

Veritas Access Appliance with Enterprise Vault

Archival Solution

Contents

Revision History	3
Introduction	4
Executive Summary	4
Scope and Target Audience	4
Solution Value	4
Solution Key Features.	5
WORM (Write Once, Read Many)	5
Storage Efficiency	6
Replication	6
Seamless Integration with NetBackup	7
Encryption	7
Shadow Copy with Versioning	7
Tight Integration with Data Insight for Better Visibility	7
Appliance AutoSupport and NetInsights Console	9
Monitoring and Detection	10
Solution Architecture	10
Enterprise Vault (EV).	10
Access Appliance	12
Solution Data Flow	14
Archival Data Flow	15
Retrieval Data Flow	15
Enterprise Vault Migrator Data Flow	16
Best Practices and Recommendations.	17
File System and Data Layout on Access Appliance.	17
Network Connectivity	17
Monitoring	18
Load Balancing.	18
NTP Server	18
Replication	18
Sizing Guidance	19
Conclusion	17
References.	18

Revision History

Version	Date	Notes
1.0	March 2018	Initial version
2.0	October 2018	Updated for Access 7.4.2 release
2.1	April 2018	Modified information on Enterprise Vault encryption
3.0	May 2021	Updated for Access 7.4.3 release
4.0	July 2022	Updated for Access 8.0 release

Introduction

Executive Summary

The exponential growth of data storage requirements for organizations leads to an increasing focus on information lifecycle management and archival strategies. Organizations seek a solution that can control ongoing costs, and meet business compliance and regulatory requirements.

Enterprise Vault, used with Access Appliance, provides a complete solution to address these challenges. Storage provided by Access Appliance is dense, highly resilient, and tightly integrated with data lifecycle management software. Enterprise Vault uses this storage to great effect, not simply for archival purposes, but also in support of data management scenarios that involve more advanced management such as data immutability or replication.

Enterprise Vault manages the archiving and retrieval of information from 80+ data sources, including native support for Microsoft® Exchange, IBM Domino®, Skype for Business, Microsoft SharePoint, and file systems, while Access Appliance is the solution's dense, on-premises storage target for data. When they are used together with other products from Veritas such as NetBackup and Data Insight, they improve the visibility of an organization's data and allow for more fine-grained data management. This reduces cost, waste, and risk, and overall makes a compelling solution for an archival use case.

Scope and Target Audience

This white paper will describe the business value and key features of an archival solution consisting of Enterprise Vault and Access Appliance. Additionally, it will provide sizing guidance and discuss solution best practices. This document is targeted for customers, partners, and Veritas field personnel interested in learning more about Access Appliance and Enterprise Vault as an archival solution.

NOTE: This document is periodically updated. The latest version can be retrieved from Veritas at the location noted in the [References](#) section.

Solution Value

Archival of data is the process of storing data that has not been referenced for a long period of time in such a way to save space or resources and still be easily accessible when it is needed. Some reasons for archiving data include controlling cost, freeing up space for incoming data, improving security, complying with legal and regulatory requirements, and classifying content for search and discovery. Archived data is best stored in a centralized storage media on-premises or in the public cloud instead of on individual laptops, desktops, or disparate storage, so it is secure, manageable and easy to locate. Data growth, management, retention, visibility, and cost are some challenges when selecting the appropriate storage platform for archived data. Archival storage on Access Appliance addresses these challenges, and provides the following key values as a target storage platform for Enterprise Vault:

- **Cost minimization:** Access Appliance provides a low-cost, disk-based solution that is easy to manage. Together with Enterprise Vault's storage-supported features such as single-instance storage and compression, Access Appliance can reduce your archival solution's overall costs while also increasing storage efficiency.
- **Management simplification:** The maintenance and management of varying secondary storage types, media, and protocols are challenges to an organization's information technology administration. A tightly integrated, single-vendor solution can improve issue resolution and simplify acquisition and long-term management.
- **Increased visibility and control:** Enterprise Vault's data insight and characterization can help with plans for future storage needs, as well as identify data that is unused, orphaned, or no longer required for regulatory or compliance purposes. Its seamless integration with Veritas Data Insight allows you to reduce resource inefficiencies, storage waste, and overall cost.

Solution Key Features

There are certain key features that companies look for in an archival solution. Primary features often include compliance, flexibility, storage efficiency, and ease of management. An archival solution using Access Appliance with Enterprise Vault provides these features, along with others, to assist in preserving an organization's valuable data.

