WHITE PAPER: TECHNICAL

Confidence in a connected world.

Enterprise Vault 8.0

Optimized Single Instance Storage (OSIS)

Alex Brown
Technical Field Enablement

White Paper: Symantec Technical

Content

If you have any comments on this Whitepaper please email EV-TFE-Feedback@Symantec.com

Introduction

This whitepaper introduces the new Enterprise Vault storage model launched in version 8.0. It will explain how the new model can help provide industry-leading archive storage footprint reduction through single instance storage and compression, as well as describe the various components and technologies employed within Enterprise Vault to do this.

This whitepaper will not provide Enterprise Vault solution sizing information or metrics; this information is available separately.

Note - This whitepaper assumes a reasonable level of knowledge of Enterprise Vault and its storage subsystem.

Terminology

Before we get started let's get some terms and definitions out of the way:

Term	Definition
OSIS	Optimized Single Instance Storage, the term used to described the new storage model used by Enterprise Vault 8.0 and later.
Archive	A logical container of archived items, for example a single user's archived email items.
Vault Store	A collection of Archives. A Vault Store consists of a SQL database for reference and one or more storage locations to store the archived items and SIS Parts.
Vault Store Group	A collection of one or more Vault Stores that may share SIS parts.
SIS Part	A shareable part of an archived item.
Fingerprint	A unique cryptographic hash based value for a shareable or unique part of an archive item.
Fingerprint Set	A collection of one or more Fingerprints for all of the parts of an archive item.
Fingerprint Database	SQL Database containing all Fingerprints for a specific Vault Store Group.
SIS Boundary	The boundary that defines the limit for sharing of SIS parts.

OSIS Explained

Optimized Single Instance Storage (OSIS) is the term given to the Enterprise Vault 8.0 sharing technology i.e. the sharing model implemented by the software in order to store archived items. The advent of OSIS has strengthened Enterprise Vault's already market-leading position to provide global, item-level, single instance storage. There are various aspects of the OSIS model that we will investigate within this whitepaper, however first let's get an idea of what OSIS can do.

When using OSIS, elements of archived items that could be considered as shareable can be identified by Enterprise Vault and be saved and stored separately as SIS parts in the Vault Store (for example, attachments on emails or entire documents from a file system). Each SIS part is potentially shareable with any subsequently archived identical SIS part. This begs the question; "How can we be sure that shared SIS parts are in fact identical?". We will answer this question in more depth a little later on, however the uniqueness of a specific SIS part is guaranteed by the process of Fingerprinting.

Benefits of OSIS

The benefits of using Enterprise Vault and its OSIS model are wide ranging:

- Enterprise Vault can share at the item level across different content sources and Vault Stores.
- Attachments to emails can be shared separately from the email itself.
- Archived content is only written to a Vault Store once, meaning that Enterprise Vault never needs to go back and update previously archived item for any reason. This makes sharing between standard (using NTFS) and secured file systems (such as WORM) possible
- Enterprise Vault can take advantage of storage level deduplication and compression from vendors such as DataDomain and NetApp to allow further reduction in the storage footprint of archived content.
- Existing SIS parts can be fully validated before sharing occurs to ensure the SIS parts really are the same.
- Existing archived items can be re-validated during retrieval to ensure that they are exactly the same as when they were stored originally.

The Sharing Process

How does Enterprise Vault know when to share?

Enterprise Vault uses Vault Store Groups to configure the sharing of items. A Vault Store Group consists of one or more Vault Stores. A Vault Store Group is also a single instance sharing boundary and as such any items stored within a Vault Store in the same Vault Store Group can potentially be eligible for sharing with other items stored within that Group. This sharing option does not apply between Vault Stores in different Vault Store Groups.

The Vault Stores contained within each Vault Store Group can each be set to one of three levels of sharing depending on storage requirements:

- Share within Group – Items are eligible for sharing with other items from any other Vault Store in the Group that is set to the same level of sharing.
- Share within Vault Store – Items are only eligible for sharing with other items archived into the same Vault Store.
- No sharing – Do not initiate any sharing at all.

