stage 1
- Stage 2

i | **NOTE:** Tags can be used with any type of container (including ownerless ones). Additionally, tags can be assigned for a mapping's source and target containers.

When you select objects and click the **Tag Assignment** button, you can create, remove, or assign tags using this window:



Tag Management Page

The **Configuration > Tag Management** page shows a list of current tags and the number of containers assigned each tag. A tag can be expanded to review the containers that have the tag.

On this page, you can:

Add: Create a new container tag.

Edit: Modify a container tag.

Delete: Remove a container tag.

Unassign containers: Unassign containers.

Assign Users (.csv Import): An imported .csv file can be used to assign one or more tags to user group containers. Click [here](#) to view instructions on creating a .csv file. The supported fields for matching the containers are:

- UserSid
- SaMAccountName
- Container Name

The format of the .csv file should be one of the following:

- SID,tagname
- samAccountName,tagname
- ContainerName,tagname

You can use the separator drop down box to change the punctuation used to separate each property (comma, semi colon, colon, vertical line). For example, using the examples above:

- SID;tagname
- samAccountName;tagname
- ContainerName;tagname

Multiple tags can be assigned to a container, for example:

- samAccountName,tag1,tag2,tag3

Assign User Groups: You can assign a tag to all members of a group using the Assign User Groups button. This allows you to select a group and assign a specific tag to all containers in that group. This can be useful when you are transitioning from using groups, to using tags.

Refresh: Update the grid.

Columns: Modify the columns in the grid.

Reset: Reset the grid to its default state.

Load: Load a previously saved .csv file for the grid.

Save: Save the grid layout as a .csv file. Click the **Shared** box to share the layout to other users.

Creating a .csv file

To create a .csv file compatible for use with Archive Shuttle, follow these steps in Microsoft Excel. This is to be used when adding multiple tags to a value.

1. Add each value into its own cell. Do not separate each value with separators such as commas, semi colons or colons, unless necessary (see below). Use the table below as an example:

Column A	Column B	Column C	Column D
ColumnName	tag1	tag2	tag3

The output of this would look like the following, using commas as the separator, and Container Name as the importing field, as defined when importing into Archive Shuttle:

ContainerName,tag1,tag2,tag3

2. Save under **CSV UTF-8 (Comma delimited) (*.csv)**.

3. Import the .csv file into the chosen grid.

If a comma is necessary in a cell, for example when separating a first name and a last name, then use the instructions below:

1. Add each value into its own cell. Use the table below as an example:

Column A	Column B	Column C	Column D
FirstName,LastName	tag1	tag2	tag3

The output of this would look like the following, using semi colons as the separator, and Container Name as the importing field, as defined when importing into Archive Shuttle:

"FirstName,LastName";tag1;tag2;tag3

2. Save under **CSV UTF-8 (Comma delimited) (*.csv)**.

3. Import the .csv file into the chosen grid.

Enterprise Vault environment

This is an overview page of all the Enterprise Vault (EV) environments that are needed for the migration project. New Enterprise Vault Directories can be added, and Archive Shuttle collects information about this Enterprise Vault directory through the assigned Archive Shuttle Enterprise Vault Collector module.

For every Enterprise Vault Directory, the table lists the following information:

- Vault Store Name
- Number of Archive in this Vault Store
- Total Item Count in this Vault Store
- Total Item Size in this Vault Store
- Total Shortcut Count in the archives' associated mailbox
- Total Legal Hold Count
- Archive Gathering Enabled
- Backup Mode enabled



NOTE: The data values will only be shown on screen when “Archive Gathering” has been enabled, and an EV Collector module has been linked (see below)

Actions to be performed on the EV Environment page

Add EV Directory: Adds an Enterprise Vault Directory to Archive Shuttle. The following information is needed in order to perform this action:

- **Module to associate:** Archive Shuttle Enterprise Vault Collector module which should be “responsible” for collecting information about this Enterprise Vault Directory.
- **Display Name:** Friendly name for this Enterprise Vault Directory. This is only a display name and can be used to identify this Enterprise Vault Directory.
- **EV SQL Server:** SQL Server name / instance where this Enterprise Vault Directory hosts its Directory database. This can be a hostname, a Fully Qualified Domain Name, an IP Address or all of the previous with a SQL Named Instance provided:
 - Example: mysqlserver.mydomain.com
 - mysqlserver
 - 192.168.0.1
 - mysqlserver.mydomain.com\instance
 - mysqlserver\instance
 - 192.168.0.1\instance

i NOTE: To configure the SQL server to use a port other than 1433, use a comma to define the port. Look at an example below to define a server using port 53179:

Without configuration: 192.168.0.1\instance

With configuration: 192.168.0.1\instance,53179

- **EV SQL Database Name:** Database name. Defaults to EnterpriseVaultDirectory. This entry is for display purposes only.

Sync all EV Environments: Queues an update request for all Enterprise Vault Directories by their respective Archive Shuttle Enterprise Vault Collector Modules.

Sync Active Directory: Queues an update request for all enabled Active Directory Domains by their respective Archive Shuttle Active Directory Collector modules.

Archive Gathering

Enable: Enable Archive Gathering for the selected Vault Store(s).

Disable: Disable Archive Gathering for the selected Vault Store(s).

Run Now: Queue an Archive Gathering request for the selected Vault Store(s).

Refresh: Refresh the information on the displayed EV Environment table.

Enterprise Vault Discovery Accelerator environment

This is an overview page of Enterprise Vault Discovery Accelerator (EV DA) customers that are needed for the migration project. New EV DA customer databases can be added here, and Archive Shuttle collects information from the environment to get information about the cases associated with that customer database through the assigned Archive Shuttle Enterprise Vault Collector module.

i NOTES:

- Enterprise Vault export and collector modules must run under the user account which has access to the Enterprise Vault and Enterprise Vault Discovery Accelerator server.
- An Enterprise Vault Discovery Accelerator environment requires dependency on an existing Enterprise Vault environment.

EV DA cases will be shown on the bulk mappings, manual mappings and existing mappings pages as an archive type. Collection of items from EV DA cases are processed automatically after a mapping is created in the Bulk Mapping wizard (when collection of items is enabled.)

All target containers are supported for migration (Office 365, Exchange, Enterprise Vault, PST and Storage Import).

For more on Enterprise Vault Discovery Accelerator, [click here](#).

User interface

Display Name	SQL Server	Database Name	Number Of Cases	Total Item Count	Total Item Size
EVD123			10	4 814 995	141,32 GiB

For every Enterprise Vault Discovery Accelerator directory, the table lists the following information:

- Display Name
- SQL Server
- Database name
- Number of Cases
- Total Item Count
- Total Item Size

The actions you can take on this page are:

- **Add:** Add an EV DA environment.
- **Run Now:** Collect cases immediately.
- **Refresh:** Refresh table to show up to date statistics.
- **Reset:** Reset the grid to default layout.
- **Columns:** Choose which columns to display.
- **Load:** Load previously saved grid layout.
- **Save:** Save current grid layout. Click the **Shared** box to share the layout to other users.
- **Export:** Export grid into a .pdf, .xls, .xlsx or .csv format.

How to add an Enterprise Vault Discovery Accelerator environment

1. On the EV DA Environment page, click **Add**.
2. Select the module to associate with this environment. This module is responsible in collecting information on the EV DA directory.

3. Enter the EV DA SQL server. This is the SQL server name/instance where this directory hosts its database. This can be a hostname, a Fully Qualified Domain Name, an IP Address or all of the previous with a SQL Named Instance provided. Examples include:
 - a. mysqlserver.mydomain.com
 - b. mysqlserver
 - c. 192.168.0.1
 - d. mysqlserver.mydomain.com\instance
 - e. mysqlserver\instance
 - f. 192.168.0.1\instance
4. Enter the EV DA customer database SQL database name. This is the name of the database associated with your EV DA customer in the Discover Accelerator user interface.
5. Enter the Link SQL server. This should be filled by default.
6. Click **Add**.

Repeat these steps for each EV DA customer you wish to add to Archive Shuttle.

Additional details

- Migration from EV DA cases supports the following filters:
 - Path
 - Item Date
 - Item Size
- Journal Transformation is currently not supported.

Enterprise Vault retention mappings

This page shows all Enterprise Vault Retention Categories that have been mapped, and allows an administrator to add or delete existing mappings.

A Retention Category mapping is necessary in cross-Enterprise Vault Site/Directory migration scenarios. This is required so that Archive Shuttle knows which Retention Category it should apply to the target item, based on the retention category of the source item.

A Retention Category mapping can also be used to change the Retention Category in intra-Enterprise Vault Site/Directory migration scenarios. This is useful if it is required to consolidate Retention Categories.

i | **NOTE:** Without Retention Category mappings in place, Archive Shuttle will export data from Enterprise Vault, but not import any items into Enterprise Vault.

Retention related mappings can be logged when interacted with. The logs will also show the retention label prior to its ingestion. This setting needs to be enabled, and the source container needs to be Enterprise Vault to be able to view the core logs.

Actions to be performed on the EV Retention Mappings page

Create Mapping: Create a new Retention Category Mapping. The following information is required, and can be selected from drop-down lists when creating the mapping:

- Source site
- Source Retention Category
- Target Site
- Target Retention Category


Create Multiple Mappings: This can be used to add multiple retention category mappings at one time.

Add Intrasite Migration Mappings: This can be used in intra-Enterprise Vault Site scenarios. It maps every Retention Category to itself, thus keeping the existing Retention Category during migration.

Delete Mapping: Deletes the selected Retention Category mapping(s).

Edit Mapping: Allows a mapping to be modified.

Refresh: Refresh the information on the displayed Retention Category table.

 **NOTE:** Unmapped retention categories can be seen on the System Health page.

SourceOne/EMX environment

This is an overview page of all the SourceOne/EmailXtender (EMX) environments that are needed for the migration project. New SourceOne/EMX environments can be added and ran here.

User interface

For every SourceOne/EMX environment, the table lists the following information:

- Environment Name
- SQL Server
- Database name
- Number of Archives
- Total Item Count
- Base Paths

The actions you can take on this page are:

- **Add:** Add a SourceOne/EMX environment.
- **Enable Organization:** Enable Exchange organization for archive collection.
- **Disable Organization:** Disable Exchange organization for archive collection.
- **Run Now:** Run archive collection now.
- **Run Organization Collection:** Run organization collection now.
- **SMTP Settings:** Set SMTP settings for archive collection. Click the checkbox to include SMTP archives, and add domains into the text box to be included into archive collection. This can be ran from the user interface, or is automatically ran once per day by using command **S1CollectSMTPArchives** in PowerShell.



NOTES:

- SourceOne does not have archives.
 - Archive Shuttle determines message ownership and associates the first instance of a message to each sender and recipient Archive Shuttle identifies, and the routes a copy of that message.
 - The command **S1CollectSMTPArchives** runs once per day and can be also ran from the user interface, on the SourceOne environment page.
- **Refresh:** Reset the grid to default layout.
 - **Columns:** Choose which columns to display.
 - **Load:** Load previously saved grid layout.
 - **Save:** Save current grid layout. Click the **Shared** box to share the layout to other users.
 - **Export:** Export grid into a .pdf, .xls, .xlsx or .csv format.

How to add an SourceOne/EMX environment


1. On the Source/EMX Environment page, click **Add**.
2. Select the module to associate with this environment. This module is responsible in collecting information on the SourceOne/EMX environment.
3. Enter the SourceOne/EMX SQL server. This is the SQL server name/instance where this directory hosts its database. This can be a hostname, a Fully Qualified Domain Name, an IP Address or all of the previous with a SQL Named Instance provided. Examples include:
 - a. mysqlserver.mydomain.com
 - b. mysqlserver
 - c. 192.168.0.1
 - d. mysqlserver.mydomain.com\instance

- e. mysqlserver\instance
 - f. 192.168.0.1\instance
4. Enter the SourceOne/EMX customer database SQL database name. This is the name of the database associated with your SourceOne/EMX customer in the SourceOne/EMX user interface.
 5. Enter the BasePath, and select whether to create an item database and/or enable archive collection.
 6. Enter the Link SQL server. This should be filled by default.
 7. Click **Add**.

Repeat these steps for each SourceOne/EMX customer you wish to add to Archive Shuttle.

Workflows for SourceOne SMTP archives

Below are workflows for SourceOne SMTP archives as a source into all supported targets in Stage 2.

 **NOTE:** These workflows do not utilize shortcut deletion.

- SourceOne SMTP to Enterprise Vault
- SourceOne SMTP to Exchange
- SourceOne SMTP to Office 365
- SourceOne SMTP to Office 365 Leavers
- SourceOne SMTP to Proofpoint
- SourceOne SMTP to PST
- SourceOne SMTP to Storage Import
- SourceOne SMTP to Storage Import (ownerless)

Links

The page shows all Links in the environment (e.g., Vault Stores, Exchange Mailbox databases) that were discovered through the Active Directory Collector Modules, and the Enterprise Vault Collector Modules. It also shows links that were created manually (e.g. Native Format Links, and Office 365 links).

For each link, you can see the following information:

Type: Enterprise Vault, Exchange, Office 365, or PST.

Name: Name of the link. This is usually the Vault Store name or the Exchange database name. For links that were created manually it will be the name that was chosen.

Computer Name: The name of the server hosting the link. This is usually the Enterprise Vault server or the Exchange server.

AS DB: If an Archive Shuttle Item database has been created for this Link. A green and white check mark will be displayed if the database exists.

Failed Item Threshold: The failed item threshold that has been configured for this link.

Number of containers: Container Count for this Link

Associated Modules: This grid will show the module to link mappings that have been defined.

Staging Area Path: Path used for export or import for this link

i | **NOTE:** The Archive Shuttle Active Directory Collector, and Enterprise Vault Collector may have already scanned Active Directory and pre-populated some of the information on this page.

Actions to be performed on the Links page

Create PST Link: Provides the ability to create a new PST link. It is possible to choose a name for the link, and to specify the PST Output Path where final PST files will be placed.

Create Office 365 Link: Provides the ability to create a new Office 365 link. It is possible to choose a name for the link.

Create Proofpoint Link: Provides the ability to create a new Proofpoint link. It is possible to choose a name for the link, and to specify the output folder.

Create Database: Create an Item database for the selected Link.

i | **NOTE:** The source Link needs to be of type "Enterprise Vault". An Item database should be created for each Enterprise Vault source Vault Store.

You need to provide the following information:

Link SQL Server: The SQL Server name / instance where you want to create the Link database. This can be a hostname, a Fully Qualified Domain Name, an IP Address or all of the previous with a SQL Named Instance provided.

Example: mysqlserver.mydomain.com

mysqlserver

192.168.0.1

mysqlserver.mydomain.com\instance

mysqlserver\instance


192.168.0.1\instance

- SQL Database Name

Default: ArchiveShuttleItem_. Read only.

Link

Map Modules: Provides the ability to allocate modules to this link.

 **NOTE:** Multiple links can be selected and modules allocated to them.

Threshold: Define an error threshold for the migration. If the count of failed messages is below this threshold, the migration is still considered successful, and Stage 2 will continue at the appropriate time.

Export Path: This option should be configured when reviewing failed items, and moving them to a secondary location. This location is where those items will be moved to from the regular staging area.

Staging Area Path: Defines a path to which items should be exported, and items ingested. Multiple links can share the same export/import path, or they can be separate in order to distribute the exported data ready for ingestion. This path can be overridden on individual links if required.

PST Output Path: Allows the configuration of a PST Output Path for links where the type is PST. The output Path can be an UNC Path or an Azure Blob Storage Account.

Temporary Path: Allows the configuration of a temporary path to store PST files while they are processed on PST as a Target. It is recommended that this is FAST disk (Preferrably SSD)

Cleanup Staging Area: Issues a command to the Archive Shuttle modules to clean up the staging area of already imported files. The EVExport, and all Import Modules will receive the command.

Enable/Disable/Run-Now Archive Collection: Allows quick access to the functionality provided on the EV Environment page. These buttons allow for archive gathering to be enabled or disabled for a link, and to issue the command to perform archive gathering now.

Enable/Disable Offline Mode: When a link is enabled for offline mode, it will perform migration based on the Offline Media that has been configured for it. The staging area, if shown on this page, will not be used.

Watermark: The high and low watermarks can be configured per link. More information on the functionality of the high and low watermarks is available in the section relating to System Health.

Office 365

Create Office 365 Link: Provides the ability to create a new Office 365 link. It is possible to choose a name for the link.

Sync Mailboxes: Synchronizes mailbox information from Office 365 to Archive Shuttle.

Sync cloud-Only Mailboxes: Synchronizes cloud only mailbox information from Office 365 to Archive Shuttle.

Sync Mailboxes Filter: Specify filters for Office 365 mailbox sync commands. Properties include:

- Display Name
- Country
- Department
- User Principal Name



NOTE: The filter is applied over full mailbox search results. The module then returns only those mailboxes that fit the filter configuration.

PST

Create PST Link: Provides the ability to create a new PST link. It is possible to choose a name for the link, and to specify the PST Output Path where final PST files will be placed.

Set Rollover Threshold: Specify the size in MB, or number of items that PST files should be rolled over.

PST Target: Set PST target path for PST links.

PST Source: Set PST source folders for PST links.

File Name Policy: Set file name policy for PST files.

Output Watermark: Set watermarks for PST output paths.

Temporary Path: Set PST temporary path for PST links.

Proofpoint

Create Proofpoint Link: Provides the ability to create a new Proofpoint link. It is possible to choose a name for the link, and to specify the output folder.

Edit Proofpoint Link: Modify the Proofpoint link.

Journal Transformation

Create Journal Transformation Link: Provides the ability to create a new Journal Transformation link. It is possible to choose a name for the link, and to specify the output folder.

Edit Journal Transformation Link: Modify the Journal Transformation link.

Storage Import

Create Storage Import Link: Provides the ability to create a new Storage Import link. It is possible to choose a name for the link, and to specify the output folder.

Edit Storage Import Link: Modify the Storage Import link.

Grid

Refresh: Refresh the contents of the grid.

Columns: Allows the selection of additional data columns that can be added to the grid.

Reset: Reset the columns and grid layout to their defaults.

Load: Load a previously saved grid format.

Save: Save the grid format. Click the **Shared** box to share the layout to other users.

Export: Export the grid in a .pdf, .xls, .xlsx or .csv format.

Advanced configuration

Naming policies

File Name Policy

The File Name Policies page in the Admin Interface allows an administrator to customize file name policies to be used in migrations where the target is PST. Tokens are used to construct the file name of the PST file when it is renamed/moved to the PST Output Path. The possible tokens can be seen above.

When creating or editing a policy a live example of a file name will be shown to help with the policy design.

Folder Name Policy

The Folder Name Policies page in the Admin Interface allows an administrator to customize folder name policies to be used in a journal archive migration. Tokens are used to construct the folder name of the place in the mailbox where items are migrated to. The possible tokens are:

Token	Description
Original SMTP address	The SMTP address of the original recipient
Mapped SMTP address	The SMTP address of the mailbox where data is being migrated to

There is also a checkbox to indicate whether the purges folder should still be used as a root, and then the constructed folder name is used as a subfolder from that.

When creating or editing a policy a live example of a name will be shown to help with the policy design.

Archive Name Policy

The Archive Name Policies page in the Admin Interface allows an administrator to customize archive name policies to be used when the RenameSourceArchive command runs in the Stage 2 part of a migration. Tokens are used to construct the name of the archive. The possible tokens are:

Token	Description
firstname	First name of the owning user
lastname	Last name of the owning user
fullname	Full name of the owning user
upn	User principal name of the owning user
SMTP address	The primary SMTP address associated with the owner
SAM Account Name	The SAM Account Name associated with the owner
Container Mapping ID	The Archive Shuttle container mapping ID
Archive Name	The name of the Enterprise Vault archive.

When creating or editing a policy a live example of will be shown to help with the policy design.

The policy can be used when editing the RenameSourceArchive command as part of a Workflow Policy.

Considerations when migrating to the recoverable items folder

The default location for active mailboxes is the following folder location **Recoverable Items\Purges** – this is a hidden folder from the end user.

i | NOTES:

- Items in the Recoverable Items folder aren't calculated toward the user's mailbox quota. In Exchange Online, the Recoverable Items folder has its own quota
- Recoverable Items\Purges folder has its own quota – soft limit 20GB, hard limit 30GB. If litigation hold is enabled the quota is automatically increased from 30GB to 100GB for the primary mailbox and unlimited for the EOA (auto expanding feature is required)
- Currently you can only ingest into the Primary Mailbox

The default location for leaver mailboxes is the folder **Journal Items**.

Changing the target folder

If you are planning to ingest into another folder, that is different from the default system policy Recoverable Items\Purges & Journal Items folder. A Folder Name Policy will be required.

Navigate to **Configuration**, then **Naming Policies** then **Folder Name Policy**.

Workflow policies

The Workflow Policies page in the Admin Interface allows an administrator to customize Workflows and create new ones. These can either be new, original workflows or copies of existing workflows with some amendments. You can create either a Stage 2 policy, or a Pre-migration policy.

Creating a new policy

A new policy can be created by clicking on “New” the Actions Bar, choosing either a Stage 2 or Pre-migration policy, and entering a name for the policy.

New policies can contain any of the commands list on the right hand side of the Workflow Policies screen.

When creating a new policy it is required to indicate the types of mapping that can use this particular workflow by selecting the check boxes at the top right of the Workflow Policies screen.

For information on the detail relating to each command, see below.

i | **NOTE:** Remember to save before leaving the screen in order to commit the changes to the Workflow Policy to the Archive Shuttle database.

Copying a policy

When a policy is being viewed it is sometimes desirable to copy the policy and then make some small alterations to it. This is achieved by using the “Copy” button in the Actions Bar. Once a copy of the policy has been made, it can be given a new name, and saved to the Archive Shuttle database.

Editing a policy

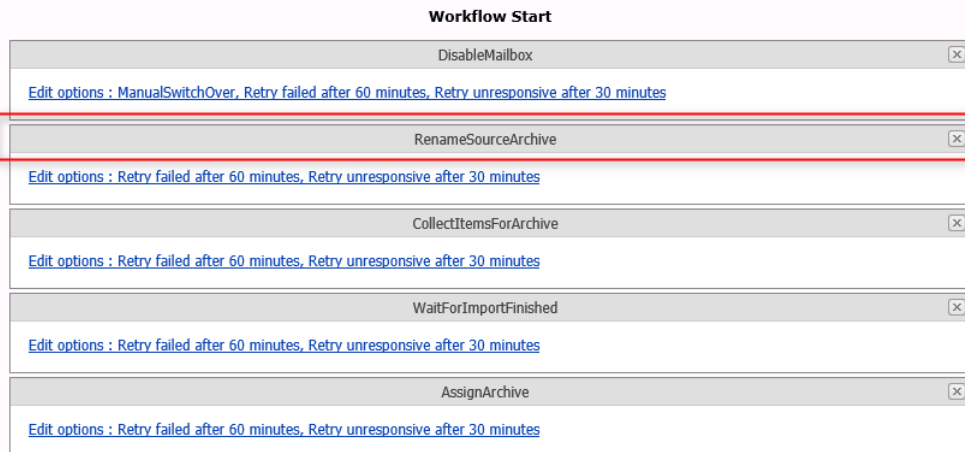
A policy can be edited by clicking on “Edit” in the Actions Bar, and selecting an existing policy from the list. If appropriate, the types of mapping that the policy can be used with can be adjusted. All changes are only committed to the Archive Shuttle database when the “Save” button is clicked.

