

Veritas Access Appliance with Enterprise Vault

Archival Solution

Contents

Contents	2
Revision History	3
Introduction	4
Executive Summary	4
Scope	4
Target Audience	4
Solution Value	4
Solution Key Features	5
WORM (Write Once Read Many)	5
Storage Efficiency	5
Replication	5
Seamless Integration with NetBackup	6
Encryption	6
Shadow Copy with Versioning	6
Tight Integration with Data Insight for Better Visibility	6
AutoSupport Feature	7
Monitoring and Detection	8
Solution Architecture	8
Enterprise Vault (EV)	9
Access Appliance	10
Solution Data Flow	12
Archival Data Flow	13
Retrieval Data Flow	13
Enterprise Vault Migrator Data Flow	14
Best Practices and Recommendations	15
File System and Data Layout on the Access Appliance	15
Network Connectivity	15
Monitoring	15
Load Balancing	16
NTP Server	16
Sizing Guidance	16
Conclusion	17
References	17
Appendix	18
Access Appliance Setup	19
Setup Access Appliance Share as a Storage Target on Enterprise Vault	29
Validation of Access Appliance as a Storage Target for Enterprise Vault	35
Setup Access Appliance as a Secondary Storage Target for Enterprise Vault Collections	42
Configuration of Episodic Replication Notes	44
Table of Figures	46

Revision History

Rev 1.0 8 Mar 2018	Initial version
Rev 2.0 1 Oct 2018	Updated for 7.4.2 Release
Rev 2.1 4 Apr 2018	Modified info on EV encryption
Rev 3.0 3 May 2021	Updated for 7.4.3 Release

Introduction

Executive Summary

With the exponential growth of data, strategies for information lifecycle management and data archival have become increasingly important for numerous organizations who are grappling to control costs and meet business compliance and regulatory requirements. Veritas Technologies Enterprise Vault with Access Appliance provides a complete solution to address these challenges. Enterprise Vault is a platform for managing information capable of archiving from an 80+ data sources including native support for Microsoft® Exchange, IBM Domino®, Skype for Business, Microsoft SharePoint and file systems. The Access Appliance acts as a dense on-premises storage target for data that has been archived using Enterprise Vault. Used together with Veritas Technologies other products, such as NetBackup and Data Insight, they improve the visibility of an organization's data and allows for better data management to reduce cost, waste, and risk. The integration of these Veritas products provides a compelling solution for the archival use case.

Scope

The purpose of this document is to provide technical details to assist in understanding the Access Appliance with Enterprise Vault as a solution for archival. It describes the components of this solution, its value, key features, sizing guidance, and some best practices. **NOTE:** This document gets updated periodically and if you downloaded a local copy of this document, please get the latest from this [link](#).

Target Audience

This document is targeted for customers, partners, and Veritas field personnel interested in learning more about the Veritas Access Appliance with the Enterprise solution for archival. It provides a technical overview of this solution, guidance in sizing, and highlights some best practices.

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 of 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. The Access Appliance addresses these challenges and provides the following key value as a target storage platform for Enterprise Vault:

- **Minimize cost** – Access Appliance provides a low-cost, disk-based solution that is easy to manage. With Enterprise Vault single instance storage and compression features which is maintained on Access Appliance reduces overall costs and enhances storage efficiency.
- **Simplify Management** – maintenance and management of varying secondary storage types, media, and protocols present challenges to IT departments. Having a single vendor to handle an organization's archival and information lifecycle management requirements improves issue resolution and simplifies management and ordering.
- **Increase visibility and control** – insight and characterizations on how data is utilized assists in planning and helps identify storage that are not being used, orphaned, and/or no longer required. Enterprise Vault's seamless

integration with Veritas Data Insight provides greater visibility to your data to reduce inefficient utilization of resources, storage waste, and overall cost.

Solution Key Features

There are certain key features that companies look for in an archival solution product. The main features that are often sought include compliance, flexibility, storage efficiency, and ease of management. The Access Appliance with Enterprise Vault provides these features plus more to assist customers in preserving their most valued data.

WORM (Write Once Read Many)

For compliance and regulatory requirements, having WORM features in a storage platform is important when selecting a storage platform for archived data. Industries such as financial and government institutions must adhere to certain data retention rules or be penalized and fined. With the WORM feature, the specified data cannot be modified or deleted until the retention period expires.

The WORM feature must be first enabled on the Access Appliance for the file system that is shared and then retention period is set at file level by Enterprise Vault. Files that are WORM enabled are protected from access by any user including root and administrative users.

NOTE: The WORM feature is only supported on an SMB share created on a CFS type file system.

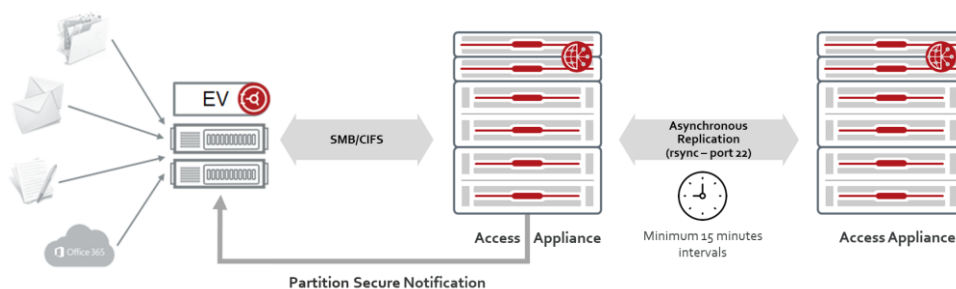
Storage Efficiency

Support for storage efficiency is one of the main factors in the decision-making during the purchase of a storage platform solution for archival. The ability to maximize storage space assists in reducing overall cost. The Access Appliance can leverage Enterprise Vault compression and SIS features. Both these features are maintained in the Access Appliance. 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

As previously mentioned, for additional data protection, Access Appliance has the ability to perform “[episodic replication](#)”. This feature can replicate data or files asynchronously to another Access Appliance or Access Software Defined Storage (SDS) cluster. The utility rsync is used to transfer data from the Access Appliance to another Access cluster via port 22 at set intervals as shown in Figure 1. After the data has been replicated, the Access Appliance acting as the source has an option to send partition secure notification to Enterprise Vault. This notification triggers Enterprise Vault to remove its safety copies.

Figure 1 - Replication Data Flow



Seamless Integration with NetBackup

Being prepared for disasters is imperative for business continuity and productivity. The Access Appliance is integrated with NetBackup such that the data archived by Enterprise Vault in Access Appliance can be further backed up to a NetBackup media server and restored in case of failures. A container-based [NetBackup client add-on package](#) is available and can be installed via the command-line shell.

Encryption

Any data encryption done by an application and archived using Enterprise Vault is maintained on the Access Appliance. The Access Appliance also has encryption capabilities in conjunction with an external Key Management System (KMS). The appliance encrypts the volume that the “file system” resides on. An external KMS such as IBM KMS is required to create the keys for the encryption.

Shadow Copy with Versioning

If customers desire to utilize the Microsoft Windows® Server Volume Shadow Copy Service with Enterprise Vault stores, then the Access Appliance supports the creation of “shadow copies” (copies of the data within a volume at a specific point in time). The advantage of enabling shadow copies on shared folders is that one can quickly recover data in case of corruption, accidental deletion or being overwritten. Each shadow copy is versioned so a specific version of the data can be recovered. The Access Appliance has support for shadow copy and it is specified as an export option when the SMB share is created. By turning on this option, Access has awareness of and the ability to store shadow copies created by Microsoft Windows Server. For more information, refer to [Storage Foundation - Quick Recovery Solutions Guide for Enterprise Vault - Windows](#) and the [Veritas Access Administrator's Guide](#).

Tight Integration with Data Insight for Better Visibility

Veritas Technologies offers Data Insight to provide better visibility of data sources within an organization’s environment. It has features in either illuminating the data either by classification to identify certain content within the file or by inspecting your entire environment across numerous locations to determine where your data is located. Data Insight is tightly integrated with Enterprise Vault to facilitate the archival of data with storage tier optimization to Access Appliance in addition to security and management of information.

[Veritas Data Insight](#) product 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. in order that administrators can identify data that can be archived or tiered to cheaper storage, enforce security, conduct data chargeback, and perform information lifecycle management and risk analysis.

With regards to this solution, Data Insight is a powerful tool to identify areas in the data sources that can be archived due to inactivity, age, or compliance. As pictured in Figure 2, data sources are scanned, analyzed and classified by Data Insight. These sources can then be manually inspected, or policies can be defined to determine the data that can be archived using Enterprise Vault and stored in the Access Appliance. Knowledge of what are in the data sources allows organizations to make more informed decisions on what to do with the data for storage optimization, security, and archival.

