

Symantec Enterprise Vault™

Virtual Vault Best Practice Guide

Who should read this paper

This Whitepaper is intended to assist customers, partners and service providers deploy the Virtual Vault feature within Enterprise Vault.

If you have any feedback or questions about this document please email them to iig-tfe@symantec.com stating the document title.

This document applies to the following version(s) of Enterprise Vault:

11.0 and later

This document is provided for informational purposes only. All warranties relating to the information in this document, either express or implied, are disclaimed to the maximum extent allowed by law. The information in this document is subject to change without notice. Copyright © 2014 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo, and the Checkmark Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

Table of Contents

Scope of Document	1
Intended Audience	1
Introduction	1
Overview	2
How Virtual Vault works	3
Initial synchronization	4
Incremental synchronization	4
Upload and archive	4
Content cache download	5
How to configure Vault Cache and Virtual Vault end-user policies	6
How to configure Vault Cache and Virtual Vault server-side operations	19
Header cache configuration options	19
Virtual Vault archiving	22
Content cache configuration options	23
Exchange Server configuration options	26
SQL server configuration options	26
Rollout Strategy	27
Before you roll out Virtual Vault	28
Server Performance Impact	28
Roll-out Strategy	28
Rollout scenarios	29
Summary	35
Appendices	
Appendix A - Virtual Vault advanced settings	
Appendix B - Vault Cache policy settings	
Appendix C - Vault Cache advanced policy settings	
Appendix D - Vault Cache client registry settings	
Appendix E - performance counters to monitor	

Document Control

Contributors

Who	Contribution
Chris Harrison	Author
Dan Strydom	Updates for EV 11.0

Revision History

Version	Date	Changes
1.0	June 2011	Published
2.0	February 2012	Updates for EV 9.0.5 (Threshold Based Sync)
3.0	May 2014	Updates for EV 11.0

Related Documents

Document Title	Document Location	Version
Virtual Vault Feature Overview	www.symantec.com/docs/TECH124851	1.0
Symantec Enterprise Vault Administrator's Guide	www.symantec.com/docs/DOC6634	11.0
Enterprise Vault 11 Performance Guide	www.symantec.com/docs/TECH125795	11.0
Enterprise Vault Compatibility List	www.symantec.com/docs/TECH38537	-
Advanced Strategies for Monitoring Enterprise Vault	www.symantec.com/docs/HOWTO74545	-

Scope of Document

This document provides best practice configuration and design details for the Virtual Vault feature of Enterprise Vault.

Virtual Vault is available for Enterprise Vault users with Microsoft Exchange archives that use Microsoft Outlook to access their email data.

There are a number of policy and server configuration options which either directly or indirectly have an impact on this feature, whether from an end-user-experience perspective, or from a server performance/resource perspective.

This best practice document aims to provide more details on these policy and server settings and their overall impact, as well as providing details of several example rollout scenarios, server settings, end user experience, performance considerations, and any other relevant information.

Intended Audience

This document is aimed at customers, partners, and support staff. It is assumed that the reader has a thorough understanding of the architecture and operational aspects of Enterprise Vault and Vault Cache. Additionally it is also assumed that the reader has read the technical whitepaper Virtual Vault Feature Overview (referenced documents).

Introduction

The majority of users use Microsoft Outlook for Windows to access both Exchange and archived email. Prior to Enterprise Vault 8.0 service pack 3, end-users were able to access archived content in Microsoft Outlook through the use of Enterprise Vault Search, Archive Explorer, and by clicking on archived shortcuts. With Enterprise Vault 8.0 service pack 3 the end-user experience has been extended to include Virtual Vault, so that a user is able to view their entire archive seamlessly within Microsoft Outlook with behaviour akin to that of a PST file.

Overview

Virtual Vault provides a seamless experience for end users accessing their Enterprise Vault archive within Microsoft Outlook. With the Enterprise Vault Add-in for Outlook, users can now see the folder structure of their archive, and access individual archived items without the need for shortcuts in their mailbox. Virtual Vault appears like a mailbox folder or personal folder in the Microsoft Outlook navigation pane but provides direct access to your archived items (see Figure 1).

Virtual vault is a component of Vault Cache which means that any administrator wishing to rollout Virtual Vault to their end-users must first enable Vault Cache. As Virtual Vault relies on the underlying Vault Cache infrastructure it is imperative that an administrator sets policy and configuration options correctly pertaining to both Virtual Vault and Vault Cache in order to achieve the desired experience for both the end-user (in Microsoft Outlook) and the administrator (with regards to the impact on the Enterprise Vault server infrastructure as a whole).

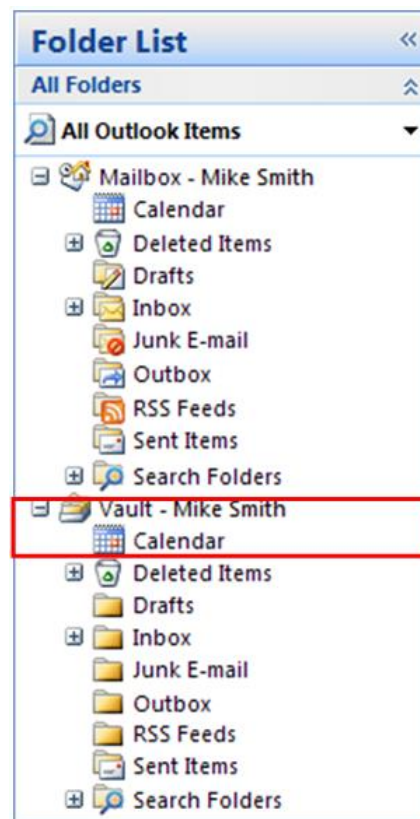


Figure 1 – Virtual Vault

This section will cover the following:

- High level introduction on how Virtual Vault works.
- Details of the policy options that control end user Virtual Vault functionality.
- Details of the policy options that control Vault Cache behaviour.
- Details of the server configuration options that control server resources applicable to Virtual Vault/Vault Cache performance.

How Virtual Vault works

Figure 2 provides an overview of the related Enterprise Vault components and how they fit into and interact with the existing Enterprise Vault Outlook Add-in.

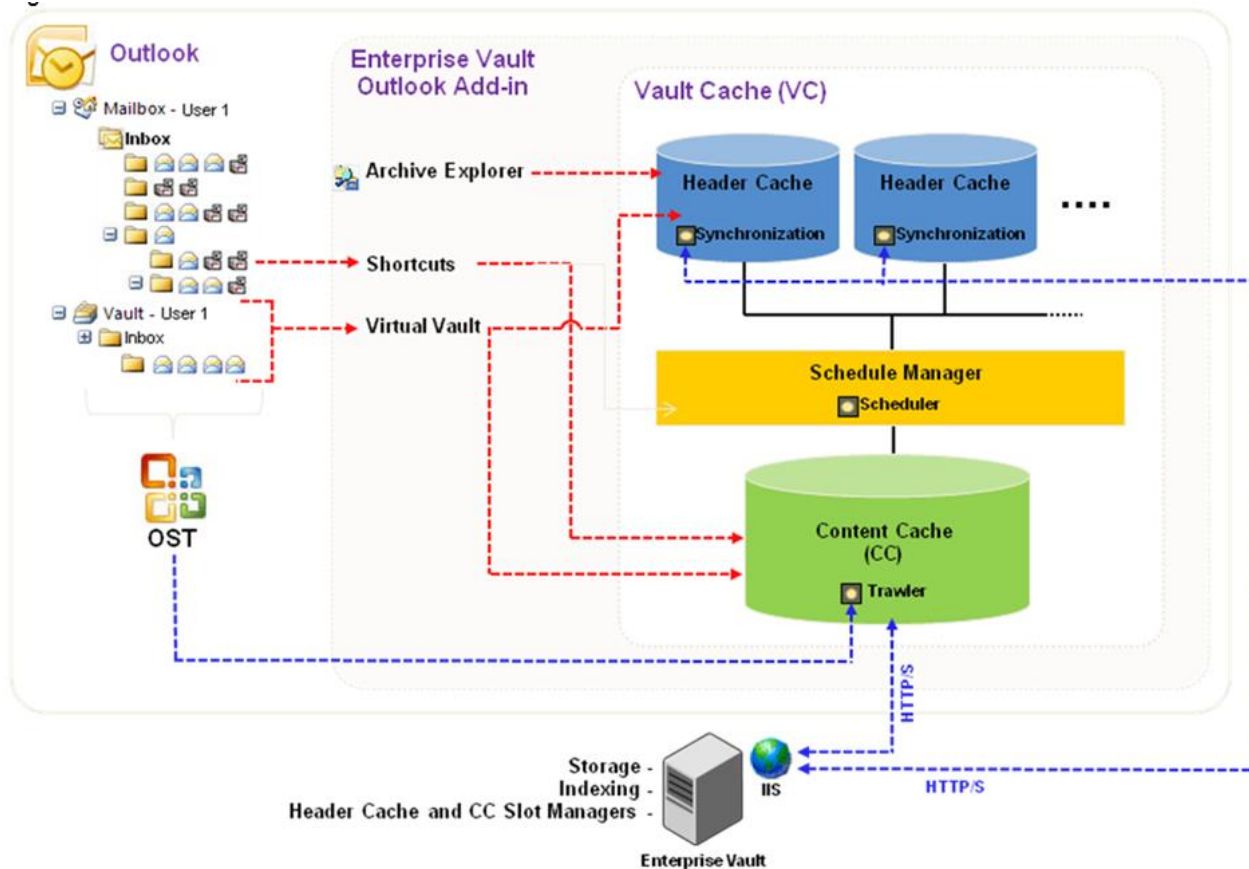


Figure 2 – Virtual Vault Overview

Virtual vault acts as the interface between the end-user and their archive by interacting with Vault Cache and its underlying components. Using the information contained within the header cache, Virtual Vault provides the end-user with a hierarchical view of their archive. Selecting a folder in the archive displays the header information for the archived content within that folder. When opening a specific item the item's full content is retrieved from the content cache and displayed to the user. The header cache consists of data such as:

- Folder hierarchy and specific folder properties (for example name, retention category, etc).
- Item specific metadata (for example recipient lists, subject, sent/received date, etc).

