Veritas Access Appliance 8.0 Solutions Guide for Enterprise Vault

January 2024

VERITAS™

Access Appliance Solutions Guide for Enterprise Vault

Last updated: 2024-01-09

Legal Notice

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

https://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

https://www.veritas.com/content/support/en_US/dpp.Appliances.html

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

APPL.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

https://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1

Introduction

This chapter includes the following topics:

About this document

This document describes how Access Appliance 8.0 can be configured as archival storage with Enterprise Vault. Veritas Enterprise Vault helps automate retention management, classification and supervision, while simplifying search and eDiscovery of unstructured data.

About Access Appliance as archival storage for Enterprise Vault

Access Appliance can be used as both primary storage and secondary storage for Enterprise Vault.

- Primary storage: Enterprise Vault can archive data directly to the storage system vault store partition.
- Secondary storage: Enterprise Vault can migrate collections of files from an NTFS or network share vault store partition to a secondary storage location.

Figure 1-1 shows how Access Appliance works as a primary archive datastore for Enterprise Vault.

Figure 1-1 Access Appliance as a primary archive datastore for Enterprise Vault

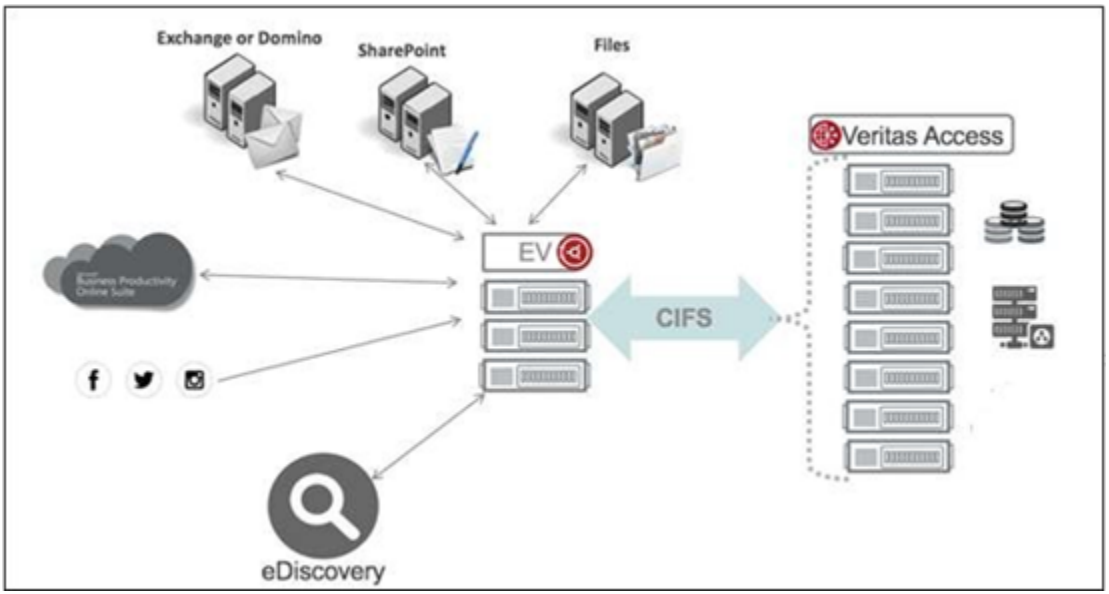

Access Appliance versions certified by Enterprise Vault

Access Appliance is certified for archival storage by Enterprise Vault.

Table 1-1 Access Appliance versions certified by Enterprise Vault

Access Appliance version	Enterprise Vault 12.x		Enterprise Vault 14.x	
	Non-WORM	WORM	Non-WORM	WORM
7.3.2	Y	Y	N	N
7.4.2	Y	Y	N	N
7.4 .3	Y	Y	Y	Y
8.0	Y	Y	Y	Y

Chapter 2

System Requirements

This chapter includes the following topics:

Server roles

Figure 2-1 shows the roles of different servers in an Enterprise Vault configuration.