Figure 2 - Data Insight with Enterprise Vault and Access Appliance


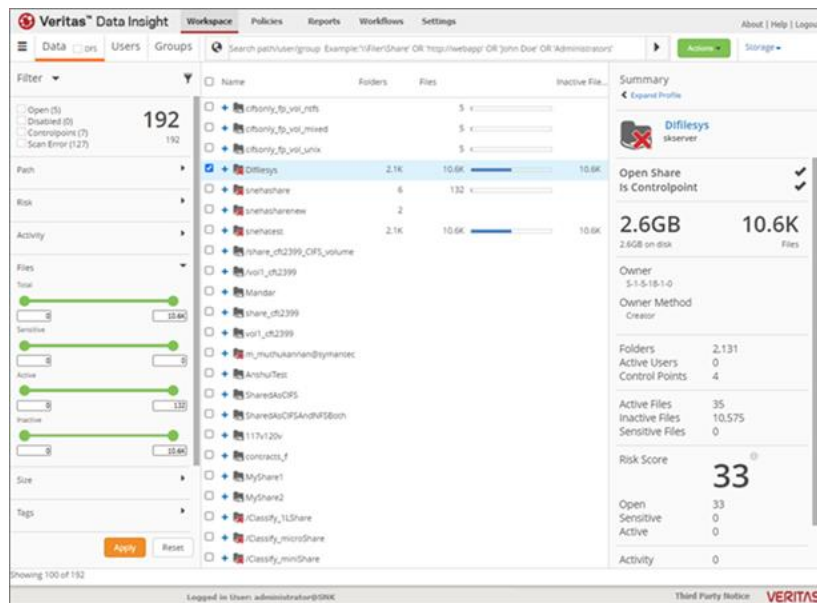


Figure 3 is a sample view of the Data Insight Console displaying a list of group shares and folders with their sizes and number of inactive files being monitored. The inactive folders or files can either be selected for deletion or archival using Enterprise Vault. For more information, refer to [Data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

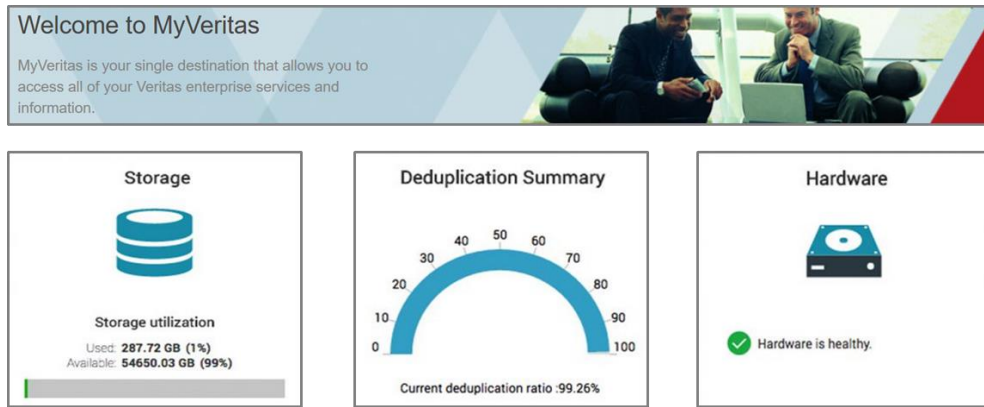
Figure 3 - Sample View of Data Insight Console - Inactive Files in Group Shares



AutoSupport Feature

For the Access Appliance, there is an AutoSupport feature that can call-home in case of hardware and software issues observed in the production deployment. The advantages of using Veritas appliances for the entire solution are the ability to automate support case management and leverage guided workflows for faster resolutions of issues and mitigation of risks. Veritas AutoSupport service provides proactive monitoring and alerting 24x7 on the health of the appliances. This feature alerts customers and/or service engineers to quickly handle the issue and reduce further risks. Enabling this feature can be done simply by registering the appliance(s) at the Veritas MyAppliance portal as shown in Figure 4 and enabling the call-home functionality. 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.

Figure 4- MyAppliance Portal View




Monitoring and Detection

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

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 file system (i.e. NTFS and UNIX). Enterprise Vault sends archive data to the Access Appliance using the Server Messaging Block/Common Internet File System (SMB/CIFS) protocol. The following sections expand on each of these components.

Figure 5 - Access Appliance with Enterprise Vault Solution High-Level Architecture



Enterprise Vault (EV)

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

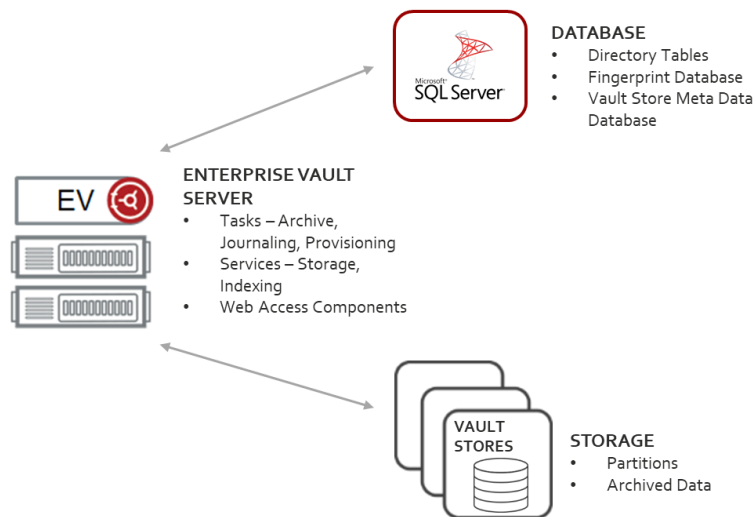
- **Compliance** – reduces risk by proactively monitoring electronic communications to comply with industry and government regulations.
- **Discovery** – allows for IT and Legal discovery with guided review to assist in reducing costs of eDiscovery, litigation and compliance demonstration.
- **Retention** – provides policy-based retention of data to keep what is important and delete waste. Data can be stored 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 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 comprising the following:
 - Archiving task which connects to target system and find items that should 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 user via a web client.
- **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.

Figure 6- Enterprise Vault Main Components




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. In addition, 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.

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

Access 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, Veritas has developed the Access Appliance for ease of acquisition, management, and support. Access Appliance is a turn-key storage solution designed for high capacity and cost optimization, making it well suited for archival use case. The Access Appliance model 3340 is comprised of two clustered nodes and one primary storage shelf and up to three additional expansion storage shelves. The appliance can scale up to 2,800 TB of usable space as shown in Figure 7.

Figure 7 - Access Appliance Rack Units



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

Table 1 - Highlights of Access Appliance Specifications

Model	CPU Processor	RAM	Ports		Capacity	Rack Units
			1 GbE	10 GbE		
3340 (2 nodes)	2 x Intel® Xeon® 4108 (1.8 GHz) per node Total: 16 core per node	384 GB per node	4 per node	2 per node	280 TB – 2800 TB (254 TiB– 2544 TiB)	Server: 2U Storage per Shelf: 5U

Note: TB - Capacity values are calculated using Base 10; TiB - Capacity values are calculated using Base 2.

The two nodes are clustered in active/active configuration such that each node can handle I/O requests, and each node is able to take over the tasks of its partner in the event of failure. Storage shelves are connected to each node 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 5 data volumes per shelf. Each volume can survive up to two simultaneous disk failures.

Access Appliance runs Red Hat Enterprise Linux (RHEL) version 7.4 or later as the operating system platform and Access software version 7.3.2 or later. The Access Appliance is a scale-up Network-Attached Storage (NAS) platform that supports multiple protocols, including NFS, SMB/CIFS, FTP, and S3. For Enterprise Vault’s purposes, data is written to the 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 the Access Appliance.

Additionally, the Access Appliance supports WORM (Write-Once-Read-Many). WORM is first enabled at the file system level, then the retention period is set at the file level. After the retention period is set, it cannot be modified or deleted by any user including root and admin until the retention period expires.

Data can also be replicated synchronously and asynchronously to another Access Appliance for additional data protection. There are two methods to replicate data to an Access Appliance:

- **Episodic Replication** – file system replication asynchronously.
- **Continuous Replication** – block level replication of volumes in synchronous or asynchronous modes.

If there is a requirement to inform Enterprise Vault to remove the safety copies, then episodic replication should be used. Episodic replication runs at regularly scheduled intervals of a minimum of 15 minutes. After replication of data to another Access Appliance, the source Access Appliance can send a partition secure notification to Enterprise Vault to remove the safety copies. This can be achieved by setting the “evpsn” flag to “yes” when creating the replication job or schedule.

NOTE: The “evpsn” option is only available for episodic replication on a CFS file system.

Refer to the Appendix section for examples of how to deploy and configure the Access Appliance with Enterprise Vault. For more information on configuration of WORM and replication, refer to the [Veritas Access Administrator's Guide](#).

Solution Data Flow

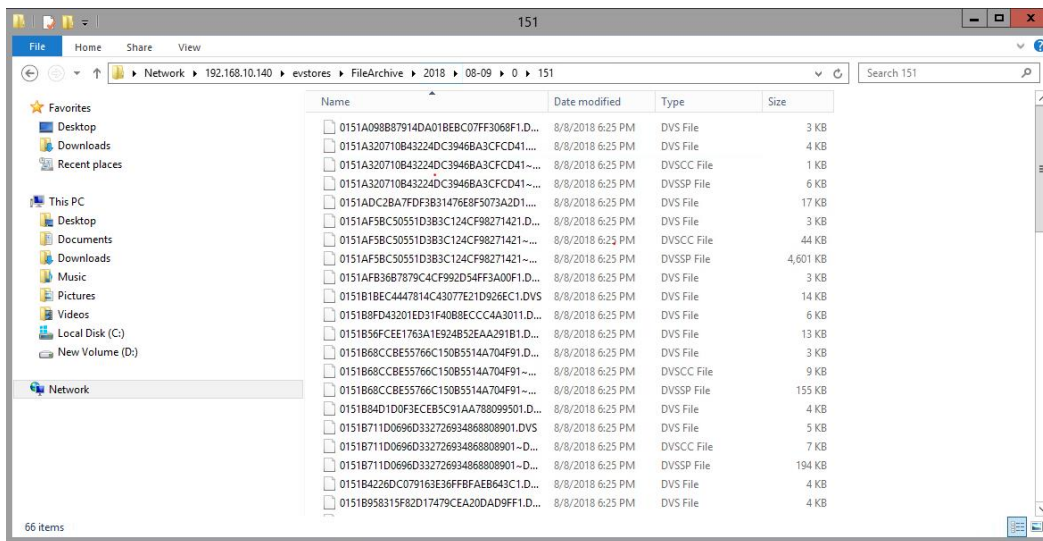
Depending on the applications or data (i.e. SharePoint, Exchange, Domino, etc.) being archived, the data flow and process within Enterprise Vault may differ. For more detailed process descriptions on the varying data flows within Enterprise Vault, refer to the [Enterprise Vault Process Diagrams](#). However, in context with the storage aspects of this solution, the 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 the Access Appliance.

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

- **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 the Access Appliance can also be a storage target. A sample view of archived data on the Access Appliance is pictured in Figure 8.

Figure 8- Sample View of EV Archived Data on Access Appliance




The next sections describe the data flow for archival and retrieval between Access Appliance and Enterprise Vault.