The header cache and content cache data are kept up to date via two automated client synchronization processes. Header cache synchronization drives the content cache download process, however they utilize different server-side components to update the end-user's workstation with the relevant data.

Header cache synchronization takes place once a day by default. The end-user's Vault Cache synchronization manager requests a slot from the header cache slot manager component on the Enterprise Vault server. The role of the header cache slot manager component on the Enterprise Vault server is to accept or refuse the end-user request for a slot, based on whether a slot is available or not. Providing that the request for a slot is accepted, the next phase will commence, which will be either an initial (full) synchronization, an incremental synchronization, or a hierarchy-only synchronization. If there is a problem with obtaining a slot, a retry mechanism with a randomized interval period is invoked that will eventually allow a slot to be successfully obtained.

Initial synchronization

An initial header cache synchronization (usually a one-time operation) involves loading an end-user's index into memory on the Enterprise Vault server which is responsible for indexing requests relating to that end-user's archive. This will consume I/O, CPU and memory on the applicable Enterprise Vault server, whilst having to share resources with other Enterprise Vault related operations on that server¹. In order to minimize the impact that this (loading large numbers of indexes) could have, the number of concurrent initial header cache synchronizations is limited by default to 10. This operation will also cause increased SQL activity that will elevate disk I/O on the disk where the databases are located.

Incremental synchronization

Once the initial header cache synchronization has taken place, further incremental synchronizations are performed against SQL tables in the end-user's vault store database. Queries are executed to determine what differences need to be synchronised down to the end-user and this data is then identified by entries in these vault store database tables. This is less costly from an Enterprise Vault server resource perspective, however there will be an additional load levied against the SQL server due to lookups being performed against the various vault store database tables.

Upload and archive

Header cache synchronization not only synchronises the folder hierarchy and item header cache data, it also has the ability to upload data archived through Virtual Vault, as well as informing the content cache what data needs to be downloaded. Allowing uploads (an archive operation) through Virtual Vault will cause additional

¹ Examples of other operations include PST migrations, end user operations (view, restore), and archiving tasks (Exchange, Domino, File System, SharePoint, Instant Messaging)

network traffic between the end-user workstations and the Enterprise Vault server. It will also consume Enterprise Vault server and SQL server resources (CPU, memory, IIS, Enterprise Vault processes, database lookups) and I/O (cache location, vault store partitions, database disk).

Content cache download

Content cache downloads are a background process and therefore “always on” unless the end-user has suspended the operation. The header cache synchronization process will notify the content cache that there are updates to the archive that require actioning. The Vault Cache synchronization manager will request a slot from the content cache slot manager component on the Enterprise Vault server. Providing a slot is available, a server-side process called MigratorServer.exe is spawned which retrieves data from the relevant vault store partition locations in order to create temporary content cache data in the cache location² specified on each Enterprise Vault server. This data exists as a file with extension .PST in a sub-folder of the Enterprise Vault server cache location called ‘VCBuilds’. Once the requisite amount of archived data has been copied into the PST file, the PST file will be downloaded to the end-user’s content cache via the Microsoft Windows Binary Intelligent Transfer Service (BITS) technology. The PST file is deleted from the VCBuilds folder once the end-user’s workstation has successfully downloaded and Vault Cache has integrated the downloaded PST file data³ into the content cache data (.DB) files⁴.

Both the synchronization and download processes will consume time and resources, depending on the amount of work there is to do. It is possible to place a limit on certain operations, or even turn them off completely, in order to ensure that the time it takes, and the resources used, are satisfactory both from an end user experience and from a server performance/resource perspective.

In order to generate the best settings for your environment it is important to understand the impact of the settings that are being modified. The next section explains some of the more important settings and the impact of changing the default values.

² The location, and other information, is discussed later on in this document under the “Additional server configuration options” section.

³ If for some reason the end-user workstation does not complete the download successfully, there are other checks that will remove this data if it should become stale.

⁴ Other mechanisms are also involved in monitoring and removing PST file data from the VCBuilds folder depending on certain conditions, these mechanisms are not covered in this document.

Full details of all the vault administration console settings listed over the next few pages are available in the appendices towards the end of this document, or in the following documents located on the Enterprise Vault server in the “...\Enterprise Vault\Documentation” folder:

- “Administrators Guide.pdf”
- “Registry Values.pdf”
- “Setting up Exchange Archiving.pdf”

How to configure Vault Cache and Virtual Vault end-user policies

Configuration settings for Vault Cache and Virtual Vault are located in the vault administration console under the container Policies > Exchange > Desktop. See Figure 3, and Appendices A-C for full details and descriptions of each policy setting. These configuration settings can also be rolled out to end-users via client-side registry settings. The exception to this rule is the policy to control how often a synchronization occurs between the end-user workstations and their Enterprise Vault server. By default a synchronization occurs every 24 hours. The only way to change the default value is to deploy a client-side registry key called OVMDCSyncFrequencyInSecs (see Appendix D).

Your company’s requirements will determine whether the settings you wish to implement can be fully configured using a number of desktop policies created through the vault administration console, or whether you will need to create registry files with the desired settings and then distribute that to the end-users (via GPO, logon script, etc).

Figure 3 Screenshots of a) the Vault Cache tab, and b) the advanced tab with Virtual Vault dropdown selected, both from a sample Exchange desktop policy.

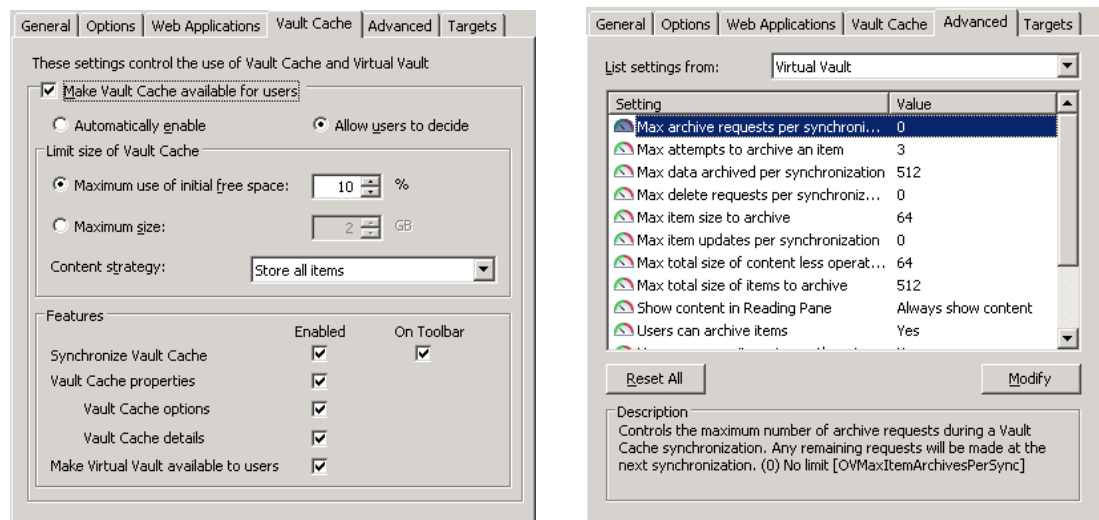


Figure 3 – Vault Cache tab with Virtual Vault Advanced settings

If you allow end-users to perform certain operations in Virtual Vault, for example the archiving or deleting of data, you may choose to change some of the other settings to tailor the behaviour of the operation to fit your company's requirements. This section covers the main operations and the relevant settings that have an effect on the behaviour of the operation in question, in the following order:

- Users archiving items
- Users deleting items
- Users reorganising items
- Users restoring items
- Users copying items
- Content strategy
- Synchronization
- Users can archive items (default = yes).

Any setting that affects the uploading / archiving process should be carefully considered.

- **Max archive requests per synchronization (default = 0, no limit).**

If you allowing archiving, do you want to throttle the number of items each end-user uploads during a synchronization? Data that is archived through Virtual Vault resides in the local header cache until a synchronization occurs and that data is then uploaded. If you do decide to throttle, whatever data is not synchronised on this attempt will reside locally until the next synchronization occurs. As archiving through Virtual Vault does not provide the safety copy functionality that scheduled archiving does⁵, this data is not secured from data loss whilst residing in the header cache if a hard drive should fail so that the header cache data could not be read/recovered.

Some companies operate on a zero data loss policy and negate the risk by backing up data in a variety of ways to minimize the possibility. Having data residing in a local cache that is not backed up would be the direct opposite of this policy. To minimize this breach of policy an unlimited threshold would enable every synchronization to upload all data that is due to be archived through Virtual Vault.

⁵ If your vault store is configured to remove safety copies after backup

If you do not throttle, consider the following:

- Can your network bandwidth cope with the additional traffic?
- Is there spare capacity on your Enterprise Vault server to handle the additional load that archiving through Virtual Vault will generate?