Figure 2-1

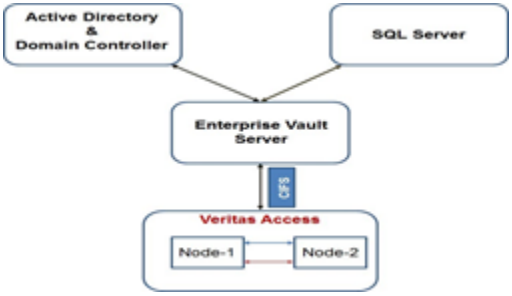

Table 2-1 gives the list of required servers and their use.

Table 2-1

Server	Use
Active Directory and Domain Controller	Required for user authentication for Enterprise Vault.

Table 2-1 *(continued)*

Server	Use
Enterprise Vault server	Required to configure Enterprise Vault services.
SQL server	Required by Enterprise Vault server to host its back-end database.
Access Appliance	Required as a storage system for archival by Enterprise Vault.

Hardware requirements

Four physical or virtual servers are required to host the following:

1. Active Directory and Domain Controller server

2. SQL server

The SQL servers should have the following minimum specifications:

- Number of CPUs: 8 (recommended)
- Power of CPUs: 2.8 GHz or more
- Memory: 16 GB (recommended)

The SQL server requires three fast local partitions (composed of one or more disks) to hold the following:

- SQL Server installation (OS Drive)
- SQL Server Data
- SQL Server Logs

3. Enterprise Vault server

The Enterprise Vault servers should have the following minimum specifications:

- Number of CPUs: 8 (recommended)
- Power of CPUs: 2.8 GHz or more
- Memory: 16 GB (recommended)

Each Enterprise Vault server requires three fast local partitions (composed of one or more disks) to hold the following:

Table 2-2

Disk use	Disk size (minimum)
OS Drive: Enterprise Vault installation Enterprise Vault Cache Enterprise Vault Storage Queue	100 GB
Enterprise Vault Indexes	75 GB
Target data for archiving This disk is only required for EVSCEV01	75 GB

4. Access Appliance cluster

See the *Access Appliance Initial Configuration Guide* for the system requirements.

Software requirements

Table 2-3 shows the software requirements for the involved software components are as follows:

Table 2-3 Software requirements

Component	Requirements
Enterprise Vault	See the *Enterprise Vault™ Installing and Configuring Guide* on SORT for the software requirements to configure Enterprise Vault server.

Table 2-4 shows the Enterprise Vault certification details for Access Appliance as a primary archival storage.

Table 2-4 Certification details

	Enterprise Vault 12	Enterprise Vault 14
Access Appliance 7.3.2 and later versions	Non-WORM, WORM	-
Access Appliance 7.4.2 and later versions	Non-WORM, WORM	-
Access Appliance 7.4.3 and later versions	-	Non-WORM, WORM

Table 2-4 Certification details *(continued)*

	Enterprise Vault 12	**Enterprise Vault 14**
Access Appliance 8.0 and later versions	Non-WORM, WORM	Non-WORM, WORM

See the *Access Appliance Release Notes* for operating system compatibility list.

Chapter 3

Installing and configuring Enterprise Vault with Access Appliance

This chapter includes the following topics:

Enterprise Vault deployment

Before you install Enterprise Vault, the Enterprise Vault Deployment scanner should be run on each of the intended Enterprise Vault servers.

If you want to configure partitions in both IPv6 and IPv4 modes, the following network settings are required on Enterprise Vault and the Active Directory (AD) server.

Network settings> **NIC Properties**> **Enable IPv6 and Enable IPv4**.

For more information, see the *Enterprise Vault™ Installing and Configuring Guide* on SORT.

Access Appliance deployment

For deploying Access Appliance, see the *Access Appliance Initial Configuration Guide*.

