



transvault
Let's move together



Enterprise Archive Solution (EAS)

Frequently Asked Questions

Date: April 2021

Version: 1.2

Transvault is a trademark of Transvault Software Ltd

This document contains proprietary and confidential information and may not be shared with any 3rd party without the express written permission of Transvault.

Microsoft 365 and Office 365 are registered trademarks of Microsoft Corporation.

IMPORTANT: The information in this document is to act as a guide only. Transvault shall not be liable for technical or editorial errors or omissions contained herein. No warranties or any kind are made, including the accuracy of the information contained herein, or whether this information reflects the latest capability of the technology referenced.

www.transvault.com

Contents

1. Why use Transvault™ for our EAS migration?.....	3
2. How does Transvault connect to EAS?	3
3. How does Transvault compare with using native EAS extraction and import tools?	4
4. Can we use EAS Storage Manager (STORM) instead of PSTs to export data out of EAS?	5
5. Will our shortcuts (stubs) work when we migrate to the new system?	5
6. How quickly can Transvault migrate archived items?	6
7. Can we migrate archived Public Folders?	7
8. How can I find out if Transvault has moved all of our data?	7
9. How does Transvault aid compliance when migrating?	7
10. Why do Exchange Journal Archives require special care?	8
11. What do you mean by 'Chain of Custody' and why is it important?	9
12. What happens if an item fails to migrate?	9
13. We are involved in a de-merger and need to split up our archive. How does this work?	10
14. If we switch from using Notes to Exchange (or vice versa), can Transvault migrate our archive?	10
15. Do I need to provide systems to run Transvault on?	11
16. Will I need to buy services?	11

1. Why use Transvault for our EAS migration?

EAS was originally developed in Canada in the late 1990's by a company called Educom. After a few years this company was bought out by Zantaz which was later bought out by Autonomy, who is now owned by HP. EAS is now supported and developed by a company called Capax.

Transvault was the first vendor to offer EAS migrations, and throughout EAS' chequered history, Transvault has maintained a deep understanding of the EAS architecture and the different challenges introduced by different (and sometimes buggy) versions.

By working with Transvault software you can be assured of a seasoned solution and support capability that can help with issues such as:

- Accessing items that have been orphaned and are not retrievable via the EAS API
- Extracting archives from an EAS service that is non-operational
- Working with an archive service that is heavily loaded
- Support for both SQL and Oracle databases
- Migration of either Notes or Exchange EAS archives
- Support for EAS configurations that use email attachment single instancing

Transvault also offers extremely high-speed migrations, advanced shortcut handling and 100% chain of custody for compliance-led migrations.

EAS journal archives can also be moved back into your preferred destination – including back into native Exchange journals, with Transvault ensuring that all vital meta-data is protected and correctly aligned to your target system.

This includes the ability to support the new Office 365 compliance model.

2. How does Transvault connect to EAS?

Transvault connects to EAS through either an API or Direct Access connector.

The API works by querying the EAS IIS Server for the message and waiting for the G-Zip archive message (which is then migrated by Transvault to your preferred destination).

The direct connector is useful for accessing *orphaned* archived content. It can also be faster than using the API as you are not required to wait for the EAS IIS Server. Instead Transvault connects straight to the archive database, locate and transfer messages directly from the file stores.

3. How does Transvault compare with using native EAS extraction and import tools?

Many archive systems, EAS included, have some form of native 'export to PST files' function, as well as the ability to import emails from PST files. This approach, however, tends to be a slow, very manually intensive process that has no error-logging, no batching or audit capability. The discrete steps of first exporting data and then re-importing it to the destination archive risk data loss and require large amounts of interim storage. There's also no shortcut conversion/management capability.

Other problems surround the use of PST files, including:

- **Large PST files are prone to corruption.** This creates problems for organizations trying to extract the contents of very large mailbox archives or Journal mailbox archives,
- **PSTs are space-inefficient.** This means you'll need lots of extra storage to pre-stage your data.
- **PST files are slow.** Added to this is the fact that native extractions are typically single threaded processes which need to be manually overseen (i.e. you could not leave several mailboxes running unattended overnight).