If you already allow manual archiving using the Enterprise Vault Outlook Add-in, there may be little or no additional impact generated as manual archiving is already in use. Ultimately there is little difference between the two operations, apart from the aforementioned safety copies generated by manually archiving an item, and the underlying mechanism used to perform the archiving operation. Archiving through Virtual Vault, and hence not generating safety copies in doing so, would have the benefit of immediately reducing the size of the end-user's mailbox from a quota perspective, which could be important for the end-user's workflow if mailbox quotas are in place and the values are low. One way of defining a threshold for this setting could be to work out the average number of items an end-user will send and receive on an average day and set the value accordingly. This approach assumes that all data sent and received will be archived through Virtual Vault. It is possible to achieve this by using rules to move all email data automatically from your Microsoft Outlook's inbox directly into a folder within your Virtual Vault. In reality, while some companies will indeed do this – basically using Virtual Vault as a direct PST replacement, – for the majority this is unlikely to be the case. Some data will be deleted by the end user within Microsoft Outlook and will never be archived. Scheduled archiving is also more than likely to be in use and archive the majority of data that meets the archiving policy that is applied to the mailbox. Archiving through Virtual Vault could therefore be considered as an ad-hoc archiving operation that compliments scheduled archiving.

Finally, you must consider PST data as a source to archive from. Customers may have already eliminated PSTs from their environment using one of several PST migration strategies available with Enterprise Vault. However if this is not the case, it is possible that end-users could archive their PST data through Virtual Vault, and as such the threshold value may need to be raised or lowered depending on the conclusions drawn from the comments above.

- **Max attempts to archive an item (default = 3).**

This setting relates to the unlikely situation whereby an item might fail to be uploaded and archived.

There are two main reasons why an item would fail to be uploaded:

- **Data integrity.** If there is an issue with the actual data itself then this could prevent either the upload of the item or the subsequent archive operation from successfully completing. In this situation it is likely that no amount of retrying will be successful.
- **Transitive environment issues.** The nature of these issues varies greatly and includes IIS related issues, network issues, resource issues, and so on.

As such, the default value for this setting (3) is there really to act as a retry mechanism for these transitive errors, which shouldn't be altogether that common as long as the environment is relatively stable for the most part. Any items which fail to archive after the retry limit is reached will appear in a search folder within Virtual Vault called "Could Not Archive". These items should be investigated with Support as to why they will not archive if the root cause for this failure is not immediately obvious.

- **Max data archived per synchronization (default = 512 MB).**

This setting is highly similar to the previously discussed *max archive requests per synchronization*. The key difference between the two is that this particular setting relates to the size of data to be uploaded rather than the number of items.

This setting is potentially more important than the earlier item-based setting as it is the size of the data uploaded that could cause larger issues. For example, consider the difference in impact of uploading 100 items each of size 10kb, versus uploading 100 items each of size 1mb. With the former, the number of items being uploaded is not problematic. With the latter, due to the size of each item it will take considerably longer to synchronise, and more resources will be used server-side in terms of IIS, memory, disk I/O, CPU and so on.

This setting therefore allows you to set a reasonable limit so that an end-user with lots of small items could successfully upload all their items in a single synchronization in a reasonable time frame. An end-user with a few large files can also do the same, while protecting the various environment resources against the end-user who has lots of large items and decides to archive them through Virtual Vault all at once.

Generally, *max data archived per synchronization* should be set to a value greater than or equal to the value of *max total size of items to archive*, so as to maximise the chance of all items being uploaded in a single synchronization.

- **Max item size to archive (default = 256 MB).**

This setting can override two other settings, namely *max data archived per synchronization* and *max total size of items to archive*. The logic behind this is that if you have specifically set this value higher than the other two, then this was done for a specific reason. The other thing to consider when setting this is that if the maximum item size to archive value is equal to the value set for *max data archived per synchronization*, a single large item could constitute the an entire synchronization, to the exclusion of any other data that requires uploading. One way of setting this value could be to duplicate the largest item size allowed to be sent or received through Microsoft Exchange⁶.

- **Max total size of items to archive (default = 512 MB).**

This setting controls the maximum size of items allowed to exist in the header cache that have not yet been uploaded for archiving. As per previous comments in this section, this value really depends on the company policy towards unsecured data. As previously mentioned, this value should be set to less than or equal to the *max data archived per synchronization* value.

- **Users can hard delete items (default = yes)^{7 8}.**

Any setting that affects deletion synchronization should be reviewed.

⁶ The assumption in this instance being that limits have been placed on the size of items allowed to be sent or received.

⁷ Deletion is only possible if the site setting *users can delete items from their archives* is enabled, and if the retention period for the item has either expired or is not enforced.

⁸ This section refers to deletion from the perspective of an end user deleting an item through Virtual Vault, and not through the server-side operation storage expiry

- **Max delete requests per synchronization (default = 0, no limit).**

Allowing deletions does not have the same end-user impact as allowing archiving through Virtual Vault. The obvious reason for this is due to the amount of data involved being so much less (no content to upload for archiving).

One of the few considerations to take into account is the effect of mass deletion of data by an end-user using Virtual Vault on the server-side components. Deletion of thousands of items will have an impact on the SQL Server (database lookups taking up CPU, memory and disk I/O) and the Enterprise Vault storage server⁹ (Enterprise Vault storage and monitoring processes, vault store partition disk I/O). You can safeguard against this by setting a threshold to a figure deemed appropriate. In all likelihood, end-user deletion through Virtual Vault will potentially occur in two different scenarios:

- When an end-user is first enabled for Virtual Vault (and now able to manipulate their archived data more easily than previous versions).
- When an unwanted item in their mailbox is archived and the end-user wants to 'clean' up their archive.

Therefore the chance of this situation (thousands of items being deleted by the end user through Virtual Vault) occurring is likely to be relatively low.

- **Users can reorganise items (default = yes).**

Any setting that affects the manipulation of items and folders within the archive should be reviewed.

- **Max item updates per synchronization (default = 0, no limit).**

An update in Virtual Vault refers to the movement of an item, or the update of an item's retention category. As with deletion synchronization, the impact of allowing end-users to update an item's location and / or retention category does not have the same level of impact that allowing archiving

⁹ A storage server is an Enterprise Vault server that has a storage service.

through Virtual Vault. In fact, changing the default value (no limit) or setting this value to a low figure may only cause confusion for end-users when some of their items are updated and others are not.

- **Users can copy items to another store (default = yes)¹⁰.**

This setting allows end-users to copy / move items out of the archive and into another location, whether that be the Microsoft Outlook mailbox, a PST file or another location. A copy / move operation out of the archive will cause an item retrieval from the Enterprise Vault server if that content does not exist in the content cache. See the *content strategy* and *max total size of content-less operations* settings below for more information.

- **Users can copy items within their archive (default = no).**

This setting mirrors the policy setting *allow shortcut copy* in the mailbox. A copy operation within the archive will cause an item retrieval from the Enterprise Vault server if that content does not exist in the content cache. See the *content strategy* and *max total size of content-less operations* settings below for more information.

- **Content strategy (default = store all items).**

Any setting that has a need for the content of the item to be available should be reviewed.

Content strategy is the setting that determines how much content is downloaded to end-user workstations that are Vault Cache enabled. The three content strategy options are:

- Do not store any items in cache. Header cache information is synchronized, however the content of archived items is not downloaded and stored in the content cache.
- Store all items. Header cache information and the content of archived items are downloaded and stored in the content cache.

¹⁰ If end users are not allowed to delete from their archives (if the site setting *users can delete items from their archives* is not enabled) a move operation will fail and the user will be prompted to copy the item instead.

- Store only items that users open. Header cache information is synchronized, but only the content of archived items that are opened in Virtual Vault are downloaded and stored in the content cache.

Having a best-fit *content strategy* policy can ensure that no – or relatively few – content-less operations need to occur, effectively removing any concerns over the load that *max total size of content-less operations* adds to the storage and SQL servers.

If your chosen strategy is *do not store* or *store only items that users open* then it is highly likely that any move / copy operation of two items or more out of the archive will require the content to be retrieved in order to process the end-user's request. You may wish therefore to set *max total size of content-less operations* to a lower value if there is a concern over resources.

If your strategy is *store all items* then it is likely that any move / copy operation of two items or more out of the archive will not initiate a server-side request for content. However, there are three main instances where this may not always be the case:

- The space allocated to Vault Cache may not be large enough to ensure all end-users are able to download their entire archive locally. In this instance the end-users can increase the size of their Vault Cache themselves.
- The *download item age limit*¹¹ setting may be configured to only download items of a certain age (for example data up to one year old), as such any older items will not have their content downloaded. Again, end-users are able to modify this value as long as the administrator has not locked this setting, and only during the initial configuration of their Vault Cache¹².
- In some situations, it is possible that the content has not yet been downloaded¹³, or pre-emptively cached, into the content cache. Pre-emptive caching is discussed in the next section below.

¹¹ Refer to the Administrators Guide for more information.

¹² Any change later on is not retrospective.

¹³ For example Vault Cache may have only just been configured and there is a large amount of content that needs to be downloaded.

- **Offline store required (default = yes).**

Previously, end-users who did not work in Microsoft Outlook cached mode were unable to use Vault Cache as Vault Cache had a dependency on the presence of an OST file. The *offline store required* setting now enables an administrator to give Vault Cache functionality to end-users that do not have an OST. However, there is a difference in behaviour depending on whether you have an OST or not.¹⁴

End-users running in Microsoft Outlook cached mode are able to utilize a trawler mechanism within the Enterprise Vault Outlook Add-in, which actively trawls the contents of the OST looking for data that is soon to be archived based on the archiving policy applied to the mailbox. This data is pre-emptively copied into the content cache before the archive operation occurs. Therefore when the data is actually archived, no subsequent download of that content from the Enterprise Vault server back down to the end-user workstation is required, thus resulting in various resource benefits¹⁵. With this configuration it often makes sense to utilize the *content strategy* of 'store all items' as this incurs little if any overhead server-side and the benefit to the end-user is a full content cache.

End-users not running in Microsoft Outlook cached mode are not able utilize the trawler mechanism. Therefore any content requiring download based on the *content strategy* policy will always have to come from the storage service, expending storage server, SQL server and bandwidth resources. With this configuration it often makes sense to utilize the *content strategy* 'do not store' or 'store only items that users open'.