Chapter 4

Access Appliance features for Enterprise Vault archival storage

This chapter includes the following topics:

Write-Once-Read-Many support

When a file is committed as Write-Once-Read-Many (WORM), the data in the file can be read but cannot be altered. The retention time for a WORM file specifies the time period for which the file must be retained after it is committed to WORM storage. The file cannot be deleted till the retention period expires. Once the retention time period has expired, the storage system allows the deletion of the file.

The retention period cannot be reduced once it is set.

The WORM property is set at the file level. The file can be committed or enabled as WORM on a file system created with 'WORM' support. A new option has been added in the `Storage> fs create` command to enable WORM support.

You can find out whether a given file system is WORM-enabled or not using the `Storage> fs worm get <fsname>` command.

See the `Storage> fs list` man page for more details.

A new option has been added in the Access Appliance GUI as a policy to enable WORM support. You can find out whether a specific file system is WORM-enabled or not using the **File Systems** tab in the Access Appliance GUI.

While creating a partition in the Enterprise Vault Administration Console, select **Device stores data in WORM mode**.

Figure 4-1 Creating a new partition

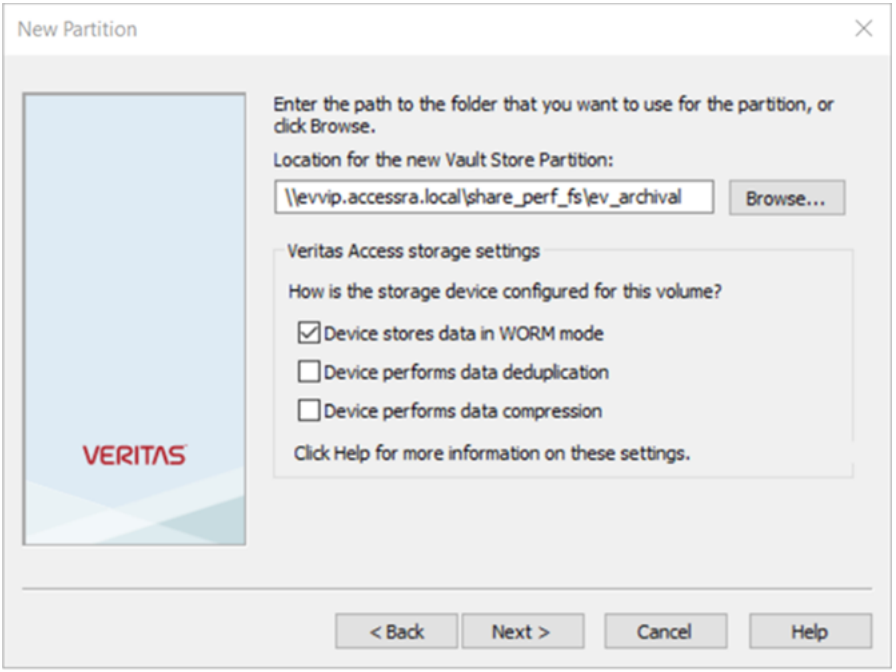

Enterprise Vault archiving processes check the target servers for items to archive at scheduled times. When an item is archived, it is automatically assigned a retention category, which defines how long it must be kept in the archives (retention time). The administrator can define different retention categories for different types of data. Enterprise Vault monitors the archives and deletes items when the retention period expires.

Enterprise Vault sets a retention time on individual files. It sets the retention time by setting the access time of the file to a date in the future. For WORM-committed files, the access time of the file indicates the retention time.

Enterprise Vault uses a CIFS share, which exports the WORM-supported file system to archive data in WORM mode. Create the CIFS share with the `full_acl` option.

Any file system and CIFS share created in this way can be used as WORM storage for Enterprise Vault's archival on WORM-enabled storage.

Note: When using Access Appliance file systems as WORM-enabled storage for Enterprise Vault, once WORM retention period is set on archived items from Enterprise Vault, it can be extended for the associated files on the Access Appliance file systems only through the Veritas Access command-line interface.