WORM (Write Once, Read Many)

An archival storage platform's WORM (or data immutability) features are important not only for regulatory and compliance reasons, but also as a bulwark against malicious modification or deletion of data by ransomware. When using the WORM feature of Access storage volumes, the specified data cannot be modified or deleted before the expiration of the retention period, satisfying data retention rules and enhancing the security of archival data at rest.

Access Appliance software has previously included a WORM setting for its data storage filesystems, and with its 8.0 software release, gains the ability to run Access Appliance in lockdown mode. An appliance running in lockdown mode runs a compliance clock to track retention that is independent of an operating system, and will not permit any operations that lead to the destruction of immutable data, whether the operation is initiated by a user, administrator, or another solution component.

When archiving data on Access Appliance using Enterprise Vault, the WORM feature must first be enabled on the target filesystem, and then the retention period is set at file level by Enterprise Vault. Files that are WORM enabled are protected from alteration or removal by any user, including root and administrative users.

NOTE: Enterprise Vault uses a Common Internet File System (CIFS) share to store data on an Access Appliance with WORM features enabled. For more information on lockdown mode operation, see the [Immutability in Access Appliance](#) section in the [Veritas Access Appliance Administrator's Guide](#).

Storage Efficiency

Storage efficiency features are often a primary factor in decision-making when selecting a platform, as increasing storage efficiency often serves to also reduce overall cost when less storage is required. Using Access Appliance as storage allows Enterprise Vault to leverage its compression and SIS (single-instance storage) features. Data compression is a mechanism that reduces the size of the file by encoding the data using fewer bits. The Enterprise Vault SIS feature allows for a single instance of a file or data storage across multiple source contents (such as email, file systems, and SharePoint). For instance, an email attachment can be sent to numerous recipients, but Enterprise Vault only maintains or archives one instance of the file in the target storage and subsequent copies are a reference to that single file.

Replication

In-depth protection of an organization's archival data often means creating additional, geographically-separate copies of that data for disaster recovery and business continuity reasons. To meet such a requirement, Access Appliance can create and maintain a data replication relationship between appliances using either scheduled or continuous replication.

An archival solution using Enterprise Vault can utilize Access Appliance's episodic replication functionality for scheduled replication. The sending appliance may also use a partition secure notification to inform Enterprise Vault that the remote copy has been successfully created, triggering the removal of unneeded local safety copies, as shown in Figure 1.




Figure 1. Episodic replication between Access Appliance instances

If removal of the safety copies is not required, Access Appliance's continuous replication facility can be used to replicate the archive. See [Replication](#) in the Best Practices section of this document.

For more information on using episodic replication, including the partition secure notification functionality, see the About Access Appliance Episodic Replication section in the Veritas Access Appliance Administrator's Guide.

Seamless Integration with NetBackup

Access Appliance can not only receive archival data from Enterprise Vault for long-term storage, but it can also transfer data to a NetBackup domain using its built-in client functionality, and subject it to NetBackup's storage lifecycle and data management policies.

For more information on Access Appliance's container-based NetBackup client functionality, see the [Configuring Access Appliance with the NetBackup client](#) section in the [Veritas Access Appliance Solutions Guide for NetBackup](#).

Encryption

Any data encryption done by an application and archived using Enterprise Vault is maintained while stored on Access Appliance, which also includes encryption capabilities in conjunction with an external Key Management System (KMS). The appliance encrypts the archival data volume and uses the KMS for key lookup and distribution.

For more information on using encryption on Access Appliance, see the [About encryption at rest](#) section in the [Veritas Access Appliance Administrator's Guide](#).

Shadow Copy with Versioning

You can use the Microsoft Windows® Server Volume Shadow Copy Service with Enterprise Vault. A shadow copy is essentially a point-in-time replica of a filesystem's contents visible to Windows file management tools. Shadow copy integration permits Windows to browse the replica data and quickly recover files in case of corruption, accidental deletion, or being overwritten. As each shadow copy is versioned, multiple shadow copies for a filesystem enables an entry point into the filesystem at different points in time.

Awareness and storage of shadow copies created by the Volume Shadow Copy Service is enabled on Access Appliance at the share level when exporting filesystems for use by CIFS/SMB.

For more information on Access Appliance and volume shadow copy service, see the Making a CIFS share shadow copy aware section in the [Veritas Access Appliance Administrator's Guide](#).

Tight Integration with Data Insight for Better Visibility

Veritas Technologies offers Data Insight to provide better visibility into an organization's data sources. Data Insight can use its classification engine to locate and classify data by searching across data sources and within individual files. This integrates with Enterprise Vault functionality to allow you to target data for archival according to a specific classification.