- **Show content in reading pane (default = when in Vault Cache).**

This setting has three values:

- Always show content. The reading pane always shows the header and content of the item that is selected in Virtual Vault.

¹⁴ For the remainder of this whitepaper, it is assumed that when cached mode is being used the OST is available, and when cached mode is not being used there is no OST available.

¹⁵ For further information regarding the relevant policy settings relating to preemptive caching, please refer to the Administrators Guide section "Advanced Exchange mailbox and desktop policy settings"

- When in Vault Cache. The reading pane shows the selected item's header. If the item is in Vault Cache, it also shows the content. If the content is not shown, a banner provides a link to open the original item. When the content strategy is *store only items that user opens*, the effect of this value is that the reading pane only shows the content of previously opened items.
- Never show content. The reading pane always shows only the selected item's header cache information. A banner provides a link to open the item.

When determining the right value, it is important to consider the content strategy being utilized. One of the major policy changes in Enterprise Vault 9.0.3 restricts certain policy settings from being used in conjunction with each other. With earlier versions of Enterprise Vault, it was possible to have *always show content* when the *content strategy* was set to *do not store any items in cache*. This combination often resulted in excessive downloading of content from the Enterprise Vault server and could affect other operations involving viewing of data. With the release of Enterprise Vault 9.0.3, *always show content* is only available if you have upgraded from an earlier release, and in that earlier release, *show content in reading pane* had the value *always show content*. For customers in this situation it is highly recommended that this setting be changed to *when in Vault Cache*.

As a rule of thumb, it is logical to show content when it is there. A setting of value *when in Vault Cache* makes sense as a way of providing full item content for those end-users who have the full content in their content cache. Other benefits include reducing network bandwidth consumption, and the Enterprise Vault server from potentially having a significant number of view requests for those users that do not have local content in their content cache. However, this could lead to a difference in behaviour in some situations. An end-user could have a full content cache, but with the space allocated to their Vault Cache being smaller than the size of their archive, resulting in some content not being inserted into the content cache. An end-user could open an item that had not been inserted into their content cache, resulting in the item's header being displayed. This behaviour could appear to the end-user as being inconsistent. To ensure a consistent behaviour, it may be worth having separate policies for end-users with content to use *always show content* and users without content to use *never show content*.

- **Max total size of content-less operations (default = 64 MB).**

This setting comes into effect when two or more items are moved or copied – effectively a restore operation – out of the archive¹⁶ and the resulting content does not exist in the content cache. This will result in an additional load on the applicable storage server servicing the request as well as the SQL server, not to mention the increase in network traffic to download the content. Setting the value too high will allow end-users to export a large amount of data and will consume higher levels of resources. Setting the value too low may result in end-users becoming frustrated when they are unable to bulk move / copy items out of their archive. However, one should question the real need that the end-user has to perform this operation.

Prior to Virtual Vault functionality it was understandable that end-users may need to restore items for one reason or another¹⁷. Now with Virtual Vault end-users are able to seamlessly access their archived data within Microsoft Outlook and perform a variety of operations on that archive as they would with regular emails in their Microsoft Outlook mailbox. The actual need therefore to restore data has disappeared and as such a low value could be workable. It may be worthwhile for already existing customers to inform end-users of the change in functionality that Virtual Vault offers so that they can understand that previous limitation such as these are no longer apparent.

- **Synchronization.**

Any way in which synchronization of data or archives that can be modified should be reviewed.

When we talk about synchronization, it is important to mention that this can relate to synchronization of data, as well as synchronization of archives. For the former, thresholds can be configured that, when met, will cause a synchronization of data to occur. For the latter, what archives a user synchronizes locally can be controlled depending on the requirements of the end-users in question, or depending on the decisions of the administrator in control of the system. In either case, the more often synchronization occurs, the more server-side resources will be used as part of this process,

¹⁶ This location could be anywhere, for example back to the Exchange store, into a PST, or onto a local or networked drive.

¹⁷ For example an end user may need to forward an item to another user, but that item's shortcut no longer exists.

therefore any change to settings that affect the frequency of synchronizations should be considered carefully.

- **OVMDCSyncFrequencyInSecs (default = 86400).**

As mentioned earlier, synchronization occurs every 24 hours by default. This can be modified via a client side registry key, OVMDCSyncFrequencyInSecs. If archiving through Virtual Vault is enabled then you may wish to reduce this value so that unsecured data does not reside in the local content cache for too long, however changing a threshold setting may be more suitable than more frequent time-based synchronization, depending on the requirements of the business.

- **OVAAllowMissedMDCSyncOnStartup (default = 0).**

By default, the Outlook client add-ins initiates a synchronization shortly after an end-user connects to the network with Outlook running if a scheduled synchronization has been missed. This particular scenario is more than likely to occur over the weekend period where most, if not all, end-users will not be working. If a large number of users all attempt to synchronize at the same time this could lead to slot contention issues, as well as impacting SQL Server resources. To mitigate against this possible scenario, a new client-side registry setting, OVAAllowMissedMDCSyncOnStartup, has been added which essentially informs the client to act as if the last scheduled synchronization was successful.

To explain, the following scenario details what would happen when using a 3 hour scheduled synchronization interval:

- 4pm – user completes a scheduled synchronization.
- 5pm – user logs off. (If the user had stayed logged in all night, their synchronization schedule would have been 7pm, 10pm, 1am, 4am, and 7am.)
- 9am – user logs in. The client behaves as if the last scheduled synchronization had taken place at 7am and a synchronization is not initiated at this stage.
- 10am – a scheduled synchronization is now performed.

- **Threshold number of items to trigger synchronization (default = 0).**

Administrators may wish to have an option other than time to control synchronization frequency. *Threshold number of items to trigger synchronization* adds the ability to control the number of items ready for synchronization before a synchronization will be triggered.

- **Threshold total size of items to trigger synchronization (default = 0).**

Similar to the above threshold setting, *threshold total size of items to trigger synchronization* adds the ability to control the size of data ready for synchronization before a synchronization will be triggered. Both of these values are clearly aimed at customers that have enabled the ability to archive through Virtual Vault, and wish to synchronize this data via a non-time specific function.

Care should be taken when configuring these values - too high and they will not be effective, too low and they may cause unnecessary synchronizations to occur. The primary method of archiving should ideally be through a scheduled archiving task whose policy has been configured according to the size of the environment and the amount / age of data that needs to be archived. Manual archiving, whether it be via the Outlook client add-in toolbar, or by archiving through Virtual Vault, should only be used for ad-hoc archiving purposes and not as a replacement for a properly configured scheduled archiving task.

- **Synchronize archive types (default = mailbox).**

This setting has three values:

- Default mailbox. Synchronize the primary mailbox archive only.
- All mailbox archives. Synchronize the primary mailbox, and any delegate mailbox archives to which the user has access.
- All mailbox and shared archives. Synchronize the primary mailbox, and any delegate or shared mailbox archives to which the user has access.

Most end-users will likely only have access to a single mailbox, and therefore will only require access to a single mailbox archive when offline. Therefore, for most administrators, *synchronize archive types* will always be left on the default value. If this is the case, it is possible to bypass some of the permissions checks that take place during a synchronization which are used to identify what other archives an end-user has access to, seeing as only the default mailbox archive is configured to synchronize down to the end user.

If this setting is changed to one of the other values, additional performance checks will be called which will cause SQL Server utilization to increase. If there are specific end-users that require access to multiple archives when offline, consider creating a separate desktop policy for these end-users.

How to configure Vault Cache and Virtual Vault server-side operations

From a server perspective this section can be split into four main parts:

- Header cache configuration options
- Content cache configuration options
- Exchange server configuration options
- SQL server configuration options

It is important to consider the environment as a whole when making changes to Enterprise Vault. This next section will provide information on changing default behaviour and thresholds of some processes. As such these changes should be made with caution and their effect on the environment should be monitored to ensure the environment can cope with the changes made.

Header cache configuration options

Initial and incremental synchronizations

As previously discussed in the section “How Virtual Vault works”, a slot mechanism is used by the Enterprise Vault server to accept or refuse an end-user request for a slot in order to perform a Vault Cache synchronization. A slot is required to perform an initial, an incremental synchronization, or a folder hierarchy only synchronization. Once Vault Cache and Virtual Vault have been deployed and end-users have completed their initial synchronizations, these slots will generally then be used only for incremental synchronizations. Incremental synchronizations are performed against SQL tables in the end-user’s vault store database. Queries are executed to determine what differences need to be synchronised down to the end-user, causing an additional load to be levied against the various vault store database tables. Depending on your SQL server configuration (see the “SQL server configuration options” section) you may wish to lower or raise the number of slots each storage server will use. The number of slots available per storage server is 100 by default. To modify this value, use the following SQL query:

```
USE EnterpriseVaultDirectory
UPDATE StorageServiceEntry
SET SyncSlots = X
WHERE ComputerEntryId = Y
```

X = the number of synchronization slots that are required.

Y = the GUID of the EV server that owns the storage service you wish to change. Use the following query to help identify the ComputerEntryId of the server you wish to change:

```
USE EnterpriseVaultDirectory
```

```
SELECT ComputerEntryId, ComputerName, ComputerNameAlternate
```

```
FROM ComputerEntry
```

A restart of the storage service is required when changing this value.

Initial synchronizations

As previously discussed in the earlier section “How Virtual Vault works”, the initial synchronization process involves loading the end-user’s index into memory on the Enterprise Vault server which is responsible for indexing requests relating to that end-user’s archive. A special type of slot called a full synchronization slot is used to provide this function.

On an Enterprise Vault server that has an indexing service configured, Indexserver.exe processes are spawned by requests for data that exists in an index. These Indexserver.exe processes are used to load indexes into memory. The larger the index/index request, the more memory that is required to service this request. As servers are equipped with finite reserves of memory, the number of Indexserver.exe processes able to be spawned is capped at 30 to guard against the possibility of too many indexing requests causing a low memory condition on the Enterprise Vault server, which could lead to loss of service.