Partition Secure Notification

You can configure Enterprise Vault to retain original items until the vault store partition in which they are archived has been backed up. During the time between archival and removal, the original items are treated as safety copies by Enterprise Vault. A Partition Secure Notification (PSN) file is created at a source partition after the successful backup of the partition at the remote site. When the vault store partition has been backed up, Enterprise Vault removes the safety copies and creates placeholders.

Veritas replication is used to take a backup of a single partition. The Partition Secure Notification feature is enabled in the Access Appliance `replication episodic job create` interface when an optional argument called `evpsn` is enabled.

Figure 4-2 shows the Enterprise Vault store settings when safety copies are removed after the partition has been backed up.

Figure 4-2

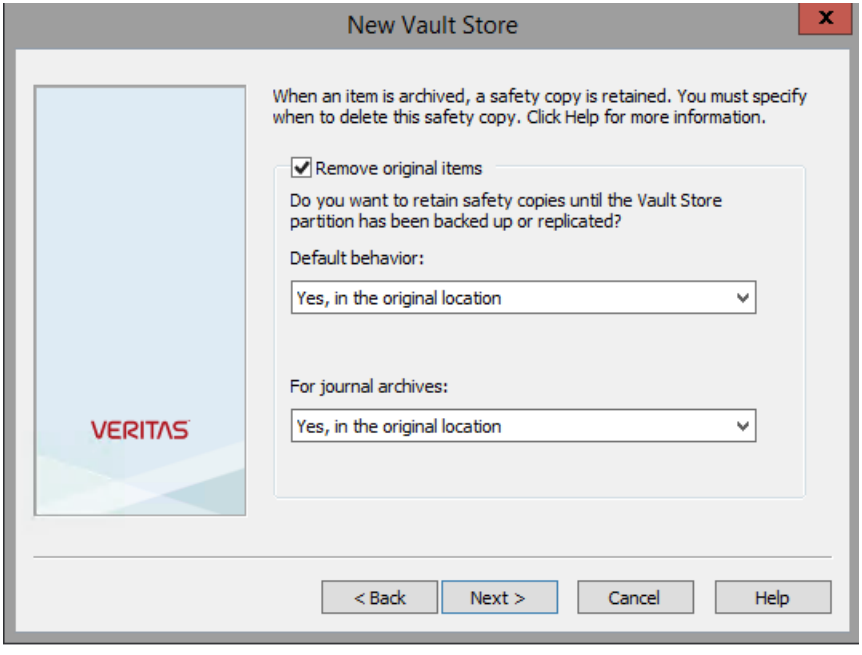

Enterprise Vault keeps the original items until the partition that contains the archived items has been backed up if **Yes, in the original location** option has been selected.

Figure 4-3 shows the Enterprise Vault store setting when a partition is created.