After Veritas Data Insight scans and analyzes unstructured data sources such as filers, SharePoint web applications, Documentum repositories, and cloud storage accounts, it classifies the data into certain categories such as ownership, age, size, activity, data access patterns, user risk, type, etc. As well as using Enterprise Vault to archive data, you can also enforce security or compliance constraints, conduct data chargeback, and perform information lifecycle management and risk analysis.




Figure 2. Data insight classification of archival data

As pictured in Figure 2, data sources are scanned, analyzed, and classified by Data Insight. The resulting scan data can then be manually inspected, or policies can be defined to determine the data that will be archived using Enterprise Vault and stored using Access Appliance. This level of in-depth knowledge of a source's contents allows organizations to make more informed decisions on what to do with the data for storage optimization, security, and archival.

For a sample view of the Data Insight Console displaying a list of group shares and folders with their sizes and number of inactive files being monitored, see Figure 3. The inactive folders or files can either be selected for deletion or archival using Enterprise Vault. For more information, refer to the [Data Insight Product Documentation](#).




Figure 3. Data Insight Console group share view

Appliance AutoSupport and NetInsights Console

Veritas appliances, including Access Appliance, can send support and telemetry data directly to Veritas for analysis and reporting using their AutoSupport functionality. The support call-home facility allows Veritas to automate the opening of support cases and dispatch of replacement parts when required. Additionally, the flow of telemetry and support data is used for the NetInsights appliance health monitoring and analytics platform available to Veritas appliances. NetInsights System Health can use this data to generate actionable steps to improve appliance reliability and further reduce risks to archival data, even communicating findings in a System Reliability Score summary. For an example summary output, see Figure 4.

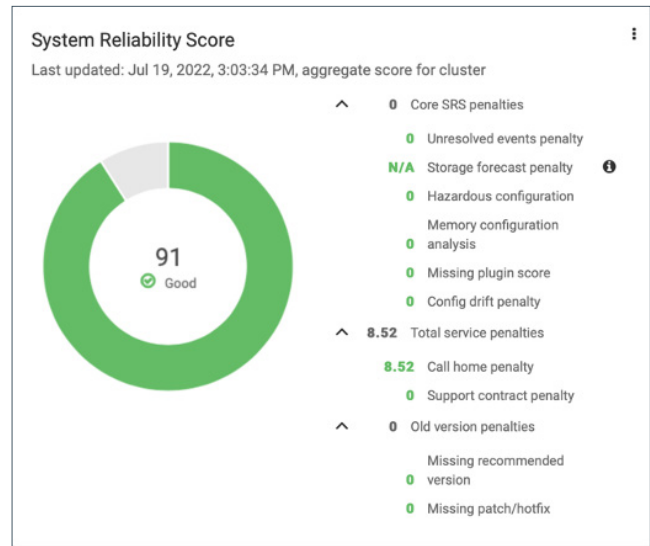


Figure 4. System Reliability Score summary

This summary can be inspected to receive an itemized and detailed list of steps that will improve system reliability and supportability. For an example of a list of recommended actions, see Figure 5.

Overall, a solution consisting of a single vendor provides end-to-end support for quicker resolution and response as opposed to having to contact multiple vendors to handle issues related to varying products and/or hardware implemented in the solution.

For more information on AutoSupport functionality, see the [Veritas Appliance Autosupport Reference Guide](#).

Monitoring and Detection

Also available on Access Appliance is Symantec Data Center Security (SDCS), an intrusion detection system. SDSCS is real-time monitoring and auditing software that offers host intrusion detection, file integrity monitoring, configuration monitoring, user access tracking and monitoring, and logging and event reports. SDSCS adds security hardening and monitoring for Access Appliance to reduce security risks and attacks.

Recommended Actions		Node Specific SRS score			
<input type="checkbox"/> Show acknowledged recommendation in the table ⓘ					
Total	Critical	Risk	Warning	Information	
2	1	0	1	0	
Export Search for UID, Category, Component, Problem Description 🔍					
Created At	UID	Severity	Category/Component	Problem Description	Recommendation
Mar 14, 2022, 9:53:25 AM	VTAS0013600	Warning	Resources Storage	Audit storage pool usage is above warning level (80%).	Please review the storage pool contents and free up space.
Apr 29, 2022, 3:16:41 AM	VTAS0013600	Critical	AutoSupport Client Transmission Service	Appliance has not called home in more than 336 hours, which is 14 day(s). Last Called Home Date and Time : 2022-03-14T15:45:34.104000	Please investigate appliance connectivity to Veritas Call Home endpoints by following the instruction below: 1. Check Last Call Home: <here> 2. Communication Test 3. Log Check 4. Certificate Check