The default number of Indexserver.exe processes able to be spawned as part of initial synchronization operations is capped at 10. This ensures that initial synchronization operations do not monopolise all Indexserver.exe processes on an Enterprise Vault server, allowing other indexing requests such as end-user search and Discovery Accelerator search to run concurrently.

With the uptake of 64bit architecture (both hardware and software) by some customers, the previous restrictions on the amount of memory on a server is not as restricted as it once was. This means that the theoretical number of concurrent Indexserver.exe processes able to be spawned safely on a server is higher

than that of their 32bit counterparts. This in turn means that the number of Indexserver.exe processes for initial synchronization operations can also theoretically be increased¹⁸.

To modify the number of Indexserver.exe processes that are able to be spawned on an Enterprise Vault server that has an indexing service configured, create the following registry key:

Name: MaxIndexServers
 Type: DWORD
 Location: HKEY_LOCAL_MACHINE\Software\KVS\Enterprise Vault\Indexing
 Default:30

To modify the number of full synchronization slots¹⁹ available – and therefore the number of Indexserver.exe processes that are able to be spawned as part of initial synchronization operations – use the following SQL query:

```
USE EnterpriseVaultDirectory  

UPDATE StorageServiceEntry  

SET FullSyncSlots = X  

WHERE ComputerEntryId = Y
```

X = the number of full synchronization slots you wish to allow.

Y = the GUID of the EV server that owns the storage service you wish to change. Use the following query to help identify the ComputerEntryId of the server you wish to change:

¹⁸ Having more available memory in a server does not necessarily mean that this is all that is required for a server to be able to service more Indexserver.exe requests without issue. Care must be taken to ensure that other hardware factors such as CPU (speed and no. of physical processors and cores), disk I/O, etc, are also at a standard that is able to cope with the increased number of Indexserver.exe processes being spawned.

¹⁹ The maximum value for full synchronization slots must be less than or equal to the number of synchronization slots configured. This is because full synchronization slots start as a 'regular' synchronization slot which is then promoted to a full synchronization slot. As such, full synchronization slots count towards the number of synchronization slots that are available.

```
USE EnterpriseVaultDirectory
SELECT ComputerEntryId, ComputerName, ComputerNameAlternate
FROM ComputerEntry
```

A restart of the storage service is required when changing this value.

Virtual Vault archiving

When archiving through Virtual Vault is enabled, end-users are able to upload data to the Enterprise Vault server to be archived. In general, large items take longer to upload and archive than smaller items.

It is generally recommended to use an appropriate archiving policy that keeps mailboxes down to size rather than rely on end users to do their own archiving. If however the administrator wants to roll out the ability to do manual archiving by dragging items into Virtual Vault, the following should be considered:

- How much additional end user training will be required to allow end users to perform manual archiving? Will end users regularly do manual archiving? What about when users are out of the office for extended periods of time, will their mailbox go over the size quota?
- Scheduled archiving performance is far greater than manual archiving. The best use of the Enterprise Vault server resources are through the use of scheduled archiving.
- If however you do want to enable minimal manual archiving, per example 10 items per user per day average, you can expect a further 5% increase in server resource use during the day (IIS, StorageOnlineOPNS, Index and SQL)
- Specifically monitor StorageOnlineOPNS as this is likely to become a bottleneck in over-use of manual archiving.

It is possible that an upload and archive operation may take longer than it should, perhaps due to a resourcing issue on the Enterprise Vault server, a network-related problem or some other issue such as large numbers of end-users simultaneously archiving data through Virtual Vault on the same storage server. A default timeout of 300 seconds for the upload and archive operation has been put in place to guard against the possibility of this operation continuing for too long a period of time. If, however, this timeout is too short a period of time – for instance because end-users are able to archive very large items through Virtual Vault – it is possible to increase this timeout value as desired.

Within the “...\Enterprise Vault\Webapp” folder, locate and open the web.config file. Locate the following line, and change the value specified (in seconds) to the desired value:

```
<add key="UploadItemExecutionTimeout" value="300"/>
```

Content cache configuration options

The majority of the configuration options for content cache are available through the vault administration console at the Enterprise Vault Servers container.

Figure 4 Server properties > Cache tab:

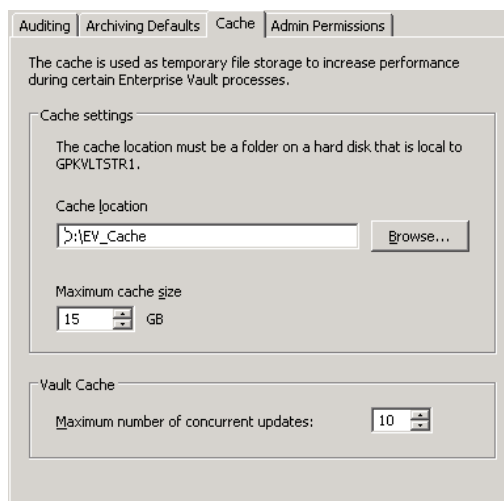


Figure 4 – Server Properties Cache tab

The 3 available settings here are:

Cache location – specifies the location that will be used for (among other things²⁰) the area where content cache data is stored temporarily after it has been built and during its transfer via BITS to the end user's workstation.

Maximum cache size – specifies the maximum size allocated for the cache.

In order to work out an appropriate maximum cache size value it is necessary to know how much data will be stored there, and for how long. Earlier in this document it was mentioned that once a slot is obtained a server-side process called MigratorServer.exe is spawned which retrieves data from the relevant vault store partition locations in order to create temporary content cache data in the cache location.

²⁰ The cache is currently used by PST Migration, FSA when the target file server is a Celerra file server, Vault Cache for both upload and download requests, and the Indexing service on Enterprise Vault 10.0.

Once an end-user has obtained a slot and a MigratorServer.exe process has been spawned, it will begin the retrieval of the most recent three months' worth of end-user data (for example January to March, April to June, etc) that needs to be downloaded. The amount of data involved can vary depending on how much has been archived for the end-user during this period. Each PST file generated by Migratorserver.exes has a maximum potential size of 500 MB. As an example, if 1.6 GB of data had been archived for an end-user over the three month period, MigratorServer.exe would create 3 x 500 MB PST files and 1 x 100 MB PST file. Once the quarter's data has been generated the slot is then released to service another end-user's request.

When sizing the cache location you therefore need to consider how much data has been archived per user over the course of the three month period, the number of concurrent updates specified (see next section), and then potentially double or triple the value depending on how fast the end-users are in having their content cache data downloaded to their workstations. Data is only downloaded when end-users are using Microsoft Outlook and are connected to the network. Therefore it may only be possible to download content during working hours only. With customers that have international sites, the size may need to be even larger given a possible twenty four / seven processing period. Backing up the environment may also require some downtime and should also be considered.

If content can be downloaded to end-user workstations faster, data can be removed from the cache location faster. As such, the cache size required could be reduced. One way to increase the throughput of the downloading of content cache is to increase the number of concurrent users able to download content cache data from the Enterprise Vault server. By default, this setting mirrors the value defined for maximum number of concurrent updates, however it is possible to change this behaviour. To modify the number of users able to download content cache data concurrently, create the following registry key:

Name: MaxCFSAllowed
Type: DWORD
Location: HKEY_LOCAL_MACHINE\Software\KVS\Enterprise Vault\AdminService
Default: 30 (minimum 1, maximum 99)

Maximum number of concurrent updates – specifies the number of concurrent content cache download slots available, default value 10, range 0 (minimum) / 50 (maximum).

The maximum number of concurrent updates value determines the number of end-users that at any one time can obtain a slot and request a MigratorServer.exe process to start building PST files ready for download to the client machine via BITS. MigratorServer.exe processes consume CPU, memory and disk I/O resources as it queries SQL for the relevant information, retrieves data back in memory, and writes that data out to disk. Therefore having a large number of MigratorServer.exe processes can significantly reduce available resources

on the Enterprise Vault server. Care also needs to be taken as to what value if used here as MigratorServer.exe processes are also used for PST migrations. A shared pool of fifty processes per storage server is used for both functions. Setting too high a value for maximum number of concurrent updates could therefore impact any PST migrations that are occurring at the same time.

If your server has sufficient resources available²¹ it is possible to increase the number of MigratorServer.exe processes (and therefore increase the maximum number of concurrent updates) allowed per storage server by using the following SQL query:

```
USE EnterpriseVaultDirectory  
UPDATE StorageServiceEntry  
SET MaxConcurrentMigrations = X  
WHERE ComputerEntryId = Y
```

X = the number of MigratorServer.exe processes you require.

Y = the GUID of the EV server that owns the storage service you wish to change. Use the following query to help identify the ComputerEntryId of the server you wish to change:

```
USE EnterpriseVaultDirectory  
SELECT ComputerEntryId, ComputerName, ComputerNameAlternate  
FROM ComputerEntry
```

A restart of the storage service is required when changing this value.

²¹ Monitor disk I/O, memory and CPU counters to determine if your server has sufficient resources available. Note, the appropriate value will vary between environments and hardware.