Figure 4-3

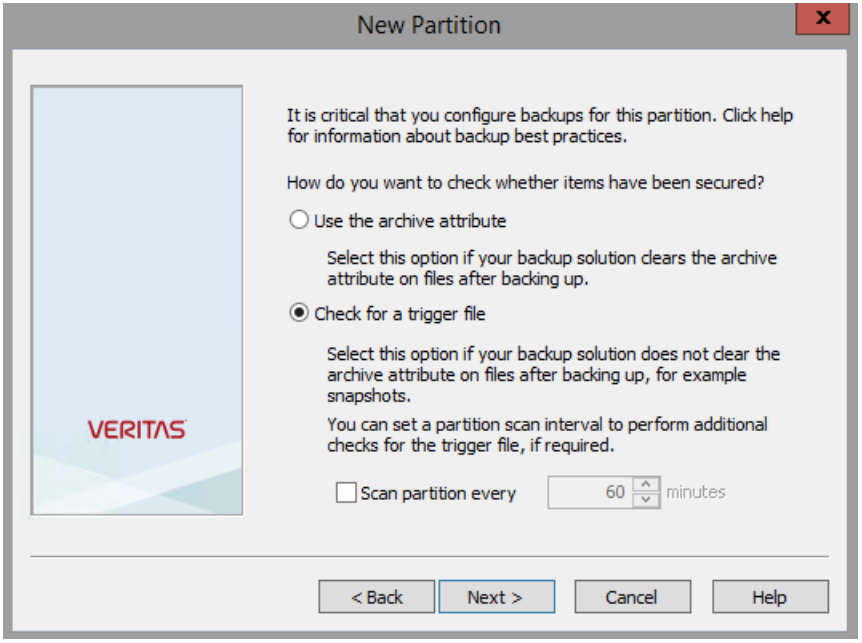

If the **Check for a trigger file** option is selected, Enterprise Vault refers to the PSN file to confirm that the archived data is secure. Accordingly, Enterprise Vault deletes safety copies and creates placeholders.

When the Enterprise Vault store is configured to delete the safety copies, it is important to configure how Enterprise Vault checks to ensure that the partition has been secured. Enterprise Vault checks for a trigger file when its Storage service starts and the backup mode is cleared.

See the `replication episodic job create` man page for more details.

See the *Configuring replication* chapter in the *Access Appliance Administration Guide* for more details on replication.

Chapter 5

Access Appliance archival policy configuration for Enterprise Vault

This chapter includes the following topics:

Configuring CIFS for the Active Directory domain mode

You have to configure Active Directory (AD) using the Access Appliance GUI.

To configure Active Directory

- Click **Settings** > **Directory service management** > **Active Directory management** to configure AD.

Access Appliance GUI policies for archival storage

The Access Appliance GUI policies provide quick and easy configuration of archival storage for Enterprise Vault. The following policies are supported:

WORM

This policy provides Write Once Read Many (WORM) storage and mirrors data across two devices, thus protecting against devices failures.

The key features of this archival policy are:

- Fault tolerance
- WORM

Prerequisites for this archival policy:

- Storage pool should be configured.

Non-WORM

This policy protects data against device failures by creating mirrored data across two devices.

The key features of this archival policy are:

- Fault tolerance

Prerequisites for this archival policy:

- Storage pool should be configured.

Non-WORM with episodic replication

This policy protects the data against device, node, and site failures. It replicates data across two clusters. Veritas recommends that these clusters should be in different data centers. This policy mirrors data across two devices and protects against device failures.

The key features of this archival policy are:

- Fault tolerance
- Replication

Prerequisites for this archival policy:

- Storage pool should be configured.
- Replication link should be set up.

WORM with episodic replication

This policy provides Write Once Read Many (WORM) storage and protects data against device, node, and site failures. It replicates data across two clusters. Veritas recommends that these clusters should be in different data centers. This policy mirrors data across two devices and protects against device failure.

The key features of this archival policy are:

- Fault tolerance
- WORM
- Replication

Prerequisites for this archival policy:

- Storage pool should be configured.

Note: To ensure seamless disaster recovery for vault stores, it is important to collect simultaneously consistent point-in-time copies of the Enterprise Vault partitions, Enterprise Vault database, and associated Enterprise Vault indices on the disaster recovery site regularly. If simultaneously consistent point-in-time copies are not collected regularly for vault stores at the disaster recovery site, the recovery points and the recovery times for those vault stores may be adversely affected.

Prerequisites for this archival policy:

- Storage pool should be configured.

Prerequisites for this archival policy:

- Storage pool should be configured.

Note: For Enterprise Vault PSN to show expected behavior with replication, it is recommended to configure replication from the Access Appliance GUI.

Veritas recommends that every vault store that is protected against disasters using replication should have backup mode set for it before replication is started for partitions in that vault store. The set of simultaneously consistent point-in-time copies should then be collected for that vault store. Once the sets of simultaneously consistent point-in-time copies have been collected, backup mode can be unset for that vault store resuming archival in the vault store.