Archival Data Flow

Enterprise Vault archives unstructured data based on a schedule. Archival policies are defined to determine what data to archive and when, retention categories, automated deletion, and other configurable parameters. Several Enterprise Vault services and tasks are involved prior to data being stored in the Access Appliance. In general, when a data is archived, Enterprise Vault services and tasks queue up requests, and process the queues as follows (illustrated in Figure 9):

1. Archive data based on scheduled policies.
2. Enterprise Vault services run the following:
 - a. Indexing services extract the text from the document, indexes the data along with the meta-data, and places them in the assigned index storage location.
 - b. Storage services check the fingerprint database to determine if the data has been archived already and if so, it is not stored again.
3. 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 contains the data and its associated information.
4. A placeholder or stub of the data replaces the file in the client view.

Figure 9 - Archival Data Flow




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

Retrieval Data Flow

Once a data is archive, Enterprise Vault presents data 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 and if user double-clicks on the archived file or email from their browser the data will be restored. As shown in Figure 10, when the client requests for the archived data the following occurs:

1. 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 data.
3. The archived parts are retrieved from the vault stores residing on the Access Appliance.
4. File is uncompressed and parts are re-constituted by the Enterprise Vault services and returned to client.

Figure 10 – Retrieval Data Flow




Enterprise Vault Migrator Data Flow

A vault store partition can be further migrated to a secondary storage target. The Access Appliance can also act as 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 11, the sequence of events includes:

1. The Enterprise Vault collector is run based on the age of files and run daily at a specified time.
2. The saveset files (DVS, DVSSP and DVSCC) from the vault store partition (not WORM enabled) stored in the Access Appliance are retrieved. 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 11 - Migration of Enterprise Vault Collections to the Access Appliance for Secondary Storage.



Best Practices and Recommendations

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

File System and Data Layout on the Access Appliance

The Access Appliance contains hardware RAID 6 controllers in its primary storage shelf, so, the simple data layout should be used. Mirroring or other software RAID layout is not necessary 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. Finally, it is not recommended to use the predefined policies for Enterprise Vault on the Access Appliance GUI since it creates a CFS file system with a mirrored layout. These predefined policies were designed to be used with Access Software Defined Storage.

Enterprise Vault requires storage for the SQL database, indexes, and vault stores (archive data). As a best practice, use Access only for the vault store partition since the database and index components of Enterprise Vault require faster storage. In addition, although the Access Appliance has a scale-out file system, Enterprise Vault has been fully certified with the CFS file system, and thus it is recommended to use the CFS file system type for Enterprise Vault.

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). The Appendix describes how to configure the clustering mode to CTDB.

The Access Appliance can be a target for multiple Enterprise Vault deployments. In these scenarios, it is best practice to not to allocate one filesystem as the target for 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. Thus, use one or more Access filesystems per Enterprise Vault deployments. Using more than one filesystem provides parallelism of reads and writes.

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 files of small files (40 KB – 50 KB) such as in email archives. **NOTE:** Maximum filesystems supported on Access Appliance is 50.

Network Connectivity

The Access Appliance has two 10 GbE uplinks per node. Each physical port maps to a virtual IP. Thus, there are four virtual IP addresses. Always present the 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 fails or becomes unreachable. Refer to Appendix section to determine how to get a listing of the virtual IPs.

Monitoring

It is important to monitor or be aware of alerts, especially storage utilization warnings and hardware critical alerts. The AutoSupport features assists in this manner, but as a best practice, it is advisable to be pro-active instead of re-active. For instance, once the storage capacity reaches 60%, it might be a good time to revisit storage utilization or plan for growth.

Load Balancing

There are two nodes on the Access Appliance configured as active/active. As a best practice, balancing the load across nodes on Access 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 more algorithms to distribute load across nodes such as least connections or weights. It also frees the Access nodes from the proxy handling and balances the network traffic between the nodes.
- **Manual load balancing** – virtual IP addresses of the nodes can be manually assigned to applications in a distributed manner. The disadvantage of this approach is that even distribution may be difficult to gauge since applications are not all equal in sense of workload.
- **DNS load balancing** – DNS includes all the virtual IP addresses of the Access nodes. DNS round-robins through the virtual IP addresses. The disadvantage of using DNS is when there are connectivity issues, the virtual IP is still in rotation until it is manually removed.

NTP Server

Connecting the Access Appliance to an NTP server is a recommended best practice in order that the hosts running Enterprise Vault, data sources, 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 user via Active Directory may fail.

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 EV version.
- **Performance** – how much workload (throughput and bandwidth) the storage platform can handle.

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

1. Archiving type (Microsoft Exchange, SharePoint, Domino, SMTP Journaling, file system, etc).
2. Based on archiving type there will be other questions including state of archiving such as steady state (on-going incremental archive), backlog (initial archiving of documents) and journal archiving. Some questions may include:
 - Microsoft Exchange/Domino Mailbox
 - i. Number of mailboxes
 - 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
 - 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 is a useful document that 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 the Access Appliance with Enterprise Vault as a target storage platform simplifies management and support, minimizes costs, and improves control and visibility. This solution is rich with features such as WORM, replication, storage efficiencies, encryption, monitoring, auto support and integration with other Veritas products such as NetBackup and Data Insight.

References

- Enterprise Vault
 - Landing Page - https://www.veritas.com/support/en_US/article.100040095
 - Performance Guides - https://www.veritas.com/support/en_US/article.100000918
- Access Appliance 3340
 - 7.3.2 - https://sort.veritas.com/documents/doc_details/AAPP/7.3.2/Appliance%203340/ProductGuides/

- 7.4.2 - https://sort.veritas.com/documents/doc_details/AAPP/7.4.2/Appliance%203340/ProductGuides/
- 7.4.3 - https://sort.veritas.com/documents/doc_details/AAPP/7.4.3/Veritas%203340/Documentation/
- Veritas Data Insight
 - https://www.veritas.com/content/support/en_US/DocumentBrowsing.html?product=Data%20Insight

Appendix


This section describes the basic steps to configure Access Appliance for Enterprise Vault. This is only a sample configuration and readers are expected to refer to the Veritas Product Documentation for Veritas [Access Appliance](#) and [Enterprise Vault](#) for definitive and specific installation, administration and configuration details.

For the configuration examples in this section, it is assumed that the following have been pre-installed and configured:

- Enterprise Vault 12.3 on Windows Server 2012 R2
 - Vault Store Group (VSG1)
 - Vault store (VS1)
- Microsoft SQL Server on Windows Server 2012 R2
- Active Directory/DNS is enabled on Access Appliance.
- NTP Server is enabled on Access Appliance.

The Enterprise Vault Server and components are networked together with the Access Appliance running Access software version 7.3.2 as shown in Figure 14. This section will go over the creation and provisioning of the Access Appliance for Enterprise Vault and configuration of the Access Appliance as target vault stores for Enterprise Vault. A directory will be created on the server hosting Enterprise Vault and data in the directory is archived during the validation section.

Figure 12 - Enterprise Vault with Access Appliance Environment Example



Access Appliance Setup

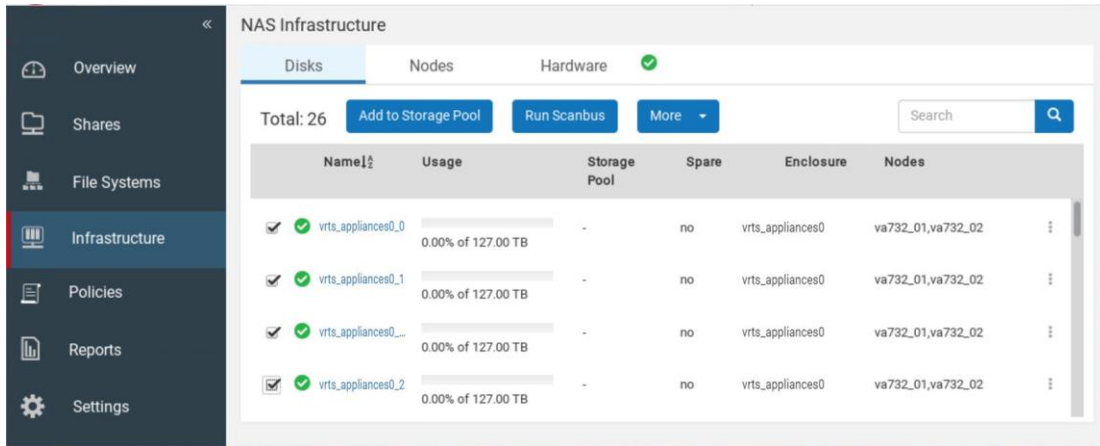
As previously mentioned, as a best practice it is not recommended to use the policies for EV. The high-level step to configure Access Appliance involves:

1. Create storage pool.
2. Create CIFS file system.
3. Set SMB/CIFS to clustering mode and enable the SMB/CIFS Server.
4. Provision storage for Enterprise Vault on Access Appliance
5. Enable WORM on file system (Optional).
6. Setup Access Appliance share as a Storage Target on Enterprise Vault.
7. Validate archival to Access Appliance


The next sections walk through the configuration of Access Appliance with Enterprise Vault.

Create Storage Pool


1. Click **Infrastructure** on left pane. **Check mark four disks** and click **Add to Storage Pool**.



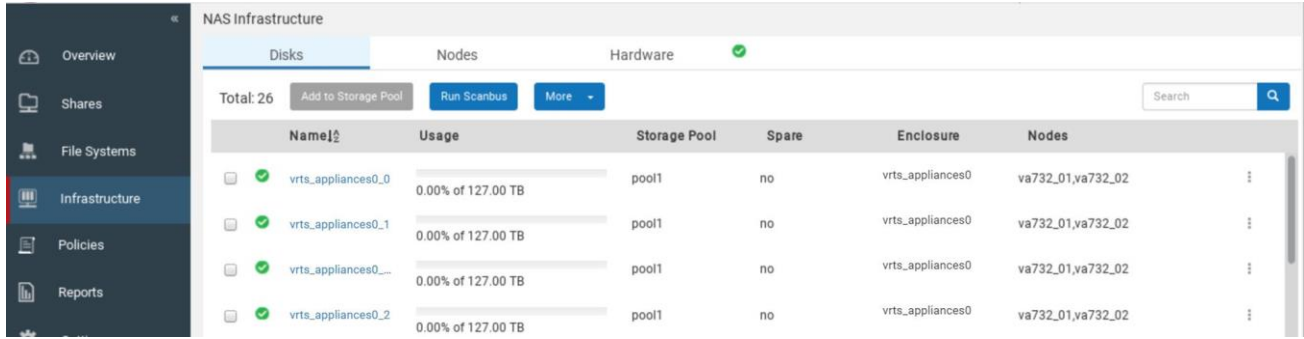
2. Select **add new storage pool**. Enter name **pool1**. Click **Next**.