Adding a new command

A new command can be added to the bottom of the list of current commands in the Workflow by single clicking on it in the list on the right hand side of the page.

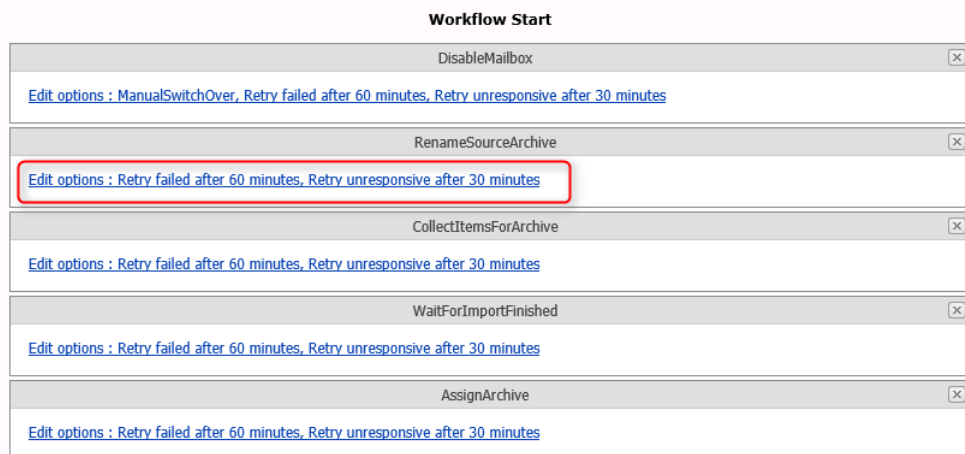
Moving a command

A command can be moved from its current position in the list by clicking on the grey title area and dragging it to a new position in the list:



Editing the Details of a Command

The details behind a command such as the number and frequency of retries can be edited by clicking on the hyper link relating to the command, as shown below:



Recommendations

- Built-in workflows should be left as-is, and copies can be made if required.
- Existing policies should be reviewed before creating new ones.
- New policies should contain the command-pair “CollectItemsForArchive” and “WaitForImportFinished”, as these perform the last-sync of the container.

Workflow commands


Below is a list of workflow commands that can be applied to workflow policies.

Name	Description
SyncEVEnvironment	Synchronize Enterprise Vault Environments vault stores and associated metadata
TransferArchiveForVaultStore	Collect archives of specified EV Vault Store
CollectItemsForArchive	Collect items of Enterprise Vault archive
CollectShortcutsFromMailbox	Collect Enterprise Vault shortcuts from Exchange On-Prem mailbox attached to Enterprise Vault archive
TransferEVContainerInfo	Collect statistics of Enterprise Vault archive
TransferDeletedItems	Collect deleted Enterprise Vault items of specified Vault Store, mark them in Item Database as deleted (DeletedDateUtc) and remove from routings for migrations which did not start yet
CollectIFoldersForArchive	Collect Icon Folders for EV Archive
CollectJournalSenderRecipients	Collect sender, recipient(s) from EV Archive item(s)
CollectDACases	Collect cases from EV Discovery Accelerator
CollectDACaseItems	Collect items from EV Discovery Accelerator case
CloseDACase	Closes EV DA Case
ExportArchive	Export items from Enterprise Vault archive
GetStagingFreeSpace	Get information about free space on Staging Area
ImportArchive	Import items into Enterprise Vault archive
ImportJournalArchive	Import items into Enterprise Vault journal archive
CheckBackupMode	Check backup mode for EV Vault Store (if BackupMode is enabled then ingest is put on hold)
CheckIndexComplete	Check index COMPLETE for EV Vault Store [SELECT COUNT_BIG(IndexCommitted) FROM dbo.JournalArchive WITH (NOLOCK) WHERE dbo.JournalArchive.IndexCommitted = 0]

Name	Description
EnableMailbox	Enable mailbox for Enterprise Vault archiving
DisableMailbox	Disable mailbox for Enterprise Vault archiving
ZapMailbox	Remove all Enterprise Vault customized archiving settings from the root of the mailbox and any folders within the mailbox
ZapArchivePermissions	Remove all Enterprise Vault archive permissions
CreateArchive	Create new Enterprise Vault archive
AssignArchive	Assign Enterprise Vault archive to mailbox by attaching archive to LegacyMailboxDN
RenameSourceArchive	Rename Enterprise Vault source archive (by default attach prefix/suffix to origin name or use custom name)
DeleteArchive	Delete Enterprise Vault source archive
RenameTargetArchive	Rename Enterprise Vault target archive (by default attach prefix/suffix to origin name or use custom name)
DisableOrphanedArchive	Set archive in Enterprise Vault system to status "Closed"
CopyPermissions	Copy Enterprise Vault Inrasite archive permissions
ResolveEVArchiveDefaultFolders	Resolve EV archive default folder icons
FixShortcuts	Fix Enterprise Vault shortcuts in the mailbox to point to the new Enterprise Vault archive
DeleteEVShortcut	Delete Enterprise Vault shortcuts from Exchange On-Prem mailbox
RevertPendingShortcuts	Revert "Pending" Enterprise Vault shortcuts back to the origin message class in the Exchange On-Prem mailbox
DeleteEVConfiguration	Remove hidden EV messages from Exchange On-Prem mailbox
RemoveVirtualVault	Remove EV Virtual Vault from Outlook for Exchange On-Prem mailbox
RemoveEVButtons	Remove (hide) buttons from Outlooks Enterprise Vault Add-In for Exchange On-Prem mailbox
RestoreEVForeignShortcuts	RestoreEVForeignShortcuts
Office365RestoreEVForeignShortcuts	Office365RestoreEVForeignShortcuts

Name	Description
EVRestoreEVForeignShortcuts	EVRestoreEVForeignShortcuts
RestoreSourceOneForeignShortcuts	Collects foreign SourceOne shortcuts information
ExchangeImportToArchiveMailbox	Import exported items into Exchange On-Prem secondary (archive) mailbox
ExchangeImportToPrimaryMailbox	Import exported items into Exchange On-Prem primary mailbox
ExchangeImportToJournalMailbox	Import exported items into Exchange On-Prem journal mailbox
TransferDomainList	Retrieve all available domains
TransferAllUsers	Retrieve all Active Directory users
TransferSingleUser	Retrieve single Active Directory user
TransferExcDBs	Retrieve Exchange databases
TransferExcServers	Retrieve Exchange servers
TransferExcDAGs	Retrieve Exchange Database Availability Groups (DAG)
TransferDeletedUsers	Retrieve deleted Active Directory users
-	Update modules via Archive Shuttle User interface
StopService	Stop Archive Shuttle service/module
StartService	Start Archive Shuttle service/module
Restart Service	Restart Archive Shuttle service/module
GetAzureCredentials	Retrieve Azure credentials stored in Credentials editor
GetPowerShellScriptCredentials	Retrieve PowershellScript credentials stored in Credentials editor
AutoRestartModule	Tries to restart unintentionally stopped module(s)
RecreateMapiProfile	Tries to re-create MAPI profile when current MAPI profile is not functioning
GetAmazonCredential	Retrieve Amazon credentials stored in Credentials editor
WaitForImportFinished	Wait until all routed items are imported
StagingAreaCleanup	Remove non-relevant or successfully imported items from the Staging area

Name	Description
SendEmailNotification	Send Email Notification
Office365ImportToArchiveMailbox	Import exported items into Office 365secondary (archive) mailbox
Office365ImportToPrimaryMailbox	Import exported items into Office 365primary mailbox
Office365DeleteShortcuts	Delete Enterprise Vault shortcuts from Office 365 mailbox
Office365RevertPendingShortcuts	Revert "Pending" Enterprise Vault shortcuts back to the origin message class in the Office 365 mailbox
Office365DeleteEVConfiguration	Remove hidden EV messages from the Office 365 mailbox profile in Outlook
Office365CollectMailboxes	Collect all available primary Office 365 mailboxes
Office365CollectArchiveMailboxes	Collect all available secondary (archive) Office 365 mailboxes
Office365CollectShortcutsFromMailbox	Collect Enterprise Vault shortcuts from Office 365 mailbox attached to Enterprise Vault archive
Office365CollectLicenses	Collect all available Office 365licenses and its statistics
Office365CreateArchiveMailbox	Creates archive/secondary mailbox for Office 365 i NOTE: The user must have an Office 365 mailbox prior to using this command.
Office365EnsureLicenseAssigned	Checks for licenses for a mailbox. The command will fail if there is no existing license. i NOTE: For pre-migration workflow policies only.
Office365RemoveLicense	Remove Office 365 license from the mailbox.
Office365ProvisionMailbox	Provision new Office 365 mailbox. i NOTE: For pre-migration workflow policies only.
Office365SetLegalHold	Set provisioned Office 365 mailbox on legal hold

Name	Description
Office365AssignLicense	Assign license to provisioned Office 365 mailbox
Office365GetMailboxInfo	Retrieve Office 365 mailbox details
Office365RemoveUser	Remove Office 365user (Stage2 command)
Office365CloudOnlyUserSync	Synchronize Office 365cloud-only users and shared mailboxes
Office365SyncSingleCloudOnlyUser	Synchronize specified Office 365cloud-only user
Office365CollectSingleMailbox	Synchronize specified primary Office 365 mailbox
Office365CollectSingleArchiveMailbox	Synchronize specified secondary (archive) Office 365 mailbox
Office365RemoveVirtualVault	Remove EV Virtual Vault from Outlook for Office 365 mailbox
Office365RemoveEVButtons	Remove (hide) buttons from Outlooks Enterprise Vault Add-In for Office 365 mailbox
Office365SwitchTarget	Switch the container from primary to secondary mailbox for Office 365 mappings <div>  NOTE: Make sure that an archive mailbox is present in the database for the switch to be successful. </div>
Office365SyncDeletedUsers	Remove deleted Office 365 users from database
Office365SyncRetentionTags	Synchronize Office 365 retention tags
NativeFormatImportToPst	Import exported items into PST file
NativeImportRename	Rename and move finished and closed Native import container
NativeImportClose	Close opened Native import container into which no more items are to be ingested
EasCollectArchives	Collect EAS archives
EasCollectItems	Collect items of EAS archive
EasExportItems	Export items from EAS archive
EasGetArchiveStatistics	Collect statistics of EAS archive
EasDeleteShortcuts	Delete EAS shortcuts from Exchange On-Prem mailbox

Name	Description
EasConvertToEvShortcuts	Convert EAS shortcuts into EV shortcuts (Create new EV shortcuts and remove matching EAS shortcuts from the mailbox)
EasDisableArchive	Disable EAS archive for archiving
EasOffice365DeleteShortcuts	Delete EAS shortcuts from Office 365 mailbox attached to EAS archive
EasCollectShortcuts	Collect EAS shortcuts from mailbox attached to EAS archive
EasGetServerInfo	Retrieve EAS server Id
ScanOfflineImport	Scan items within offline media location for import
ExportFailedItems	Move permanently failed items from the Staging Area to a location defined on the target link
ExportFailedItemsToFTP	Copy permanently failed items from the Staging Area to a FTP location
ProofPointImportUser	Import/Convert data to ProofPoint RawXML format
ProofPointImportJournal	Import/Convert journal data to ProofPoint RawXML format
S1ResolveFolders	Collect folders from mailbox attached to SourceOne archive
S1CollectArchives	Collect SourceOne archives
S1CollectItemsForArchive	Collect items of SourceOne archive
S1ExportArchive	Export items from SourceOne archive
S1ConvertToEvShortcuts	Convert SourceOne shortcuts into EV shortcuts (Create new EV shortcuts and remove matching SourceOne shortcuts from the mailbox)
S1DeleteShortcuts	Delete SourceOne shortcuts from Exchange On-Prem mailbox attached to SourceOne archive
S1Office365DeleteShortcuts	Delete SourceOne shortcuts from Office 365 mailbox attached to SourceOne archive
S1CollectOrganizations	Retrieve Exchange Organizations from SourceOne/EMX
S1SyncArchives	Collect statistics for SourceOne archives
S1SyncPstArchives	Collect owners for SourceOne PST archives

Name	Description
	i NOTE: If this command does not run, ensure that PST archives were collected from SourceOne via the S1CollectArchives command.
DellArchiveManagerCollectArchives	Collect archives of Dell Archive Manager
DellArchiveManagerCollectItemsForArchive	Collect Items of Dell Archive Manager archive
DellArchiveManagerExportArchive	Export items from Dell Archive Manager archive
DellArchiveManagerDeleteShortcuts	Delete Dell Archive Manager shortcuts from Exchange On-Prem mailbox attached to archive
DellArchiveManagerConvertToEvShortcuts	Convert Dell Archive Manager shortcuts into EV shortcuts (Create new EV shortcuts and remove matching Dell Archive Manager shortcuts from the mailbox)
DellArchiveManagerDisableArchive	Disable Dell Archive Manager archive for archiving
DellArchiveManagerOffice365DeleteShortcuts	Delete Dell Archive Manager shortcuts from the Office 365 mailbox attached to archive
DellArchiveManagerCollectShortcutsForArchive	Collect Dell Archive Manager shortcuts from mailbox attached to archive
DellArchiveManagerSyncArchives	Collect statistics for Dell Archive Manager archives
PstExportScanDirectory	Scan specified directory for PST source files
PstExportScanPst	Collect items of PST file
PstExportExportItems	Export items from PST file
PSTExportReexportPst	Initiates reexporting of items for corrupted PSTs and removes the corrupted PST file.
PowerShellScriptExecutionRunScript	Execute custom PowerShell script
JournalExplosionCleanupImportedItems	Remove successfully imported Journal Explosion items from the Staging area
JournalExplosionCheckQuota	Check current usage of recoverable items quota for Journal Explosion
GetExistingFilesInOutputPath	Retrieve content of OutputPath area to check if PST files are still present (based on this information watermark statistics are updated)
UpdateModulesViaUpdater	Update Archive Shuttle modules via dedicated Updater module

Name	Description
StagingAreaUsedSpaceForLink	Starts command that returns the used space in bytes for staging area per link
MetalogixCollectArchives	Collect Metalogix archives
MetalogixCollectItems	Collect items of Metalogix archive
MetalogixExportArchive	Export items from Metalogix archive
MetalogixCollectShortcuts	Collect Metalogix shortcuts
MetalogixDeleteShortcuts	Delete Metalogix shortcuts
MetalogixConvertToEvShortcuts	Convert Metalogix shortcuts to EV
MetalogixDisableArchive	Disables Metalogix archive
MetalogixOffice365DeleteShortcuts	Metalogix Office 365Delete Shortcuts
StorageImportToUNC	Import items to UNC
StorageImportToAzure	Import items to Azure
StorageImportToAmazon	Import items to Amazon
StorageImportSpaceCheck	Check free/used space for Storage Import output path
StorageImportUNCspaceCheck	Command for UNC space check
StorageImportAzureSpaceCheck	Command for Azure blob space check
StorageImportAmazonSpaceCheck	Command for Amazon bucket space check

Migration filters

Migration filters allow an administrator to construct filters that can narrow the data during migration.

The grid contains the following information.

- **ID:** The identification number associated with the policy.
- **Policy Name:** The title of the policy.
- **Definitions Count:** The number of filter conditions added to the policy.

Actions to be performed on the Migration filters page

- **New Policy:** Create a new policy to add migration filters to.
- **New Filter Condition:** Add a migration filter to the policy.
- **Edit:** Edit a migration filter in a policy.

- **Delete:** Delete a migration filter from a policy.
- **Refresh:** Refresh the current screen.
- **Auto refresh:** Choose time parameters to refresh the current screen automatically.
- **Columns:** Choose which columns to display in the grid.
- **Reset:** Reset the grid to its default state.
- **Load:** Load grid layout.
- **Save:** Save grid layout. Click the **Shared** box to share the layout to other users.
- **Export:** Export the grid in a .pdf, .xls, .xlsx or .csv format.

How to create a migration filter policy

1. Click **New** under the Policy section.
2. Enter a name for the migration filter policy.
3. On the newly created policy, select the checkbox, and click **Edit** under Filter Condition.
4. Select a policy from the drop down list.
5. Select a filter value to filter the data by.
6. Select an operator.
7. Input a value and select the data size from the drop down menu.
8. Click **Save**.

Filter policies

The Filter Policies page in the user interface allows an administrator to customize filters relating to the data migration and create new ones.

In order to filter by 'Path', the 'Collect Extended Metadata' must be enabled in System Configuration.

Filters should be given a name to easily identify them, and can be based on:

Field	Description	Source
ArchivedDate	The date the item was archived	EV

Field	Description	Source
Path	The path to the archived item	All
RetentionCategory	The retention category of the archived item	EV
HasShortcut	The item has message properties implying it was archived	All
ItemDate	The date/time of the item itself.	All
ItemSize	The size of the archived item	All
HasLegalHold	Whether the items is on legal hold	EV

Filter conditions are logically ANDed together. For example, it is possible to migrate data from a particular path AND below a particular size.

i | NOTE: The HasShortcut filter should not be auto-enabled.

Example filters

The following section gives some examples of filters that can be created. Once they have been created in the Admin Interface, they can then be used when performing container mapping operations.

Only migrate items which have shortcuts.

Create New Filter

- Give the filter a name, e.g., “With Shortcut”

Add a filter condition

- Click “New” in the “Filter Condition” section.
- Choose the “Policy” which was created previously.
- Choose “HasShortcut” from the “Filter By” selection
- Choose “Yes” in the “Value” selection
- Click Add

Only migrate items which belong to a particular retention category

Create New Filter

- Give the filter a name, e.g., “3 Year Retention Category”

Add a filter condition

- Click on “New” in the “Filter Condition” section.
- Choose the “Policy” which was created previously.
- Choose “RetentionCategory” from the “Filter By” selection
- Select the appropriate retention category from the drop-down list of categories from the source environment
- Click Add

Only migrate data with ItemDate younger than 2011-12-31

Create New Filter

- Give the filter a name, e.g., “Newer than 2011”

Add a filter condition

- Click on “New” in the “Filter Condition” section.
- Choose the “Policy” which was created previously.
- Choose “ItemDate” from the “Filter By” selection
- Select the Operator “YoungerThan” from the Operator drop down
- Select the date 31st December 2011 from the date picker
- Click Add

Changing a mapping's filter

You may need to adjust filters for a mapping that is currently running. The results of changing filters depends on the stage in which it is changed in the mapping:

- Prior to collection: items not yet collected will use the filter assigned to the mapping
- Prior to routing: items not yet routed will use the filter assigned to the mapping
- After routing: the core service must be restarted. Once this is done, routing will also be restarted. Filters will be applied from this point; any items routed before the restart will not be affected by new filters. If previously routed items need to be re-filtered, they will need to be deleted from the ItemRouting table if necessary

System configuration

The System Configuration page in the Archive Shuttle Admin Interface contains many settings that can be used to customize the migration and environment. Some of these settings will affect the workflow of the migration while others will affect throughput and performance.

Changing these configuration options should be done after careful analysis of the current workflow, environment and throughput.

Changes to system settings take place after a few minutes when the appropriate module checks in for new work items.

Schedule settings

On the System Configuration page different settings/configurations can be applied according to a schedule.

When you first enter the System Configuration, the default schedule is shown. Configuring specific schedules consists of the following three steps:

1. Create and name a new schedule.
2. Make the required configuration changes for this schedule, above, or below the default schedule.
3. Select the times of day, or days of week that this schedule applies.

Date/Time is based on the date and time of the Core.

Example: More export parallelism on Saturday/Sunday

In this example, you will configure higher Enterprise Vault export parallelism for Saturday and Sunday. The steps to follow are:

1. Navigate to the System Configuration page.
2. Review the current configuration, this is the default configuration when no other configuration is set.
3. Via the button on the toolbar, add a new schedule. Give the schedule a name (e.g. Weekend) and choose a color from the available list.
4. Change the EV Export module parallelism settings to the new, higher values that you wish to use. When changes are made the entry will become bold, and a checkbox will appear underneath **Custom**.
5. Click on **Save**, in the toolbar, to commit those changes.
6. From the toolbar click on **Time Assignment** in order to review the current scheduled configurations.

7. Click on the new schedule on the right hand side, and then highlight all of Saturday and all of Sunday.
8. Click on **Save** on the schedule view in order to commit those changes.

General settings

This section contains general settings relating to Archive Shuttle:

Item	Description
General	
Do not re-export on File Not Found	Normally if an import module files that a required file is not present on the staging area it will report this to the Archive ShuttleCore and the Core will instruct the export module to re-export the item. If this behavior is not required, select this checkbox
Turn off the post-processing of secondary archives	This will disable the post processing modules from performing actions against an Exchange (or Office 365) Personal Archive
Autoenable Stage 1	Automatically enable Stage 1 when Pre-migration is complete.
Stage 2 Default Priority	The normal priority which will be applied to users when they enter Stage 2.
Autoenable Stage 2	<p>With this option enabled when a container gets to 100% exported and 100% imported, it will automatically move into Stage 2.</p> <p>Without this option, set containers will remain in Stage 1 until Stage 2 is enabled for them manually.</p>
Delete Shortcut Directly after Import	<p>With this option selected, successfully ingested items in to Exchange, or Office 365 will have the shortcut removed from the target container straight away. Without this option selected, the shortcuts will only be deleted once the Stage 2 command is executed.</p> <p>This option can greatly enhance the user experience when migrating data back to existing containers (e.g., primary mailbox)</p>

Item	Description
Archive Migration Strategy	<p>When a batch of users has been selected for migration and given a specific priority, within that, a migration strategy can be used to also govern the order that the migrations will take place. The options here are:</p> <ul style="list-style-type: none"> • Smallest archive first (based on archive size) • Largest archive first (based on archive size) • Oldest first (based on the archived date/time of items) • Youngest first (based on the archived date/time of items) <p>The default strategy equates to a random selection within the batch that was chosen</p>
Do not transmit non-essential Active Directory Data	Stops the Active Directory Collector module from returning metadata about users which is not required for the product to function.
Use HOTS format	Instructs the Core and Modules to expect and use HOTS format for items on the staging area.
Clear staging area files older than [hours]	When the staging area cleanup process runs it will only remove files older than this number of hours.
Disable automatic staging area cleanup	Stops the automatic cleanup of the staging area.
Enable Auto Restart of Modules	Causes modules to automatically restart.
Number of retries for Auto Restart of Modules	Specify the number of times modules attempt to restart.
Logging	
Delete ItemRoutingErrorLogs after successful Export/Import	This option allows the system to remove any reference to issues/errors during export or import once an individual has been successfully exported or imported.
Clear Module Performance Logs	By default module performance logs will be kept for 30 days, but this can be changed in this

Item	Description
	setting.
Item Database	
Default size [MiB]	Shows the size that new item databases will be created.
Default log file size [MiB]	Shows the default size of new item database log files.
Item database update retry timeout [hours]	Shows the number of hours that will elapse before upgrades of the item databases will be retried.
User Interface	
Global Timezone	By default the system operates in UTC
My Timezone Override	Allows a user specific timezone to be specified
Item size scale	Can be used to control whether item sizes in the user interface are displayed as bytes, Mib, GiB, TiB, or displayed in automatic scale
Options 'All' on grids will be disabled when this threshold is exceeded	If a value is specified here it will stop Archive Shuttle giving the option to display 'all' values on a grid, when there are a large number of items to display.
Folder	
Folder name for folder-less items	Specify a name to use for items which do not have a folder.
Treat all messages as folderless	If enabled this will ignore the folder which was obtained during the export of an item, and just ingest all items in to the folder specified as the 'folderless items' folder
Manage Folder Translations	Manage translation of default mailbox folders.
Chain of Custody	

Item	Description
Re-export on Chain of Custody error	Instructs the import module to report a Chain of Custody error back to the Core, and for the Core to then queue work for the export module to re-export the item.
Enable Chain of Custody for Extraction	Causes the export modules to generate a hash of items (to be stored in the Item Databases) when items are exported
Enable Chain of Custody for Ingestion	Causes the import modules to validate the Chain of Custody information when items are ingested. If an item fails this check it will not be ingested.
Reliability	
Allow import of messages with empty recipient e-mail	If a message is found without recipient information it will normally fail the ingestion process. This can be overridden with this option
Allow import of corrupt messages	Items which fail the MSG file specification may not be ingested into the target container. Some of these failures can be overridden with this option checked.
Enable item level watermarks	Archive Shuttle will stamp items that get imported with a Watermark, specifying details about the Item like Source Environment, Source Id, Archive Shuttle Version, ItemRoutingId, LinkId.
Fix Missing Timestamps	If the message delivery time is missing on a message the messaging library will generate it, from other properties already present on the message
Missing timestamp date fallback	The date to be used in case of a missing timestamp

FTP/SMTP settings

This section contains settings relating to the FTP/SMTP configuration (for sending out reports):

Item	Description
SMTP Server	The FQDN of an SMTP server which can be used to send emails
SMTP Port	The port to be used when connecting to the SMTP server
Use SSL	A flag to indicate that SSL should be used to connect
Use Authentication	A flag to indicate whether the SMTP server supports anonymous connections, or requires authentication
SMTP Username	The username to be used for authentication
SMTP Password	The password to be used for authentication
SMTP From Email Address	The 'from' address to be used in the outbound mail
SMTP From Name	The display name to be used in the outbound mail

FTP/SFTP settings

This section contains settings relating to the FTP or SFTP configuration:

Item	Description
URL	The URL or IP address to used to connect to the server
Username	The username to be used when connecting to the server
Password	The password to be used when connecting to the server
Use SFTP	If ticked, Archive Shuttle will connect and use SFTP. Otherwise, Archive Shuttle will use FTP

Journal Transformation settings

This section contains settings relating to journal transformations.