Configuring the replication job

You have to configure replication to activate replication-related policies. As part of this configuration, you have to set up the replication VIP and replication link.

Archival policies use the Veritas replication feature.

See the *Access Appliance Administration Guide* for more information on configuring replication.

To configure replication

1. Log on to the Access Appliance GUI.
2. Click **Settings** -> **Replication Management**

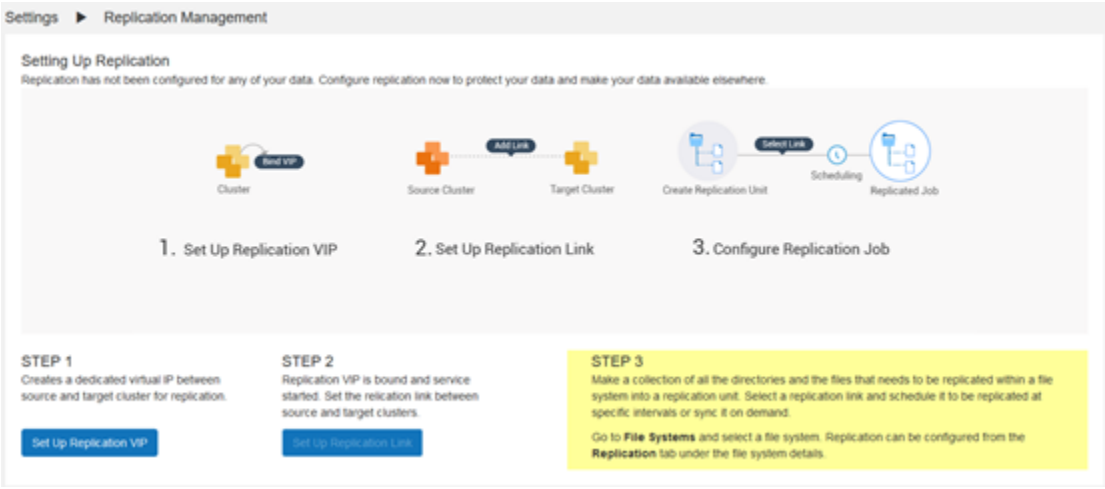

3 Enter the required information for setting up the replication VIP.

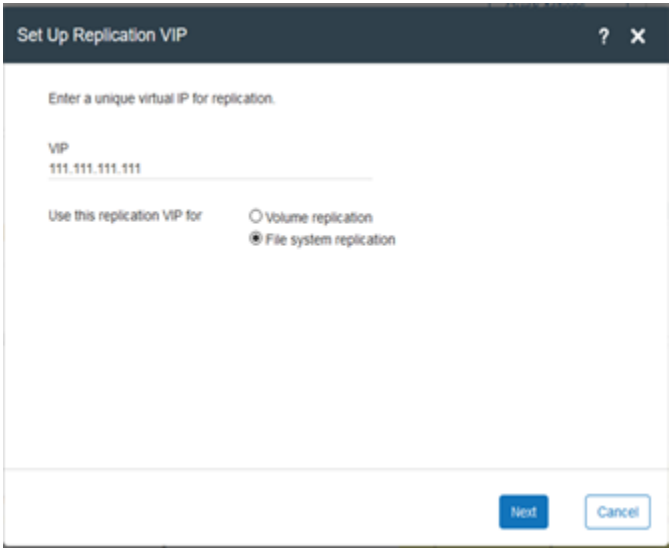

4 Enter the required information to set up the replication link.

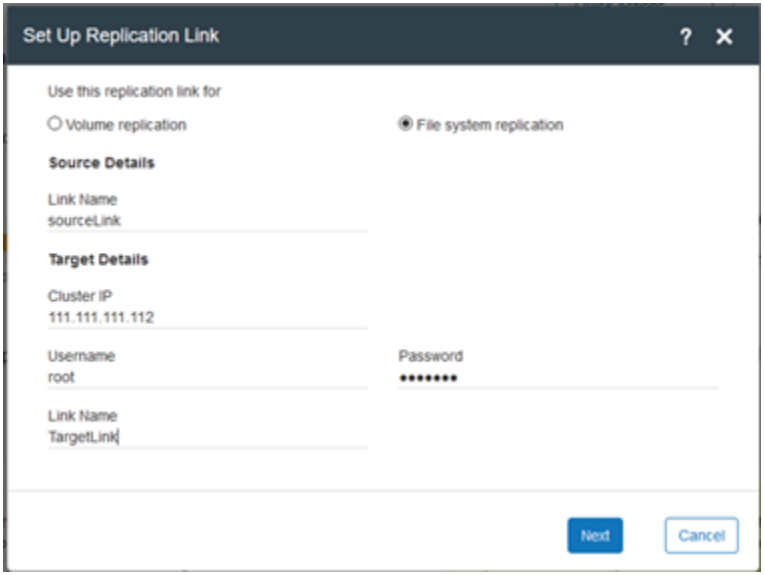

Configuring the archival policy using GUI

Once policy prerequisites are completed (like configuring storage pool and replication set up), you have to activate the archival policy.