3. Click clock icon on top to check the status of this activity. Click **Show All Recent Activities**. Wait for the operation to show success.



4. After success, pool1 is created.

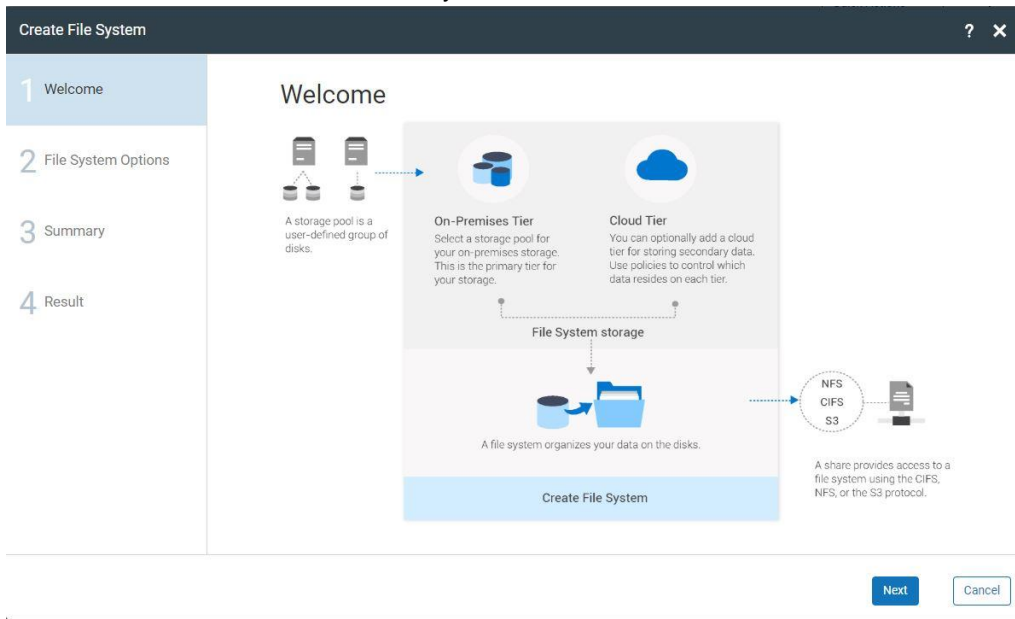


Creating File System

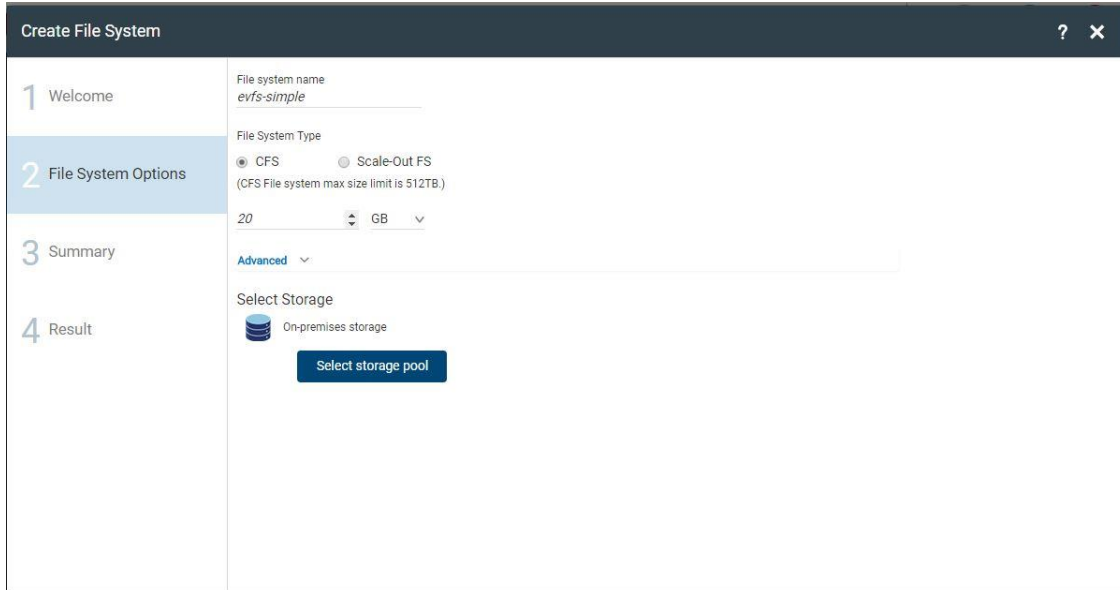
1. Click **Quick Actions** at top and select **Create File System**.




2. Follow the wizard to create the File System. Click **Next**.




- Set **name** of file system, file system **type** to CFS and **size** of file system with a **Simple** layout and **Block Size** (default: 8192). For this example, name is **evfs-simple**, size is **20 GB**. Click on **Select storage pool**.



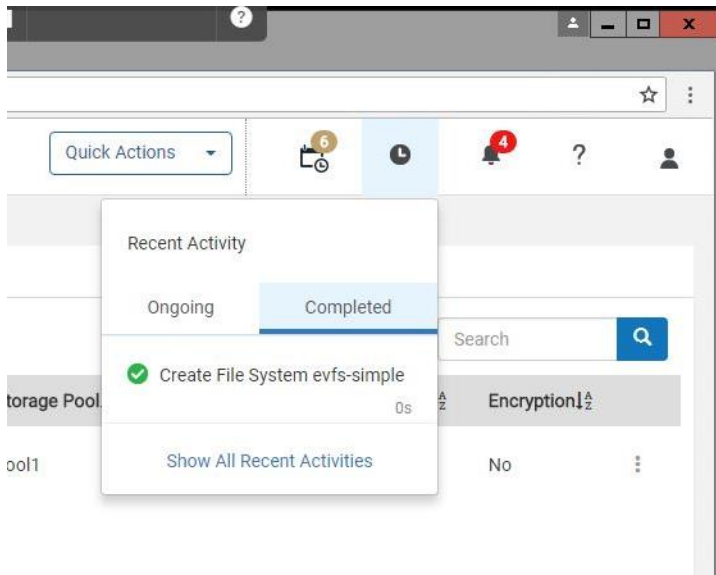
- Check mark the desired storage pool: **pool1**




5. Review Summary and click **Next** and **Finish**.



6. Click the **Clock icon** and check the activity of creation.



7. Click **File Systems** on left pane and validate that the **evfs-simple** file system has been created.



Set SMB/CIFS Clustering Mode and Enable the SMB/CIFS Server

1. Before enabling the SMB/CIFS services, set the mode of the SMB/CIFS share to CTDB. Enter the Access CLISH via ssh. Setting the SMB/CIFS mode involves the following steps to be done from the CIFS prompt:
 - a. **server status** – check the server status. If status indicates that the clustering mode is already in CTDB mode, then resume to step 2.
 - b. **server stop** - stop the SMB/CIFS service.
 - c. **set clustering_mode ctdb** - set the clustering mode to CTDB
 - d. **server start** - start the SMB/CIFS service.
 - e. **server status** - check the server status to verify that clustering mode has been set to ctdb.

Sample output is shown below.

```
va732> cifs
Entering CIFS share mode...
ajva.CIFS> server status
CIFS Status on va732_01 : ONLINE
CIFS Status on va732_02 : ONLINE

Homedirfs           :
Security             : user
Clustering Mode     : normal

va732.CIFS> server stop
Stopping CIFS Server.....Success.

va732.CIFS> set clustering_mode ctdb
Global option updated. Note: Restart the CIFS server.

va732.CIFS> server start
Uninstalling 'normal' Clustering Mode....Success.
Installing 'ctdb' Clustering Mode.....Success.

Starting CIFS Server.....Success.
va732.CIFS> server status
CIFS Status on va732_01 : ONLINE
CIFS Status on va732_02 : ONLINE

Homedirfs           :
Security             : user
Clustering Mode     : ctdb
```


Enable Worm on File System (Optional)

1. Log on to the Access Appliance CLISH and enable WORM on files ystem created using “**storage fs worm set evfs-simple**” command. Check that WORM has been set by entering “storage fs list evfs-simple”.

```
va732> storage fs worm set evfs-simple
ACCESS fs SUCCESS V-493-10-2189 Enabled WORM for evfs-simple file
system.

va732> storage fs list evfs-simple
General Info:
=====
Block Size:          8192 Bytes
Version:             Version 13
Volume Encrypted:   No
Max IOPS:           0
va732_01:           online
va732_02:           online

Primary Tier
=====
Size:                20.00G
Use%:                3%
Used:                257.58M
Layout:              simple
Mirrors:             -
Columns:             -
Stripe Unit:        0.00 K
Meta Data:           metaOk
FastResync:          Disabled

1. Mirror 01:
List of pools:       pool1
List of disks:       vrts_appliances0_1

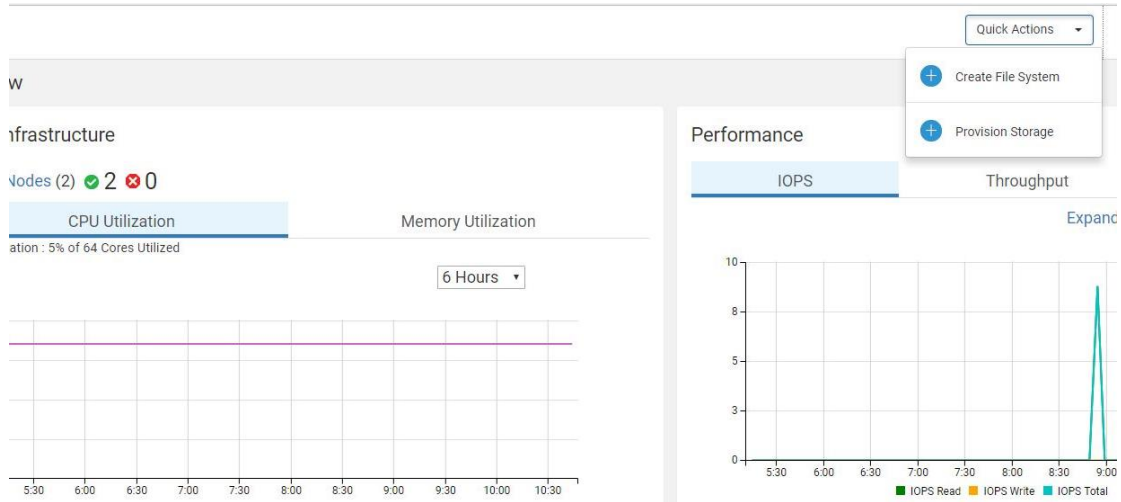
FS Type:             Normal

Defrag Status:       Not Running
Fullfsck Status:     Not Running
Resync Status:       Not Running
Rollsync Status:     Not Running
Relayout Status:     Not Running

WORM Enabled:       Yes
```