Other issues associated with using 'native' EAS tools include:

- **There's no error management.** If a mailbox extraction fails owing to a corrupted email, the extraction will stop with no indication of which item was corrupted and no ability to pick up where the extraction left off.
- **You lose important compliance data.** If you plan to migrate the contents of Journal archives, it's likely you will have captured vital BCC'd and distribution list recipients. *PSTs cannot preserve such information when you move to a new archive.*
- **No records.** There is no audit table of what data has been moved
- **No shortcut support.** Options for giving users access to migrated data post-migration can be limited (e.g. it may not be possible to recreate shortcuts in the destination archive, at best giving users access via a separate search portal).

By comparison Transvault offers multi-threaded, high-speed extraction with full error-logging and chain-of-custody. It also offers the unique ability to convert shortcuts and synchronize folders ensuring that users retain transparent access to their emails post migration.

Performance comparison: A project to extract 10.5 million emails from a leading on-premise solution using the supplied tool took 5.5 man months to run (and re-run) single mailbox

extractions. Transvault Migrator is benchmarked to migrate the same amount of data in 24 hours. Note that some customer environments have towards a billion items archived.

Cost comparison: Manually exporting the contents of a single mailbox into a PST file can take 1.5-3 hours (depending on size), with additional time required for checking completion, updating manual logs, fixing any corruptions etc. A recent costing exercise based on hiring a temp (at \$18 per hour) to manually migrate PST files, and an average mailbox migration time of 3 hours indicated a cost of around \$54 per mailbox. *Bearing in mind this is just part of the process (the contents of PST files then need to be ingested into the target archive) manual migration costs can reach over \$100 per mailbox.*

4. Can we use EAS Storage Manager (STORM) instead of PSTs to export data out of EAS?

Whilst it is possible to extract larger amounts of data more easily using the EAS Storage Manager, this process loses all original folder information. This means that end users will no longer be able to navigate and locate emails using their familiar folder structures in your new archive platform.

5. Will our shortcuts (stubs) work when we migrate to the new system?

Yes. Transvault provides a comprehensive shortcut management service that ensures users have a seamless experience when they migrate.

As your data is moved, the corresponding EAS stub or 'shortcut' is converted to point to the new archive. This activity can take place whilst users are online – there's no need for any downtime or for users to logout.

In the situation where shortcuts aren't supported by the target archive (for example, if you're migrating from EAS back into Microsoft Exchange or Office 365), legacy shortcuts are replaced (rehydrated) with the original item.

Another key feature designed to ensure a great user experience is that Transvault synchronizes with the current status of users' shortcuts as they exist in their mailbox at the time of the migration.

For example, where users have re-folded their shortcuts 'post-archiving', Transvault ensures the corresponding items end up in the right folders post-migration.

Similarly, where users have deleted their shortcuts, you can configure Transvault not to migrate the corresponding item, as this is confusing and causes concern for end users when deleted items reappear.

This service is vital, as not all archive platforms keep track of shortcut re-folding and deletion activity.

NB - EAS itself does a good job of synchronizing the location of shortcuts, however there are some scenarios where deletions might not be picked up by EAS. *An example is where users drag their shortcuts into the Outlook deleted items folder, rather than highlighting them and using the delete option that links with EAS.*

6. How quickly can Transvault migrate archived items?

Transvault offers the fastest and safest migration times in the industry. This is confirmed by the many Partners we have that have direct, project-based experience of Transvault vs other migration solutions.

Transvault's multi-threaded, multi-server capability means that multiple extraction and ingestion pipes can be set up between your source and target system, thus driving your environment to capacity. You can also process a single mailbox using *multiple* migration threads – ideal for processing large archives and journals in the fastest possible times.

For example, Transvault's CloudStream service, combines the latest Microsoft ingestion technologies with sophisticated algorithms, such as traffic flow analysis, to guarantee unrivalled ingestion speeds into Office 365.