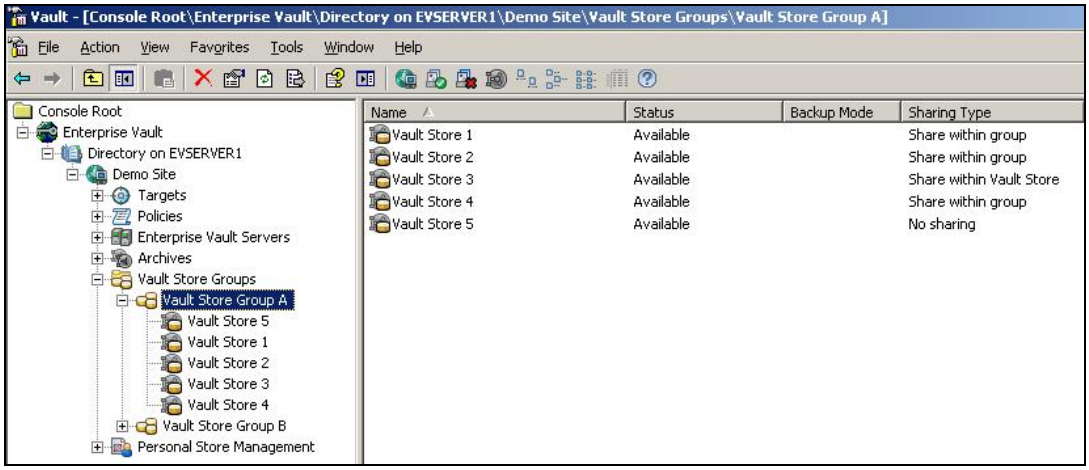

Figure 1.Example Vault Store Group in the Enterprise Vault Administration Console

Figure 1 shows an example Vault Store Group and note the “Sharing Type” column. In this case Vault Stores 1, 2 and 4 are set to “Share within group” and will attempt to share items between these three Vault Stores. Vault Store 3 is set to “Share within Vault Store” and will therefore only share items within its own partitions and finally Vault Store 5 is set to “No sharing” and will therefore not attempt to share any items at all.

The sharing level of Vault Stores within a Group can be changed at any point however note that the new sharing level will only apply to items archived after this change.

Using Vault Store Groups it is possible to design Enterprise Vault archiving solutions that provide global or data centre level sharing. Global SIS is not necessarily always the best option for all customers as external network bandwidths (such as those used by a WAN) can affect performance. Vault Store Groups can also be used to create Chinese walls between certain archived content to ensure the content is not shared.

What is Fingerprinting?

Fingerprinting is the process by which the Enterprise Vault OSIS Fingerprint engine generates a unique identifier “Fingerprint” for each part of an archived item. Even apparently similar items will generate a completely different identifier if they contain different content. Therefore, these Fingerprints are the mechanism by which Enterprise Vault can determine whether it has previously archived all or part of an item. There are two types of Fingerprints:

- Individual Fingerprint is generated for each part of a archived item (shareable and non-shareable). All the individual Fingerprints for each part of an archived item comprise a Fingerprint set.
- Cumulative Fingerprint is generated on the fly using a items Fingerprint set as a reference. Is used to uniquely identify, verify or re-Fingerprint an entire archived item during various storage operations to ensure authenticity.

A Fingerprint itself is a string of hexadecimal characters which are generated by the Fingerprint engine using a combination of item type (Exchange email, Domino email, File, SMTP email), a SHA256 hash of the original item and the original size. Below is an example Fingerprint:

1M46157AF50408DA53A15FC9516457E2AFE52584F2BE352C7491CD255BF493CBB3s2AC

Each Fingerprint generated will be stored in a Fingerprint database. Each Vault Store Group has a dedicated Fingerprint database for storing Fingerprints. For this reason the SIS boundary within Enterprise Vault is at the Vault Store Group level. The Fingerprint Database consists of a number of tables[1] which hold and catalog these Fingerprints for use when Enterprise Vault archives another item and needs to check whether the item (or part of it) has already been stored.