Provision Storage for Enterprise Vault on Access Appliance


1. Click **Quick Actions**, select **Provision Storage** and follow the wizard.




2. Select the Provisioning type as **Storage for Enterprise Vault**.

The screenshot shows the 'Provision Storage' wizard. The title bar reads 'Provision Storage' with a help icon and a close icon. The wizard has two steps: '1 Provisioning Option' and '2 Storage Options'. Under '1 Provisioning Option', there is a section titled 'Select Provisioning type' with four radio button options: 'Storage for NFS', 'Storage for CIFS', 'S3 Storage for NetBackup', and 'Storage for Enterprise Vault'. The 'Storage for Enterprise Vault' option is selected. At the bottom right, there are 'Next' and 'Cancel' buttons.


3. Select **File System** as Storage Type and select **evfs-simple** as the file system to create the shares on.




4. Enter **name** of Share Name to be evstores. Select the **Access Type** and **Export Options**. In this example, the settings are **rw**, **owner is lab\levsvc** and **file system mode 1777**. Click on **Set** to set the export options.



5. Review the options and click **Next**.



6. Click on **Shares** on the left pane and check that the evstores share has been created successfully.




Setup Access Appliance Share as a Storage Target on Enterprise Vault

1. On the Access Appliance CLISH, execute the “**network ip addr show**” command and use one of the **virtual IPs** for the Enterprise Vault configuration of vault store partition.


```
va732> network ip addr show
```

IP	Netmask/Prefix	Device	Node	Type	Status
--	-----	-----	----	----	-----
192.168.10.125	255.255.255.0	pubeth0	va732_01	Physical	
192.168.10.126	255.255.255.0	pubeth1	va732_01	Physical	
192.168.10.127	255.255.255.0	pubeth0	va732_02	Physical	
192.168.10.128	255.255.255.0	pubeth1	va732_02	Physical	
192.168.10.130	255.255.255.0	pubeth0	va732_01	Virtual	ONLINE
(Con IP)					
192.168.10.140	255.255.255.0	pubeth0	va732_02	Virtual	ONLINE
192.168.10.141	255.255.255.0	pubeth0	va732_01	Virtual	ONLINE
192.168.10.142	255.255.255.0	pubeth1	va732_02	Virtual	ONLINE
192.168.10.143	255.255.255.0	pubeth1	va732_01	Virtual	ONLINE

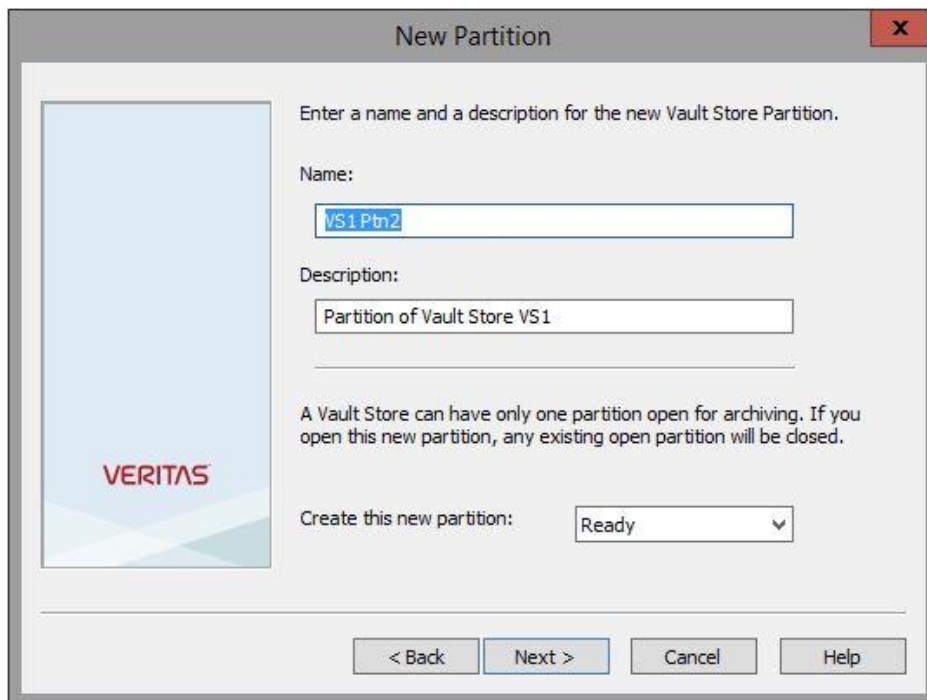
2. On the Enterprise Vault server, connect to the **SMB/CIFS share** ([\\<IP address>evstores](#)) on the Access Appliance using the virtual IP noted in previous step and **create a directory, FileArchive**.



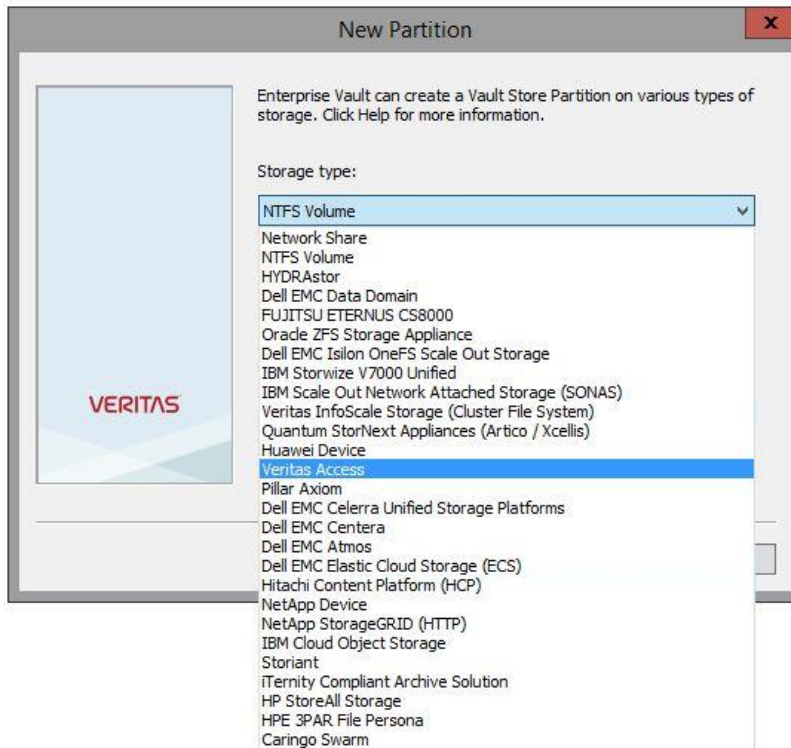
3. **Start** the Enterprise Vault management console and traverse to Directory → Site (i.e. Lab) → Vault Store Groups → (VSG1) → VS1 → Partitions. **Right-click and click New → Partition** to start the partition creation wizard. Users can also create Smart Partition similarly if classification tags are desired. In this example, a regular partition is created.



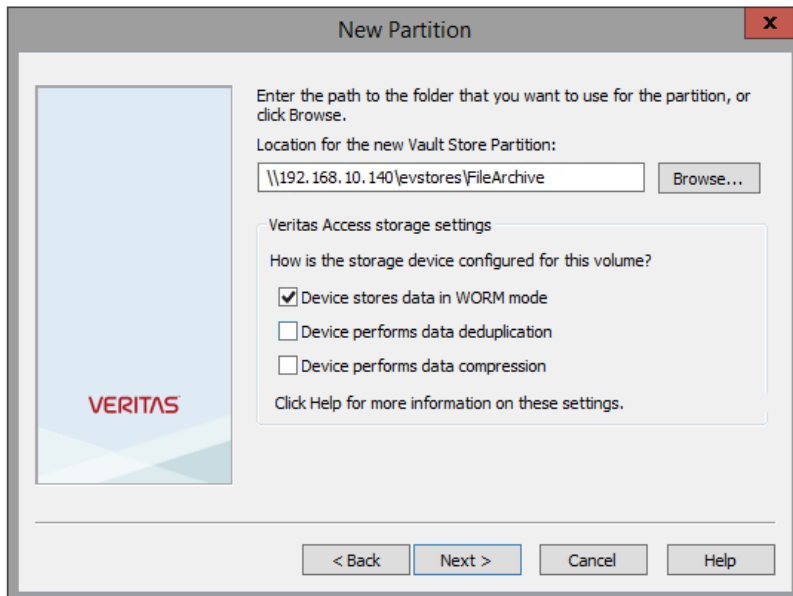
4. Enter **name** of new partition as VS1 Ptn2, and **description** as Partition of Vault Store VS1.