Figure 5. NetInsights recommended actions detail display


Solution Architecture

At a high level, the basic components of this solution consist of data sources to archive, Enterprise Vault, and Access Appliance as shown in Figure 5. Enterprise Vault archives diverse sources such as Microsoft Exchange, Microsoft SharePoint, IBM Domino, SMTP messages, and unstructured file data (i.e. NTFS or UNIX/Linux filesystems). Enterprise Vault sends archive data to Access Appliance using the Server Messaging Block/Common Internet File System (SMB/CIFS) protocol. The following sections expand on each of these components.

Enterprise Vault (EV)

Enterprise Vault software is a scalable archiving platform and well known for managing information in the following featured areas:

- **Compliance:** Reduces risk by proactively monitoring electronic communications to comply with industry and government regulations.
- **Discovery:** Allows for IT and Legal discovery with guided review to assist in reducing costs of eDiscovery, litigation, and compliance demonstration.
- **Retention:** Provides policy-based retention of data to keep what is important and delete waste. Data can be stored with appliance-enforced immutability via automatic or manual classification.
- **Optimization:** Reduces storage with the single instance storage (SIS) feature. With SIS, if a file has already been found within the sharing boundary, then another copy of the file is not stored. Data is also compressed prior to sending to the target storage platform.



Enterprise Vault main components as shown in Figure 6 include:

- **Enterprise Vault (EV) Server**
Runs several tasks and services:
 - Archiving task which connects to target system to discover items to be archived
 - Storage services responsible for storing the items in Vault Store partitions (i.e. folder in storage)
 - Indexing services that index any text, document, text of email, etc. for fast searching and retrieval
 - Web access components to enable viewing, searching, and restoring archived data by a user accessing the server with a web browser
- **Microsoft® SQL Server**
Contains numerous database tables relating to configuration information of Enterprise Vault, the hashes or fingerprint of every single item archived, monitoring and reporting data, and the vault store meta-data.
- **Vault Store Partitions**
Storage for the data archives.

The components of Enterprise Vault can be run on a single large system or distributed over several servers. For instance, multiple Enterprise Vault servers can be managed by a single administration console with each server handling different sources to archive, and running various tasks and services.

There is a graphical user interface (GUI) for Enterprise Vault that is responsible for administration, configuration, and management of archival targets and storage, along with the ability to change the settings relating to retention of data, monitoring, and reporting. Additionally, there are add-ons that are bundled with Enterprise Vault to support archiving the various data sources such as email exchange, file systems, SharePoint data, SMTP messages, etc. Extensions developed or co-developed with partners that extend the functionality of Enterprise Vault are available in the [Veritas Technology Partner Program](#). Also, refer to the [Enterprise Vault Compatibility Charts](#) for more information on third party integrations.

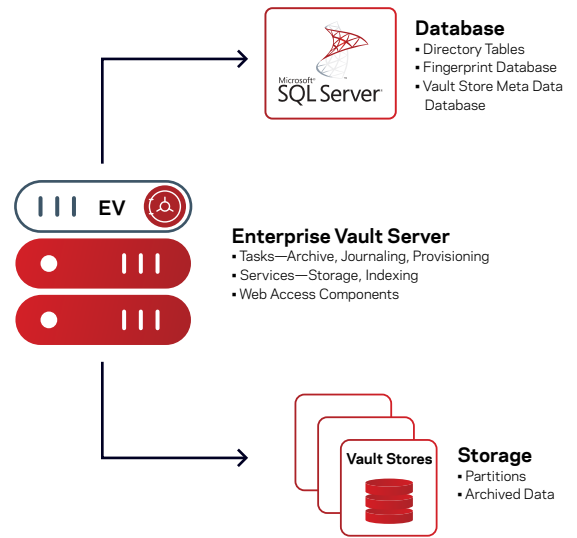


Figure 6. Enterprise Vault architectural components

Components such as EV Cache, SMTP holding folder, PST holding folder, etc. are beyond the scope of this document. However, for more information on Enterprise Vault, refer to [Veritas Enterprise Vault Product Documentation](#).