The Storage process

Now that we've discussed the mechanism by which Enterprise Vault's OSIS can achieve single instance storage let's look at a brief example of how an item is stored:

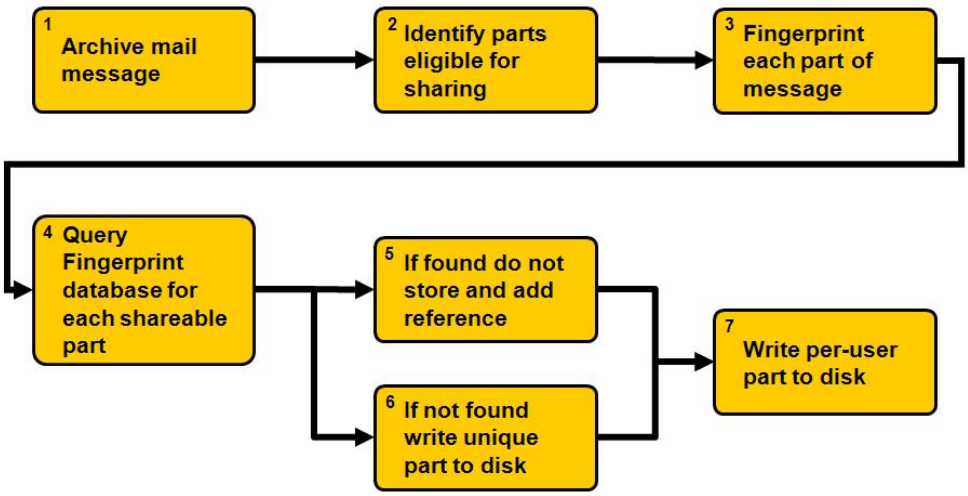

Figure 2. OSIS Archiving Workflow

We will assume we are archiving an email item (Domino or Exchange):

1) An email is archived from a users mailbox via the normal archiving mechanism[2].

[1] There are actually 256 member tables per Fingerprint database, and with each member table able to reasonably store up to around 4 billion Fingerprints, plus the ability for the entire Fingerprint database to roll over to a new instance up to 32,000 times, means that each Fingerprint database can store and reference approximately 32,000,000,000,000,000 (Quadrillion) Individual Fingerprints.

[2] We also assume that the item is being archived into a Vault Store that is enabled for sharing within its Vault Store Group.

2) Enterprise Vault identifies any part of that message that could be eligible for sharing e.g. attachments, and splits them from the primary item creating SIS Parts. Other sharable properties such as Message body & the Recipient List are also split into another SIS Part. All remaining properties are classified as per-user information (non-shareable) and will be stored in a Per User part[3].
3) Each of these parts are now Fingerprinted via the Fingerprint engine, generating a set of unique identifiers.
4) Enterprise Vault then queries the fingerprint database for each fingerprint and finds out if any of the designated SIS parts have already been stored.
5) If the SIS part has been stored before then Enterprise Vault will not store the SIS part to disk but instead increments the reference to that SIS part in the database.
6) If the SIS part has not been stored before, Enterprise Vault will write the unique SIS part to disk (including the index conversion copy which we'll talk about later) and creates a new Fingerprint reference in the Fingerprint database.
7) Enterprise Vault then always writes the Per User part to disk regardless of the sharing that has taken place.

Naming Conventions

Enterprise Vault stores each archived item and its associated SIS parts to the currently open Vault Store partition for the archive into which the item is being archived into. There are 3 main types of files used in the archive:

- *.DVS – This file contains all un-sharable (per user) properties relating to the archived item.
- *.DVSSP – This file or files each contain an individual SIS part.
- *. DVSCC – This file contains the HTML content conversion of the SIS part[4].

[3] Per User parts are never eligible for sharing.

[4] The content conversion is an HTML representation of the item. Enterprise Vault uses Oracle/Stellent's Inside Out content converters to generate this copy. The content conversion copy is used for indexing the item (AltaVista indexes this copy and not the original) and also optionally for displaying the item in HTML rather than original format if the user so wishes.

Figure 3 shows a screenshot of the files created when archiving an email item with attachments. Here we can see that Enterprise Vault determined that the message had 3 shareable parts (SIS parts). The Vault Store partition therefore contains a single DVS file for the non-shareable (Per User) properties of the email message, 3 DVSSP files (SIS parts) and 2 DVSCC files[5].