5. Select **Veritas Access** as the storage type to use.



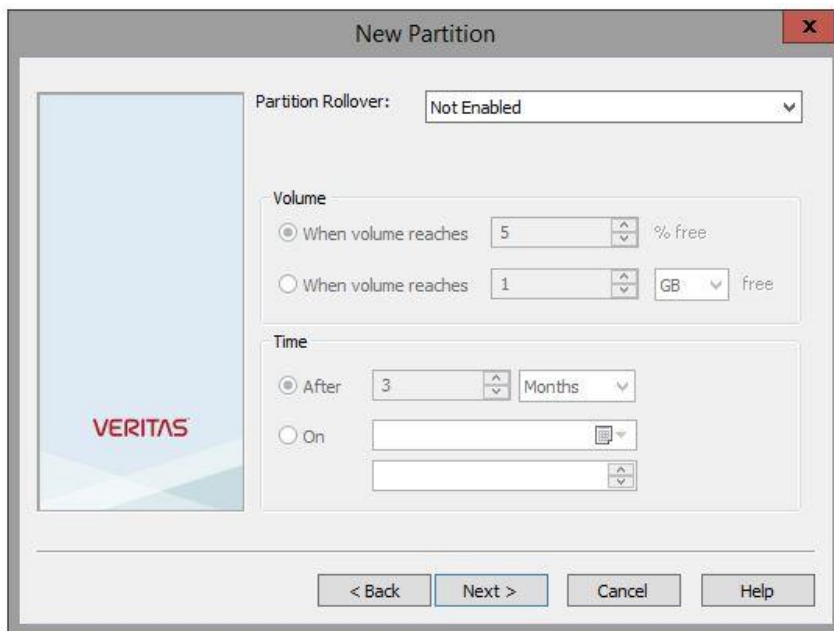
6. Enter the **path** of the Veritas Access **shares**. In this example, Exchange Archive is a directory that was created in step 2. Place a checkmark on **Device stores data in WORM mode** (optional) since the Veritas Access file system for the share has been WORM enabled.



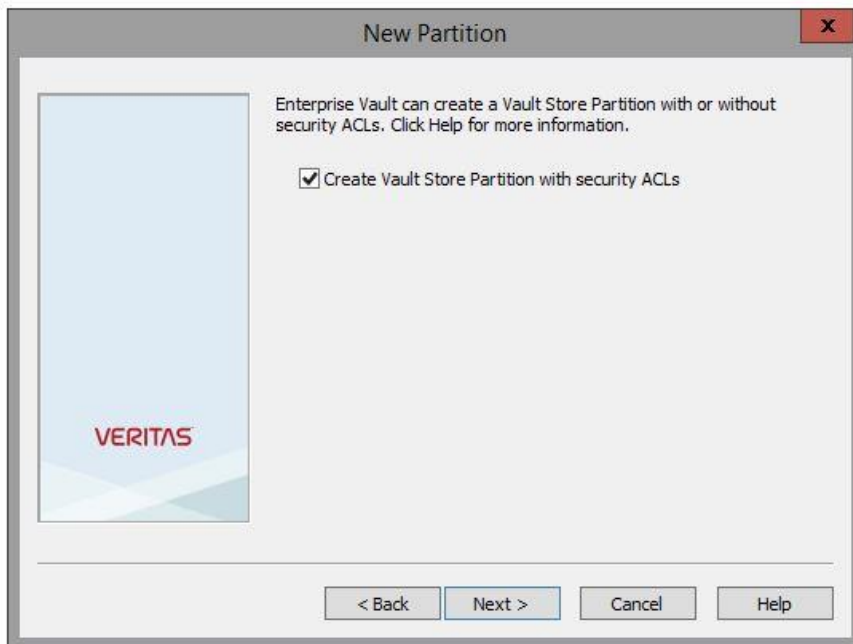
- Click **Run Test** to test connectivity.




- Use the **default Partition Rollover**. In this example, it is set to **Not Enabled**.




9. Use the **default setting** to create vault store partition with security ACLs.




10. Use the **default archive attribute**.



11. Review the configuration and **click Finish**.




12. Validate that the **partition VS1 Ptn2** is created.




Validation of Access Appliance as a Storage Target for Enterprise Vault


1. Change the VS1 Ptn2 (Veritas Access) partition created from previous section to **Open** such that the archives will be directed to this partition. If there is another partition that is open, close those first.



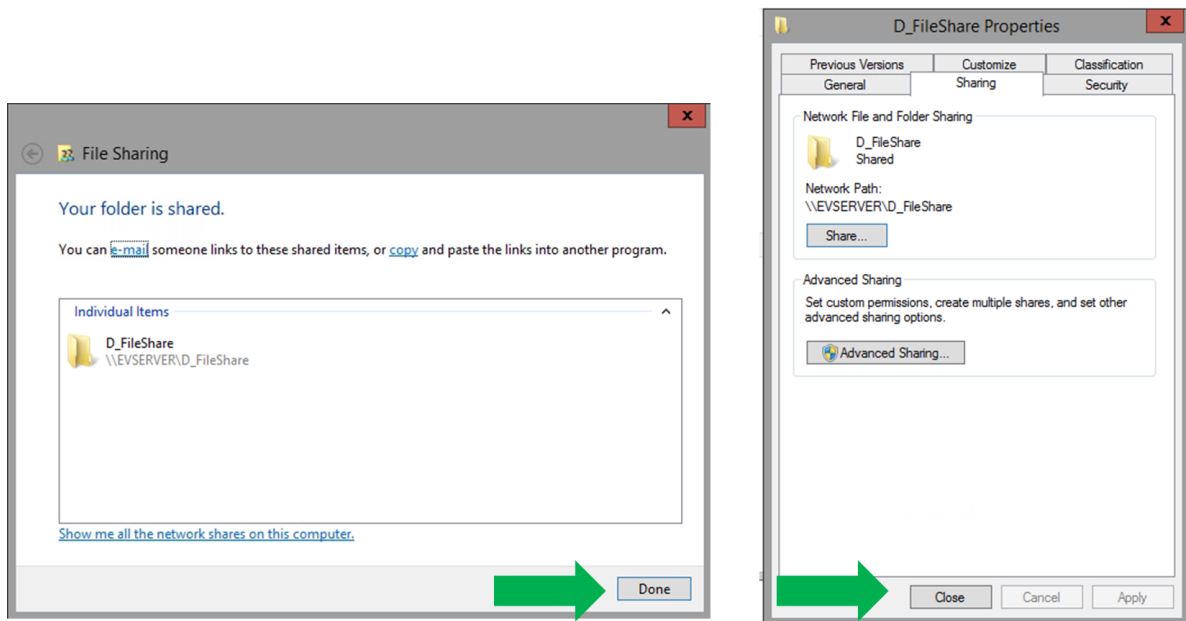
2. Create a file share to archive. **Create a directory** on server.




3. **Right-click** and select the **properties** of the D_Share folder. Click on **Share** tab. Select the **people** to share with and click **Share**.



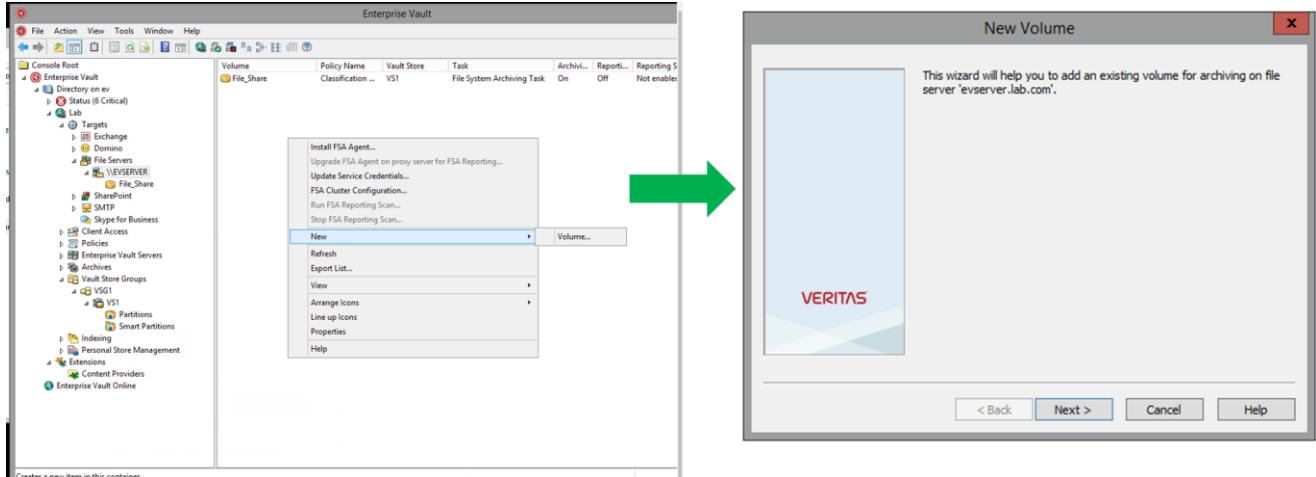
4. Click **Done** and **Close**.