Item	Description	Notes
Delete items from staging are after initial export	If enabled items are removed from the staging area after initial export.	It is recommended to enable this if there is a small subset of users to be migrated. It reduces the requirements on the size of the staging area.
Delete items without journal archive migration import routings	Items without journal archive migration import routings will be deleted from the staging area. It is recommended to enable this if all user mappings are already enabled for import.	It's recommended to turn this on when all user mappings are enabled for import.
Time to expire inactive items (hours)	Set value to '0' to disable this setting. Can only be applied if setting 'Disable automatic staging area cleanup' is inactive.	
Mailbox Quote Exceeded Threshold (%)	Ingest of items will stop once mailbox size has reached this quota. Ingesting will restart after quote is increased or size has decreased to below quota threshold.	Ingest will resume if the quota is increased, or the mailbox size is reduced.
Ingest users with most items first	Ingest users with higher amount of item counts first.	
Skip items which do not belong to user mapping	Items which are older than the user will not be ingested into the target. This setting isn't applied for leavers, re-created users and manually mapped users.	This setting does not apply to leavers, re-created users and manually mapped users.
Export items for users enabled for ingestion only	Only sender recipients collected via EV index which are enabled for ingestion are	

Item	Description	Notes
	exported. When this setting is disabled, all collected items are exported.	
Enable collection of sender recipients on export	EV export module will start to collect sender recipient data during export of the journal archive. 'Export items for users enabled for ingestion only' setting has to be enabled per mapping to apply this setting.	
Export items with permanently failed collection of sender recipients	EV export module will export item for which collection of sender recipients failed permanently.	
Enable all SMTP addresses discovered upon exploding	All SMTP addresses discovered during explosion process will be automatically enabled for further processing. The setting is ignored if 'list of accepted domains' is specified.	
List of accepted domains	All exploded SMTP addresses are stored in the Archive Shuttle system, however only addresses of domains listed here will be automatically enabled for further processing. All other addresses will be filtered out and not available for user interaction. This setting only applied during the explosion process, and not for already exploded addresses.	If this list is altered during a migration, it will only affect future explosion of items, not the existing items.
Enable automatic 'Merge by preview'	All leaver mapping(s) that are available for merging by preview will be automatically enabled.	

Item	Description	Notes
Check mailbox quota only for importing users	The quota check command will be executed only for users that are actively importing.	

EV Collector module

This section contains settings relating to the EV Collector Module:

Item	Description
Collector Parallelism	Defines how many archives will be collected in parallel
Collect Extended Metadata	Reads legal hold and path of items using the Enterprise Vault API for filtering purposes. It is recommended that this is only enabled when filtering by folder
Use the BillingOwner on Archives which would be otherwise ownerless	This uses the owner set as "Billing Owner" in Enterprise Vault as the Owner of the archive instead of trying to use the entries in the Exchange Mailbox Entry table. This is useful where an Active Directory account relating to an archive has been deleted, for an employee who has left the company for example. This would normally show as Ownerless in Archive Shuttle, but with this switch enabled the Enterprise Vault Collector module will attempt to use the "Billing Owner".
Ignore LegacyExchangeDN when matching EV users	With this option enabled the ownership detection for EV archives is modified so that the LegacyExchangeDN is not used.
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place

EV export module

This section contains settings relating to the EV Export Module:

Item	Description
General	
EV Export Archive Parallelism	Defines how many archives will be exported in parallel. Total thread count = EV Export Archive Parallelism multiplied by EV Export Item Parallelism.
EV Export Item Parallelism	Defines how many items should be exported in parallel per archive. Total thread count = EV Export Parallelism multiplied by EV Export Item Parallelism.
EV Export Storage	<p>If using Azure for the staging area storage (or you are migrating a journal archive for an older version of EV where Archive Shuttle is doing the envelope reconstruction) ensure this option is set to Memory, otherwise File System or Memory with File System Fallback can be selected. In either of the described situations if 'File' is chosen, an error will be reported in the export module log file and export will not proceed.</p> <p>This setting can be adjusted if there are problems exporting very large items.</p>
Export Provider Priority	Specifies the order in which EV export mechanisms will be tried.
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place
Failures handling	
Export Messages with invalid Message ID from Journal	When enabled it will mean that items from a journal that require envelope reconstruction will still be processed (and a P1 message generated) even if the Original Message ID attribute can not

Item	Description
	<p>be found in the item that was retrieved from EV (meaning that an EV Index lookup can not be performed).</p> <p>Note: This setting will have an impact in that it may mean BCC information is not added to the P1 envelope (since it can not be obtained from EV)</p>
Prevent exporting of items if envelope reconstruction fails	This will prevent Archive Shuttle from providing an item from EV if the envelope reconstruction fails. With this setting disabled some items may be provided without an appropriate P1 (Envelope) message.
Fail items permanently on specified errors	<p>Indicates whether Archive Shuttle should mark certain items as permanently failed, even on the first failure.</p> <p>i NOTE: A full or partial error message can be entered.</p>
Error message(s) to permanently fail items on	A list of error messages which will cause items to be marked as permanently failed (if the previous setting is enabled)
Journal Transformation	
Process messages without P1 header	Data for journal transformations are parsed from reconstructed P1 headers. BCC recipients are not available for the journal transformation user mappings.

EV import module

This section contains settings relating to the EV Import Module:

Item	Description
Offline Scan Parallelism	Indicates the number of threads that should be used when scanning offline media
EV Default Retention Category ID	When a non-Enterprise Vault source is used, and the target of the migration is an Enterprise Vault environment, indicate here the

Item	Description
	retention category ID to apply to the data which is ingested
EV Import Archive Parallelism	<p>Defines how many archives will be imported in parallel.</p> <p>Total thread count = EV Import Archive Parallelism multiplied by EV Import Item Parallelism.</p>
EV Import Item Parallelism	<p>Defines how many items should be imported in parallel per archive.</p> <p>Total thread count = EV Import Parallelism multiplied by EV Import Item Parallelism.</p>
Import Journal Archive Through Exchange	Imports a journal archive through Exchange instead of through the Enterprise Vault API. Elements from the staging area will be added to Exchange (for an appropriate Enterprise Vault task to process) rather than directly into Enterprise Vault.
Journal Mailbox Threshold	If using the 'Import Journal Archive Through Exchange' option then this setting can be used to limit when ingest will be stopped while the appropriate task processing the mailbox catches up.
Suspend imports while EV is archiving	Disables import module while Enterprise Vault is in its archiving schedule.
Ingest Provider Priority	Indicate the type of ingest provider to use.
Read file to memory	Allow reading of files to system memory before ingestion
Read file to memory threshold (bytes)	Items below this size will be read into memory (to speed up ingestion), whereas items above this size won't be read into memory
Fail items permanently on specified errors	If enabled items which encounter one of the specified errors below will be marked as permanently failed, regardless of the failed item threshold.

Item	Description
Error message(s) to permanently fail items on	A list of errors
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place

EV provisioning module

This section contains settings relating to the EV Provisioning Module:

Item	Description
Convert orphaned into shared archives	Enterprise Vault 11 onwards cannot handle Folder structures in shared archives. Only to use on advice of Quadrotech support.
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place

EWS is the preferred method of shortcut processing and should be used wherever reasonable to support.

Shortcut process module

Item	Description
Shortcut Process Parallelism	Specifies how many mailboxes are post-processed at the same time.
Shortcut Process Item Parallelism	Specifies how many shortcuts per mailbox are processed in parallel.
Delete Shortcuts options	There are three options when deleting shortcuts:

Item	Description
	<ul style="list-style-type: none"> • Shortcuts related to migrated items: shortcuts linked to migrated items will be deleted only. This settings is set by default. • Shortcuts not related to migrated items: shortcuts with no association to migrated items will be deleted only. • All shortcuts: shortcuts are removed without an ownership check. <p>i NOTE: This is a hard deletion of the shortcut(s).</p>
Delete messages with EV properties but without proper shortcut Message Class	Deletes messages which are not proper EV shortcuts, but have EV properties such as SavesetId or ArchiveId. (Setting ignored for migrations to PST).
Use EWS for Processing	Use EWS as access method for communication with Exchange/Office 365 system. For Office 365 and Exchange 2019 and above, this setting has to be enabled.
Shortcut deletion maximum batch count	Specifies maximum number of batches used for shortcut deletion in parallel.
Process shortcuts from primary and secondary mailbox	Collect and restore Shortcuts for both primary and secondary mailboxes. Hybrid scenarios are also supported.
Convert EV foreign shortcuts to regular items	Converts EV foreign shortcuts to regular messages on target mailbox. Conversion is done only when EV source archive for these shortcuts has not been found on the EV source server.
Limit stored results	Enables limit to store results in local database based on threshold value.
Threshold database limit	Defines maximum records stored in local data base when module stops asking for work items. This setting is applied only when setting 'limit stored results' is enabled.

Exchange import module

This section contains settings relating to the Exchange Import Module:

Item	Description
General	
Use per server EWS Url for Exchange import	If this is enabled then the import module will use the EWS Url configured in Active Directory on the Exchange Server object rather than a general Url for all ingest requests
Import Root Folder	When ingesting data in to Exchange mailboxes or personal archives it is sometimes required to ingest the archived data into a top level subfolder (and then the archive folders beneath that). Specify the name of that top level folder here.
Ingest Provider Priority	Specifies the order of the ingestion providers that should be used e.g. AIP first, EWS Batch second etc.
Maximum Batch Size Bytes	The maximum size that a batch of items can be, which is then sent in one go to Exchange
Maximum Batch Size Items	Maximum number of items in a batch
Exchange Timeout Seconds	Timeout in seconds until Archive Shuttle aborts the ingest (i.e. upload/processing)
Disable reminders for appointments in the past	This will remove the MAPI properties relating to whether a reminder has been sent/fired or not as the item is ingested into the target. If this is not enabled reminders may appear for long overdue archived items.
Mark migrated items as read	If this is enabled all migrated will be marked-as-read by the import module
Disable migration on specified error(s)	If enabled items which encounter one of the specified errors below cause the mapping to be disabled.
Error message(s) to disable migration on	List of errors

Item	Description
Fail items permanently on specified errors	If enabled items which encounter one of the specified errors below will be marked as permanently failed, regardless of the failed item threshold.
Error message(s) to permanently fail items on	A list of errors
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place
Threading/Parallelism	
Offline Scan Parallelism	Number of threads that will be used for scanning offline media
Exchange Mailbox Parallelism	Defines how many Exchange mailbox imports will be ingested to in parallel.
Exchange Batch Parallelism	Defines how many batches will be ingested per mailbox in parallel
Exchange Item Parallelism	Defines how many items will be ingested per mailbox in parallel
Connectivity	
Exchange Version	Specify the version of Exchange which is in use.
Disable Certificate check	Disable the certificate validity when connecting to Exchange
Exchange Connection URL	Specify an Autodiscover URL if the default one does not work
Use Service Credentials for Logon	Authenticate to Exchange with the credentials which the Exchange Import Module Server is running as.

Office 365 module

This section contains settings relating to the Office 365 Module:

Item	Description
General	
Number of fastest servers to use	Determines how many servers from the list of fastest are used. They will be picked randomly by the module
Import Root Folder	When ingesting data in to Office 365 mailboxes or personal archives it is sometimes required to ingest the archived data into a top level subfolder (and then the archive folders beneath that). Specify the name of that top level folder here.
Ingest Provider Priority	Determines which ingestion methods will be used to ingest data into Office 365 and in which order
Office 365 Batch Size Bytes	The maximum size that a batch of items can be, which is then sent in one go to Office 365
Office 365 Batch Size Items	Maximum number of items in a batch
Office 365 Timeout Seconds	Timeout in seconds until Archive Shuttle aborts the ingest (i.e. upload/processing)
Check mailbox quota only for importing users	Mailbox quota only used for users that are actively being imported.
Mailbox Quota Exceeding Threshold (%)	Item ingest will be stopped once mailbox size has reached the defined threshold of the mailbox quota. The default value is set at 95%.
Disable reminders for appointments in the past	This will remove the MAPI properties relating to whether a reminder has been sent/fired or not as the item is ingested into the target. If this is not enabled reminders may appear for long overdue archived items.
Mark migrated items as read	If this is enabled all migrated will be marked-as-read by the import module

Item	Description
Convert journal messages to Office 365 journaling format	If this option is enabled information in the P1 envelope gets added to an attribute called GERP, and added to the message as it is ingested in to Office 365. This makes those items Office 365 journal-format messages. (See additional information here)
Allow synchronization of un-associated mailboxes	Mailboxes with no association to AD User present in AS database will be also synchronized and stored (matching is based on UPN/Email address)
Enable auto-recreate deleted Office 365 Users as Leavers	Deleted Office 365 Users which are marked as Enabled to Recreate will be recreated as Leavers once task 'RecreateDeletedOffice 365Users' will run
Collect retention tags	<p>Collect retention tags from Office 365.</p> <p>Disabled: No retention tags are collected.</p> <p>Active tags and labels only: active labels and tags are collected.</p> <p>All tags and labels: all tags and labels are collected.</p>
Allow ingest of items without Office 365 retention mapping	If the collection of retention tags is active, but an explicit mapping of retention tags hasn't been performed then some items may not ingest into the target. If this option is enabled, if a mapping doesn't exist, the item will still be ingested.
Preferred domain for ambiguous user match	Specify a domain (with or without the @) to use when trying to match users with ambiguous names.
Disable migration on specified error(s)	If enabled items which encounter one of the specified errors below cause the mapping to be disabled.
Error message(s) to disable migration on	List of errors

Item	Description
Fail items permanently on specified errors	If enabled items which encounter one of the specified errors below will be marked as permanently failed, regardless of the failed item threshold.
Error message(s) to permanently fail items on	A list of errors.
Disable ingest account on specified error	If enabled items which encounter one of the specified errors below will result in the ingest account being disabled from further ingest requests.
Error message(s) to set ingest account as disabled	A list of errors.
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place
Remove Invalid Recipients	Recipients with empty PR_DISPLAY_NAME and PR_EMAIL_ADDRESS properties failing on error [CorruptData, LID: 40618 - State machine stalled.] are removed from a message during ingest.
Virtual Journal	
Virtual Journal Item Count Limit	The maximum number of items in a virtual journal mapping before a new mapping will be created.
Virtual Journal Size Limit	The maximum size of a virtual journal mapping before a new mapping will be created.
Threading/Parallelism	
Offline Scan Parallelism	Number of threads that will be used for scanning offline media

Item	Description
Office 365 Mailbox Parallelism	Defines how many items will be ingested per mailbox in parallel.
Office 365 Item Parallelism	Defines how many items will be ingested per mailbox in parallel.
Office 365 Batch Parallelism	Defines how many batches will be ingested per mailbox in parallel
Connectivity	
Use fastest server (round-trip)	If enabled, from time to time the Office 365 module will get a list of servers responding to Office 365 ingest requests and use only those for ingestion.
Office 365 Exchange Version	Specify the Office 365 Exchange version
Disable certificate check	Disable the certificate validity check when connecting to Office 365
Use Multiple IP from DNS	When Office 365 returns multiple IP address entries for it's ingestion service this setting will allow the ingest module to communicate to all of those IP addresses instead of just one. For this to work, the 'Disable certificate check' option must be enabled.
Exchange Connection URL	Specify an Autodiscover URL if the default one does not work.
OAuth concurrent batches	By default, Archive Shuttle use 5 batches in parallel. This value can be adjusted.

PST Export module

This section contains settings relating to the PST Export Module:

Item	Description
General	

Item	Description
File Parallelism	The number of PST files to ingest from simultaneously
PST item collection file parallelism	The number of PST files to scan simultaneously
Limit stored results	When the module is working, should the number of items it is tracking and not sent the result to the Core be limited?
Threshold database limit	The number of items allowed to be stored locally, and not sent to the Core, before the module stops requesting more work.

Native Import module

This section contains settings relating to the Native Import Module:

Item	Description
Stamps a header to imported messages for ProofPoint to identify message source	When enabled a message-header is added to each item as it is added to a PST file. The header is called x-proofpoint-source-id and has the itemid/item routing id, as the value. For example: x-proofpoint-source-id: 91016abe-51e3-bdd6-132f-fb6763ecc751/2865103
Native Import File Parallelism	Defines how many PST files will be imported to in parallel.
Native Import Item Parallelism	Defines how many items will be ingested in parallel per PST file.
Finalize finished PSTs in Stage 1	With this option enabled finished/full PST files will be moved to the output area whilst the mapping is still in Stage 1. This will only happen on PSTs which are complete, i.e. those that have split at the predefined threshold. For migrations lower than the threshold, which therefore have just a single PST, this PST will not be moved in Stage 1.

Item	Description
Verify closed PST file consistency	Every closed PST will be checked for its correct items count and if the file itself is not corrupted.
Mark PST as corrupted on specified errors	If checked, PST(s) which item(s) failed on specified error will be marked as corrupted and not processed anymore.
Error message(s) to mark PST as corrupted	Based on listed error message(s), PST(s) will be marked as corrupted immediately.
Fail items permanently on specified errors	If enabled items which encounter one of the specified errors below will be marked as permanently failed, regardless of the failed item threshold.
Error message(s) to permanently fail items on	A list of errors.
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place
Re-create MAPI profile handling	
Enable re-creation of MAPI profile	If enabled, MAPI profile can be recreated with manual action from Modules page, or automatically based on error messages listed in setting below.
Error messages to cause re-creation of MAPI profile	Based on listed error message(s), module will initiate re-creation of MAPI profile.
Time of module inactivity before re-creating MAPI profile	The time of inactivity of NFI module after which the MAPI profile will be re-created.

EAS Zantaz module

This section contains settings relating to the EAS Zantaz module:

Item	Description
EAS Archive Parallelism	The number of archives to process in parallel
EAS Item Parallelism	The number of items to process in parallel per archive
Limit stored results	Stops the module trying to get more work, if there is a backlog of transmissions to be sent to the Core.
Threshold database limit	The number of items in the backlog before limiting takes place

SourceOne module

This section contains settings relating to the SourceOne module:

Item	Description
SourceOne Container Parallelism	The number of archives to process in parallel
SourceOne Collect Items Container Parallelism	The number of items to process per archive in parallel
SourceOne Sync Archives Batch Size	Defines how many archives will be synced in one batch.
Ignore Missing Transport Headers	Ignores missing transport headings in case of export from Journal export.
Email address used in case 'sender is null' error	Specify an email address to be used if the module reports a 'sender is null' error
Export messages failing on 'SAFETY: Mismatched Attachment Counts from native message!' error.	Use this setting to export messages which fails on mismatched attachment count error.
Limit Stored Results	Select this option to limit stored results.
Threshold Database Limit	Limit of records in local database, when module stops asking for work items.

Dell Archive Manager module

This section contains settings relating to the Dell Archive Manager module:

Item	Description
Container Parallelism	Defines how many archives will be exported in parallel.
Collect Item Parallelism	Defines how many items should be collected in parallel per archive.
Limit Stored Results	Select this option to limit stored results.
Threshold Database Limit	Limit of records in local database, when module stops asking for work items.
Collect size of archives	If enabled the module will also collect the overall size of archives

PowerShell Script Execution module

This section contains settings relating to the PowerShell Script Execution module:

Item	Description
Limit stored results	If enabled helps protect the module from some failures where backlog of results can't be sent to the Core.
Threshold database limit	The number of results that will stored locally before the module stops processing more work. (Combined with previous setting).

Monitoring

Day to day administration of migration activity

The tasks that need to be performed from day to day during an Archive Shuttle based migration vary from project to project; however, aspects that should always be performed are:

Monitoring Performance

Monitoring Workflows

Monitoring Module-Level Activity

Adjusting Priority of Migrations

Monitoring Migrations

Pausing, Resuming, and Skipping Workflow Steps

Changing Workflows

Setting Users Up for Migration

Configuring, Loading, and Saving Filters

Monitoring Performance

Performance of the migration can be viewed in various ways within the Archive Shuttle Web Interface and can be exported to a variety of formats. It is recommended to review the following:

Progress & Performance

Over a period of time, the number of containers (users) processed should start to increase. In fact, all the statistics should rise towards 100% during the migration. It is also possible to view the information displayed on this page both for the overall progress of the migration and an individual link level. To get a fine level of detail on the progress of a migration, it is also possible to show the graphs on the page as items/second rather than items/hour.

The data which is displayed and the order it is displayed in can be customized to meet the needs of the particular migration. These views can be stored, and can be switched between when required. Each chart or data grid is a widget that can be moved around the screen, or removed by clicking on the small 'x' at the top right of the widget. Once removed, widgets can be re-added if required.

Link Statistics

Ensure each link is progressing towards 100% over time

Performance

Once steady-state migration has been reached the number of items being exported and imported per hour should be fairly consistent.

System Health

Ensure that the Free Space amounts do not reach a warning or critical level.

Also ensure that modules are not deactivated or otherwise not running.

Finally ensure that the link 'used' space is within normal parameters.

Monitoring Workflows

Monitoring workflows involves reviewing the Stage 1 Status screen for hangs and failures. By default, the number of export errors and the number of import errors are shown as data columns on the screen. Additional data columns can be added (by clicking on **Columns** in the Actions Bar and dragging additional columns on to the data grid).