Exchange Server configuration options

Before data has been uploaded and archived, the act of archiving data through Virtual Vault (dragging and dropping data into Virtual Vault from an end-user mailbox) immediately removes data (a hard delete) from Microsoft Outlook (and therefore Microsoft Exchange), and inserts it into end-users' local header cache, where it resides unsecured prior to being uploaded and archived. The window of opportunity for data loss, whilst being minimal, is still very much a concern for customers. To offset this, customers can make use of native Microsoft Exchange functionality called DumpsterAlwaysOn. This allows end-users the ability to retrieve data that has been hard deleted from any folder in their mailbox. More information regarding this setting can be found here:

<http://support.microsoft.com/?kbid=246153>

SQL server configuration options

The SQL server is the heart of the Enterprise Vault system and as such requires careful management. The Enterprise Vault Performance Guide (referenced documents) makes some suggestions around appropriate hardware sizing for particular scenarios. Additionally, it recommends that Enterprise Vault databases should have a scheduled weekly plan that performs the following actions:

- Rebuild indexes
- Update statistics
- Shrink databases

These operations are vital to ensure operations such as incremental synchronizations can perform well.

As previously mentioned, incremental synchronizations will elevate disk I/O of the disk where Enterprise Vault databases are located. It is therefore important to monitor I/O response times. "Average Disk Seconds Per Transfer" for the same I/O path will help give an indication, along with looking at the SQL Server sysprocesses and waitstats for buffer based I/O wait information. These, coupled with "Physical Disk: Avg. Disk Queue Length" exceeding 2 for continuous periods of 10 minutes or more during a regular 24 hour period, and/or "Physical Disk: % Disk Time" exceeding 55% for continuous periods of 10 minutes or more during a regular 24 hour period, could indicate an I/O bottleneck. To avoid the potential of disk 'thrashing' it may be worth considering splitting your Enterprise Vault databases across 2 or more independent disk spindles to spread the disk I/O load.

Enterprise Vault also makes use of the SQL tempdb database, and as such it is recommended to follow Microsoft best practices.

“Tempdb supports only one data filegroup and one log filegroup. By default, the number of files is set to 1. Multiple files can be created for each filegroup. Adding more data files may help to solve potential performance problems that are due to I/O operations. Increasing the number of files helps to avoid a latch contention on allocation pages (manifested as a UP-latch). The recommended solution is to set the number of files to match the number of CPUs that are configured for the instance. This is only a recommendation—the number of files might not be the same as the number of CPUs.”

More information regarding best practice for Tempdb usage and recommended practices around optimisation and I/O monitoring can be found here:

<http://technet.microsoft.com/en-gb/library/cc966545.aspx#ECAA>

Rollout Strategy

There are several strategies that can be used to rollout Vault Cache and Virtual Vault functionality.

The first stage may very likely involve a documentation and training phase. No matter how seamless the integration with Microsoft Outlook, most companies will likely want to write an end-user guide or some training materials to explain to their end-users the new functionality they have, and how to use it. Administrators will need to tailor these guides and materials in relation to the policy settings they have configured.

The next phase may be a decision on who to enable first, second, and so on. Depending on the scale of the environment, administrators may find it easier to rollout based on region, business unit, provisioning group, and so on, or alternatively have a small implementation that could allow all end-users to be enabled all at once. Ideally, a "big bang" bulk enablement approach in larger deployments should be avoided and a more considered rollout should be adopted so that the impact on the environment can be more easily identified and understood. For end-users that already have Vault Cache functionality, enabling Virtual Vault will have no server-side impact; these end-users therefore are the easiest to enable.

The final phase should involve the actual rollout/enabling of the end-users, while monitoring the environment to ensure it is performing within acceptable limits.

Before you roll out Virtual Vault

Before you roll out Virtual Vault you need to know:

- How will Virtual Vault impact end user workflow? Are they used to small mailboxes and forced to drag items to PST, and VV will replace their PST?
- Or do they have large mailboxes and prefer to search?
- Are the Outlook clients in cache mode?
- How many other archives are users synching and how often?
- What is the Sync Policy going to be – age (up to x years, size (xGB)?
- What other operations occur on the EV server during the day? PST migrations (client or server side), journaling, public folder archiving – always good to build a 24hr schedule of what happens when

Server Performance Impact

When you are designing for Virtual Vault, expect a 10% utilisation increase during the day when Outlook is open and calling for slots. The 10% applies to IIS, StorageOnline, Index and SQL Databases.

- This 10% increase should not affect the number of mailboxes the server can support with Mailbox Archiving – provided you do your scheduled archiving during the night-time hours as per best practice.
- Expect a single, appropriately specified Enterprise Vault server to support up to and above 10,000 Virtual Vault users, provided the scheduled archiving and Virtual Vault sync times do not clash
- Size the Enterprise Vault Server Cache storage appropriately – in large environments it should be capable of supporting 1,000 IOPS.

Roll-out Strategy

Always ensure you use the latest Enterprise Vault client, where compatible with the server version. The Enterprise Vault slot manager will handle bulk enablement without overloading the server. It is however recommended to stage the roll-out as you would for any large scale deployment of archiving – Baseline the server performance (see Advanced Monitoring Strategies whitepaper, referenced documents) and start with a small group of pilot users.

Enterprise Vault 10.0.2 introduced a new feature known as “Vault Cache Diagnostics”, and this feature can assist administrators determine the status of Vault Cache builds in the wild, thereby confirming the status of a group of users rolled out before progressing with the next group of users. Note that the minimum client version compatible with this feature is 9.0.4 and 10.0.2, earlier Outlook Add-in will not report the diagnostic information to the EV server.

The following best practice recommendations should be followed where possible:

- Initially enable all users for contentless cache (“Do not store any items in cache” / “Store only items that user opens”). This allows all users to use Virtual Vault as soon as possible.
- Enable with-content cache (“Store All Items”) for only those users who do not always have access to their online archives. However, for better performance, the initial synchronization should be done when the users have access to their archive over a fast network.
- Limit the size or time range of items in the with-content cache. Users typically require access to the most recent items rather than their entire archives.
- If manual archiving is enabled, enable Microsoft Exchange functionality **DumpsterAlwaysOn**.

Rollout scenarios

A Virtual Vault rollout encompasses a whole host of considerations ranging from specific end-user functionality policies, network bandwidth utilization, server disk space/performance, and other Enterprise Vault-related processes such as archiving, indexing and PST Migrations.

Some examples of questions to consider when determining the best fit scenario are:

- What problems do I have today that Virtual Vault will help solve?
- Who should I roll Virtual Vault out to first (for example pilot users, IT, etc)?
- How do I roll Virtual Vault out to the rest of the end-user population?
- What operations do you currently expose to end-users now and do you want the same operations available with Virtual Vault (for example manual archive, restore)?
- Do I need different policies to cater for different types of user (for example desktop/office bound, laptop, mobile, OWA, thin clients)?
- Are there any known or potential bottlenecks (network related, server) that may affect optimum performance for end-users and/or Enterprise Vault services?
- Are there any other operations occurring during the working day that I need to take into account?

The way in which Vault Cache and Virtual Vault functionality is configured will depend on what best suits your company’s needs and therefore this configuration will vary as requirements for one organisation will be different to another.

The following scenarios in this final section will detail common business requirements and describe the associated best-fit policy settings to achieve those requirements²².

Scenario 1

Company A has recently purchased and installed Enterprise Vault. They have rolled out the Enterprise Vault Outlook Add-in to all 500 end-users in preparation for using scheduled mailbox archiving, although no archiving has been done at this stage.

The high level solution is as follows:

- Laptop and desktop users both utilize Microsoft Outlook cached mode as part of company policy to ensure access to email even when Microsoft Exchange services are unavailable.
- Laptop and desktop users both need a local copy of all their archived data to facilitate offline access whilst avoiding unnecessary network traffic.
- Access to OWA is not available.
- Scheduled archiving is to be used as the only archiving mechanism – manual archiving will not be utilized.
- Virtual vault is to be made available to all end-users so as to streamline end user training.

This is potentially the simplest of scenarios to create a policy for. Having the same end-user experience for laptop and desktop users simplifies the number of policies and the defining of the appropriate provisioning groups. If OWA is not available for end-users then there is a better case for expiring shortcuts from mailboxes sooner. Potentially, a shortcut-less policy implementation could be considered.

The following configuration options should be considered:

- Maximum concurrent updates – can either leave this value at the default value (10) or set a little higher (20) as content cache downloading should not be excessive for two main reasons:

²² Note, these scenarios are sample scenarios that provide suggestions on how an administrator could configure policy settings to provide the best experience for their end-users. The size/number of seats given as an example in each scenario does not suggest that that scenario is the best solution for the specified number of seats.

- Content cache downloads are at a steady state (no backlog to process) due to no archived data already existing.
- Most archived data will be automatically inserted into the content cache locally by the Enterprise Vault Outlook add-in prior to it being archived (via the trawler mechanism) so little or no content cache downloading should need to occur. (Note: inactivity period and inactivity period units²³ should be configured so that items are pre-cached by the Enterprise Vault Outlook Add-in prior to being archived.)
- Users can archive items should be set to no.
- Content strategy should be set to store all items.
- Store in vault²⁴ option should also be disabled.
- Shortcut deletion can be set low (for example 7 days) as end-users have been trained to access their archived data through Virtual Vault and do not have a (legacy) reliance on using shortcuts, archive explorer or search for access.
- Administrators may also wish to consider what toolbar buttons (if any) in general may be of use for their end-users.

Scenario 2

Company B has been a long-standing customer of Enterprise Vault, with many of their 1000 end-users having mature multi-gigabyte archives. Prior to their latest upgrade they were running with Enterprise Vault 2007 and had configured offline vault²⁵ for all end-users.

²³ Refer to the Administrators Guide for more information.

²⁴ Exchange Desktop policy option

²⁵ Offline vault was the first implementation of offline archive access, replaced in Enterprise Vault 8 with Vault Cache, a redesigned implementation of offline archive access providing better scalability and performance.

The high level solution is as follows:

- Laptop and desktop users both utilize Microsoft Outlook cached mode as part of company policy to ensure access to email even when Microsoft Exchange services are unavailable.
- Laptop and desktop users both utilize offline vault to provide a local copy of their archived data and will want the same experience after upgrade.
- OWA is available both internally and externally for desktop users only.
- Separate mailbox and desktop policies already exist for laptop and desktop users.
- All end-users are able to manually archive and restore.
- All end-users are to be able to move items and folders within their archive.
- Shortcut deletion is enabled and configured to remove shortcuts sixty days after items have been archived.