Your chosen archive migration partner can establish a proof of concept (POC) to establish likely throughput rates in your specific environment. *They will also be able to give you guideline speeds seen at other customer sites for a similar migration path.*

It is important to note achieving best performance is always dependent on environmental factors such as:

- Available network bandwidth
- Speed of the storage subsystem on which the legacy archive sits as well as the destination storage

- The ingestion performance of the target archive system (typical speeds may range from 5-10 GB/hour)
 - The scheduling of other project elements such the commissioning of the target environment.
- Also note that migration speeds we refer to are end-to-end migration times, and not just *extraction* speeds into an interim format. *Extraction speeds are misleading as this is only half of the 'story' (as the extracted data then needs to be ingested into the target archive).*

7. Can we migrate archived Public Folders?

Yes. Legacy Public Folders can be migrated out of EAS for Exchange into:

- Public Folder archives in Veritas Enterprise Vault for Exchange
- Microsoft Exchange Public Folders
- Office 365 Public Folders (Modern Public Folders)
- PST files (for ingestion into any target platform that supports this format)

You can also use Transvault to migrate Public Folders still hosted on Exchange into Office 365. Additional Public Folder platforms will be added in future releases of Transvault. Please contact us to check.

8. How can I find out if Transvault has moved all of our data?

Transvault includes comprehensive auditing that tracks the migration of each individual item. This is further bolstered by that fact that wherever possible, migrations are made in one transaction, direct from source to destination and with no interim steps that could lead to data being tampered with or going missing 'in transit'. In the event of a future eDiscovery, the audit enables you to prove that you have not lost items during the move and that chain-of-custody has been maintained.

If the migration of an item fails (e.g. owing to data corruption or intermittent network issues) tasks can be re-run to target *just* these items, adding them into the new archive alongside items that are already migrated.

Reporting is also available to highlight how much data has been moved for each 'mailbox' in the archive.

9. How does Transvault aid compliance when migrating?

Migration methods that rely on interim PST or EML files are subject to human error and have no tracking or auditing mechanism to prove that a migration was 100% successful.

As described earlier in this document, Transvault eliminates opportunity for human error: your data is automatically moved in 1 step, direct from the source to destination using banking-style transactions. Each item moved is checked for integrity to ensure your data will be viable in the new environment.

There's also complete auditing of the migration process, with detailed reports that show 1:1 mappings of the ID of the item in the source archive and the ID of the new item as it is moved to the destination archive, enabling you to demonstrate of a complete 'Chain of Custody' for your data while undergoing migration.

See also next the following questions relating to compliance concerns.

10. Why do Exchange Journal Archives require special care?

First, Journal archives tend to be extremely large, making a manual extraction approach slow and subject to size-related problems that can occur when using PST files as an interim store.

Transvault allows journal and large mailboxes to be split into a number of separately handled virtual mailboxes of a user-defined size. This allows multiple processing threads to be applied to the migration of a single journal mailbox, significantly speeding up the migration task.

Secondly, most of the organizations that capture emails into Exchange Journal mailboxes for compliance reasons use the Envelope Journaling feature. This feature was developed by Microsoft as a way to preserve vital header information including BCC'd recipients and the expanded members of any distribution lists.

From a compliance perspective, this data must be preserved and available for access when performing eDiscovery. Losing this vital recipient data would jeopardize the validity of any future eDiscovery case, as not all the people that received any given email would be included in an investigation.

The important thing to note is that different archive platforms store envelope information in different ways.

This creates a very special requirement at migration time that neither manual methodologies nor vendor extraction tools can support.

For example, with Transvault you can move EAS journal archives into EV journal archives. To do this Transvault reconstructs EAS journals into the format required by EV.

And, if you're moving data from EAS into another platform, such as Microsoft Office 365,

Transvault is able to migrate the legacy journal format into the new compliance model that has been introduced by Microsoft, ensuring all your data is intact, fully discoverable and stored in accordance with Microsoft's licencing policies.

11. What do you mean by 'Chain of Custody' and why is it important?

Chain of custody refers to the reliable recording of processes and procedures that occur while evidence (physical or electronic) is being captured, held, transferred or disposed of.

It is vital that the evidence in question has remained free from alteration and that secure handling has been provided at all stages. This is important since any later data investigation or discovery against migrated emails will seek to validate that the data has been handled correctly and that nothing was missed out.

If at all possible, transferring your archives in 1 step, end-to-end, is the best way to preserve chain-of custody. This approach avoids any chance of manual intervention or loss of tight control. Other migration techniques that involve extraction to interim storage areas and file formats (such as PST or EML files) risk your data being inadvertently lost, maliciously tampered with, or indeed corrupted while they are waiting to be ingested into the target archive. It is therefore difficult to vouch for the security of your data 'in transit'.