The Stage 1 Status screen can also be filtered and sorted to show (for example) all containers with less than 100% imported. This can be done as follows:

1. Navigate to the **Stage 1 (Sync Data)** screen.
2. Click Reset on the Actions Bar in order to return the page to the default view
3. In the text box under **Imported Percentage** enter **100**.
4. Click on the small filter icon to the right of the text box, and choose "is less than" from the pop-up list
5. Click on the Apply button at the right of the filter row.

The page will now show those containers that have been enabled for Stage 1 that have not reached the 100% import stage.

Monitoring Module-level Activity

Valuable information about module-level responsiveness can be gained from the **Module Dashboard** in Archive Shuttle. Using the dashboard it is possible to see detailed information relating to the operations being performed by particular modules on particular servers involved in the migration. For example, if new mappings have just been created, and yet one or more Enterprise Vault Export modules are showing no items being exported, it may indicate an area that should be investigated further.

Adjusting Priority of Migrations

When reviewing the data synchronization on Stage 1, it might be necessary from time to time to adjust the priority order that they are running. A good way to do this is to ensure that all new

mappings are created with a priority of 10, then, other priorities can be assigned on the **Stage 1** screen in order to raise or lower the priority of particular mappings.

i | NOTE: The lower the priority number, the higher the priority.

Monitoring Migrations

Monitoring migrations involves reviewing the **Stage 2 Status** screen for failures. Each command, performed by the appropriate module, will report back results to the Archive Shuttle Core.

This screen can also be filtered to show (for example) all containers that have not yet finished the migration completely. This can be done as follows:

1. Navigate to the **Stage 2 (Switch User)** screen.
2. Click **Reset** on the Actions Bar in order to return the page to the default view
3. In the drop down list under **Finished** select **No**.
4. Click on the **Apply** button at the right of the filter row.

The page will now show those containers that have been enabled for Stage 2, but have not yet completed the migration.

Pausing, Resuming, and Skipping Workflow Steps

Particular steps in a Workflow can be paused, resumed or even skipped. This can be done as follows:

1. Navigate to the **Stage 2 Status** screen.
2. Click **Reset** on the Actions Bar in order to return the page to the default view
3. Select one or more containers.
4. Click the appropriate action button, e.g., **Pause**, **Suspend**, **Resume**.

Changing Workflows

At any time a particular container mapping can have the workflow changed to a new one. In order to do this, perform the following:

1. Navigate to the **Stage 2 (Switch User)** screen.
2. Click **Change Workflow Policy** on the Actions Bar
3. Select a new workflow for this mapping.
4. Click **[Save]** to commit the change.

Once a new workflow has been selected, all previous workflow steps will be removed for the mapping and the new workflow will begin with the first command in that workflow.

This can also be used to re-run a chosen workflow associated with a container mapping.

Setting Users Up for Migration

Most archive migrations are performed by selecting groups or batches of user archives (containers) to process through Stage 1 and Stage 2. The mapping and selection is performed on the **Bulk Mapping** screen in the Archive Shuttle Web Interface. This screen can be filtered and sorted with the current list of data columns, and additional data columns can be added to help facilitate the selection of containers.

There are many ways that the selection can be defined. Below is an example of selecting users based on the source Vault Store:

1. Navigate to the **Bulk Mapping** or **Manual Mapping** screen.
2. Click **Reset** on the Actions Bar in order to return the page to the default view
3. In the text box under **Link Name** enter the name of one of the source Vault Stores, and press enter or click into a different text box
4. Click the **Apply** button at the right of the filter row.

The page will now refresh to show all archives (containers) in that particular Vault Store (link).

This selection can be further refined, before selecting some or all of the containers and performing the **Add Mapping** function from the Actions Bar.

Configuring, Loading and Saving Filters

Both the **Stage 1 (Sync Data)** and the **Stage 2 (Switch User)** pages in the **Admin Interface** allow for additional, useful data columns to be added to the screen. These can be re-arranged and sorted as required.

Both of these pages also allow for these filters, and column configurations to be saved under a friendly name. Previously saved filters can be loaded, allowing an administrator to jump between slightly different views of the migration.

i | **NOTE:** The filters are saved per user and can be used on any machine that accesses the Admin Interface from the same Windows account.

Monitoring a migration


It's important to monitor Archive Shuttle during a migration. This article gives advice on the things to monitor.

Active systems

Archive Shuttle works best at scale. Migrating a few hundred items is not what it is intended to do. Fully utilized systems ensure that all aspects of the system, the workflow and processes are busy.

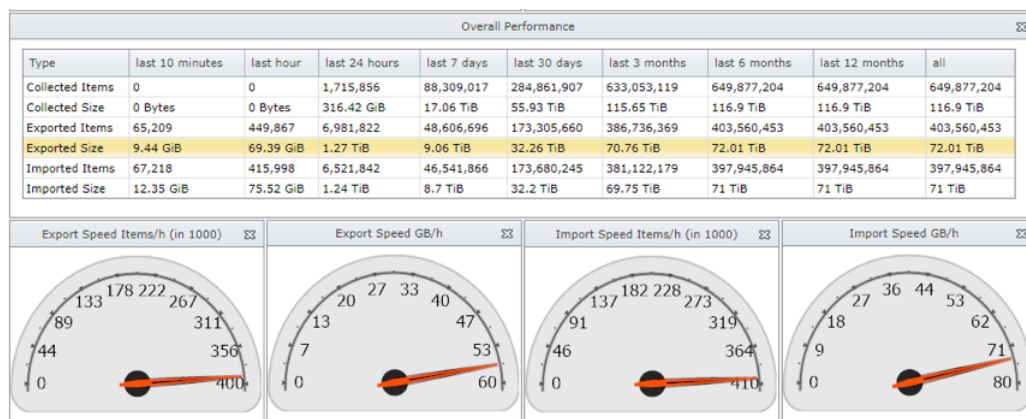
Progress & Performance

Take a look at the Progress & Performance screen in the user interface.



This dashboard is modular, meaning you can re-arrange widgets, remove ones which aren't needed and store that layout to use later. The page can be configured to automatically refresh (for example every 60 seconds).

The dashboard gives a great overview of the overall health of your migration.



In the graphic above we ask the question whether there is any potential issue with the highlighted statistics.

Red box

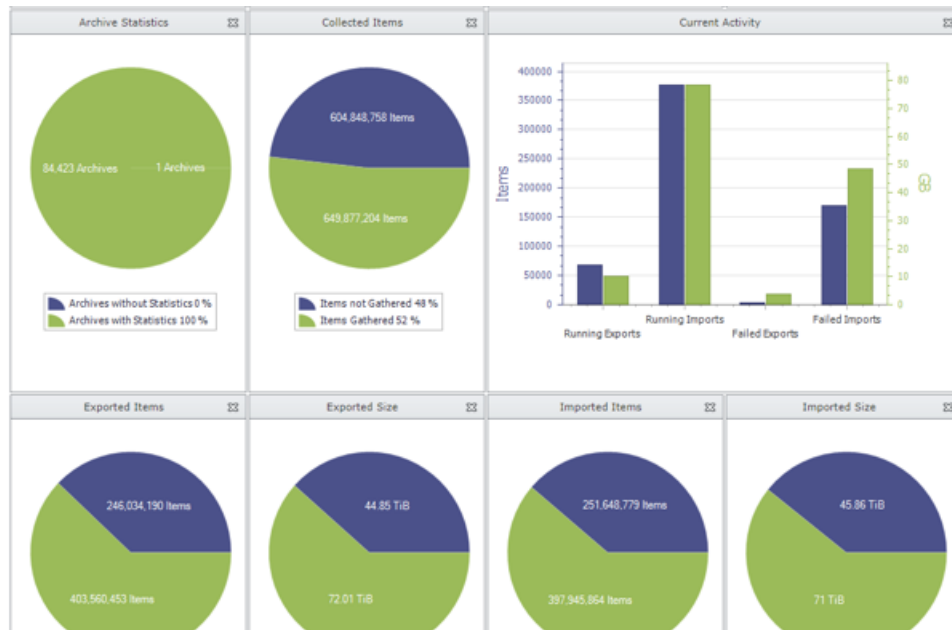
Overall things look healthy, however export is slightly slower than import. This means that eventually the system will run out of work to do.

Green box

Collection of archived item metadata is currently not happening. That could be completely normal.

Blue box

This again is showing that import is currently faster than export, and that may mean that the system will run out of work to do.



Lower down on the Progress & Performance page are more indicators about the overall health of the migration.

Stage 1

Stage 1 does all the heavy lifting in a migration, and if Stage 2 is not enabled, Stage 1 will continue topping up the migration (with any additional archived items from the source environment).

Conditional formatting is applied to results:

- Containers with no items are shaded blue
- Containers with rows in red are potential issues

More examples of the Stage 1 page showing a journal migration in progress:

Container	Full Name	Tags	Has Owner	Routed Count	Exported Count	Retriable Errors	Permanent Errors	Errors	Percentage	Imported Count	Retriable Errors	Permanent Errors	Errors	Percentage	Priority	Item Collection Enabled	Migration Enabled	Has Sub-Container	Stage 2 Enabled	Stage 2 Finished	Failed Item Threshold	Needs Provision	Apply Filter
13					5 000	5 000	0	0	0	100%	13 498	0	0	0	2%								
Primary Smtip Address																							
qtjt_ajeyi.fo@qtqa.onmicrosoft.com					0	0	2 500	0	0	0	0	0	0	0	100%								2 500
qtjt_alexander.ri@qtqa.onmicrosoft.com					0	0	2 500	0	0	0	0	0	0	0	100%								2 500
qtjt_althagafi.o@qtqa.onmicrosoft.com					0	0	2 500	0	0	0	0	0	0	0	100%								2 500
qtjt_anastasopoulos.k@qtqa.onmicrosoft.com					0	370	2 130	0	0	0	0	0	0	0	85%								2 500
qtjt_andrews.w@qtqa.onmicrosoft.com					0	146	2 354	0	0	0	0	0	0	0	94%								2 500
qtjt_annunziata.ne@qtqa.onmicrosoft.com					0	2 154	346	0	0	0	0	0	0	0	13%								2 500
qtjt_ayagi.m@qtqa.onmicrosoft.com					0	1 332	1 168	0	0	0	0	0	0	0	46%								2 500
qtjt_aronstein.ws@qtqa.onmicrosoft.com					0	2 500	0	0	0	0	0	0	0	0	0%								2 500
qtjt_asher.l@qtqa.onmicrosoft.com					0	2 500	0	0	0	0	0	0	0	0	0%								2 500
qtjt_balazs.g@qtqa.onmicrosoft.com					0	0	0	2 500	0	0	0	0	0	0	0%								2 500

The screenshot shows the 'Remedy Management' page in the Archive Shuttle application. The interface includes a sidebar with navigation options like 'Dashboards', 'Manage', 'Reports', and 'Support'. The main area displays a table of remediation activities.

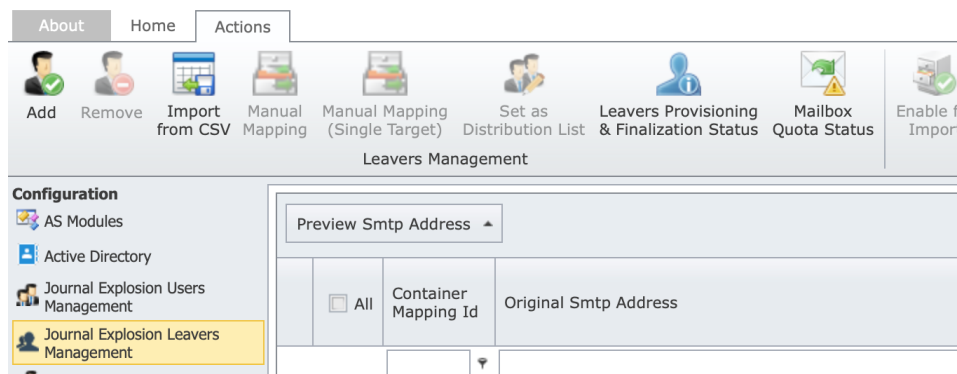
	Journal Explosion Remediation Id	Journal Explosion User Mapping Id	Journal Explosion Remediation Type Id	Remedy Data	Created Date Utc	Remedy Start Date Utc	Remedy End Date Utc	Remedy Failed Date Utc	Remedy Error	Remedy State	Apply Filter
<input type="checkbox"/>	1	1	Delete User Mapping	More Info...	2020-08-04 08:03:46 AM	2020-08-04 08:03:46 AM	2020-08-04 08:03:46 AM	2020-08-04 08:03:46 AM	dsadasdas	Finished	
<input type="checkbox"/>	2	2	Delete User Mapping	More Info...	2020-08-04 08:14:48 AM	2020-08-04 08:14:48 AM	2020-08-04 08:14:48 AM	2020-08-04 08:14:48 AM	dsadasdas	Finished	

Page 1 of 1 (2 items) | Page size: 30

Here you can see the status of any remediation activities which have taken place on user mappings.

Remediation status

If you're processing leavers, then reviewing the leaver provisioning status page is essential. You can access it from here:



Again this page can be sorted and filtered to try to pick out any problematic users.

Archive Shuttle

rob.wilcox@quadrotech-it.com

Finalize Retry Convert Rollover Mailbox Quota Status Refresh Reset Columns Load Save Export

Dashboard Progress & Performance System Health User Information Module Dashboard Manage Bulk Mapping Manual Mapping Existing Mappings Offline Media Stage 1 (Sync Data) Stage 2 (Switch User) Reports Link Statistics Performance Failed Items License Scheduled Reports Support Support Bundles

	User Sid	Primary Smtg Address	Link Name	Status	Processing Error	Is Finalization Postponed	Is Mailbox Quota Exceeded	Is Recoverable Items Quota Exceeded	Apply Filter
	S-1-5-21-1234567890-100000000-100000000-11344		0365 JE - JournalStore	Finalized					
	S-1-5-21-1234567890-100000000-100000000-20313		0365 JE - Other01	Licensing Started					
	S-1-5-21-1234567890-100000000-100000000-21639		0365 JE - JournalStore	Finalized					
	S-1-5-21-1234567890-100000000-100000000-41061		0365 JE -	Finalized					
	S-1-5-21-1234567890-100000000-100000000-41630		0365 JE -	Finalized					
	S-1-5-21-1234567890-100000000-100000000-42112		0365 JE -	Finalized					
	S-1-5-21-1234567890-100000000-100000000-42394		0365 JE -	Licensing Started					
	S-1-5-21-1234567890-100000000-100000000-43371		0365 JE -	Finalized					
	S-1-5-21-1234567890-100000000-100000000-4461		0365 JE -	Finalized					
	S-1-5-21-1234567890-100000000-100000000-48066		0365 JE -	Licensing Started					
	S-1-5-21-1234567890-100000000-100000000-48091		0365 JE -	Finalized					
	S-1-5-21-1234567890-100000000-100000000-48743		0365 JE -	Licensing Started					
	S-1-5-21-1234567890-100000000-100000000-48863		0365 JE -	Licensing Started					
	S-1-5-21-1234567890-100000000-100000000-48901		0365 JE -	Licensing Started					
	S-1-5-21-1234567890-100000000-100000000-48933		0365 JE -	Licensing Started					
	S-1-5-21-1234567890-100000000-100000000-49425		0365 JE -	Finalized					

Base system

Monitor the base or basic Journal Commander System will ensure fundamental operational activity. One of the ways to proactively monitor the system itself is to use the HealthCheck API.

Completing a journal migration

In this section we'll review some of the things that need to take place at the end of a migration using Archive Shuttle.

Progress & Performance


The summary information at the top of the Progress & Performance gives a good indicator of what has progressed in the migration.

User finished migration:	100%
Stage 1 (Synced):	99.93%
Stage 2 (Switched):	100%
Basic statistics	
Users	5 075
Mailboxes	112
Archives	3 810
Archived Items	203 161 244
Archived Items Size (Compressed)	23,72 TB
Mappings	3 774
Throttled Mappings Count	0
Enabled for Sync (Stage 1)	3 774
Enabled for Switch (Stage 2)	3 774
Finished Migration	3 774
Routed Items	203 024 847

Failures

Failures during the migration are inevitable.

Assuming that the source environment is fairly healthy the actual result in most migration is just a tiny fraction of a percentage, but a small fraction can still appear to be tens of thousands of items:



In this example 140,000 failed items is actually 0.069% within the migration.

Further, with this example, most items failed to export from the source meaning that source integrity was a possible issue.

Logging

Archive Shuttle performs logging for activities in two locations, both of which can be useful for troubleshooting purposes. Logging can be configured to record data at a number of different levels: Info, Debug, Trace, Warn, Error, and Fatal.

Module-level logging

Each module records log information at the INFO level though that can be overridden if required. When the modules are installed, the location of these log files can be chosen. By default, the location is inside the program files folder in a subfolder called Logs.

i **NOTE:** When the module logs (or default core log from System Configuration) for particular modules are set to a different logging level than 'Info', the logging level is set back to 'Info' after a specific period of time (default 12 hours).

The current location for the log files can be seen on the Modules page in the Archive Shuttle Admin Interface.

In addition to logging data locally, each module also transmits this log information to the Archive Shuttle Core, using a Web Service.

Each module can log extended information about item-level migration details; for example, on successful migration of an item an EV Import module may log the Archive Shuttle reference for an item, the source item ID and the target item ID.

Core-level logging

The Archive Shuttle Core logs information as well as the modules. The location where the log files will be written can be configured during the installation of the Archive Shuttle Core. In addition, to the core operations the Archive Shuttle Core will also log user interface actions to a file on the Archive Shuttle Core.

The Archive Shuttle Core also receives logging from each module. These files have .Client in their filename. This means it is not normally necessary to get log files from the servers where modules are installed.

Logging levels

There are several logging levels in log files, that vary in severity.

Solution

Here are listed all allowed logging levels in order from most severe to least severe along with explanation and example when to use it:

Fatal

- The highest log level, only the most important information is logged
- Minimum entries in the log file

Error

- High log level, application crashes/exceptions are logged
- Average number of entries in the log file

Warn

- High log level, incorrect behavior but the application can continue
- Average number of entries in the log file

Info

- Default log level, logging contains all above, used for normal behavior like item export
- Average number of entries in the log file

Debug

- Low log level, executed queries, user authenticated, session expired, etc
- Average number of entries in the log file

Trace

- The lowest log level, everything is logged, begin method X, end method X, etc.
- Maximal number of entries in the log file

Reporting

i | **NOTE:** The reports are in PDF format.

Each report can be configured to email one or more recipients, and is controlled via separate schedules, or, optionally triggered immediately using the 'Send Now' button.

i | **NOTE:** In order to send out reports it is necessary to use the System Configuration page and enter information in the SMTP Settings page.

This is controlled on the Scheduled Reports screen in the Archive Shuttle User Interface.

Link statistics

This page shows an overview of all containers that have been switched to target and shows the status in the Stage 2 workflow.

i | **NOTE:** Additional columns are available if the 'Advanced view' is enabled (under the 'Views' tab) though this makes the user interface slower if there are a large number of mappings.

The following basic information about each container is visible in the grid view:

- Source User Name
- Group
- Source Container Type
- Link (for the source)
- Target User Name
- Group
- Target Container Type
- Link (for the target)
- Command

The current command which is being ran by Archive Shuttle.

- Status: The status of the command
- Error: Information relating to errors

- Date/Time
- Next Command: The next command to be executed by Archive Shuttle
- Finished

Actions to be performed on the Workflow Status page

These actions can be performed on this page:

Refresh: Refreshes the tables data.

Columns: Allows the selection of additional data columns, which might then be used to select groups of containers.

Reset: Resets the grid to the default columns, and removes any filters.

Load: Loads a saved filter/page-layout.

Save: Saves the current filter/page-layout. Click the **Shared** checkbox to share the layout with other users.

Add to Group: One or more containers can be added to an existing or new group. This group membership can then be used for filtering and batching of users for migration.

Reset Workflow Status: Can be used to retry the current command.

Skip: Skip the current command in the workflow.

Pause: Pause the workflow.

Resume: Continue from a paused workflow.

Change Policy: Allows a new policy to be chosen and the workflow restarted with that policy.

Export to .pdf, .xls, .xlsx, .csv: Allow data to be exported to several formats.

Performance dashboard

The Performance Dashboard provides information about the migrations being performed on an hour by hour basis. The screen allows an administrator to select information relating to:

- The last day
- The last week
- The last month
- A custom time frame

Information about the migration is then displayed in a tabular format as follows:

Item	Description
Hour	Date/Time
Exported Count	Number of items exported
Exported Size	Size of all data exported
Imported Count	Number of items imported
Imported Size	Size of all data imported

It is also possible to view the information on a per-link basis, and each set of data is also displayed as a graph at the right hand side of the page.

Optionally, you can configure the page's widgets to display compressed sizes.

Scheduled reports

Archive Shuttle includes a number of pre-built reports can be scheduled to be delivered to different recipients. This is controlled on the Scheduled Reports screen in the Archive Shuttle User Interface.

The current reports which are available:

Report Name	Description
Performance	<p>This report shows export and import performance in a bar graph both as a summary and per source link for:</p> <ul style="list-style-type: none"> • The last day • The last week • The last 30 days • The last 3 months • The last 12 months • A custom time frame can also be chosen
Stage2Reports	<p>This report shows information relating to Stage 2, including:</p>

Report Name	Description
	<ul style="list-style-type: none"> • Containers that have finished Stage 2 in the last 24 hours • Containers still in Stage 2 • Containers which are in Stage 2 and have errors • Containers in Stage 2 which may be stuck/hung • Containers in Stage 2 which have been paused.
FailedItemReport	<p>This report shows items which have failed migration. The report is divided up per link, and shows:</p> <ul style="list-style-type: none"> • Items which have failed to export • Items which have failed to import

i | **NOTE:** The reports are in PDF format.

Each report can be configured to email one or more recipients, and is controlled via separate schedules, or, optionally triggered immediately using the **Send Now** button.

i | **NOTE:** In order to send out reports it is necessary to use the **System Configuration** page and enter information in the **SMTP Settings** page.

Failed items

This page shows an overview of any per-item failures that Archive Shuttle has encountered during the migration. When the screen is first entered a particular link should be chosen from the drop down list. This will filter the Failed Items screen to show only items relating to that link.

The following basic information about each failure is logged in the grid view:

- Container Mapping ID: The Archive Shuttle ID for the container mapping – reference only
- Item Routing ID: The Archive Shuttle ID for the item which has failed – reference only
- Is Exported Failed?: An indicator field to show whether it is an export failure
- Is Import Failed?: An indicator field to show whether it is an import failure

- **Is Failed?:** An indicator field to show whether it is concerned that the error is permanent (e.g., file not found on exporting an item)
- **Item Routing Error Count:** The current error count for the item
- **Error Text:** The error which was reported to Archive Shuttle from the underlying module, or from Exchange/Enterprise Vault
- **Last Date Time UTC:** The last date/time where this was reported by Archive Shuttle
- **In FTP:** An indicator to show whether the item was uploaded to the Quadrotech FTP Server.
- **Download Link:** A hyperlink to the failed item.

The Failed Items page also allows an administrator to submit failed items for reprocessing. This is achieved by selecting one or more items and selecting the button on the Actions Bar. It is also possible to reprocess all failed items. When reprocessing is selected, this will have the effect of resubmitting the item to the appropriate module. For example, if the issue is that the item could not be exported, then that item will be resubmitted to the export module. Likewise, if the issue is that the item could not be ingested into Office 365, then the item will be resubmitted to the Office 365 module.



NOTES:

- Reprocessing failed items will reset the Item Routing Error Count for those items.
- As of Archive Shuttle 9.5, FTP settings (URL, username, password) are no longer hardcoded for the ExportFailedItemsToFTP command and use the settings from System Configuration.
- As of Archive Shuttle 10.4, there is also the ability to permanently fail permanently fail certain export or import failures to other sources after the first attempt.