Some thought needs to be given for OWA access as Virtual Vault is only available through Microsoft Outlook; the only access to Enterprise Vault data would be through shortcuts in the mailbox, archive explorer and search²⁶.

The following configuration options should be considered:

- Maximum concurrent updates – can either leave this value at the default value (10) or set a little higher (20) as content cache downloading should not be excessive for two main reasons:
- Content cache downloads are at a steady state (no backlog to process) due to offline vault already existing on end-user workstations and being fully populated/up to date²⁷.
- Newly archived data will be automatically inserted into the content cache locally by the Enterprise Vault Outlook Add-in prior to it being archived (via the trawler mechanism) so little or no content cache downloading should need to occur. (Note: inactivity period and inactivity period units should be configured so that items are pre-cached by the Enterprise Vault Outlook Add-in prior to being archived.)
- Users can archive items should be set to yes.

²⁶ Assuming the Enterprise Vault OWA Extensions have been installed on the Exchange servers.

²⁷ Upgrading from offline vault to Vault Cache does not require a re-download of locally cached data.

- Users can copy items to another store should be set to yes²⁸.
- Store in vault and restore from vault should also be enabled.
- Users can reorganize items should be set to yes.
- The current shortcut deletion setting does not necessarily need changing. However the need for shortcuts in a mailbox is lessened with Virtual Vault functionality as archived data is more easily accessible. The new value does however need some consideration. Shortcut deletion for the desktop users should be configured to a reasonable value due to their dependency on shortcuts for (easy) access to archived data when using OWA (externally). This value could differ for laptop users who do not use OWA, which would therefore require an additional mailbox policy to cater for the differing values. A low shortcut deletion value (for example 7 days) or even a shortcut-less policy for laptop users could be considered, whilst a higher (for example 30 days, 60 days, etc) shortcut deletion value will still be required for desktop users that have a reliance on using shortcuts, Archive Explorer or Search for access through OWA.

Scenario 3

Company C has been a long-standing customer of Enterprise Vault, with many of their 5000 end-users having mature multi-gigabyte archives. Prior to their upgrade they had configured Vault Cache for their 1000 laptop users but not for the remaining desktop users.

The high level solution is as follows:

- Laptop users generally utilize Microsoft Outlook cached mode to ensure access to email when not in the office.
- Desktop users do not use Microsoft Outlook cached mode.
- Separate desktop policies already exist for the laptop and desktop users.
- All end-users are able to manually archive and restore.

With desktop users not using Microsoft Outlook cached mode, some thought needs to be given to the *content strategy* utilized and any other settings related to this.

The following configuration options should be considered:

²⁸ Review other options associated with “Users can copy items to another store”

- Maximum concurrent updates – can either leave this value at the default value (10) or set a little higher (20) as content cache downloading should not be excessive for two main reasons:
- Content cache downloads are at a steady state (no backlog to process) due to offline vault already existing on end-user laptops and being fully populated/up to date.
- For laptop users, newly archived data will be automatically inserted into the content cache locally by the Enterprise Vault Outlook Add-in prior to it being archived (via the trawler mechanism) so little or no content cache downloading should need to occur. (Note: inactivity period and inactivity period units should be configured so that items are pre-cached by the Enterprise Vault Outlook Add-in prior to being archived.)
- Offline store required should be set to no otherwise desktop users would be unable to enable Vault Cache.
- Content strategy should be configured differently for laptop and desktop users due the former using Microsoft Outlook cached mode, whilst the latter do not. For the former you should use a value of Store all items, and for the latter consider a value of do not store or store only items that users open.
- Show content in the reading pane will potentially be different for laptop and desktop users based on the content strategy used.
- Users can archive items should be set to yes.
- Users can copy items to another store should be set to yes.
- Store in vault and restore from vault should also be enabled.

Scenario 4

Company D has been a long-standing customer of Enterprise Vault, with many of their 35,000 end-users having mature multi-gigabyte archives. Prior to their upgrade they had configured Vault Cache for their 1000 laptop users but not for the remaining desktop users. There is also a significant percentage of end-users who access network resources from home via server farms running Microsoft Terminal Services and Citrix server.

Citrix environments are supported if the following conditions are met, as outlined in the Enterprise Vault Compatibility Chart (referenced documents):

- Vault Cache and Virtual Vault must be configured in header-only mode (no content caching)
- The network connecting the Citrix server and file server hosting the Vault Cache and Virtual Vault files must offer high speed and low latency
- Outlook 2010 or later is required
- Servers must be Windows 2008 R2 or later

- Enterprise Vault Outlook Add-In 9.0.5 or later service pack, or 10.0.4 or later service pack, or 11.0 is required.

Summary

In summary, this document has described in detail the various policy and configuration settings that an administrator can change when implementing Vault Cache and Virtual Vault for an environment, whilst having an awareness of the implications making these changes, both from an end-user perspective and from a server environment perspective.

Appendix A - Virtual Vault advanced settings

Table 1

Advanced setting	Description	Impact
Max archive requests per synchronization	Controls the maximum number of archive requests during a Vault Cache synchronization.	<p>When a user stores unarchived items in Virtual Vault, the archive operation does not take place until after the next Vault Cache header synchronization.</p> <p>No limit or a high value can increase the time that is required to complete a Vault Cache synchronization. This effect is a consideration if the additional load affects the Enterprise Vault server.</p> <p>Also, until the items that a user has stored in Virtual Vault are archived in the online archive, there are no backup items.</p>
Max attempts to archive an item	Specifies how many times Enterprise Vault tries to archive an item.	<p>Specifies how many times Enterprise Vault tries to archive an item.</p> <p>The archive operation is tried this number of times before the item is listed in the Virtual Vault search folder named 'Could Not Archive'.</p>
Max data archived per synchronization	Controls the maximum amount of data in megabytes that can be uploaded during a Vault Cache synchronization.	<p>Controls the maximum amount of data in megabytes that can be uploaded during a Vault Cache synchronization. Any remaining data is uploaded at the next synchronization.</p> <p>No limit or a high value can increase the time that is required to complete a Vault Cache synchronization. This effect is a consideration if the additional load affects the Enterprise Vault server.</p> <p>Also, until the items that the user stores in Virtual Vault have been archived in the online archive, there are no backup items.</p>

Advanced setting	Description	Impact
		The value of this setting must be greater than or equal to the value of Max item size to archive . If not, the value of Max item size to archive is used.
Max delete requests per synchronization	Controls the maximum number of delete requests during a Vault Cache synchronization. Any remaining requests are made at the next synchronization.	Controls the maximum number of delete requests during a Vault Cache synchronization. Any remaining requests are made at the next synchronization. Deletion requests use relatively few resources on the Enterprise Vault server.
Max item size to archive	Controls the maximum size in megabytes of an item that can be moved or copied into Virtual Vault.	Controls the maximum size in megabytes of an item that can be moved or copied into Virtual Vault. If this value is similar to the value of Max total size of items to archive , a full synchronization can consist of one item. The Max item size to archive value may be used automatically for Max data archived per synchronization or Max total size of items to archive . It is used if the value of those settings is less than the Max item size to archive value.
Max item updates per synchronization	Controls the maximum number of property change requests during a Vault Cache synchronization. Any remaining requests are made at the next synchronization.	Controls the maximum number of property change requests during a Vault Cache synchronization. Any remaining requests are made at the next synchronization.
Max total size of content-less operations	Controls the maximum total size in megabytes of copy and move operations when items have no content in Vault Cache.	Controls the maximum total size in megabytes of copy and move operations when items have no content in Vault Cache. This setting only applies when two or more standard Microsoft

Advanced setting	Description	Impact
		Outlook items ²⁹ with no content are involved in the operation. Retrieval of one item is allowed regardless of its size.
Max total size of items to archive	Controls the maximum total size in megabytes of pending archive data in Vault Cache.	Controls the maximum total size in megabytes of pending archive data in Vault Cache. Pending archive data consists of items that the user has moved or copied into Virtual Vault. These items are pending archive until Vault Cache synchronization has successfully uploaded and archived them. The value of this setting must be greater than or equal to the value of Max item size to archive . If not, the value of Max item size to archive is used.
Show content in reading pane	Controls whether content is shown in the Microsoft Outlook reading pane.	This setting is useful if the Vault Cache content strategy is Do not store any items in cache or Store only items that user opens . The setting can prevent Enterprise Vault from downloading the content of every item that the user selects in Virtual Vault. Instead, a banner provides a link to open the item. If the item itself is a document, it is not displayed in the reading pane. A message in the reading pane advises the user to open the item to read the item's contents.
Threshold number of items to trigger synchronization	Specifies the total number of pending archive items in Virtual Vault that triggers automatic Vault Cache synchronization.	Pending archive data consists of items that the user has moved or copied into Virtual Vault. These items are pending archive until Vault

²⁹ For example mail items, calendar items, tasks and contacts.