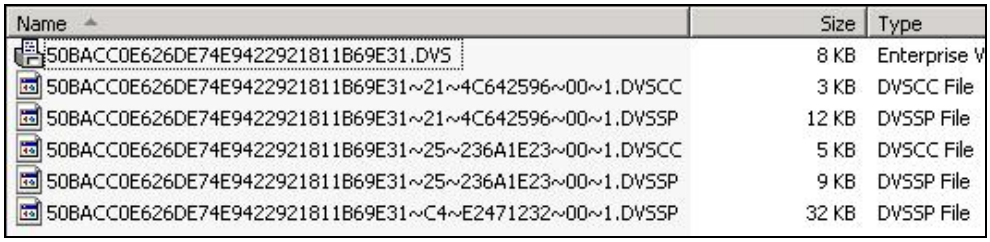

Figure 3. Example Archived Item

The names for each of these files and the path to the items within the Vault Store partition are generated from the Enterprise Vault Transaction ID[6] of the item, the current date and a number of other attributes.

In the example shown the path to these items would be as shown in Figure 4. The first level folder under the Vault Store Partition root is named after the current year, the next level the current month and day (hyphenated), the next level is the first character of the items Transaction ID, and the final level is the next 3 characters of the Transaction ID.

The actual names of the DVS, DVSSP and DVSCC files start with the full Transaction ID, in this case – 50BACC0E626DE74E9422921811B69E31:

Figure 4. Example Vault Store Partition Folder Path

[5] Why only 2 DVSCC files ? In this case one of the SIS parts (the shareable part of the email message) was already in HTML format and therefore Enterprise Vault is intelligent enough to see this and not generate a duplicate HTML copy, using the original for indexing purposes instead.

[6] A Enterprise Vault Transaction ID is a unique internal identifier that is assigned to an item when it is archived and used in subsequent operations for that item.

Storing items in this manner brings advantages, namely that each day Enterprise Vault will only store items into a single month-day folder. This means that previous month-day folders are no longer updated and Enterprise Vault backups can therefore be optimized.

Compression and deduplication

Enterprise Vault is usually configured to compress all archived items to further reduce the overall storage footprint of the archive. This is not always desirable however and for this reason compression can be configured on or off to suit the storage device to which Enterprise Vault is writing to. Take the following examples:

- If Enterprise Vault is writing to a storage device which already performs storage level compression then compressing the files first, before writing them to disk, is a waste of resources. Compression can therefore be turned off.
- If Enterprise Vault is writing to a storage device which performs block or segment level deduplication, in order to allow the items written to disk by Enterprise Vault to deduplicate with other items on the storage device, compression must be disabled[7]. In this case Enterprise Vault can continue to compress the non-shareable parts of an archived item but will not compress the shareable SIS parts.

With these examples in mind, each Vault Store partition within Enterprise Vault can be configured to individually take advantage of any compression or de-duplication features of the storage device to which they are storing to.

[7] Otherwise compression in this instance will act like encryption and not allow any further deduplication of the files.

Figure 5. Vault Store Partition – Compression and deduplication options

Figure 5 shows these options when creating a new Vault Store Partition in the Enterprise Vault Administration Console.

Collections

Enterprise Vault collections remain very similar in Enterprise Vault 8.0 in that they still provide a useful way to reduce the number of objects stored on a file system, and so decrease backup windows and increases storage performance. They also remain the enabler for further migration of archived data onto lower tiers of storage via the data migrator function.

Since items are not changed once they are written to disk, it is now possible in Enterprise Vault 8.0 to collect archived items much sooner than was possible before. Collecting after 1 day for example would be perfectly reasonable as this would not have any affect on any further potential sharing of those items[8].

[8] As the sharing process does not require modifying the archived item on disk.

When gathering archived items into collection files, Enterprise Vault tries to ensure that all relevant DVS, DVSSP and DVSCC files for the same archived item are collected into the same CAB file. However over time and as more sharing occurs this becomes more difficult as a highly shared SIS parts may possibly only exist in one collection file.