This action of reprocessing failed items can take a few minutes before the appropriate modules receive the command, perform the action, and feedback the results.

An additional option on the Failed Items screen is to 'Move Selected Items'. This is specifically for the situations where items have failed to be imported to the target, and never can be. By selecting one or more of these items, they can be moved from the staging area to a different area specified on the target link (via an option on the Links Page).



NOTE: Permanently failed items might include items which are too large to be ingested into the target (e.g., in an EV to Exchange migration where message size limits are in place).

In the case of a missing setting for an operation, a notification will be shown.

Index fragmentation

Archive Shuttle has a reliance on a performant SQL Server in order to achieve high throughput of items during migration. If SQL is struggling to handle the load, performance will drop. The Index Fragmentation page shows some key metrics about the SQL Indices associated with each link

database. These can get fragmented over time during the Enterprise Vault Collection and export/import processes. The information on the index fragmentation page is updated hourly.

The screen highlights tables/indices that are of particular concern according to this:

Row Highlighting	Reason
No highlighting	Fragmentation is not significant, or the number of pages in the index is not over 1000.
Yellow	Fragmentation is between 10 and 30 % and the page count in the index is more than 1000.
Red	Fragmentation is over 30 % and the page count in the index is more than 1000. Accessing data associated with this table/index will not be performant.

The recommended actions to take are as follows:

Row Highlighting	Action
No highlighting	No action required
Yellow	Perform an “Index Reorganization”
Red	Perform an “Index Rebuild”

More information specific to SQL Server and management of the Archive Shuttle databases can be found in the *Quadrotech Archive Shuttle SQL Best Practices Guide*.

License information

The License page in the Admin Interface gives an overview of the usage to-date of Archive Shuttle against the purchased amount of licenses and license volume. It also shows key information about when the license will expire.

If required, a license can be extended and then added to the system. See Adding a New License.

Data on this page is refreshed hourly and shows exported data quantities at that time. User licenses are used when a user enters Stage 2.

A new license can be uploaded, if required.

License expiry notifications may be seen from time to time as the migration progresses. The notifications will appear as follows:

Notification	Description
Date	When there is 7 days left before the license expires
Volume	When there is 3% left of the volume license
Users	When there is 5% of the user license count remaining

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. Connectors for PST as a source, and PST or Exchange as target are available through preview via Product Management.

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 Explosion tab and create a new journal archive migration 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 Explosion 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 Explosion 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 Explosion 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.

Check/confirm general location for ingest items

In the System Configuration section of the user interface, go to the **Journal Explosion Settings** tab.

On this page you can set a general folder to be used when ingesting the items into the target container. This can be overridden on an individual user mapping. If one is not specified then the **Recoverable Items Folder** is used in the target container.

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 Explosion, 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.

Export Starts

Journal archive migration is done during export. As the item is exported, the P1 information is reviewed for SMTP addresses. These addresses then need to be mapped, as described in the next section.

Journal Transformation user management

As items export from the source container and are analyzed for recipient information a list is built up that needs to be reviewed, mapped, and enabled for migration. This is done on the **Journal Explosion 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 **Archive Shuttle > System Configuration** and select Journal Explosion User Management.
2. From the Select Journal Explosion Name drop-down, select All or the required Journal Explosion container mapping.
3. Click **Add**.
4. In the Enable Users for Journal Explosion 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 Mappings 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 Explosion User Management.
2. From the Select Journal Explosion Name drop-down, select All or the required Journal Explosion Container Mapping.
3. Click Import from .csv.

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 Explosion and is added to the User Mappings 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 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 Mappings** 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 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

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)

<input type="checkbox"/> All	Container Mapping	Original Smtip Address	Mapped Smtip Address	Is Valid Preview Address	Group Name	Folder Name Policy	Has Owner	Is Distributor	User Mapping Associate Status	Is Existing Preview User	Is Staged/Finished	Item Count	Total Size	Appl. File
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Preview Smtip Address: mirrye_asqata-leaver-01@mirryquadrtech.onmicrosoft.com (Count=1)														
<input type="checkbox"/>	1	asqata-leaver-01@hotmail.com	mirrye_asqata-leaver-01@mirryquadrtech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			1	91 KIB	
Total:												1	91 KIB	
Preview Smtip Address: mirrye_asqata-leaver02@mirryquadrtech.onmicrosoft.com (Count=4)														
<input type="checkbox"/>	1	asqata-leaver02@gmail.com	mirrye_asqata-leaver02@mirryquadrtech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			1	90 KIB	
<input type="checkbox"/>	1	asqata-leaver02@outlook.com	mirrye_asqata-leaver02@mirryquadrtech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			1	91 KIB	
<input type="checkbox"/>	1	asqata-leaver-01@outlook.sk	mirrye_asqata-leaver02@mirryquadrtech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			1	91 KIB	
<input type="checkbox"/>	1	asqata-leaver02@hotmail.com	mirrye_asqata-leaver02@mirryquadrtech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			2	186 KIB	
Total:												5	458 KIB	
Preview Smtip Address: mirrye_asqata-leaver03@mirryquadrtech.onmicrosoft.com (Count=1)														
<input type="checkbox"/>	1	asqata-leaver03@outlook.com	mirrye_asqata-leaver03@mirryquadrtech.onmicrosoft.com	✓		Leavers (Default)	✓		✓			1	92 KIB	
Total:												1	92 KIB	
Preview Smtip Address: mirrye_asqata-qax10JE-recyle@mirryquadrtech.onmicrosoft.com (Count=1)														
<input type="checkbox"/>	1	asqata-qax10JE-recyle@qa.quadrtech-it.com	mirrye_asqata-qax10JE-recyle@mirryquadrtech.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 Quadrotech 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 Quadrotech 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.](#)

Selective migration based on recipients collected from EV Index

Recipients can be collected separately. It has been seen to improve performance, and reduces the staging area space requirements.

A Journal Transformation setting enables this feature.

Journal Transformation Settings

- EV Collector Module
- EV Export Module
- EV Import Module
- EV Provisioning Module
- Shortcut Process Module
- Exchange Import Module
- Office 365 Module
- PST Export Module
- Native Import Module
- EAS Zantaz Module
- SourceOne Module
- Dell Archive Manager Module

☐ Delete items from staging area after initial export
☐ Delete items without journal transformation import routings
 Time to expire inactive items (hours)
 Mailbox Quota Exceeding Threshold (%)
☒ Ingest users with most items first
☐ Skip items which not belong to user mapping
☒ Export items for users enabled for ingestion only
☒ Enable All SMTP Addresses Discovered Upon Exploding

Once it is enabled, when the bulk mapping wizard is used part way through the wizard you will be asked if you want to 'export items for users enabled for migration only'. This will now default to true, but you can disable it on that mapping if you wish, and the old way will be used.

On the summary screen, you will see confirmation that Archive Shuttle will use the EV Index.

Items are collected. Once some item data has been collected, the collection for sender and recipient information is handled via the EV Collector, and it will query the EV Index. All of the recipients will be collected and sent to the core. This collection can be retried if required.

The 'Existing Mapping page' has an advanced view option. If this is selected the mappings will show 'export based on ev index' if we've used this feature. It can be disabled on a mapping by mapping basis as well. This is dynamic and can be done on a mapping which is part way through being processed.

Journal Transformation

Change method actively used for export JT times.

Column is available only when Advanced view is triggered

Mapping ID	Source	Target	WF Policy	Filter Policy	Priority	Item Collect Enabled	Migrate Enabled	Stage 2 Enabled	Stage 2 Finished	Needs Provision	Export Based on EV Index	Apply Filter
1	mirjournalproduct	QA EV10 01 VSG1	Ownerless / Unass	QATA_JournalExpl Journal Transform Default (No Filter)		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	
2	BLPEJ01 (Owner)	QA EV10 01 VSG1	Ownerless / Unass	QATA_JournalExpl Journal Transform Default (No Filter)		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	
3	JEPproductionPST (VSG1V51	Ownerless / Unass	QATA_JournalExpl Journal Transform Default (No Filter)		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	
4	Archive Shuttle - r	QA EV10 01 VSG1	Ownerless / Unass	QATA_JournalExpl Journal Transform Default (No Filter)		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	

Selective migration is also visible on the Stage 2 screen, after adding additional columns:

- Export based on EV index
- Collecting count
- Collected count
- Retryable errors (Sender recipient collection)
- Permanent errors (Sender recipient collection)

The layout can be saved, if required.

If there are errors during collection, the mapping won't be marked in 'red' as they're only failures during the collection itself. Sender / Recipient collection can be retried from the Journal Transformation tab.

User mappings can be created based on this collected sender/recipient information, and the mappings enabled. When this is done, the EV Export module will then begin to export the data

for those mappings. This reduces the storage area space needed for a Journal Transformation migration.

Leavers are supported in this scenario as well.

The EV Collector must run under a service account which has access to query the index for the archives mapping for migration.

Finally the failed item page has been enhanced to introduce a new tab 'Sender Recipient Collection'. This will show any issues with collection of sender/recipient information from the EV Index. These errors can be retried, if required.

There is a fallback mechanism for collection. Collection will retry 10 times, before being marked as permanently failed. These items are exported in the old way, in order to get the sender/recipient information from the item itself, and the sender/recipient information gathered from that.

Journal Transformation detail grid

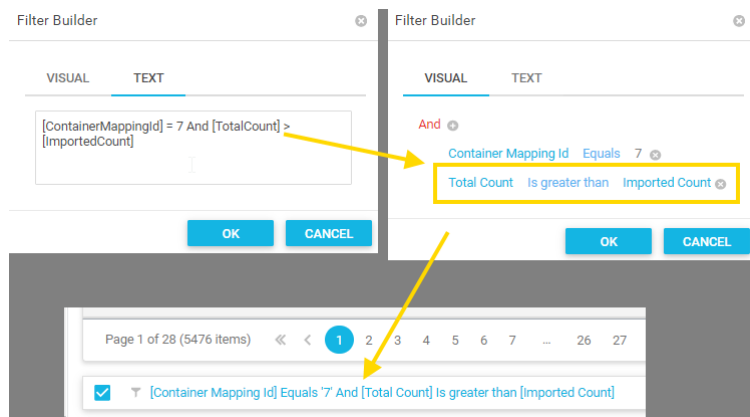
The Journal Transformation detail grid on the Stage 1 page of the user interface contains many details relating to the item level detail surrounding a mapping which is in migration.

Here is an example of how the screen might look:

Mapped SMTP Address	Container Mapping Id	User Id	Original SMTP Address	Folder Name Policy	Has Mailbox	Is Ready For Import	Enabled For Import	Is Deleted	Imported Count	Imported Size	Ready For Import Count	Ready For Import Size	Total Count	Total Size	Apply Filter
Total															
Mapped SMTP Address: asqata-qao365JE-rec7@qa.quadrotech-it.com (Count=1)															
		2: 5-1-5-21-2153452352-22	asqata-qao365JE-rec7@qa.quadrotech-it.com	Leavers (Default)		Ready				N/A	N/A	N/A	N/A	N/A	N/A
Total															
Mapped SMTP Address: asqata-qao365JE-rec1@qa.quadrotech-it.com (Count=1)															
		2: 5-1-5-21-2153452352-22	asqata-qao365JE-rec1@qa.quadrotech-it.com	Leavers (Default)		Ready			7	651 KiB	0	0 Bytes	7	651 KiB	
Total															
Mapped SMTP Address: asqata-qao365JE-rec2@qa.quadrotech-it.com (Count=1)															
		2: 5-1-5-21-2153452352-22	asqata-qao365JE-rec2@qa.quadrotech-it.com	Leavers (Default)		Ready			5	467 KiB	0	0 Bytes	5	467 KiB	
Total															
Mapped SMTP Address: asqata-qao365JE-rec3@qa.quadrotech-it.com (Count=1)															
		2: 5-1-5-21-2153452352-22	asqata-qao365JE-rec3@qa.quadrotech-it.com	Leavers (Default)		Ready			7	650 KiB	0	0 Bytes	7	650 KiB	
Total															
Mapped SMTP Address: asqata-qao365JE-rec4@qa.quadrotech-it.com (Count=1)															
		2: 5-1-5-21-2153452352-22	asqata-qao365JE-rec4@qa.quadrotech-it.com	Active Users (Default)		Ready				N/A	N/A	N/A	N/A	N/A	N/A
Total															
Mapped SMTP Address: minjye_asqataleaver02@otqa.omicronsoft.com (Count=3)															
		2: 5-1-5-21-1234567890-10	asqataleaver02@gmail.c	Leavers (Default)		Ready			1	90 KiB	0	0 Bytes	1	90 KiB	
		2: 5-1-5-21-1234567890-10	asqataleaver02@gmail.c	Leavers (Default)		Ready			1	91 KiB	0	0 Bytes	1	91 KiB	
		2: 5-1-5-21-1234567890-10	asqataleaver02@gmail.c	Leavers (Default)		Ready			2	186 KiB	0	0 Bytes	2	186 KiB	
Total summary for mapped JT user mappings/JT mapping depends on grouping															
Total															

As shown in the screenshot, it is possible to perform grouping to aid with management of the migration.

In addition the filtering possibilities have the ability to enter free text:



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](#).

To PST

This can be found in the *Quadrotech Archive Shuttle Migration Guide*.

Which Office 365 license to use for leavers or virtual journal

Archive Shuttle has ability to provision mailbox in Office 365 and set flags like “Legal Hold”, “Hide from GAL”, etc for “Leavers”, or “Virtual Journal” type of migrations.

There are several differences between each Office 365 Plan/License options. This can be checked [here](#) or [here](#).

The biggest difference from the migration perspective is the mailbox size/storage limits as the data will be migrated there:

Feature	Microsoft 365 Business Basic and Standard	Microsoft 365 Business Premium	Microsoft 365 Enterprise E3/E5	Office 365 Enterprise E1	Office 365 Enterprise E3/E5	Office 365 Enterprise F3
User mailboxes	50 GB	50 GB	100 GB	50 GB	100 GB	2 GB
Archive mailboxes	50 GB	1.5 TB	1.5 TB	50 GB	1.5 TB ¹	Not available
Shared mailboxes	50 GB	50 GB	50/100 GB	50 GB ²	50/100 GB	50 GB
Resource mailboxes	50 GB	50 GB ³	50 GB	50 GB	50 GB	50 GB ³
Public folder mailboxes	100 GB	100 GB	100 GB	100 GB	100 GB	100 GB
Group mailboxes	50 GB	50 GB	50 GB	50 GB	50 GB	50 GB

[Click here](#) for more information on storage limits.

Therefore, Archive Shuttle recommends to use Enterprise E3 or E5 licenses for migration as “Leavers”. If the mailboxes will not exceed 50GB, use an E1 license.

Enterprise E1 and all Microsoft 365 Business licenses include Exchange Plan 1 licenses. Enterprise E3 and E5 licenses include Exchange Plan 2 licenses. [Click here](#) for more information.

| NOTES:

- Use of an E1 license does not include data loss protection. Microsoft recommends to use an additional Exchange Online archiving service. For more information, [click here](#). You may also need to confer with Microsoft to determine whether the terms of your licensing agreement are impacted by the legal hold feature.
- Review this associated article.
- Archive Shuttle also supports the government equivalent of each of these licenses mentioned.

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.

Journal envelope processing

Prior to Enterprise Vault 10.0.3 journal envelope information was not returned by the Enterprise Vault API. In versions of Enterprise Vault prior to 10.0.3 a process called 'Envelope Reconstruction' is used to rebuild a journal message, where the API-retrieved item can be added, thus giving the entire P1 & P2 message.

Envelope Reconstruction does have a performance impact but is essential for maintaining data integrity and Chain of Custody for journal archive migrations.

In testing the performance impact it was observed that if extract performance was N with an Enterprise Vault 9.0.4 server doing the journal extract (and envelope reconstruction) and that machine was upgraded to 10.0.4, the extract performance would be between 4x and 6x faster than N.

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 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.

EV collector journal collection logs

Occasionally, it may be necessary to check the EV Collector log files if there are any issues present. With large journal archives, the collection can take some time, and it might be important to check for collection issues.

First, Archive Shuttle will need to find and open the log files. They are stored on:

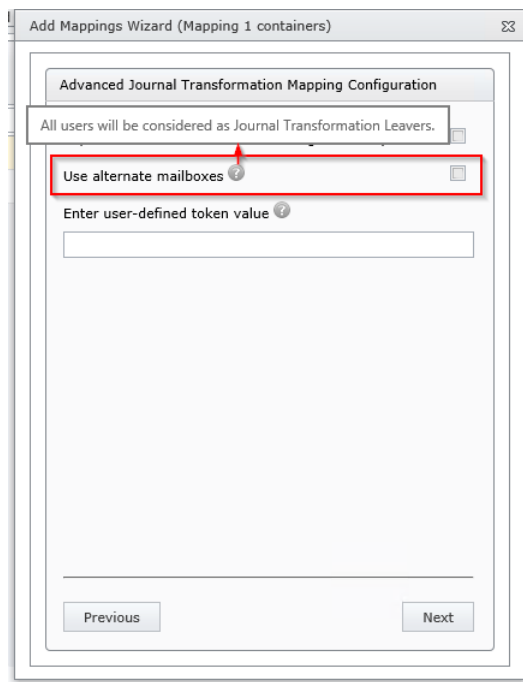
- a) Machine where are the modules installed (ie: C:\Program Files (x86)\QUADROtech\Logs), in many cases it is the Source EV server. The log file name is "AS_EVCollector_Log.txt"
- b) Machine where is core installed (ie: C:\Program Files\QUADROtech\Logs). You will find there two different log files for EV collector, first is the Client (Module) log file with name "ArchiveShuttle.WebServicesLog.HOSTNAME.EVCollector.Client.txt", and the second one is the core log ArchiveShuttle.WebServicesLog.HOSTNAME.EVCollector.txt

Second, the log level may need to be changed to TRACE to collect detailed information.

Alternate mailbox support for discovered addresses associated with active users

In a Journal Transformation migration you might want to migrate active user content into inactive alternate mailboxes. This may simplify the migration and reduce project complexity. You may also want to do this if active users recently departed the organization or re-join after leaving.

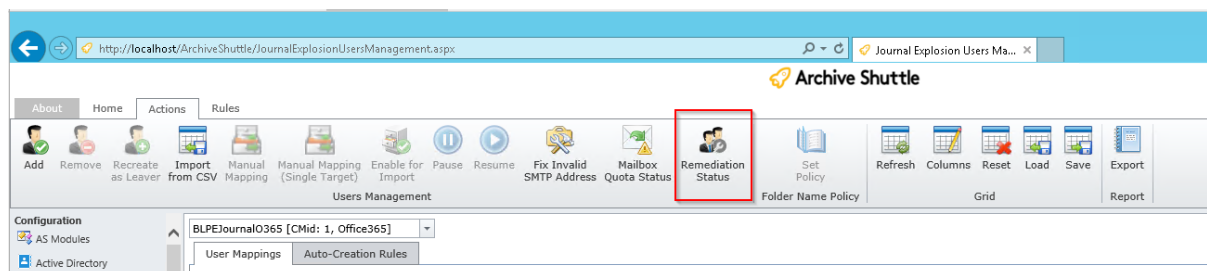
In the Bulk Mapping wizard, when you select the Journal Transformation options there is an option called 'Use alternate mailboxes':



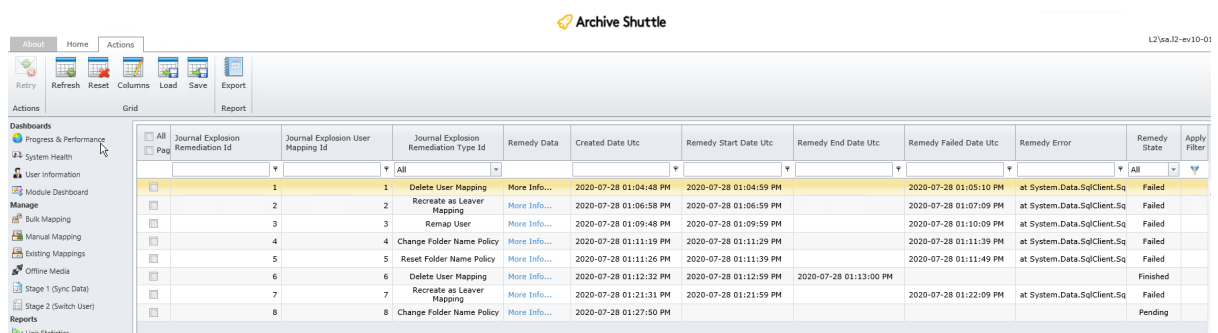
If the option is enabled all collected Journal Transformation users will be treated as leavers after consolidation. This means that the users will be managed from the Journal Transformation Leavers Management page in the user interface.

Retry failed remediations

When performing a Journal Transformation it is possible to retry any remediation actions which fail.



To start this process click on the button in the user interface, and a new page will be displayed:




More information can be obtained on this screen for each action, and each action can be retried. If a retry operation is performed it's executed immediately.

Support handling of invalid or changed SMTP addresses

Archive Shuttle has the option to handle invalid or changed addresses. A changed address can occur when a primary SMTP address is changed on a user object.

When Archive Shuttle performs an Office 365 Sync operation, the user mapping in Journal Commander will be marked as 'changed address'.



The 'Repair changed address' button can be used to update/replace the Primary SMTP Address in all the appropriate Archive Shuttle database tables for the selected user mappings.

If a user mapping is not enabled for import the change of address will be immediate, as there is only a small update required. For a user mapping which is enabled for import the change of address may take a few minutes because there could be a large number of updates required.

Providing accepted domains for a migration

Accepted domains are used in a Journal Transformation configuration to limit which domains are enabled for migration.

Consider a large journal archive where there might be hundreds of SMTP domains contained within messages. If you know you are only interested in:

domain1.com

domain2.com

and no others.

You can set the Accepted Domains in the System Configuration to:

domain1.com

domain2.com

and only those will be eligible for enabling for migration.

Why legal hold needs to be applied for data retention after migration

As part of any Archive Shuttle migration, all active user mailboxes must be put on litigation hold prior to any migration activity. This is not done by the product. It should be applied before migration begins and remain in place until the end of the migration project. By default Archive Shuttle does not allow migration to active user accounts during Journal Transformation if litigation hold is not applied.

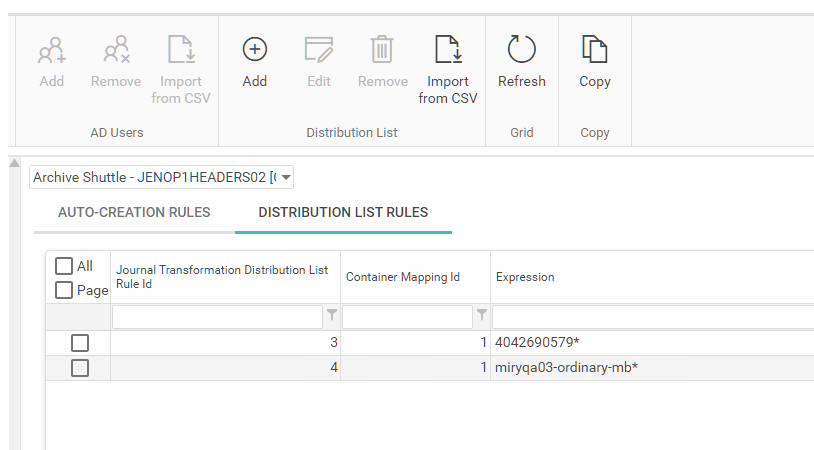
- Preserve mailbox items deleted by users or automatic deletion processes such as MRMs
- Preserve items indefinitely or for a specific duration
- Enable In-Place eDiscovery searches of items placed on hold

Once litigation hold is enabled the storage quota for the Recoverable Items folder is automatically increased from 30 GB to 100 GB for the primary mailbox and unlimited for the EOA (Auto expanding feature is required)

Automatically identify known distribution lists and mark them as such

Under Journal Transformation, in the user interface, there is an option called 'Auto-creation Rules Management'. This option allows an operator to create rules around email addresses to identify certain email addresses as distribution lists.