Advanced setting	Description	Impact
		Cache synchronization has successfully uploaded and archived them.
Threshold total size of items to trigger synchronization	Specifies the total size in megabytes of pending archive items in Virtual Vault that triggers automatic Vault Cache synchronization.	<p>Pending archive data consists of items that the user has moved or copied into Virtual Vault. These items are pending archive until Vault</p> <p>Cache synchronization has successfully uploaded and archived them.</p>
Users can archive items	Controls whether users can archive items manually using Virtual Vault.	<p>Controls whether users can archive items manually by adding new items to Virtual Vault using standard Outlook actions. Examples of these standard Outlook actions are drag and drop, move and copy, and Rules.</p> <p>If you disable this setting, users can still create folders if Users can reorganize items is enabled.</p>
Users can copy items to another store	Controls whether users can copy and move items from a Virtual Vault to another message store.	<p>Controls whether users can copy and move items from a Virtual Vault to another message store. If users can copy or move items out of Virtual Vault and the content is available in Vault Cache, the items are retrieved from Vault Cache.</p> <p>If the Vault Cache content strategy is Do not store any items in cache, the items are retrieved from the online archive. In this case, use the Virtual Vault advanced setting Max total size of content-less operations to control the maximum total size of view, copy, and move operations.</p>
Users can copy items within their archive	Controls whether users can copy items within their archive.	If users can copy items within their archive and the content is available in Vault Cache, the items are retrieved from Vault Cache.

Advanced setting	Description	Impact
		<p>If the Vault Cache content strategy is Do not store any items in cache, the items are retrieved from the online archive. In this case, use the Virtual Vault advanced setting Max total size of content-less operations to control the maximum total size of view, copy, and move operations.</p>
Users can hard delete items	Controls whether users can hard delete items from Virtual Vault.	<p>For this setting to take effect, the option Users can delete items from their archives must be enabled on the Archive Settings tab in the Site Properties dialog box.</p> <p>If you disable this setting, users can still move items to the Deleted Items folder if Users can reorganize items is enabled.</p>
Users can reorganize items	Controls whether users can reorganize items in Virtual Vault.	<p>This setting can enable users to move items between folders and to create, move, rename, or delete folders.</p> <p>If folders still exist in the mailbox, users cannot move, rename, or delete them.</p> <p>Users can hard delete only empty folders, unless Users can hard delete items is enabled.</p>

Appendix B - Vault Cache policy settings

Table 2

Setting	Description	Impact
Make Virtual Vault available to users	Controls whether users are enabled for Virtual Vault	<p>This setting enables Virtual Vault for users that have Vault Cache.</p> <p>For existing Vault Cache users it may be worthwhile informing them of this new functionality prior to it becoming available for use.</p>
Content Strategy	Controls the content strategy for Vault Cache	<p>If all items are stored locally in users' content cache then any request to view/restore items will all be performed using their content cache. If the setting has been changed to "Do not store" or "Store only items that users open" then a view/restore request will initiate a server-side request for the applicable content (unless that content has previously been cached due to an earlier request). This could have an adverse affect on the storage server facilitating these requests if it is unable to cope with the additional loading.</p>

Appendix C - Vault Cache advanced policy settings

Table 3

Advanced setting	Description	Impact
Offline store required	Controls whether Vault Cache can be enabled when no offline store is present.	<p>Users have offline store (OST) files if Outlook Cached Exchange Mode is enabled. If a user does not have an OST file, Enterprise Vault cannot perform preemptive caching.</p> <p>If there is no preemptive caching, there is an increased load on Vault Cache content synchronization for newly archived items. The increased load is only a consideration if the Vault Cache content strategy is Store all items.</p>
Synchronize archive types	Controls what is synchronized by Vault Cache.	<p>If this is policy is left at the default value of default mailbox this will reduce the number of required permissions checks when end-users synchronize. If there are specific end-users who require multiple archives to be synchronized offline a separate desktop policy should be created for these end-users.</p>

Appendix D - Vault Cache client registry settings

Table 4

Advanced setting	Description	Impact
<p>OVMDCSyncFrequencyInSecs</p> <p>(Refer to the Registry Values documentation for details on the location to configure this registry key)</p>	<p>Number of seconds between scheduled synchronization.</p>	<p>Each time a user synchronises there will be additional load placed on IIS to obtain a slot, SQL, to work out where the slot should be given, what permissions the user has, what items need to be uploaded, updated, downloaded, deleted, any folder hierarchy changes, retention category changes, and Enterprise Vault storage server, that has to interact with both IIS and SQL, whilst also servicing any requests for content resulting from this, and any other regular operations outside of this feature such as journal archiving, user retrieval/view requests, etc.</p>
<p>OAllowMissedMDCSyncOnStart up</p> <p>(Refer to the Registry Values documentation for details on the location to configure this registry key)</p>	<p>Enables or disables the behavior of scheduled synchronization when the last synchronization is missed.</p>	<p>When a user starts Outlook, a Vault Cache synchronization is initiated if a scheduled synchronization has been missed. Synchronization slot and SQL Server contention issues can occur if a large number of users, who have missed a scheduled synchronization, all start Outlook within a short timeframe.</p> <p>OAllowMissedMDCSyncOnStart up allows you to configure clients so that missed scheduled Vault Cache synchronizations are ignored.</p> <p>A Vault Cache synchronization occurs at the next scheduled synchronization time .</p> <p>If the default value (0) is specified, or the setting does not exist, then a Vault Cache synchronization is initiated when a user starts Outlook, if a scheduled synchronization has been missed.</p>

IIS log diagnostics

Statistic and diagnostic information regarding end-users' Vault Caches and Virtual Vaults can be obtained from Enterprise Vault server IIS logs. A string of information is appended to IIS requests made to /EnterpriseVault/slot.aspx after each synchronization has finished. Below is an example of what the entry looks like:

2009-11-11 12:26:26 W3SVC1 10.12.28.10 GET /EnterpriseVault/Slot.aspx

ArchiveID=1CF3100D14F76FE41BC944F1F0367B74D1110000ETG25EV.example.com&Slot=9496e3c6-bc60-428a-80c8-c8ff1f5e342d&release=true&MSt=0&CSt=3&Ls=2009-07-06T12:26:24&Ci=1&Td=0 80

USER\TestUser30 192.168.0.1 EnterpriseVaultOutlookExt-V8.0.3.0 200 0 0

The diagnostic information is logged in the following format:

MSt=0&CSt=3&Ls=2009-07-06T12:26:24&Ci=1&Td=0

Table 5 details the different sections and their possible states/corresponding values.

Table 5

Abbreviation	State	Value
MSt: header cache status	mdcSyncStatusFailedEmptyArchive	-9
	mdcSyncStatusFailedServerSyncing	-8
	mdcSyncStatusPrevented	-7
	mdcSyncStatusFailedInsufficientDisk	-6
	mdcSyncStatusFailedSlot	-5
	mdcSyncStatusFailedCouldNotConnectToE VWebServerVirtDir	-4
	mdcSyncStatusFailedOffline	-3
	mdcSyncStatusFailedDirty	-2
	mdcSyncStatusFailed	-1
	mdcSyncStatusSuccess	0
	mdcSyncStatusNone	10

Abbreviation	State	Value
	mdcSyncStatusPending	19
	mdcSyncStatusInProgress	20
	mdcSyncStatusInProgressAcquireSlot	30
	mdcSyncStatusInProgressAcquireSlotWaiting	31
	mdcSyncStatusReset	2147483645
	mdcSyncStatusSuspended	2147483646
	mdcSyncStatusShutdown	2147483647
CSt: content cache status	CC_INITIALIZING	0
	CC_NEVER_DOWNLOADED	CC_INITIALIZING + 1
	CC_NODOWNLOAD	CC_NEVER_DOWNLOADED + 1
	CC_COMPLETE	CC_NODOWNLOAD + 1
	CC_DBBUILDING	CC_COMPLETE + 1
	CC_DBDOWNLOADING	CC_DBBUILDING + 1
	CC_RETRY_WEBSERVER	100
	CC_BITS_NOT_AVAILABLE	200
	CC_FAILED_FIRST_DBBUILD	300
	CC_FAILED_NEXT_DBBUILD	CC_FAILED_FIRST_DBBUILD + 1
	CC_FAILED_BITS_DOWNLOAD	CC_FAILED_NEXT_DBBUILD + 1
	CC_FAILED_PROCESS_NEXT_ARCHIVE	CC_FAILED_BITS_DOWNLOAD + 1
	CC_FAILED_ADDING_ARCHIVES	CC_FAILED_PROCESS_NEXT_ARCHIVE + 1
	CC_FAILED_WEBSERVER	CC_FAILED_ADDING_ARCHIVES + 1
	CC_FAILED_GENERIC	400
Ls: last sync time		YYYY-MM-DDTHH:MM:SS
CI: number of items in the content cache		Integer

Abbreviation	State	Value
Td: total number of items to download		Integer

Appendix E - performance counters to monitor

Below are the list of counters that should be monitored when performing either baseline testing of your Vault Cache and Virtual Vault settings, or when making a change to one of the settings.

Operating System counters

.NET CLR Interop
.NET CLR Jit
.NET CLR LocksAndThreads
.NET CLR Memory
.NET CLR Networking
.NET CLR Remoting
ASP.NET Apps v2.0.50727
ASP.NET v2.0.50727
Cache
LogicalDisk
Memory
MSMQ Queue
Network Interface
Paging File
PhysicalDisk
Process
Processor
Server
System
Thread

SQL server counters

SQLServer:Access Methods
SQLServer:Buffer Manager
SQLServer:Buffer Partition

SQLServer:Cache Manager
SQLServer:CLR
SQLServer:Cursor Manager by Type
SQLServer:Cursor Manager Total
SQLServer:Databases
SQLServer:General Statistics
SQLServer:Latches
SQLServer:Locks
SQLServer:Memory Manager
SQLServer:Plan Cache
SQLServer:SQL Statistics
SQLServer:Transactions
SQLServer:Wait Statistic

About Symantec

Symantec is a global leader in providing security, storage, and systems management solutions to help consumers and organizations secure and manage their information-driven world. Our software and services protect against more risks at more points, more completely and efficiently, enabling confidence wherever information is used or stored. Headquartered in Mountain View, 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 website.

Symantec World Headquarters
350 Ellis St.
Mountain View, CA 94043 USA
+1 (650) 527 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2014 Symantec Corporation.
All rights reserved. Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.