Specification	Access Appliance 3340	Access Appliance 3350
Number of nodes	2	2
CPU	Intel Xeon 4108	Intel Xeon Silver 4210
Memory (per node)	384, 768, 1152, or 1536 GB	
Rack units	Node: 2U Storage shelf: 5U Fully expanded system: 24U	
Expansion shelves	1-4 per appliance	
Shelf storage	82 drives per shelf (all 4TB or all 10TB drives)	
10/25Gb network ports	2	2 or 4
Storage total	280 TB-2800 TB usable per appliance	

Enterprise Vault can send archive data to varying storage types (disk, tape, and cloud). For those seeking an on-premises disk-based solution for faster recovery times, fine-grained control and/or greater simplicity when compared to tape or cloud, Access Appliance boasts ease of acquisition, simplified management, and tight integration with other Veritas appliances and software as an archival solution. Access Appliance is a turn-key storage solution designed for high capacity, resilience, and tight integration with Veritas archival and cost optimization, making it well suited for an archival use case. Access Appliance (model 3340 or 3350) comprises two clustered nodes and one primary storage shelf at minimum, and up to three additional expansion storage shelves, allowing a total of 2800 TB of archival storage space per appliance.

Highlights of Access Appliance specifications are shown in Table 1. Refer to the [Access Appliance datasheet](#) for more detailed information.

In general, Access Appliance is meant to be general-purpose network-attached storage with the high-capacity, resilience, and density required for archival storage solutions. As well as common network file protocols such as NFS and SMB/CIFS, it can serve object data using the S3 protocol, and accept data from NetBackup media services as a storage server providing a deduplicated data pool.

Access Appliance nodes are clustered in an active/active configuration such that each node can handle I/O requests as well as taking over the tasks of its partner node in the event of failure. Storage shelves are connected to both nodes in parallel and configured with dynamic multi-pathing capabilities so I/O can be sent to either node for performance and availability purposes. The redundant hardware RAID controllers in the primary shelf aggregate its storage into RAID 6 volumes with two parity disks for every 14 data disks, in five data volumes per shelf. Each volume can survive up to two simultaneous disk failures.

For Enterprise Vault's purposes, data is written to Access Appliance using the SMB/CIFS protocol. An SMB share is exported and maps to a single file system of type clustered file system (CFS). The size of a CFS file system can scale up to 2800 TB on Access Appliance.

For more information on configuration of WORM and replication, refer to the [Veritas Access Administrator's Guide](#).

Solution Data Flow

Depending on the applications or data (i.e., SharePoint, Exchange, Domino, etc.) being archived, the data flow and process within Enterprise Vault may differ. For more detailed process descriptions on the varying data flows within Enterprise Vault, refer to the [Enterprise Vault Process Diagrams](#). However, in context with the storage aspects of this solution, Access Appliance acts as a SMB/CIFS target for Enterprise Vault store partitions and/or a secondary storage location where files can be migrated from the vault store partition. When data is archived from Enterprise Vault, a shortcut or stub of the data is created on the client side and the contents of the data are moved to Access Appliance using the SMB/CIFS protocol to free space on the client. When the archived data is once again accessed, a restore is initiated from Enterprise Vault to retrieve the data from its archive on Access Appliance.

In Enterprise Vault 8.0 and later, an archived item is stored in [several proprietary formats](#) on Access Appliance, including:

- **DVS (saveset)**—message header information of data. In the case of email, this refers to the date sent, senders, recipients, and main portion of the message body.
- **DVSSP (saveset shared part)**—shared part of the data (i.e. attachments in email).
- **DVSCC (saveset converted content)**—converted content of the attachment into HTML, text or raw text. The DVSCC file is what is used by the indexing services.

If collections are enabled, the files are stored as **CAB** (Microsoft Windows® Cabinet) files. An example where collections are created is when the data within Enterprise Vault partition is collected and migrated to secondary storage in which Access Appliance can also be a storage target. A sample view of archived data on Access Appliance is pictured in Figure 8.




Figure 7. File view of Enterprise Vault data on Access Appliance

Archival Data Flow

Enterprise Vault archives unstructured data based on a schedule. Archival policies are defined to determine what data to archive and when, along with retention categories, automated deletion, and other configurable parameters. Several Enterprise Vault services and tasks are involved prior to data being stored in Access Appliance. In general, when any data is archived, Enterprise Vault's services and tasks queue up requests, and processes them as follows (illustrated in Figure 8):

- First, archive data based on scheduled policies
- Second, Enterprise Vault services:
 - Extract the text from the document, create and store an index of the extracted data along with the metadata, and place them in the assigned index storage location
 - Check the fingerprint database to determine if the data has already been archived and can be referenced instead of stored again
- Once the data has been indexed and checked, the storage service places a reference of the data in the vault store database and the single instance parts of the data are compressed and converted into several proprietary files (DVS, DVSSP, DVSCC) that contain the data and associated information
- A placeholder or stub of the data replaces the file in the client view

NOTE: The indexes and database are not recommended for storage in Access Appliance. Only the archived data segments in the vault store partitions are stored on Access Appliance.