Operators can add a rule to automatically mark ‘mailboxes’ matching the expression as distribution lists. Recipients are evaluated frequently and compared against a list of known distribution lists for their organization. When a match is made that recipient is marked as ‘set as distribution list’ in the solution automatically.



Enable logging

In this article we'll discover information about log files and logging that can be performed by Archive Shuttle.

Logging Overview

All log files are centralised on the core server under the logging directory specified during installation.

It is rare to need to access module servers in order to review log files.

Logs are:

- Generated per module
- Roll over automatically (10 rollover files are kept)
- Default to 'info' level reporting

Important Log Files

Several important log files include:

WEBUI.TXT



This is the log file for all user interface interactions

ArchiveShuttle.WebServicesLog.txt

This is the web services log file. It contains server side information about module to core communication. It can show which tasks or commands were queued and executed.

ArchiveShuttle.Service.Log.txt

This contains information about execution of database updates, and scheduled tasks.

 JournalCommander.WebServicesLog.I2-jomu-84.EVImport.Client.txt
 JournalCommander.WebServicesLog.I2-jomu-84.EVImport.Core.txt

Module core logs

These have a *.ModuleName.Core.txt suffix. They contain communication between the core and applicable module.

Module client logs

These have a *.ModuleName.Client.txt suffix.



NOTE: All log files from a specific machine include the machine name in the filename, for easy identification.

Logging Levels

- Current processing statistics
- Summary of actions taken
- Duration of actions taken

Here is an example:

```
OnPerformanceTimerElapsedImpl|Last 2 minute(s) statistics: Count: [0] | Items/s: [0.00] | kb/s: [0] | ACPU: [1] | MaxCPU: [6] | MinCPU: [0] | TSize: [0 kb] | ASize: [0 kb]
OnPerformanceTimerElapsedImpl|Last 2 minute(s) statistics: Total Physical Memory: [31 GB] | Free memory: [29057 MB]/[88.7 %] | Avg Free Memory: [29056 MB] | Process Memory: [0 MB]
```

Errors, warning, exceptions

Log levels are used to clearly identify exceptions. All warnings are logged at WARN logging level. All exceptions are ERROR or FATAL log levels.

How to identify what log to look in

Many issues are highly reproducible. Try to identify the conditions resulting in an issue. Being able to reproduce an issue on demand helps issue resolution and investigation.

Determining which log file to look into often comes with experience.

- webui – user interface errors/issues
- services – database upgrades and scheduled task execution
- *client* – problems with an actual module
- *core* – problems with retrieving and processing of commands

How to adjust logging levels

Go to **Configuration > Modules**, then select the desired module and choose the ‘Set Log Level’ button. Changes take place almost instantly and there is no need for modules to be restarted.

Core logging level is set via System **Configuration > General**.

Identifying the point of failure

At the highest level, Archive Shuttle recommends to:


- Enable trace level logging on the Core (if appropriate) and modules which may be involved in the issue
- Try to repeat the activity of action which resulted in an error (for example retrying failed items, or retrying a Stage 2 workflow command)
- Search the log files for ERROR
- Check the thread ID of that error message and search in the logs for everything that the thread has done around about the time of the failure.

Shared mailbox over quota rollover

If a mailbox is over quota, a new one can be created automatically and any remaining data for this transformation is ingested into the new mailbox.

The new mailbox has a _1 suffix added, eg John.doe_1@domain.com. Additional rollovers for this container will be _2, _3 and so on.

There is a button in the user interface so that an operator can perform this action:



This is an additional possibility in a migration. Previously an over quota shared mailbox could be converted to a normal mailbox, but that consumes an additional Microsoft license. By creating a new, shared mailbox additional licenses are not consumed. It is also shown to be faster ingesting into this type of mailbox compared with moving data from a primary mailbox to a secondary/personal archive mailbox.

Re-exporting a PST for a certain user

When migrating to PST, you might encounter corrupt PST files and need to re-export items for a certain user. Here's how to do it:

1. On the Stage 1 (Sync Data) page, expand the line that contains JC PST mappings.
2. Right-click on the Primary Smtp Address of the mapping that contains PSTs you want to view/re-import and select Show list of PSTs. A window opens showing associated PSTs.
3. Select the desired item(s) and click the Enable For Re-Export button.

NOTE: If you notice corrupt PST files that you want to re-export, but Stage 2 is complete, you'll need to manually reset the workflow before you can re-export the PST files.

Additional information

In step 2 above, when the window opens and displays the user's PSTs, you'll notice the window's URL includes the user, item database ID, and container mapping ID.

Limitations

A maximum of 10 PST files are processed (deleted/reset) for a Journal Explosion re-export. If needed, you can re-export remaining files after the first batch of 10 are ingested.

Adding Sender-Recipient SMTP addresses

Manually add an active Sender-Recipient SMTP address

1. Go to Configuration.
2. Select **Journal Transformation User Management** page.
3. From the **Select Journal Transformation Name** drop down, select **All** or the required Journal Transformation Container Mapping.
4. Click Add toolbar button.
5. In the **Enable Users for Journal Transformation** pop out window, select the required active Sender-Recipient SMTP and click **Save & Close** button.

The selected Sender-Recipient SMTP address is now enabled for use in Journal Transformation and is added to the **User Mappings** screen.

Add Sender-Recipient SMTP address using .csv file

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

The .csv file requires the format to be: **OriginalSMTPAddress,PrimarySMTPAddress(optional)**

1. Go to Configuration
2. Select the 'Journal Transformation User Management' page
3. From the 'Select Journal Transformation Name' drop down, select 'All' or the required Journal Transformation Container Mapping
4. Click Import from CSV toolbar button
5. Click Browse and select your .csv file
6. Ensure the option, 'Update existing user mappings' is NOT selected
7. Click Import

The selected Sender-Recipient SMTP address is now enabled for use in Journal Transformation and is added to the 'User Mappings' screen.

Manually enable an active address for import

To manually enable a Sender-Recipient SMTP address:

1. In the 'User Mappings' screen, Select the required Sender-Recipient SMTP address
2. Ensure the selected Sender-Recipient SMTP address status is:
 - a. Has Mailbox = Yes

- b. Is Ready for Import = Ready
 - c. Enabled for Import = No
 - d. Is Deleted = No
3. Then click the 'Enable for Import' toolbar button
 4. Ensure the 'Enabled for Import' status changes to 'Yes' (Enabled) for the selected Sender-Recipient SMTP address

Monitor progress of an 'Enabled for Import' Sender-Recipient SMTP address

To monitor the progress of the Enabled Sender-Recipient SMTP address Journal Transformation mapping:

1. Go to Manage & Operations > Stage 1
2. Select the required Journal Transformation container mapping
3. Expand the Journal Transformation container mapping to view the enabled Sender-Recipient SMTP sub-containers.

Troubleshooting

Sender-Recipient SMTP address does not appear

If the required Sender-Recipient SMTP does not appear in the 'Enable Users for Journal Transformation' pop out window when you attempt to add a user for Journal Transformation.

You will need to confirm the SMTP address is correct and that the Stage1 Journal Transformation mapping has completed exporting/identifying sender-recipient SMTP addresses. As only exploded sender-recipient SMTP addresses appear.

If you are adding a Sender-Recipient SMTP using csv file, ensure the file is formatted correctly and contains valid Sender-Recipient SMTP address.

If the export has completed for all Journal Transformation Container Mapping. Further investigation of the logs and database will be required by Archive Shuttle.

CSV Import Fails

Check the WebUI log for the tenant and you will see the UI activity, example below:

```
2020-02-10 14:41:43Z | 11405 | chs-as-01 | 53 | TRACE | ReadUsersFromFile | File [E:
\WebHotel\cola\UI\AppData\UploadTemp\2d4adde5-a376-4664-88a3-0d34e9cb8711.csv] was
processed, [1] lines were read with total of [0] user mappings
```

Monitoring a migration

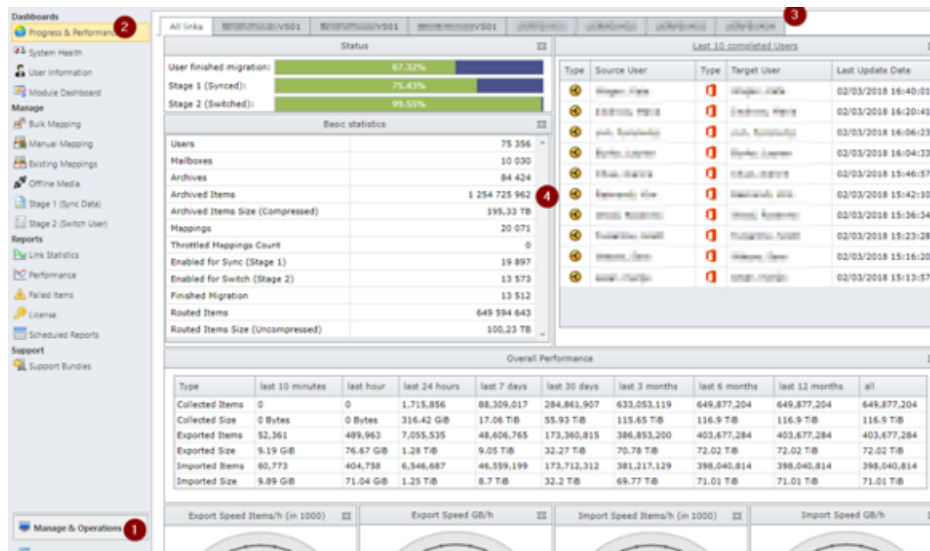
It's important to monitor Archive Shuttle during a migration. This article gives advice on the things to monitor.

Active systems

Archive Shuttle works best at scale. Migrating a few hundred items is not what it is intended to do. Fully utilized systems ensure that all aspects of the system, the workflow and processes are busy.

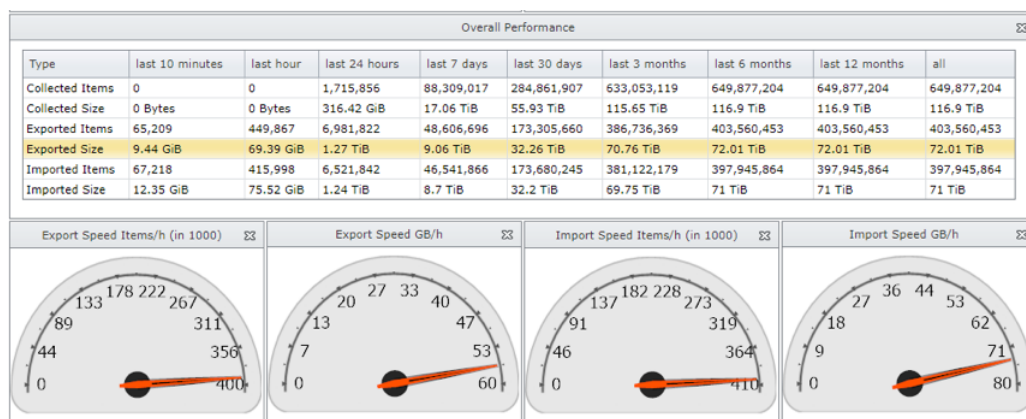
Progress & Performance

Take a look at the Progress & Performance screen in the user interface.



This dashboard is modular, meaning you can re-arrange widgets, remove ones which aren't needed and store that layout to use later. The page can be configured to automatically refresh (for example every 60 seconds).

The dashboard gives a great overview of the overall health of your migration.



In the graphic above we ask the question whether there is any potential issue with the highlighted statistics.

Red box

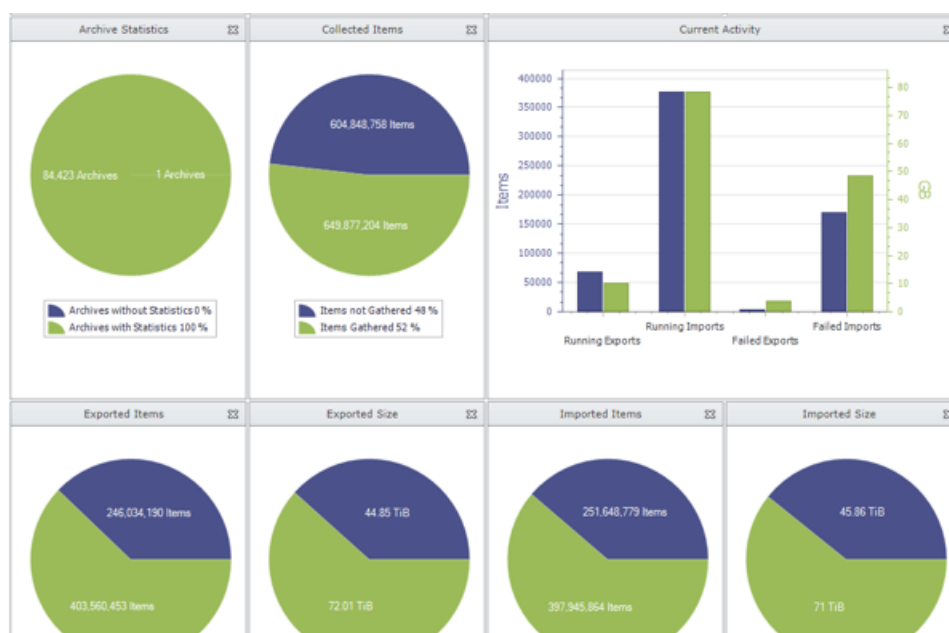
Overall things look healthy, however export is slightly slower than import. This means that eventually the system will run out of work to do.

Green box

Collection of archived item metadata is currently not happening. That could be completely normal.

Blue box

This again is showing that import is currently faster than export, and that may mean that the system will run out of work to do.



Lower down on the Progress & Performance page are more indicators about the overall health of the migration.

Stage 1

Stage 1 does all the heavy lifting in a migration, and if Stage 2 is not enabled, Stage 1 will continue topping up the migration (with any additional archived items from the source environment).

Conditional formatting is applied to results:

- Containers with no items are shaded blue
- Containers with rows in red are potential issues

More examples of the Stage 1 page showing a journal migration in progress:

<input checked="" type="checkbox"/> All	<input checked="" type="checkbox"/> Pag	Source	Export	Import	Priority	Item Collects Enabled	Migration Enabled	Has Sub-Container	Stage 2 Enabled	Stage 2 Finished	Failed Item Threshold	Needs Provision	Apply Filter
Full Name	Tags	Has Owner	Routed Count	Exported Count	Retryable Errors	Permanent Errors	Errors	Percentage	Imported Count	Retryable Errors	Permanent Errors	Errors	Percentage
13			5 000	5 000	0	0	0	100%	13 498	0	0	0	2%
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Primary Smtip Address	Not Export	Importing	Imported	Not Imported	Retryable Errors	Permanent Errors	Errors	Import Skipped	Percentage	Total Count	
qtjt_xjayi.fb@qtqa.onmicrosoft.com			0	0	2 500	0	0	0	0	0	100%	2 500	
qtjt_alexander.tl@qtqa.onmicrosoft.com			0	0	2 500	0	0	0	0	0	100%	2 500	
qtjt_althagafi.o@qtqa.onmicrosoft.com			0	0	2 500	0	0	0	0	0	100%	2 500	
qtjt_anastasopoulou.k@qtqa.onmicrosoft.com			0	370	2 130	0	0	0	0	0	83%	2 500	
qtjt_andrews.w@qtqa.onmicrosoft.com			0	146	2 354	0	0	0	0	0	34%	2 500	
qtjt_annunziata.ne@qtqa.onmicrosoft.com			0	2 154	346	0	0	0	0	0	13%	2 500	
qtjt_ayagi.m@qtqa.onmicrosoft.com			0	1 332	1 168	0	0	0	0	0	46%	2 500	
qtjt_aronstein.ws@qtqa.onmicrosoft.com			0	2 500	0	0	0	0	0	0	0%	2 500	
qtjt_asher.l@qtqa.onmicrosoft.com			0	2 500	0	0	0	0	0	0	0%	2 500	
qtjt_balan.g@qtqa.onmicrosoft.com			0	0	0	2 500	0	0	0	0	0%	2 500	

Page 1 of 21 (206 items) 1 2 3 4 5 6 7 ... 19 20 21

Page 1 of 1 (1 items) 1 2 3 4 5 6 7 ... 19 20 21

Page size: 50

Container Mapping Id Equals '13'

Activate Windows

Clear

Office 365 Mailbox Quota Status

On this page, regular mappings in which Office 365 is the target is displayed. To access the Mailbox Quota Status page:

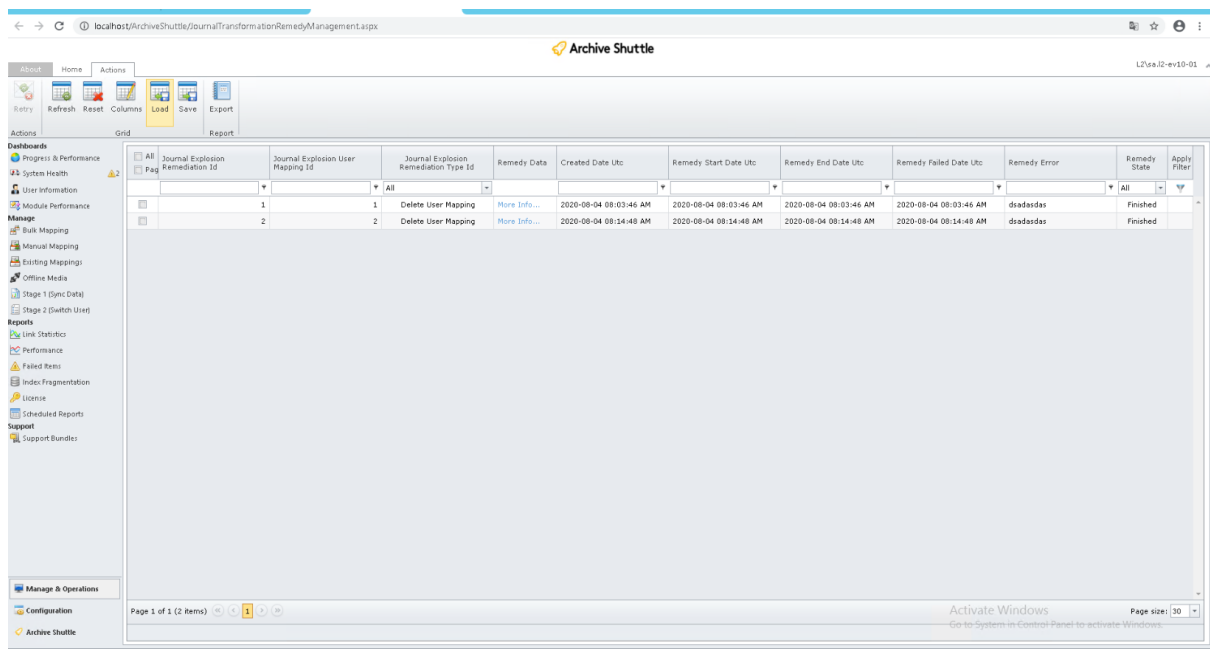
1. Go to **Stage 1 (Sync Data)**.
2. Click the Office 365 tab.
3. Click **Mailbox Quota Status**.

The page can be sorted by the quota used, and it is worth checking the recoverable items folder quota as well as the regular mailbox quota.

If the mailbox quota is reached for an Office 365 target mailbox, the migration will be disabled automatically. All mappings will be disabled if the quota is exceeded. You can check the threshold for quota by going to **System Configuration > Office 365 Module > Mailbox Quota Exceeding Threshold (%)** (the default quota value is 95%). This is indicated with a red check mark in the Is Quota Exceeded column in the grid. You can see if a migration has been disabled by going to the **Stage 1 (Sync Data)** page, and viewing the Migration Enabled column in the grid.

Leaver status

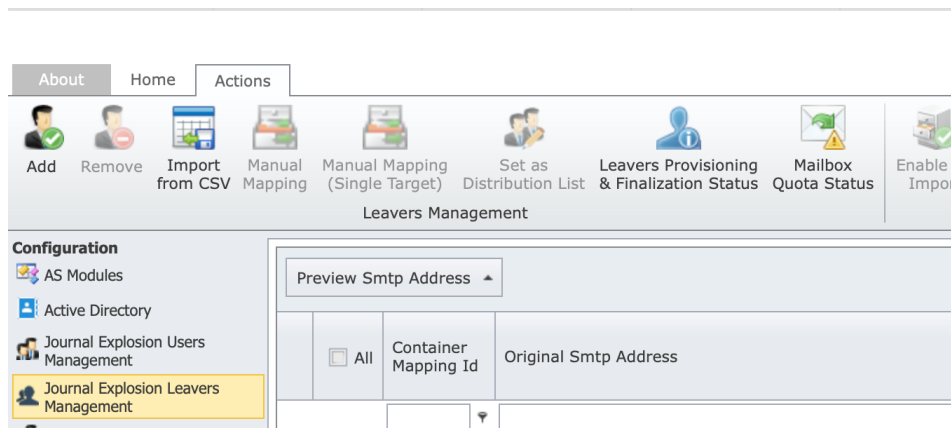
From the User Management page, you can access the remediation page.



Here you can see the status of any remediation activities which have taken place on user mappings.

Remediation status

If you're processing leavers, then reviewing the leaver provisioning status page is essential. You can access it from here:



Again this page can be sorted and filtered to try to pick out any problematic users.

Archive Shuttle

rob.wilcox@quadrotech-it.com

Finalize Retry Convert Rollback Mailbox Quota Status Refresh Reset Columns Load Save Export

Dashboard Progress & Performance System Health User Information Module Dashboard Manage Bulk Mapping Manual Mapping Existing Mappings Offline Media Stage 1 (Sync Data) Stage 2 (Switch User) Reports Link Statistics Performance Failed Items License Scheduled Reports Support Support Bundles

User Sid	Primary Smtg Address	Link Name	Status	Processing Error	Is Finalization Postponed	Is Mailbox Quota Exceeded	Is Recoverable Items Quota Exceeded	Apply Filter
S-1-5-21-1234567890-100000000-100000000-11344		0365 JE - JournalStore	Finalized					
S-1-5-21-1234567890-100000000-100000000-20313		0365 JE - Other01	Licensing Started					
S-1-5-21-1234567890-100000000-100000000-21639		0365 JE - JournalStore	Finalized					
S-1-5-21-1234567890-100000000-100000000-41061		0365 JE -	Finalized					
S-1-5-21-1234567890-100000000-100000000-41630		0365 JE -	Finalized					
S-1-5-21-1234567890-100000000-100000000-42112		0365 JE -	Finalized					
S-1-5-21-1234567890-100000000-100000000-42394		0365 JE -	Licensing Started					
S-1-5-21-1234567890-100000000-100000000-43371		0365 JE -	Finalized					
S-1-5-21-1234567890-100000000-100000000-4661		0365 JE -	Finalized					
S-1-5-21-1234567890-100000000-100000000-48066		0365 JE -	Licensing Started					
S-1-5-21-1234567890-100000000-100000000-48091		0365 JE -	Finalized					
S-1-5-21-1234567890-100000000-100000000-48743		0365 JE -	Licensing Started					
S-1-5-21-1234567890-100000000-100000000-48863		0365 JE -	Licensing Started					
S-1-5-21-1234567890-100000000-100000000-48901		0365 JE -	Licensing Started					
S-1-5-21-1234567890-100000000-100000000-48933		0365 JE -	Licensing Started					
S-1-5-21-1234567890-100000000-100000000-49425		0365 JE -	Finalized					

Base system

Monitor the base or basic Journal Commander System will ensure fundamental operational activity. One of the ways to proactively monitor the system itself is to use the HealthCheck API.