5. **Copy some files into the D_Share folder to archive.**



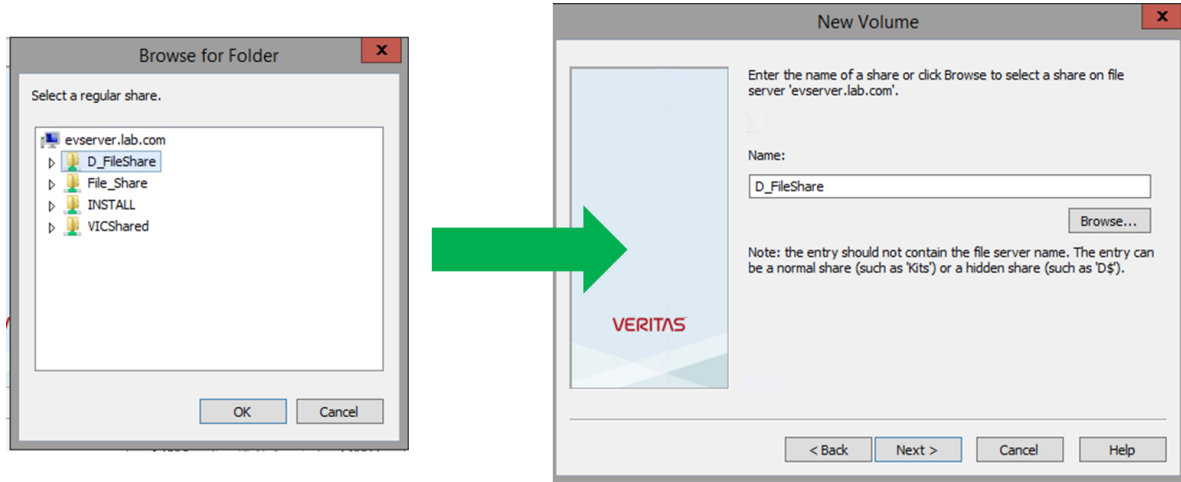
6. **Add the file share (directory) created, as a target on Enterprise Vault. On Enterprise Vault, expand Target → File Servers. Click on EVSERVER and on right pane, do a right-click and select New → Volume. Click Next.**



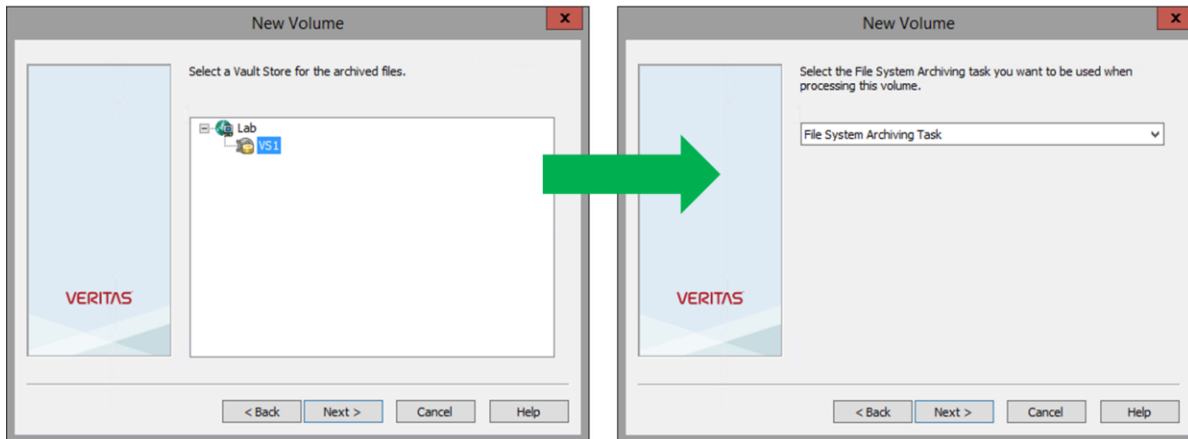
7. Click on **Browse** and select **Regular** on share type.



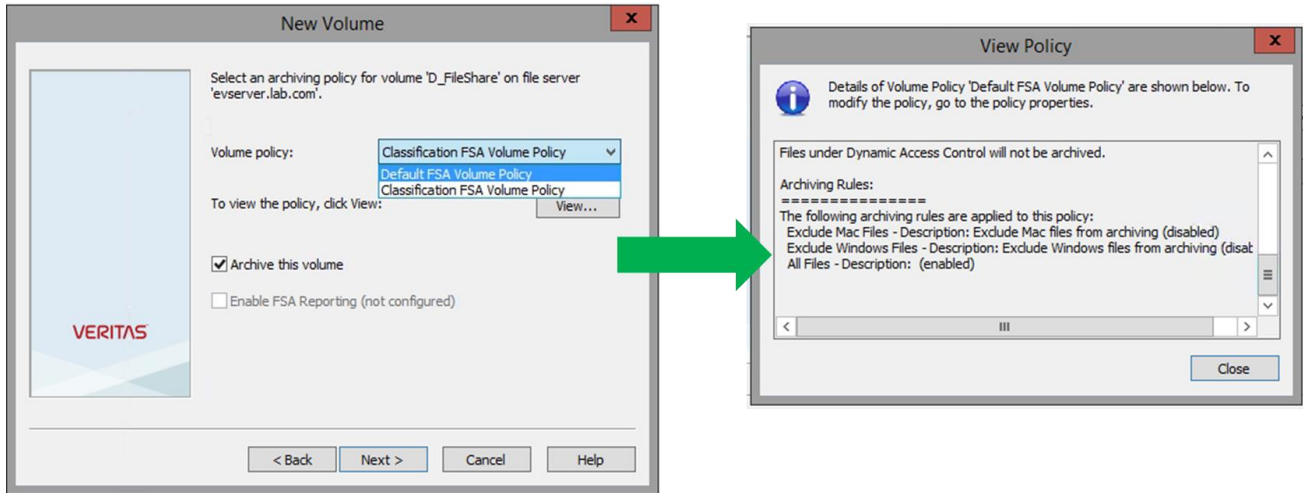
8. Select **D_FileShare** on next pane. Click **Next**.




9. Select **VS1** as the Vault Group which contains the Veritas Access partition and click **OK**. Then, select **File System Archiving Task** to do the archive. Click **Next**.




- Select the Volume Policy to be **Default FSA Volume Policy**. Click on **View** and the policy indicates the Archive Rules. In this example, all files are archived except for certain windows and MAC files. Click **Next, Finish** and **Close**.



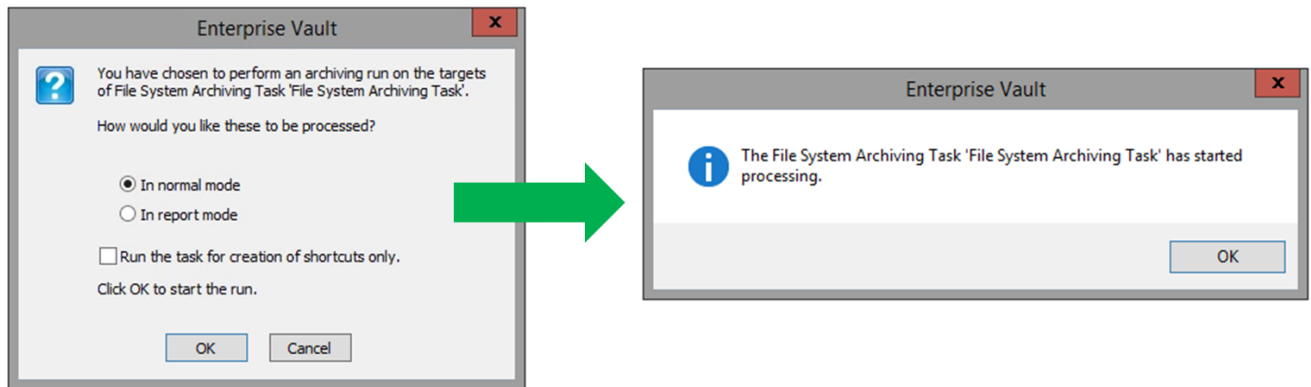
- Click and expand the **Enterprise Vault Servers → *.lab.com → Tasks**. Right-click on the right pane and select **Start** to start the File System Archiving Task.




12. After the File System Archiving Task is in Running status, right-click on task again and select **Run Now** to start the file system archiving task on the target file shares (e.g. D_FileShare).



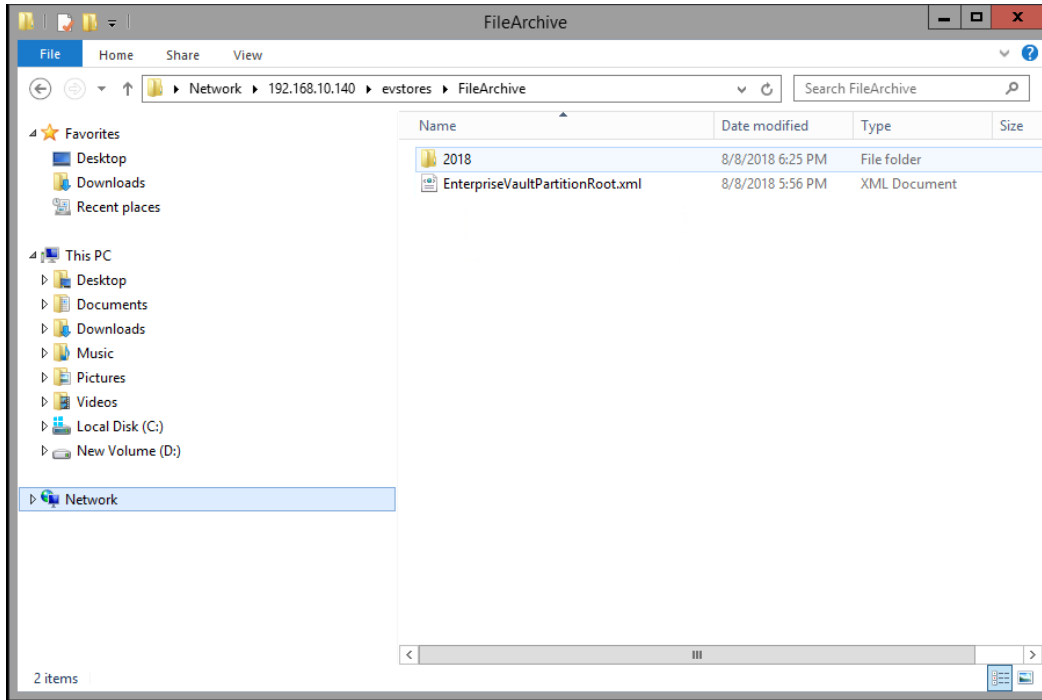
13. On next pane, select **In normal mode** and click **OK**.




14. The status will change from **Running** to **Processing**.



- After the status moves from processing to running again. Go to the Veritas Access SMB/CIFS share (**\\<IP address>levstores\FileArchive**) and check that the file share has been archived. There will be directory that would contain the archives.



- Traversing the EV archives directory, file types *.DVS, *.DVSCC, and *.DVSSP can be seen.