Figure 8 Enterprise Vault archival data flow

Retrieval Data Flow

Once a data has been archived, Enterprise Vault presents an entry listing to the client as a shortcut and the end user can seamlessly access the file as if the file was not archived. For instance, the original extension or file type, icon, and size of file can be seen. If a user double-clicks on the archived file or email from their browser, the data will be restored. As shown in Figure 9, when the client requests the archived data, the following occurs:

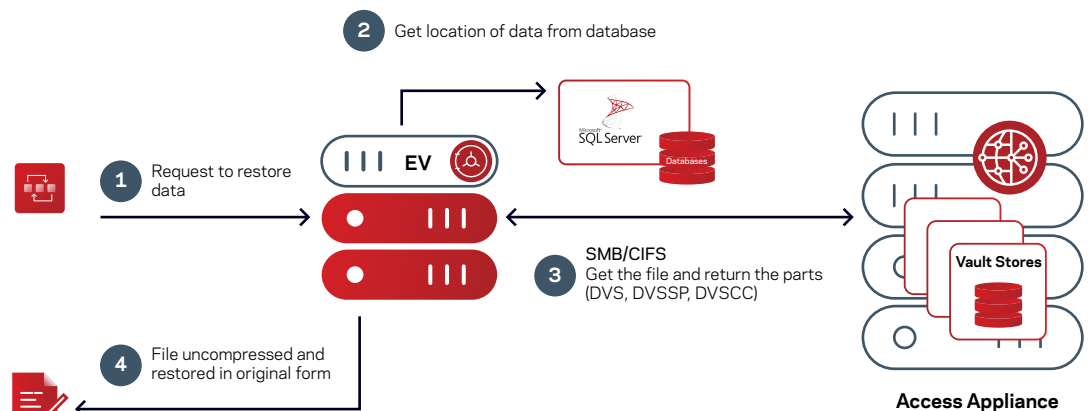


Figure 9. Retrieval data flow

1. The request is sent to EV and goes through the web server (Microsoft IIS) running on the EV server to handle the request
2. The storage services query the SQL database regarding the location of the archived data
3. The archived parts are retrieved from the vault stores residing on Access Appliance
4. The file is uncompressed and parts are re-constituted by Enterprise Vault and returned to the client.

Enterprise Vault Migrator Data Flow

A vault store partition can be further migrated to a secondary storage target. Access Appliance can also act as a secondary storage target for the Enterprise Vault migrator. The data archived in a partition is migrated as a collection file (CAB). Collections are migrated based on age, or according to a specified schedule. As illustrated in Figure 10, the sequence of events includes:

1. The Enterprise Vault collector is run daily at a specified time, based on the age of files
2. The saveset files (DVS, DVSSP and DVSCC) from the vault store partition (not WORM enabled) stored in Access Appliance are retrieved, and a collection (CAB) file is generated
3. Collection files are migrated based on age
4. Collection files are then migrated to an SMB/CIFS share on another Access Appliance




Figure 10. Enterprise Vault migrator data flow

Best Practices and Recommendations

Following best practices is important in creating an optimal deployment. This section covers some best practices relating to Access Appliance as an archival storage solution for Enterprise Vault.

File System and Data Layout on Access Appliance

Access Appliance uses hardware RAID controllers in its primary storage shelf to provide redundant RAID 6 storage volumes. Therefore, the default simple layout therefore is sufficient for data protection and availability purposes. The simple layout also makes it easier to grow the volume later without having to be concerned about matching the stripe volume size for a striped layout or consuming additional storage as in a mirrored layout.

Enterprise Vault requires storage for its SQL database, indexes, and vault stores (archive data). As a best practice, use Access Appliance only for the vault store partition, since the database and index components of Enterprise Vault require faster storage.

By default, the SMB/CIFS share is in normal clustering mode. In this mode, only one node is responsible for servicing the requests. For simultaneous servicing of a share in which either node can service the requests, the SMB/CIFS share should be configured in Clustered Trivial Database mode (CTDB). See the [About CIFS clustering modes](#) section in the [Veritas Access Appliance Administrator's Guide](#).