Completing a migration

In this section we'll review some of the things that need to take place at the end of a migration using Archive Shuttle.

Progress & Performance

The summary information at the top of the Progress & Performance gives a good indicator of what has progressed in the migration.


User finished migration:	100%
Stage 1 (Synced):	99.93%
Stage 2 (Switched):	100%

Basic statistics	
Users	5 075
Mailboxes	112
Archives	3 810
Archived Items	203 161 244
Archived Items Size (Compressed)	23,72 TB
Mappings	3 774
Throttled Mappings Count	0
Enabled for Sync (Stage 1)	3 774
Enabled for Switch (Stage 2)	3 774
Finished Migration	3 774
Routed Items	203 024 847

Failures

Failures during the migration are inevitable.

Assuming that the source environment is fairly healthy the actual result in most migration is just a tiny fraction of a percentage, but a small fraction can still appear to be tens of thousands of items:



In this example 140,000 failed items is actually 0.069% within the migration.

Further, with this example, most items failed to export from the source meaning that source integrity was a possible issue.

PowerShell guide

The PowerShell module allows you to manage Archive Shuttle migrations using a command line interface approach. Using PowerShell means that much of the administration work can be configured to be executed remotely, and it can be part of other scripted solutions (such as AD user migration, or Exchange Mailbox migration).

PowerShell module requirements

Before installing the PowerShell Module the following requirements need to be met:

- .Net Framework 4.0
- Exchange Online v3 PowerShell module

Connecting to Archive Shuttle Core

To issue any commands via PowerShell, the first step is to connect to the Archive Shuttle Core. This can be done as follows:

```
$Core = Connect-ASCore -ComputerName "Name of Server" -Credential "some domain\user"
```

You will then be prompted for the password for the account.

The \$Core variable can then be used for additional commands, as shown in the example in the next section.

To connect to a Cloud Core, use:

```
$Core = Connect-ASCore -ComputerName cloudtenantname.archiveshuttle.com -  
AuthenticationType Basic -UseHTTPS -Credential "cloud account"
```

i | NOTE: PowerShell should be launched as an administrator.

Performing a Command

Once a connection has been made to the Archive Shuttle Core, commands can be executed. For example, a list of current mappings can be shown as follows:

```
Get-ASMappings -Core $Core
```

PowerShell paging functionality

The standard PowerShell paging functionality commands can be used with the Archive Shuttle cmdlets. For example:

```
Get-ASContainers -Core $core -first 10 -skip 0
```

List of PowerShell commands and descriptions

A list of available commands can be obtained by running the following PowerShell command:

```
Get-Command -Module ArchiveShuttle.PowerShell
```

The commands currently available are covered in the next section.

Name	Description
Add-ASArchiveNamePolicy	Creates a new archive name policy
Add-ASContainer	Adds a container of a given type.
Add-ASContainerMapping	Creates a container mapping for a given source container ID. It will apply a given Workflow and Filter based on the ID's passed to the command.
Add-ASContainerToTag	Adds a container to a particular tag.
Add-ContainerToUser	Assigns user to chosen container.
Add-ASEnvironment	Creates a new environment.
Add-ASEVRetentionCategoryMapping	Creates a new retention category mapping for Enterprise Vault to Enterprise Vault migrations.
Add-ASFileNamePolicy	Adds a new file name policy.
Add-ASFilterCondition	Adds a condition to an existing filter.
Add-ASFilterPolicy	This can be used as input for the Add-ASContainerMapping command
Add-ASGroup	Creates a new group.
Add-ASItemDatabase	Creates an Item Database for a given LinkID (or Link Name).

Name	Description
Add-ASLink	Creates a new link.
Add-ASMappingConfigurationTemplate	Adds a new mapping configuration template.
Add-ASOffice365LeaversConfiguration	This command creates leavers settings.
Add-ASPstSourcePath	Adds PST sources for a PST link.
Add-ASSettingSchedule	Add a Schedule to be used for Settings.
Add-ASTag	Creates a new tag.
Add-ASUserToGroup	Adds a user (or users) to a group.
Add-ASWorkflowPolicy	This command creates new workflow policy.
Add-ASWorkflowPolicyStep	This command adds a step to a given workflow policy.
Connect-ASCore	Connects to an Archive Shuttle Core using the name given. If no credentials are supplied, then you will be prompted to input them.
Disable-ASActiveDirectorySync	Disables Active Directory synchronization for a particular domain.
Disable-ASArchiveGathering	Disables archive gathering.
Disable-ASCollection	Disables Item Gathering for the specified Container Mapping ID.
Disable-ASMigration	Disables Migration for the specified Container Mapping ID.
Disable-ASModule	Disables a module in Archive Shuttle.
Enable-ASActiveDirectorySync	Enables Active Directory synhronization for a particular domain.
Enable-ASArchiveGathering	Enables archive gathering.
Enable-ASCollection	Enables Item Gathering for the specified Container Mapping ID.

Name	Description
Enable-ASModule	Enables a module in Archive Shuttle.
Enable-ASMigration	Enables Migration for the specified Container Mapping ID.
Enable-ASStage2	Enables Stage 2 for the specified Container Mapping ID.
Get-ASArchiveNamePolicy	Gets information about a particular policy.
Get-ASADDomain	Returns a list of Active Directory domains.
Get-ASADGroup	Returns Active Directory groups.
Get-ASADGroupMember	Returns the members of a specific AD Group.
Get-ASArchive	Returns information about archives collected by Archive Shuttle.
Get-ASCommandQueueStatus	Retrieves an entry from the CommandQueue table.
Get-ASConfiguration	Retrieves information about a particular entry in the SettingDefinition table.
Get-ASContainers	Retrieves a list of containers from Archive Shuttle.
Get-ASCoreVersion	Displays version information for Archive Shuttle.
Get-ASEnvironment	Display information about a specific environment which has been added to Archive Shuttle.
Get-ASEVRetentionCategoryMapping	Gets information about a retention category mapping.
Get-ASExchangeDatabase	Retrieves a list of Exchange Databases.
Get-ASExchangeServer	Retrieves a list of Exchange Servers in a particular Active Directory Domain.

Name	Description
Get-ASFailedItemThreshold	Fetches the failed item threshold.
Get-ASFileNamePolicy	Gets information about a file name policy.
Get-ASFilterPolicy	Get a list of Filter Policies, or a specific policy.
Get-ASFolderNamePolicy	Retrieves information about a folder name policy.
Get-ASGroup	Retrieves a list of groups. It is also possible to retrieve a particular group by name (in order to get the ID) or by ID (in order to get the name).
Get-ASHealthStatus	Returns health status. Parameters can be set to return status on specific components.
Get-ASLicense	Returns information about licenses.
Get-ASLink	Returns information about a specific link.
Get-ASLinkHealth	Returns information about links.
Get-ASLinkModule	Returns information about modules assigned to a link.
Get-ASMailbox	Returns information about mailboxes.
Get-ASMailboxQuotaStatus	Retrieves mailbox quota status.
Get-ASMappingConfigurationTemplate	Retrieves information about a mapping configuration template.
Get-ASMappings	Retrieves a list of container mappings that have already been made in Archive Shuttle.
Get-ASModule	Returns a list of modules.
Get-ASModuleCredential	Returns module credentials.
Get-ASModulePerformance	Gets some statistics relating to module performance.

Name	Description
Get-ASOffice365LeaversConfiguration	Gets information about a leavers configuration.
Get-ASOffice365License	Gets information about licenses collected from Office 365.
Get-ASPremigrationStatus	Check pre-migration status.
Get-ASProgressStatistics	Gets information about the overall progress and statistics relating to a migration, as shown on the Progress & Performance screen in the user interface.
Get-ASRetentionCategory	Gets a retention category.
Get-ASScheduledTask	Gets information about all scheduled tasks, or a specific task based on name or ID.
Get-ASSettingSchedule	Get a Schedule to be used for Settings.
Get-ASStage1Statistics	Retrieves Stage 1 statistics for a specified container mapping ID.
Get-ASStage2Status	Retrieves Stage 2 status information for a specified container mapping ID.
Get-ASTag	Returns information relating to a Tag.
Get-ASUser	Retrieves details about a specific user from within Archive Shuttle.
Get-ASUserSyncStatus	Returns information about when a user object was last synchronized.
Get-ASVaultStore	Retrieves a list of Vault Stores.
Get-ASWatermark	Retrieves the high and low watermark for a link or for all links (if a LinkID is not specified).
Get-ASWorkflowPolicy	Get a list of Workflow Policies, or a specific policy.

Name	Description
Get-ASWorkflowPolicyStep	Gets information about a step in a workflow policy.
Remove-ASContainerMapping	Removes an existing mapping from Archive Shuttle.
Remove-ASLinkDatabase	Removes a specific link database.
Remove-ASMappingConfigurationTemplate	Removes a mapping configuration template.
Remove-ASPstSourcePath	Removes a specific PST Source Path.
Restart-ASModule	Issues a command to restart a specified Archive Shuttle module.
Set-ASArchiveNamePolicy	Updates an archive name policy.
Set-ASCommandIntervalLastExecutedDate	Sets the last executed date on an entry in the CommandIntervalExecutionStatus table.
Set-ASConfiguration	Updates a configuration option in the SettingDefinition table.
Set-ASContainerMapping	This command sets various values of a container mapping, After the command is executed the mapping can be checked in the Archive Shuttle user interface. All parameters in the cmdlet are mandatory except FilterPolicyId.
Set-ASEnvironment	Defines an Environment.
Set-ASFailedItemThreshold	Sets the failed item threshold.
Set-ASFileNamePolicy	Updates a file name policy.
Set-ASFolderNamePolicy	Updates a folder name policy.
Set-ASGroup	Modifies the name of an existing group.
Set-ASLinkModule	Assigns a module to a link.
Set-ASMappingConfigurationTemplate	Updates a mapping configuration template.

Name	Description
Set-ASModule	Sets a module.
Set-ASModuleCredential	Set module credentials.
Set-ASModuleLogLevel	Sets the logging level of a particular module.
Set-ASOffice 365LeaversConfiguration	Updates parameters in an Office 365 leavers configuration.
Set-ASPremigrationAction	User can change, pause, resume, retry or skip pre-migration workflow policy.
Set-ASPstSourcePath	Adds a PST Source path for a PST link.
Set-ASPstTargetPath	Sets the PST Output Path for a PST link.
Set-ASPstTemporaryPath	Sets the temporary path to be used when creating PSTs.
Set-ASScheduledTask	Modify a scheduled task
Set-ASSettingSchedule	Modify a Schedule to be used for Settings.
Set-ASStage2Status	Changes values for a mapping in Stage 2.
Set-ASStagingAreaPath	Sets a staging are for a particular link.
Set-ASTag	Updates a specific tag name.
Set-ASWatermark	Sets the high and low watermark for a particular link.
Set-ASWorkflowPolicy	This command updates a workflow policy.
Set-ASWorkflowPolicyStep	Updates a given step in a workflow policy.
Set-ASStage2Status	Sets the status for container mapping in Stage 2.
Start-ASAdSyncSingleUser	Starts synchronization of a specific user from Active Directory.
Start-ASArchiveGathering	Starts archive gathering.

Name	Description
Start-ASCommand	Starts a specific command.
Start-ASModule	Starts a specified module.
Start-ASOffice 365SyncMailboxes	Starts the procedure to synchronize all mailbox metadata from Office 365 into Archive Shuttle.
Start-ASOffice 365SyncSingleUser	Starts the procedure to synchronize a single users' metadata from Office 365 into Archive Shuttle.
Start-ASRetryFailedItems	This command retries failed (retryable, permanent) or hanging items for a specified container mapping. This actions can be defined for export, import or both. In case there it is not specified what should be retried, then everything available is retried for the whole container mapping.
Start-ASStagingAreaCleanup	Starts the clean-up process for the staging area.
Start-ASSyncADDomains	Starts synchronization of metadata about AD Domains.
Start-ASSyncADUsers	Starts synchronization of metadata from a particular Active Directory domain.
Stop-ASModule	Stops a specified module.
Update-ASMailboxSizeAndQuota	Updates mailbox size and quota status.

Examples

In the following examples, the Connect-ASCore command has already been run, using this syntax:

```
$Core = Connect-ASCore -ComputerName "Name of Server" -Credential "some domain\user"
```

List containers which are less than 90% imported

```
Get-ASStage1Statistics -Core $Core | Where-Object  
{$_ .ImportStatistics.Percentage -le 90}
```


ContainerMappingId	SourceFullName	RoutedCount	ExportStatistics.ProcessedItems	ExportStatistics.ErrorsCount	ExportStatistics.PercentAge	ImportStatistics.ProcessedItems	ImportStatistics.ErrorsCount	ImportStatistics.PercentAge
6	rowi09	102015	29874	126	29 %	7904	0	7 %
7	rowi08	129100	29857	143	23 %	4419	0	3 %

List people in Stage 1 who have not started Stage 2

```
Get-ASStage1Statistics -Core $Core | Where-Object {$_.Stage2Enabled -eq $false}
```

ContainerMappingId	SourceFullName	RoutedCount	ExportStatistics.ProcessedItems	ExportStatistics.ErrorsCount	ExportStatistics.PercentAge	ImportStatistics.ProcessedItems	ImportStatistics.ErrorsCount	ImportStatistics.PercentAge
6	rowi09	22062	17	0	0 %	0	0	0 %
7	rowi08	22100	18	0	0 %	0	0	0 %

List all the mappings with the tag 'Batch1' which are not yet mapped

```
Get-ASContainers -Core $Core -TagName 'Batch1' | Where-Object {$_.IsMapped -eq $false}
```

ContainerId	Name	LinkId	ItemCount	ItemSize	EnvironmentName
f2977992-747e-46...	rowi-2k7-11	426b4f3c-ae15-4f...	50000	783291	ev
12784114-05f2-46...	rowi-2k7-10	426b4f3c-ae15-4f...	50000	784933	ev
d9340387-b6be-46...	rowi11	426b4f3c-ae15-4f...	308	9594	ev
c75be693-b377-49...	yrowi-2k10-03	426b4f3c-ae15-4f...	9999	255191	ev
45c9553b-c4d3-4a...	rowi-2k7-05	426b4f3c-ae15-4f...	50000	786736	ev

Output Results to file

It might be required to output the results of PowerShell commands to a file, an example of how to output to a file is shown below:

```
Get-ASContainers -core $core | Out-File C:\Temp\ContainersList.txt
```

Get Command-Level Help

Additional help on the syntax of the commands can be obtained from the PowerShell help, for example:

```
Help Get-ASContainers
```

Exchange Online module commands

Below are a list of commands that you can use with the Exchange Online module in PowerShell v3. This should be used once PowerShell v3 has been installed. For more on the Exchange Online module, including requirements and setup, go to the *Quadrotech Archive Shuttle Planning Guide*.

Visit this article from Microsoft for more about the module.

PowerShell v3 command	Exchange Online command
Get-Mailbox	Get-EXOMailbox
Get-MailboxFolderStatistics	Get-EXOMailboxFolderStatistics

PowerShell v3 command	Exchange Online command
Get-MailboxStatistics	Get-EXOMailboxStatistics

System health events with PowerShell

PowerShell commands enable you to return health statuses for your tenant. You can get a health status report for your tenant as a whole, or within certain parameters, such as the UI status, license status etc.

Below, you can find the following commands and parameters, along with screen examples, to retrieve health status updates. If a search returns 'IsHealthy:true', there have been no errors found.

i | NOTE: You must be connected to Archive Shuttle Core.

Get-ASHealthStatus

Retrieve health status across the tenant.

Example: Get-ASHealthStatus / ConvertTo-Json -Depth 4

```
PS C:\Windows\system32> get-ashealthstatus | ConvertTo-Json -Depth 4
{
  "IsHealthy": true,
  "Errors": null,
  "UIResponse": {
    "ProductName": "Quadrotech Archive Shuttle",
    "ProductVersion": "10.5.41217.4",
    "IsActive": true,
    "Exception": null
  },
  "DBResponse": {
    "DBInfo": [
      {
        "DatabaseName": "QA-PEHD-130_ArchiveShuttleDirectory",
        "SchemaVersion": "10.5.0.0",
        "IsActive": true,
        "Exception": null,
        "DatabaseCrudResult": [
          "ArchiveShuttle.Common.Resources.PowerShell.Commands.DatabaseCru
          ArchiveShuttle.Common.Resources.PowerShell.Commands.DatabaseCru
          ArchiveShuttle.Common.Resources.PowerShell.Commands.DatabaseCru
          ArchiveShuttle.Common.Resources.PowerShell.Commands.DatabaseCru
        ]
      }
    ],
    "Exception": null
  },
  "WSResponse": {
    "Version": "10.5.41217.4",
    "IsActive": true,
    "Exception": null,
    "Urls": [
      {
        "Url": "asws/ArchiveShuttle.svc",
        "IsActive": true,
        "Exception": null
      },
      {
        "Url": "asws/ArchiveShuttleLogging.svc",
        "IsActive": true,
        "Exception": null
      },
      {
        "Url": "asws/ArchiveShuttleTransfer.svc",
        "IsActive": true,
        "Exception": null
      },
      {
        "Url": "asws/ArchiveShuttleUpdateService.svc",
        "IsActive": true,
        "Exception": null
      }
    ]
  }
}
```

UIStatus

Retrieve health status on the user interface.

Example: Get-ASHealthStatus -UIStatus | ConvertTo-Json -Depth 1

```
PS C:\Windows\system32> get-ashealthstatus -UIStatus | ConvertTo-Json -Depth 1
{
  "IsHealthy": true,
  "Errors": null,
  "UIResponse": {
    "ProductName": "Quadrotech Archive Shuttle",
    "ProductVersion": "10.5.41217.4",
    "IsActive": true,
    "Exception": null
  },
  "DBResponse": null,
  "WSResponse": null,
  "LicenseResponse": null,
  "Status": 0,
  "Error": null
}
```

WSStatus

Example: Get-ASHealthStatus -WSStatus | ConvertTo-Json -Depth 3

```
PS C:\Windows\system32> get-ashealthstatus -WSStatus | ConvertTo-Json -Depth 3
{
  "IsHealthy": true,
  "Errors": null,
  "UIResponse": null,
  "DBResponse": null,
  "WSResponse": {
    "Version": "10.5.41217.4",
    "IsActive": true,
    "Exception": null,
    "Urls": [
      {
        "Url": "asws/ArchiveShuttle.svc",
        "IsActive": true,
        "Exception": null
      },
      {
        "Url": "asws/ArchiveShuttleLogging.svc",
        "IsActive": true,
        "Exception": null
      },
      {
        "Url": "asws/ArchiveShuttleTransfer.svc",
        "IsActive": true,
        "Exception": null
      },
      {
        "Url": "asws/ArchiveShuttleUpdateService.svc",
        "IsActive": true,
        "Exception": null
      },
      {
        "Url": "asws/ArchiveShuttleWebUI.svc",
        "IsActive": true,
        "Exception": null
      }
    ]
  },
  "LicenseResponse": null,
  "Status": 0,
  "Error": null
}
```

DBStatus

Retrieve health status on the database.

Example: Get-ASHealthStatus -DBStatus | ConvertTo-Json -Depth 6

```
PS C:\Windows\system32> get-ashealthstatus -DBStatus | ConvertTo-Json -Depth 6
{
  "IsHealthy": true,
  "Errors": null,
  "UIResponse": null,
  "DBResponse": {
    "DBInfo": [
      {
        "DatabaseName": "QA-PEHO-130_ArchiveShuttleDirectory",
        "SchemaVersion": "10.5.0.0",
        "IsActive": true,
        "Exception": null,
        "DatabaseCrudResult": [
          {
            "CrudOperation": "Create",
            "IsCrudOperationAvailable": true
          },
          {
            "CrudOperation": "Insert",
            "IsCrudOperationAvailable": true
          },
          {
            "CrudOperation": "Update",
            "IsCrudOperationAvailable": true
          },
          {
            "CrudOperation": "Delete",
            "IsCrudOperationAvailable": true
          }
        ]
      }
    ],
    "Exception": null
  },
  "WSResponse": null,
  "LicenseResponse": null,
  "Status": 0,
  "Error": null
}
```

DBStatus with ItemDatabases status

Example: Get-ASHealthStatus -DBStatus -IncludeItemDbs | ConvertTo-Json -Depth 6

```
PS C:\Windows\system32> Get-ASHealthStatus -DBStatus -IncludeItemDbs | ConvertTo-Json -Depth 6
{
  "IsHealthy": true,
  "Errors": null,
  "UIResponse": null,
  "DBResponse": {
    "DBInfo": [
      {
        "DatabaseName": "QA-PEHO-130_ArchiveShuttleDirectory",
        "SchemaVersion": "10.5.0.0",
        "IsActive": true,
        "Exception": null,
        "DatabaseCrudResult": [
          {
            "CrudOperation": "Create",
            "IsCrudOperationAvailable": true
          },
          {
            "CrudOperation": "Insert",
            "IsCrudOperationAvailable": true
          },
          {
            "CrudOperation": "Update",
            "IsCrudOperationAvailable": true
          },
          {
            "CrudOperation": "Delete",
            "IsCrudOperationAvailable": true
          }
        ]
      },
      {
        "DatabaseName": "QA-PEHO-130_ASIDB_f3d76e06-00f6-eb11-83c4-005056982247",
        "SchemaVersion": "10.5.0.0",
        "IsActive": true,
        "Exception": null,
        "DatabaseCrudResult": [
          {
            "CrudOperation": "Create",
            "IsCrudOperationAvailable": true
          }
        ]
      }
    ]
  }
}
```

LicenseStatus

Retrieve health status on the database along with status update on licenses.