Large Files

Large files, in Enterprise Vault 8.0, are items being archived that are larger than 50 MB in size and these are processed differently from smaller files. To begin with the archived item will not be compressed (regardless of the compression setting) and it will not be eligible for collection. If Enterprise vault is configured to migrate data then the SIS parts and content collections are migrated directly rather being placed into a CAB file first. Large files will also have a content conversion copy (DVSCC) created but will only be indexed if this content conversion copy is less than 30 MB in size.

File extensions for large files are also slightly different:

- If sharing is not enabled large files are stored as *.DVF files.
- If sharing is enabled large files are stored as *.DVFSP files.
- Large file converted content files are stored as *.DVFCC files.

Retrieval of archived items

Now that we understand how items are stored, retrieving them is relatively simple:

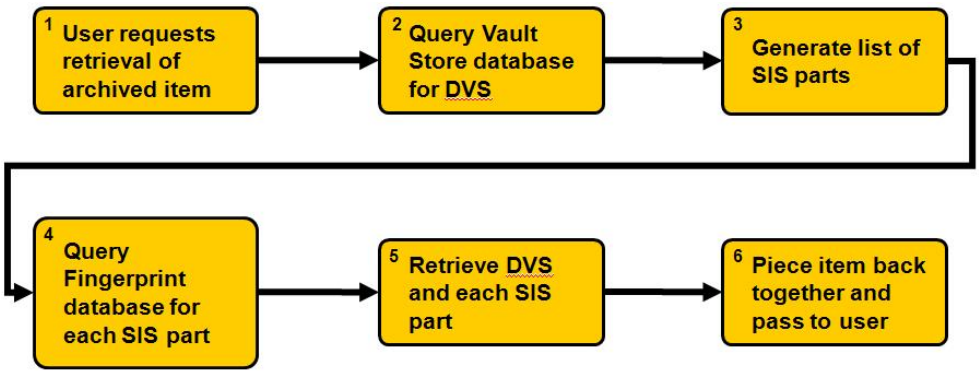

Figure 6. OSIS Retrieval Workflow

1) A user requests an archived item from Enterprise Vault.

2) Enterprise Vault looks up the location of the DVS file in the relevant Vault Store database.
3) A list of associated SIS parts that comprise this item is generated.
4) For each SIS part Enterprise Vault uses the Fingerprint database to find its location.
5) These various parts are then retrieved
6) And put back together to create the original item which is passed to the user.

The benefit here is that Enterprise Vault doesn't need to retrieve all parts of an item if the user only specified an exact part (such as an attachment). In this case Enterprise Vault will just return the part of the item that was requested rather than wasting resources and bandwidth returning the entire item.

Sharing Exceptions and Limitations

There are several notable exceptions to sharing that are worth mentioning.

Microsoft Office 2003 or earlier documents that have been sent as an attachment to an email via Microsoft Outlook might not share as expected. This is because Outlook 2003 (or earlier) makes some minor changes to the document metadata which results in the document having a different fingerprint to the original. Therefore this document would not share with the same document that has not been emailed, but it will share with the same document sent or received on another Microsoft Outlook message. For this same reason office documents that have been printed may not share as printed time is stored and updated in the document metadata. This limitation does not apply to Office 2007 documents or non-Office file types such as PDF or JPG.

Similarly, document attachments sent/received via a Lotus Notes email message will not share with the original documents as Lotus exports attachments in proprietary format which again results in a different fingerprint. They will however share with the same documents sent/received on other Notes email messages.

Conclusion

We have covered a lot of ground in this whitepaper however here are some key points to remember and take away:

- Enterprise Vault OSIS can share at the item level across Vault Stores.
- Sharing of items across different content sources and storage devices is possible using OSIS.
- Attachments to emails can be shared separately from the email itself.
- Storage level deduplication technologies offered by vendors such as DataDomain and NetApp can be leveraged when using OSIS.
- User access to archived items has been optimized to only return the content that they have requested.

About Symantec

Symantec is a global leader in providing security, storage and systems management solutions to help businesses and consumers secure and manage their information. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com