Transvault's wide range of platform-specific connectors avoids the need to use interim formats, plus its optimized performance makes it viable to move directly across the network without making copies of your archive.

Additionally Transvault's forensic 1:1 auditing provides evidence that your data has been successfully transferred to a new environment, and provides both a source ID and target ID.

12. What happens if an item fails to migrate?

As your data is moved, Transvault carries out a series of integrity checks to ensure your email records will be viable post-migration.

Any items that fail to migrate are automatically re-processed a specified number of times and/or at a different time of day.

Failures to migrate an item may be temporary, for example, owing to environmental issues such as poor network bandwidth or high loading on the legacy archive.

'Permanent' failures tend to be low – typically .001% of the overall email quantity. They are usually attributable to pre-existing problems in the source archive (i.e. not caused by the migration process). As such, it is likely that these items would NOT have been picked up by any audit or eDiscovery exercise. In the event of a permanent failure, a full log of the item(s) in question is produced to enable investigation. You can attempt a manual retrieval of any failed messages *directly* from the log, a feature which massively reduces troubleshooting overheads.

If your organization requires further investigative work to be carried out on failed items, data remediation services are available to help and where possible, fix the problem to the satisfaction of your legal team.

13. We are involved in a de-merger and need to split up our archive. How does this work?

Transvault can filter items by user, groups, folders and dates, thus enabling data to be selectively and incisively migrated to one or more different locations.

Transvault can also re-write internal email addresses in sender and recipient fields so that the email is usable with any new domain naming or recipient-addressing conventions.

Using TransVault CTM feature, journal archive content can be split based upon email addresses and domains so that a single journal can become two or more separated journals each containing emails only relevant to the specific firms.

14. If we switch from using Notes to Exchange (or vice versa), can Transvault migrate our archive?

Yes. Transvault offers migration and conversion capability between Exchange and Notes archives. E.g.

- If you're using EAS for Notes you can migrate their contents to EAS for Exchange, using Transvault to perform the migration along with on-the-fly email format conversion.
- An organization using EAS for Notes can use Transvault to migrate (and convert) its archived emails into say, an Enterprise Vault for Exchange archive.
- You can also migrate .NSF files into an Exchange-based platform.

In most cases, Transvault can convert shortcuts so they work with the new archive system.

Where necessary, Transvault can re-write internal email addresses so that legacy emails can be successfully replied to. *This is useful in the event of a merger or other scenario where the email address may have changed.*

Transvault also includes the ability to decrypt Notes emails.

Additionally organizations can reduce the time it takes to migrate their 'live' email system between Exchange and Notes by ramping up their archiving activity to minimize the size of the live mailboxes that need to be migrated.

Benchmark testing has further shown that Transvault has superior conversion performance and message format fidelity when compared with leading 3rd-party mailbox migration applications.

15. Do I need to provide systems to run Transvault on?

Working with Transvault solutions gives you the choice of making your move on-premises or in the Cloud. If you have existing server bandwidth and your migration is within the same location or domain, then a fully on-prem approach, with the option of local or remote control, may be the way to go. Alternatively you can make your move in Microsoft Azure.

Using Transvault in Azure means there's no hardware to set up in advance. You also have the added benefits of built-in resilience, load balancing and distribution for geographically distributed locations. Plus, if you're moving to Office 365, network latency is inherently reduced, resulting in faster data transfers.

Whichever route you take, your move can be managed remotely, and you can check on migration progress remotely and 'on the go' using Transvault's ToolBox App from a smart phone or a tablet.

16. Will I need to buy services?

Transvault Services have two tailored offerings: Managed Migration Services is a fully managed end-to-end migration service or Assisted Migration Services a service to allow knowledge transfer to internal teams for the delivery of your own migration with the support of Transvault throughout the migration.

Alternatively, you can sit back and relax with a complete managed migration service that includes expert advice on best practices, compliance know-how, project management, archive-specific expertise, expert troubleshooting and comprehensive management reporting.

You can also elect to run your migration on-premises or in the cloud (e.g. Microsoft Azure).

One of our many specialist, Transvault certified partners across the globe will be able to advise on and deliver the appropriate Transvault products and services for your organisation.

To contact a reseller or partner, please see <http://www.transvault.com/> for more information.