Example: Get-ASHealthStatus -LicenseStatus | ConvertTo-Json -Depth 6

```
PS C:\Windows\system32> get-ashealthstatus -LicenseStatus | ConvertTo-Json -Depth 6
{
  "IsHealthy": true,
  "Errors": null,
  "UIResponse": null,
  "DBResponse": null,
  "WSResponse": null,
  "LicenseResponse": {
    "IsValid": true,
    "Failures": [
    ],
    "Warnings": [
    ],
    "Exception": null
  },
  "Status": 0,
  "Error": null
}
```

Users will not be able to see invalid licenses in the parameter as the user cannot connect to the Core. This will show in the error message below.

```
ComputerName: qa-peno-130
connect-ascore : Core no available because license is invalid or exceeded.
At line:1 char:1
+ connect-ascore
+ ~~~~~
+ CategoryInfo          : LimitsExceeded: (:) [Connect-ASCore], Exception
+ FullyQualifiedErrorId : ArchiveShuttle.PowerShellModule.PowerShellCommands.Core.ConnectC
ore
```

Running PowerShell with Archive Shuttle

It is possible to add customized steps into a Stage 2 workflow and to execute PowerShell scripts. These can perform customer or environment specific actions. There are essentially two steps which need to be performed:

Addition of scripts

Customization of the Workflow

Addition of scripts

This section of the user interface allows an administrator to add PowerShell scripts to perform various tasks.

Scripts can be added, edited and deleted. When adding a script it should be given a friendly name, and description. The script itself can be added directly, or it can be uploaded from a file that you already have.

Customization of workflow

PowerShell scripts can be added to an existing workflow. It is also possible to add multiple scripts, if required. To add a script edit the desired workflow, and from the right hand side of the screen click on the **PowerShellScriptExecutionRunScript** command. This will add that command to the bottom of the workflow.

The command can be moved to the appropriate position in the workflow.

The command in the workflow needs to be edited so that you can provide the name of the script that you want to execute.

Variables that can be used in PowerShell execution

The table below identifies the PowerShell variables/token that can be used to run a PowerShell script in Archive Shuttle.

Name	Description
{UserSID}	User Security identifier
{Email}	Email address of the user
{LegacyExchangeDN}	Legacy Exchange DN of the mailbox
{UPN}	User principal name
{sAMAccountName}	sAMAccountName of the user
{PrimarySMTPAddress}	Primary SMTP Address of the user
{Identity}	Identity (eg alias/name) of the user
{UserDisplayName}	Display name / user name
{UserFirstName}	User's first name
{UserLastName}	User's last name
{ContainerMappingId}	Container Mapping Id in Archive Shuttle
{SourceContainerId}	Source Container Id for created/or future mapping
{TargetContainerId}	Target Container Id for created/or future mapping
{EVArchiveId}	User's Archive Id in Enterprise Vault

Troubleshooting

Archive Shuttle Troubleshooting can be broken down in to a number of different areas, as outlined below.

EV export issues

EV exports not working

If Archive Shuttle is unable to export data from Enterprise Vault, the following steps and information will help Support investigate the issue:

General information

- EV version
- Location of the EV Export Module (is it on an EV Server)
- Archive Shuttle Core version
- Archive Shuttle modules version

Steps to follow

Is it all items for all mappings? Yes/No. If Yes, go to step 1. If no, go to step 2.

Step 1

1. Check the staging area is configured for the links which are required.
 - a. Some links may not have the default staging area.
 - b. Check the staging area is accessible from the location where the export modules are installed
 - c. Collect a screenshot
2. Check modules are associated with the links
 - a. Some links may not have had any modules configured for them
 - b. Collect a screenshot
3. Check the modules are enabled
 - a. This is on the 'Modules' page

- b. Collect a screenshot
- 4. Check the modules are scheduled appropriately
 - a. This is on the 'Modules' page
 - b. Collect a screenshot

Step 2

Is it some items for all mappings? Yes/No. If Yes continue with this step. If No, go to Step 3.

1. Check the Failed Items screen for the mappings
 - a. Remember to select the link when entering the Failed Items screen for the first time
 - b. Collect a screenshot
2. Check the error that has been presented on the Failed Items Screen
 - a. Often the error will indicate what the problem is, or where to check
 - b. Collect a screenshot
3. Check that the items can be downloaded
 - a. Click on the 'Download' link. This will call download.asp for the item, from Enterprise Vault
 - b. Does the item download successfully? Yes/No
 - c. Does the item open correctly in Outlook? Yes/No
4. Try reprocessing the items from the failed items screen which will submit them to the export module again.
 - a. It may take a few minutes for the items to actually get reprocessed
 - b. Check the Failed items screen afterwards
 - c. Indicate whether the errors are the same and collect a screenshot. Same error? Yes/No
5. Try reprocessing the items by clicking on 'Failed Exports' on the Stage 1 screen
 - a. It may take a few minutes for the items to actually get reprocessed
 - b. Check the Failed Items screen afterwards
 - c. Indicate whether the errors are the same and collect a screenshot. Same error? Yes/No
6. Check the license has not expired
 - a. Take a screenshot of the license page. Note: The license is month/day/year format.

Step 3

Did export start, and then just stop? Yes/No. If Yes continue with this step.

1. Check that the module doesn't have a scheduled defined

- a. This is on the 'Modules' page
 - b. Collect a screenshot
2. Check the Health Dashboard and ensure that the link is not used up (ie watermarks and 'exported but not imported')
 - a. Collect a screenshot
3. Check the disk space thresholds on the Health Dashboard
 - a. Collect a screenshot
4. Select 'Hanging Exports' on the Stage 1 screen
 - a. Describe what happens
 - b. Wait 1-2 minutes, and refresh the Stage 1 screen to see if the exports start again
5. Check the ItemRoutingErrorCount in the ItemRouting table in SQL
 - a. If the error count is 10 or higher, the items will not be automatically reprocessed
 - b. They can be manually reprocessed by changing the Error Count in SQL, or from the Failed Items screen
6. Check the license has not expired
 - a. Take a screenshot of the license page. Note: The license is month/day/year format.

Step 4

1. Gather logs
 - a. Ensure that the module is configured for TRACE level logging
 - b. This is done on the Modules page
 - c. Reprocess the item(s), or click on 'Failed Exports', or click on 'Hanging Exports' on Stage1 screen
 - d. Wait one to five minutes
 - e. Generate a 'Support Bundle' for review by Support

EV import issues

EV Imports not working

If Archive Shuttle is unable to import data into Enterprise Vault, the following steps and information will help Support investigate the issue.

General information

- EV version
- Location of the EV Import Module (is it on an EV Server)
- Archive Shuttle Core version
- Archive Shuttle modules version

Steps to follow

Is it all items for all mappings? Yes/No. If yes, go to step 1, if no, go to step 2.

Step 1

1. Check that the retention categories are mapped.
 - a. Unmapped categories can be seen on the System Health page
 - b. This is required for intra or inter site migrations
 - c. Take a screenshot of the Retention Category mappings
2. Check that the EV Import modules are scheduled
 - a. This is on the Modules page
 - b. Take a screenshot of the schedule
3. Check that the EV Import module is enabled, and the Core is able to communicate with it
 - a. The module should not show as 'red' on the Modules page
 - b. Take a screenshot of the Modules page
4. Check that the link containing the target vault store has an Import Module which is mapped
 - a. This is on the Links Page
 - b. Take a screenshot, and highlight the link

Step 2

Is it some items for all mappings? Yes/No. If yes, continue with this step, if no, go to step 3.

1. Check the disk space is not exhausted
 - a. Collect a screenshot
2. Check that the exported-but-not-imported figure has not exceeded 100%
 - a. Take a screenshot of the System Health page
3. Review any errors on the Failed Items screen
 - a. Take a screenshot of the errors

Step 3

Did the import start, and then just stop? Yes/No. If yes, continue with this step, if no go to step 4.

1. Check that the module doesn't have a scheduled defined
 - a. This is on the modules page
 - b. Take a screenshot of the schedule
2. Check the System Health page to ensure the target vault store/link is not in backup mode
 - a. Take a screenshot of the System Health Page
3. Check the 'Index Incomplete' values on the System Health page
 - a. Take a screenshot of the System Health Page
4. Check the disk space is not exhausted
 - a. Collect a screenshot
5. Check that the exported-but-not-imported figure has not exceeded 100%
 - a. Take a screenshot of the System Health page
6. Review any errors on the Failed Items screen
 - a. Take a screenshot of the errors

Step 4

1. Gather logs:
 - a. Ensure that the module is configured for TRACE level logging
 - b. This is done on the modules page
 - c. Reprocess the item(s) or click on 'Failed Imports', or click on 'Hanging Imports'
 - d. Wait one to five minutes
 - e. Generate a 'Support Bundle' for review by Support.

EV collector issues

EV Collector not working

If Archive Shuttle is unable to collect data from Enterprise Vault, the following steps and information will help Support investigate the issue:

General Information

- EV version

- Location of the EV Collector Module (is it on an EV Server?)
- Archive Shuttle Core version
- Archive Shuttle modules version
- Mailbox Archive Item collection or Journal Archive Item collection
 - Verifying the type of EV Collector Module

Steps to follow

Is it all items for all mappings? Yes/No. If Yes, go to step 1, if no, go to step 2.

Step 1

1. Check modules are associated with the links
 - a. Some links may not have had any modules configured for them
 - b. Collect a screenshot
2. Check the modules are enabled
 - a. This is on the 'Modules' page
 - b. Collect a screenshot
3. Check the modules are scheduled appropriately
 - a. This is on the 'Modules' page
 - b. Collect a screenshot
4. Check the Link database is created for the Source Link
 - a. Source Link should have a Link database created
 - b. Collect a screenshot
5. Check the EV vaults stores are collected
 - a. All vault stores should be showed on EV Environments screen
 - b. Collect a screenshot

Step 2

Is it some items for all mappings? Yes/No. If Yes continue with this step, if No go to Step 3.

1. Check if all Mappings are showing on Stage 1 and Existing Mapping screens
 - a. Include the source and target link name column
 - b. Collect a screenshot
2. Try re-collecting the items

- a. Trigger the item collection from Existing mapping screen
 - b. Check the log files for the collection
 - c. Are there any errors? Yes/No
3. Check the license has not expired
 - a. Take a screenshot of the license page. Note: The license is month/day/year format.

Step 3

Did Collection start, and then just stop? Yes/No. If Yes, continue with this step.

1. Check that the module doesn't have a scheduled defined
 - a. This is on the 'Modules' page
 - b. Collect a screenshot

Step 4

1. Gather logs
 - a. Ensure that the module is configured for TRACE level logging
 - b. This is done on the Modules page
 - c. Reprocess the item(s), or click on 'Failed Exports', or click on 'Hanging Exports'
 - d. Wait one to five minutes
 - e. Generate a 'Support Bundle' for review by Support (Ensure each Collector Module is selected in the Bundle)

Disk space reducing unexpectedly

Administrators may notice that free disk space reduces unexpectedly on closed Enterprise Vault partitions containing Collection (.CAB) files.

The Enterprise Vault Export module may log the following error:

```
2017-02-10 08:43:18Z|5240| 53|ERROR| Process|Error occurred. ArchiveID:
[xxxxxxxxxxxxxxxxxxxxxxxx_vault] | ItemID: [694xxxxx-xxxx-xxxx-xxxx-
xxxxxxxx28] Void Get(Int32)
An internal failure occurred. Internal Error: 'Failed to extract CAB file.
[0xc0041b0f]'.
at KVS.EnterpriseVault.Interop.IItem.Get(Int32 DetailLevel)
at
ArchiveShuttle.Module.EVExportModule.ItemContentRetrievers.ContentRetriever
Base.FetchItem(IItem item, Int32 detailLevel)
at
ArchiveShuttle.Module.EVExportModule.ItemContentRetrievers.StandardContentR
etriever.FetchItem(IItem item)
```

```

at
ArchiveShuttle.Module.EVExportModule.ItemContentRetrievers.ContentRetriever
Base.ProcessAsFile(StagingArea stagingArea, EVExportResult item, String
fileName)
at
ArchiveShuttle.Module.EVExportModule.ItemContentRetrievers.ContentRetriever
Base.Process(EVArchive archive, Int32 commandId, Int32 mappingId,
EVExportResult item, StagingArea stagingArea)

```

On NTFS partitions where Collections have been enabled, archived data are collected into CAB files after a period defined by the administrator. Once this data is collected, subsequent requests to retrieve or restore it, result in this data being extracted from its corresponding CAB file to the closed partition. This data is restored in the form of temporary files with the extensions .archdvs, .archdvssp and .archdvscs.

These temporary files should then be deleted from the partition by Enterprise Vault the next time the Collection process runs if these files have not been accessed in the last 24 hours. The collection process schedule is defined on the Collections tab in the properties of each partition.

When performing a migration with Archive Shuttle it will, of course, need to retrieve the items inside the CAB files, and it does this by using the Enterprise Vault API. The Enterprise Vault API extracts the items from the CAB files, as described in the previous section. These temporary files may then remain on the partition disk.

If it is determined that there is a significant number of .archdvs, .archdvssp and .archdvscs files on the closed partition's disk location (one way to determine this is by running the command **dir *.archdvs, *.archdvscs, *.archdvssp /S > extractedfiles.txt** from the root of the partition's disk location), then follow these steps:

1. Verify that the Vault Store Database that hosts this partition is not in Backup Mode:
 - a. On the Enterprise Vault console, click the Vault Store Group that contains the Vault Store
 - b. Check the Backup Mode column in the right-hand pane of the Vault Administration Console
 - c. If Backup Mode is set for the Vault Store containing the affected partition and there is no backup currently in progress, clear the Backup Mode by right-clicking on the Vault Store and selecting Clear Backup Mode
2. Once Backup Mode is not set, manually run Collections:
 - a. Open the properties of the affected partition and select the Collections tab
 - b. Verify that the daily file collection schedule is set correctly and then click the Run Now button. An event should then be logged in the Enterprise Vault event log indicating that a Collection run has started. When this run completes an event will also be logged confirming that the Collection run has completed
 - c. See the Veritas KB article 'Free disk space reduces unexpectedly on closed NTFS partitions'

It is also possible to delete the *.archdvs files manually. There are few commands that can be used to clean the *.archdvs files:

1. CMD – needs to be run on the folder:

```
del *.archdvs* /s
```

2. PowerShell – path need to change. Below is just an example:

```
Remove-Item -path C:\FolderXY\Folder2XY\*.archdvs*
```

3. Quadrotech Archive Shuttle has also produced a tool to handle ARCHDVS files to assist in this situation:

See 'Handling ARCHDVS files whilst migrating' for more details.

Handling ARCHDVS files whilst migrating

If Enterprise Vault was configured to use Enterprise Vault Collections, when migrating a large number of ARCHDVS files, they might be placed on the file system as the Enterprise Vault API extracts the item from the collection file (CAB file). These files can consume considerable space, which may lead to low disk space situations.


Quadrotech Archive Shuttle have produced a tool to assist in this situation. While the issue is one caused by Enterprise Vault, this tool may help the situation.

The utility is able to automatically delete temp files using a few options:

- Using a lightweight **Filesystem watcher**, where the Application gets notified by the OS that a new file has been created
- Using a periodic **scanner**, which scans the filesystem each XX minutes
- Using a low disk space threshold
- On program startup

All these options also have a Timeout option, which is a minimum age that a file must have in order for it to be eligible for deletion.

You can register multiple paths, each with its own settings.

 **NOTE:** Use with care, and ideally in a test environment first.

We advise that you use the tool using Log only mode first (you will see this option in the configuration dialog), and then check the log file to see if it is performing as expected.

Usage:

1. Double-click the .exe, and you'll get a configuration dialog (Configuration mode)

Quadrotech Enterprise Vault Tempfile Cleaner

2. You can Add, Edit, or Delete directories.
 - a. Each can have its own settings
 - b. Click Save after editing or Add after adding a new directory
3. Start the .exe from the command line with a command line argument > Execution mode
 - a. For example, Quadrotech.EnterpriseVault.TempFileDeleter.exe EXECUTE

Contact Quadrotech Archive Shuttle support if you have any feedback regarding this tool.

This tool can be downloaded [here](#).

Office 365 import issues

Office 365 Imports not working

If Archive Shuttle is unable to import data into Office365, the following steps and information will help Support investigate the issue:

General Information

- Location of the Office 365Module (is it on the Archive Shuttle Core Server)
- Archive Shuttle Core version
- Archive Shuttle modules version

Steps to follow

Are all mappings affected? Yes/No. If yes, go to step 1, if no, go to step 2.

Step 1

1. Check that the Office 365modules are scheduled
 - a. This is on the Modules page
 - b. Take a screenshot of the schedule
2. Check that the Office 365module is enabled, and the Core is able to communicate with it
 - a. The module should not show as 'red' on the Modules page
 - b. Take a screenshot of the Modules page
3. Check that the Office 365link has the Office 365module mapped
 - a. This is on the Links Page
 - b. Take a screenshot, and highlight the link
4. Check that the link has the same staging area as where the export was performed to
 - a. This is on the Links page
 - b. Take a screenshot, and highlight the link
5. Check that the Office 365 Credential Editor has been provided Admin Credentials
 - a. You can see this in the Office 365 Credential Editor on the Server where the module is installed
 - b. Take a screenshot of the Credential Editor
6. Check that the Office 365Service Accounts have been granted the Application Impersonation role:
 - a. You can see this PowerShell. More information is available in this article.
 - b. Take a screenshot of the role assignment.
7. Review errors on the Failed Items screen
 - a. If the autodiscover service couldn't be located, you will have to change that setting in the module.exe.config.

Step 2

Is it some items for all mappings? Yes/No. If yes, continue with this step, if no, go to step 3.

1. Check the storage quota is not exhausted
 - a. Take a screenshot of the mailbox properties
2. Check that the exported-but-not-imported figure has not exceeded 100%
 - a. Take a screenshot of the System Health page
3. Review any errors on the Failed Items screen
 - a. Take a screenshot of the errors

Step 3

Did the import start, and then just stop? Yes/No. If yes, continue with this step, if no go to step 4.

1. Check that the module doesn't have a schedule defined
 - a. This is on the modules page
 - b. Take a screenshot of the schedule
2. Check the storage quota is not exhausted
 - a. Take a screenshot of the mailbox properties
3. Check that the exported-but-not-imported figure has not exceeded 100%
 - a. Take a screenshot of the System Health page
4. Review any errors on the Failed Items screen
 - a. Take a screenshot of the errors

Step 4

1. Gather logs:
 - a. Ensure that the module is configured for TRACE level logging
 - b. This is done on the modules page
 - c. Reprocess the item(s) or click on 'Failed Imports', or click on 'Hanging Imports'
 - d. Wait one to five minutes
 - e. Generate a 'Support Bundle' for review by Support

Exchange import issues

Exchange Imports not working

If Archive Shuttle is unable to import data into Exchange, the following steps and information will help Support investigate the issue:

General Information

- Exchange Version information
- Location of the Exchange Import Module (is it on the Archive Shuttle Core Server)
- Archive Shuttle Core version
- Archive Shuttle modules version
- Outlook version

Steps to follow

Is it all items for all mappings? Yes/No. If yes, go to step 1, if no, go to step 2

Step 1

1. Check that the Exchange Import modules are scheduled.
 - a. This is on the Modules page
 - b. Take a screenshot of the schedule
2. Check that the Exchange Import module is enabled, and the Core is able to communicate with it.
 - a. The module should not show as red on the Modules page
 - b. Take a screenshot of the Modules page
3. Check that the link containing the target Exchange Information store has an Exchange Import Module that is mapped.
 - a. This is on the Links page
 - b. Take a screenshot, and highlight the link
4. Check that the link has the same staging area as the export was performed to.
 - a. This is on the Links page
 - b. Take a screenshot, and highlight the link
5. Check the version of Outlook that is installed where the module is located.

- a. Ensure that the default profile is configured for an account that has full control on all mailboxes
- b. Ensure that the default profile is configured to launch straight into Outlook, and that it is an online profile

Step 2

Is it all items for all mappings? Yes/No. If yes, go to step 1, if no, go to step 2.

1. Check that the disk space is not exhausted.
 - a. Collect a screenshot
2. Check that the exported-but-not-imported figure has not exceeded 100%.
 - a. Take a screenshot of the System Health page
3. Review any errors on the Failed Items page.
 - a. Take a screenshot of the errors
4. Check that EWS Impersonation has been configured.
 - a. This is a setting in the module's .exe.config file.

Step 3

Did the import start, and then just stop? Yes/No. If yes, continue with this step, if no go to step 4

1. Check that the module doesn't have a schedule defined.
 - a. This is on the modules page
 - b. Take a screenshot of the schedule
2. Check the disk space is not exhausted.
 - a. Collect a screenshot
3. Check that the exported-but-not-imported figure has not exceeded 100%.
 - a. Take a screenshot of the System Health page
4. Review any errors on the Failed Items page.
 - a. Take a screenshot of the errors
5. Check that EWS Impersonation has been configured.
 - a. This is a setting in the modules .exe.config file.

Step 4

1. Gather logs:

- a. Ensure that the module is configured for TRACE level logging
- b. This is done on the modules page
- c. Reprocess the item(s) or click on 'Failed Imports', or click on 'Hanging Imports'
- d. Wait 1-5 minutes
- e. Generate a 'Support Bundle' for review by Support.

Native Format Import issues

Native Format Imports not working

If Archive Shuttle is unable to import data into Native Format, the following steps and information will help Support investigate the issue:

General Information

- Location of the Native Format Import Module (is it on the Archive Shuttle Core Server)
- Archive Shuttle Core version
- Archive Shuttle modules version
- Outlook version

Steps to follow

Is it all items for all mappings? Yes/No. If yes, go to step 1, if no, go to step 2.

Step 1

1. Check that the Native Format Import Module are scheduled
 - a. This is on the Modules page
 - b. Take a screenshot of the schedule
2. Check that the Native Format Import Module is enabled, and the Core is able to communicate with it
 - a. The module should not show as 'red' on the Modules page
 - b. Take a screenshot of the Modules page
3. Check that the Native Format Import link does have the Office 365 module mapped
 - a. This is on the Links Page
 - b. Take a screenshot, and highlight the link
4. Check that the link has the same staging area as the export was performed to
 - a. This is on the Links page

- b. Take a screenshot, and highlight the link
5. Check the version of Outlook which is installed where the module is located
 - a. Ensure that the default profile is configured for an account which has full control on all mailboxes
 - b. Ensure that the default profile is configured to launch straight into Outlook, and that it is an online profile
6. Review errors on the Failed Items screen
 - a. If the autodiscover service couldn't be located you will have to change that setting in the module.exe.config.

Step 2

Is it some items for all mappings? Yes/No. If yes, continue with this step, if no, go to step 3.

1. Check the storage quota is not exhausted
 - a. Collect a screenshot
2. Check that the exported-but-not-imported figure has not exceeded 100%
 - a. Take a screenshot of the System Health page
3. Review any errors on the Failed Items screen
 - a. Take a screenshot of the errors.

Step 3

Did the import start, and then just stop? Yes/No. If yes, continue with this step, if no go to step 4

1. Check that the module doesn't have a schedule defined
 - a. This is on the modules page
 - b. Take a screenshot of the schedule
2. Check that the exported-but-not-imported figure has not exceeded 100%
 - a. Take a screenshot of the System Health page
3. Review any errors on the Failed Items screen
 - a. Take a screenshot of the errors

Step 4

1. Gather logs:
 - a. Ensure that the module is configured for TRACE level logging
 - b. This is done on the modules page

- c. Reprocess the item(s) or click on 'Failed Imports', or click on 'Hanging Imports'
- d. Wait one to five minutes
- e. Generate a 'Support Bundle' for review by Support.

Stage 2 command issues

Archive Shuttle Stage 2 Command Troubleshooting can be broken down in to a number of different commands, as outlined below:

CollectItemsForArchive Issues

EnableMailbox Issues

DisableMailbox Issues

ZapArchivePermissions Issues

AssignArchive Issues

RenameSourceArchive Issues

DisableOrphanedArchive Issues

FixShortcuts Issues

DeleteEVShortcut Issues

RevertPendingShortcuts Issues

WaitForImportFinished Issue - WaitForImportFinished command may show failed

WaitForImportFinished Issue - Problem with Stage 2 in WaitForImportFinished staying in Processing state

Office365DeleteShortcuts Issue - Stage 2 Office365DeleteShortcuts command fails with "Too many concurrent connections opened"

Office365DeleteShortcuts Issue - Stage 2 Office365DeleteShortcuts command fails with "Autodiscover service could not be located.."