An Access Appliance instance can be a target for multiple Enterprise Vault deployments. In these scenarios, it is a best practice to not target a single filesystem with multiple Enterprise Vault deployments. Each Enterprise Vault deployment does not have knowledge that the filesystem is being used by other Enterprise Vault partitions and will perceive that it has full use of the filesystem capacity. Therefore, you should use one or more Access Appliance filesystems per Enterprise Vault deployment. Using more than one filesystem also aids with the parallelism of I/O to underlying disks and volumes.

Since some operations such as filesystem check and NetBackup client backups are done at the filesystem level, it is recommended to create file systems not more than 5 TB in size and use partition rollover especially for archives that have a lot of small files (40 KB-50 KB) such as in email archives. NOTE: For availability reasons, the maximum number of filesystems supported on Access Appliance is 50.

Network Connectivity

Access Appliance can be configured with either two or four 10/25 GbE network connections. Each physical port maps to a virtual IP. Thus, there are four virtual IP addresses. Always present a virtual IP to clients or client applications so that they will automatically transition to the other node if one node fails or the physical links on one node fail or become unreachable.

Monitoring

It is important to monitor warnings and alerts, especially storage utilization warnings and hardware critical alerts. AutoSupport and call-home telemetry assists in providing timely alerts for this issue, but a best practice is to be proactive instead of reactive about a storage growth ceiling. For instance, once an appliance's storage capacity reaches 60 percent, it's a good time to revisit storage utilization and plan for growth.

Load Balancing

There are two nodes in an Access Appliance instance, active/active nodes capable of partner takeover in case of failure. As a best practice, balancing the load across nodes is recommended. Load balancing can be achieved using any of the following techniques:

- **External load balancing:** Using an external load balancer such as HAProxy or F5 allows for its algorithms to distribute load across nodes such as least connections or weights. This method also frees the Access Appliance nodes from the proxy handling and balances the network traffic between the nodes.
- **Manual load balancing:** Creating additional virtual IP addresses on Access Appliance nodes means they can be manually assigned to applications in a distributed manner. A disadvantage of this approach is that even distribution may be difficult to gauge, since applications are not necessarily equal with regard to workload.
- **DNS load balancing:** Access Appliance's DNS record includes all the virtual IP addresses of the Access Appliance nodes. In response to lookups, DNS round-robins through the virtual IP addresses. One disadvantage of using DNS is that the virtual IP will remain with its associated DNS record, even if there are connectivity problems, unless it is manually removed.

NTP Server

Connecting Access Appliance to an NTP server is a recommended best practice so that the hosts running Enterprise Vault, data sources, and Active Directory are time synchronized. If date and time are not synchronized between the hosts and Access Appliance, issues may arise. For instance, communication to authenticate a user via Active Directory may fail.

Replication

Deploying Access Appliance to multiple locations is a common way to add site-level resiliency to archival storage. As well as supporting solution-specific data replication methods such as NetBackup AIR, Access Appliance has episodic and continuous replication facilities designed to update a remote copy of an appliance filesystem.

For replicating Enterprise Vault archival data, it is important to consider how original copies/safety copies will factor into the flow of data when choosing a replication scheme.

If Enterprise Vault is configured to remove original items only after data has been archived in Access Appliance, and subsequently replicated so that there are always two copies of the archived data:

- Access Appliance should use episodic replication for the replication job and set the `evpsn` argument to **yes** when creating the episodic replication job. In the [Veritas Access Appliance Solutions Guide for Enterprise Vault](#), see the [Partition Secure Notification section](#).
- Enterprise Vault should set the vault store partition properties to check for a trigger file.

If Enterprise Vault is configured not to remove the original items, or if replication using Access Appliance isn't what determines if there are sufficient copies:

- Access Appliance should either use episodic replication and set the `evpsn` argument to **no** when creating the episodic replication job, or use continuous replication. In the Access Appliance Administrator's Guide, see the [Configuring continuous replication section](#).
- In Enterprise Vault, it is recommended to set the vault store partition properties to check for a trigger file, rather than use the archive attribute. See [figure 4-3](#) in the Veritas Access Appliance Solutions Guide for Enterprise Vault.

Sizing Guidance

Access Appliance is used as a storage target for Enterprise Vault stores partitions. In planning for the vault store partitions for Enterprise Vault, there are two considerations:

- **Capacity:** How much archive data can be stored. As previously mentioned, Enterprise Vault archives items in several proprietary formats such as DVS, DVSSP, and DVSCC, single DVS file, or CAB file, depending on the software version.
- **Performance:** How much workload (throughput and bandwidth) the storage platform can handle.