Setup Access Appliance as a Secondary Storage Target for Enterprise Vault Collections

Provision Storage on the Access Appliance


1. Similarly, as described in the previous sections, create a file system (**ev-secondary**) and provision storage (**evfs-secondcopy – CIFS**) on the Access Appliance for Enterprise Vault.



The screenshot shows the 'Shares' page in the Veritas Access Appliance interface. It displays a table of provisioned storage entries. The table has columns for Path, Share Name, Access Protocol, File System, and File System Status. There are three entries listed, all with a status of 'online'.

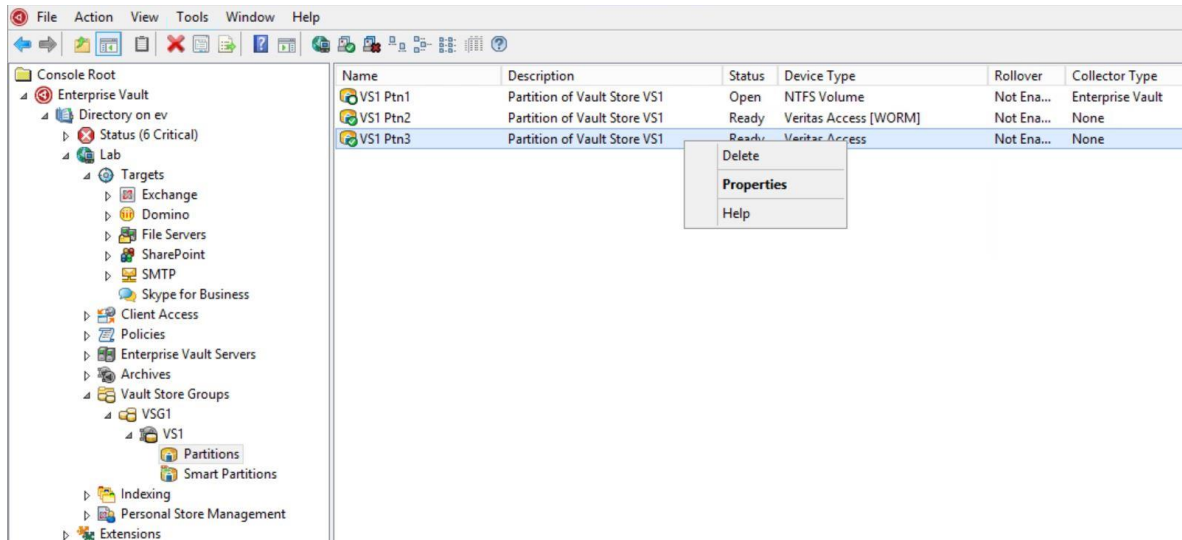
Path	Share Name	Access Protocol	File System	File System Status
/vx/ev-secondary	evfsa-secondcopy	CIFS	ev-secondary	online
/vx/ev-fsa	evfsa-store	CIFS	ev-fsa	online
/vx/evfs-simple	evstores	CIFS	evfs-simple	online

2. Create a directory (**FSA-SecondCopy**) on the host where Enterprise Vault is running for the Access Appliance share that acts as the secondary storage target (<\\<IP Address>evfsa-secondcopy>).

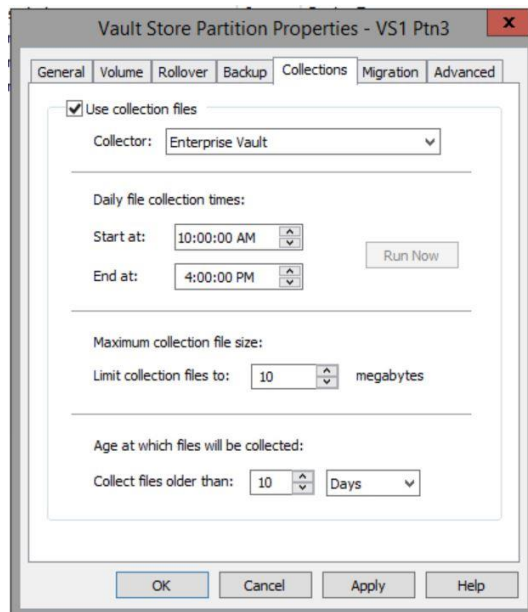


Configure EV Migrator with Access Appliance as a Secondary Storage Target.


1. On Enterprise Vault, **right-click on vault store partition (VS1 Ptn3)** and select **Properties**. In this example, vault store partition named VS1 Ptn3 is a non-WORM Veritas Access device type that will be migrated to another CIFS share on the Access Appliance.



2. Collections would need to be first enabled on vault store partition such that the collections are migrated to the secondary storage target. Thus, go to **Collections Tab** and **check mark Use collection files** and select **Collector to be Enterprise Vault**. Set the **start** and **end** of the daily collection files and the **maximum collection file size** and the **age** at which files will be collected. Click **Apply** and then select **Yes** to use the file collection software.



3. Click on **Migration** tab. Check mark the box next to **Migrate files** and select **Enterprise Vault** as Migrator. Click **OK** when on the message indicating that some secondary storage affects the response time of Enterprise Vault. Enter the **age** at which collection files will be migrated and **age** at which migrated collection files will be removed from primary storage. Then enter the secondary storage location: [\\<IP address>\evfsa-secondcopy\FSA-SecondCopy](#). Click **Apply** and then **Yes** to use file migrator software.



Configuration of Episodic Replication Notes

For configuration of episodic replication, please follow the procedure outlined in the [Access Administrator Guide](#). For Access 7.3.2 release, use the Access CLISH to configure the replication for both source and destination Access cluster.

NOTE the following:

- Both source and destination Access cluster need to be on the same version.
- An un-used virtual IP is required on both source and destination Access cluster.
- Both the source and destination would need to export and import keys in order to establish communication. If NOT using the URL path to specify the key location during the import keys, enter the command “episodic config import keys” first and hit carriage return. Then you are prompted to enter the key. Refer to below figure for an example where the Access cluster destination is doing an import of the key generated from the source Access cluster:

```
tmeaccess1.Replication> episodic config import_keys
```

```
Enter replication key of remote cluster::ssh-rsa
```

```
AAAAB3NzaC1yc2EAAAADAQABAAQDdKDr2vjDO0z/CPre74rS6Io1KmoY0gs+
4Cx6Fb4UXCshk397EKV1gGBPAef0s32IT64pcpi69Yp0Qz6xgiN8cMCNPJGf5ow
1M+Dk48XjXa60HI8DoNn9PknC2SR++GyEC1AVj2mkMe2q8P2sUvwBwDPy2z/YDg
JMVgsXjqxL0bGeTe5jXm13peWkesLTnku5Z3AdJdDi9PiH8hkLZdNQZmpmi7hU
VNpYv6nPNROxLCCeh1EP//l+JlUCsjOVoxz1N3KOLNQvW4sZR1nVoC14ndZp5NH
eC/79VN7iiFKZIk7XHupr7uruekgFdUu7d7cRy7S9mVegZxorWQzMKv4b
root@va732_01::va732_192.168.10.130
```

- d) In Access Appliance, the “evpsn” flag is set when creating an episodic replication job.
- e) In Enterprise Vault, set the vault store Partitions properties to check for a trigger file. Set the partition scan interval.

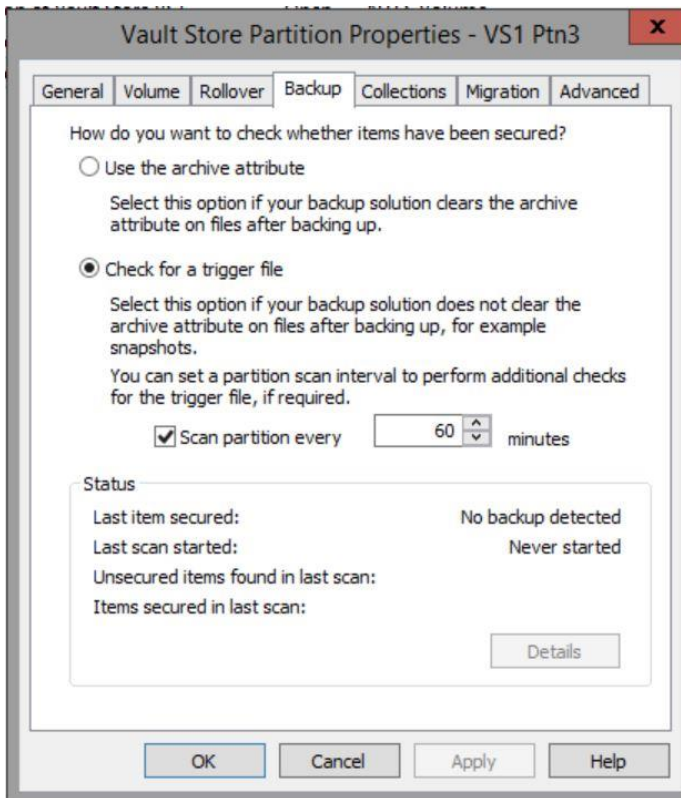


Table of Figures

Figure 1 - Replication Data Flow	5
Figure 2 - Data Insight with Enterprise Vault and Access Appliance.....	7
Figure 3 - Sample View of Data Insight Console - Inactive Files in Group Shares.....	7
Figure 4- MyAppliance Portal View	8
Figure 5 - Access Appliance with Enterprise Vault Solution High-Level Architecture.....	8
Figure 6- Enterprise Vault Main Components	9
Figure 7 - Access Appliance Rack Units	10
Figure 8- Sample View of EV Archived Data on Access Appliance	12
Figure 9 - Archival Data Flow	13
Figure 10 – Retrieval Data Flow.....	14
Figure 11 - Migration of Enterprise Vault Collections to the Access Appliance for Secondary Storage.....	14
Figure 12 - Enterprise Vault with Access Appliance Environment Example	19

ABOUT VERITAS

Veritas Technologies is a global leader in data protection and availability. Over 50,000 enterprises—including 99 of the Fortune 100—rely on us to abstract IT complexity and simplify data management. 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 supports more than 500 data sources and over 150 storage targets, including 60 clouds. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

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

For specific country offices
 and contact numbers,
 please visit our website.