A Veritas account team will assist in the sizing of Access Appliance based on your requirements.. Some additional parameters that might enter in the equation when estimating archival storage requirements include:

1. Archive type (Microsoft Exchange, SharePoint, Domino, SMTP Journaling, file system, etc.)
2. Based on archiving type, additional questions may include state of archiving such as steady state (on-going incremental archive), backlog (initial archiving of documents) and journal archiving. This includes:
 - Microsoft Exchange/Domino Mailbox
 - v. Number of mailboxes
 - vi. Average size of messages and attachments
 - vii. Estimated percentage of messages with attachments
 - viii. Average number of messages sent daily/yearly per mailbox
 - ix. Average number of messages received internally and externally, daily/yearly per mailbox
 - x. Annual growth in number of mailboxes archived, messages, and average size of messages and attachments
 - SharePoint
 - i. Number of documents to archive daily
 - ii. Typical average size of documents (Office documents, images, PDF files, etc.)
 - iii. Percentage of documents greater than 20 KB
 - iv. Annual percentage growth of number and average size of files
 - v. Typical ingest rate
 - SMTP Journaling
 - i. Initial number of SMTP journal messages daily
 - ii. Average size of messages and attachments
 - iii. Estimated percentage of messages with attachments
 - iv. Average number of messages sent daily/yearly per mailbox
 - v. Average number of messages received internally and externally, daily/yearly per mailbox
 - vi. Annual growth in number of mailboxes archived, messages, and average size of messages and attachments
 - File system
 - i. Number of files
 - ii. Typical compressed size of file in percentage or average number of duplicates of each file
 - iii. Average size of files
 - iv. Typical ingest rate
 - v. Annual growth of number and average size of files

- PST Migration
 - i. Number of messages and attachments in PST file
 - ii. Total size of messages in PST files
 - iii. Average message and attachment sizes
 - iv. Percentage of messages with attachments
- 3. Performance and/or service level requirements.

Also refer to the [Enterprise Vault Performance Guide](#), which describes how to calculate the estimated disk space for Enterprise Vault storage (indexes, database, and vault store partitions), performance (EV hourly ingest rate, rules of thumb for IOPS), and other considerations for each archive type.

Conclusion

Veritas Enterprise Vault with Access Appliance offers an end-to-end solution for information cycle management and data archival. Implementing Access Appliance as a dense, flexible, and resilient storage target platform with Enterprise Vault simplifies management and support, minimizes costs, and improves control and visibility. Furthermore, Access Appliance deepens the solution's capabilities with WORM and immutability features, solution-integrated episodic replication, storage efficiency, encryption, monitoring, AutoSupport and telemetry, and integration with other Veritas products such as NetBackup and Data Insight.

References

- Veritas Access Appliance with Enterprise Vault Archival Solution
https://www.veritas.com/content/dam/Veritas/docs/white-papers/WP-Access_Appliance_with_EV_Solution-EN.pdf
- Veritas Appliance AutoSupport Reference Guide
<https://sort.veritas.com/DocPortal/pdf/118608033-130448966-1>
- Veritas Access Appliance Solutions Guide for Enterprise Vault
<https://sort.veritas.com/DocPortal/pdf/146116601-146116608-1>
- How Data Insight Works
https://www.veritas.com/support/en_US/article.100039166
- Data Insight User's Guide
<https://sort.veritas.com/DocPortal/pdf/140216462-155484145-1>
- Veritas Access Appliance Administrator's Guide
<https://sort.veritas.com/DocPortal/pdf/146126550-151725528-1>
- Veritas Access 3340 Appliance Product Description
<https://sort.veritas.com/DocPortal/pdf/125460431-134247411-1>
- Veritas Access 3350 Appliance Product Description
<https://sort.veritas.com/DocPortal/pdf/125460431-154188696-1>
- Veritas Access Appliance Solutions Guide for NetBackup
<https://sort.veritas.com/DocPortal/pdf/146127092-151725537-1>

About Veritas

Veritas Technologies is a global leader in data protection and availability. Over 80,000 customers—including 87 percent of the Fortune Global 500—rely on us to abstract IT complexity and simplify data management. The Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas Enterprise Data Services Platform supports more than 800 different data sources, over 100 different operating systems, more than 1,400 storage targets, and more than 60 different cloud platforms. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

VERITAS™

2625 Augustine Drive
Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For global contact
information visit:
veritas.com/company/contact