Note: It is recommended that you keep the size of the file system such that there are a maximum of 100 million files on the file system.

For example, if the average size of file object from Enterprise Vault is 50 KB, then the size of the file system should be ~ 100 M x 50K = 5TB.

Note: To use the archival policy, ensure that you have five disks in the configured pool as the policy creates five-column striped file system layout. Otherwise, you can perform similar operations using the Veritas Access CLISH. See "Configuration of CIFS shares for archival using Veritas Access CLISH" on page 26.

To configure the archival policy using GUI

1. Log on to the Access Appliance GUI.
2. Click on **Policies & Select Archival Policy**. Activate the required policy by clicking **Activate** on the policy.

   The figure below shows activation of Enterprise Vault WORM policy.

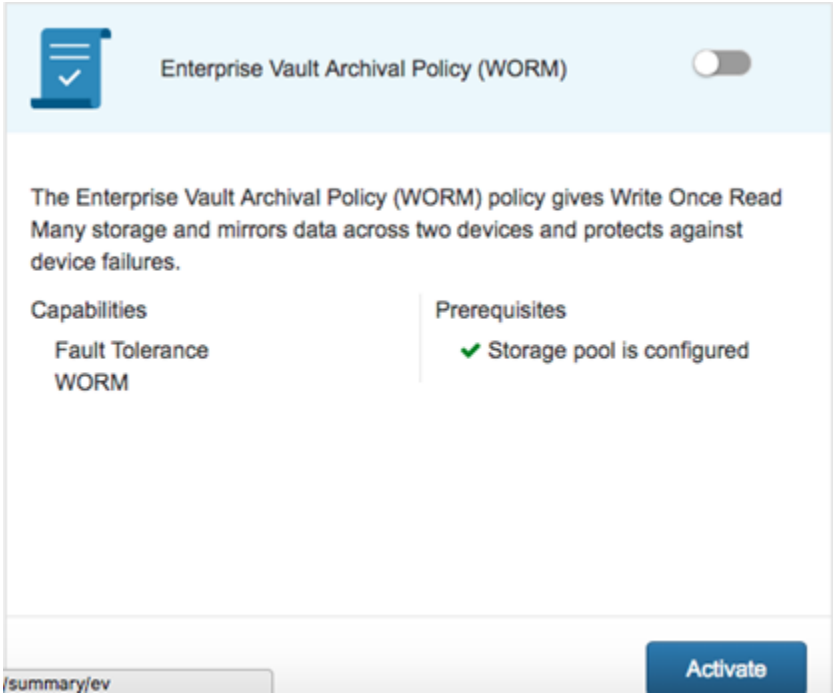

3 Select the storage pool during policy activation.

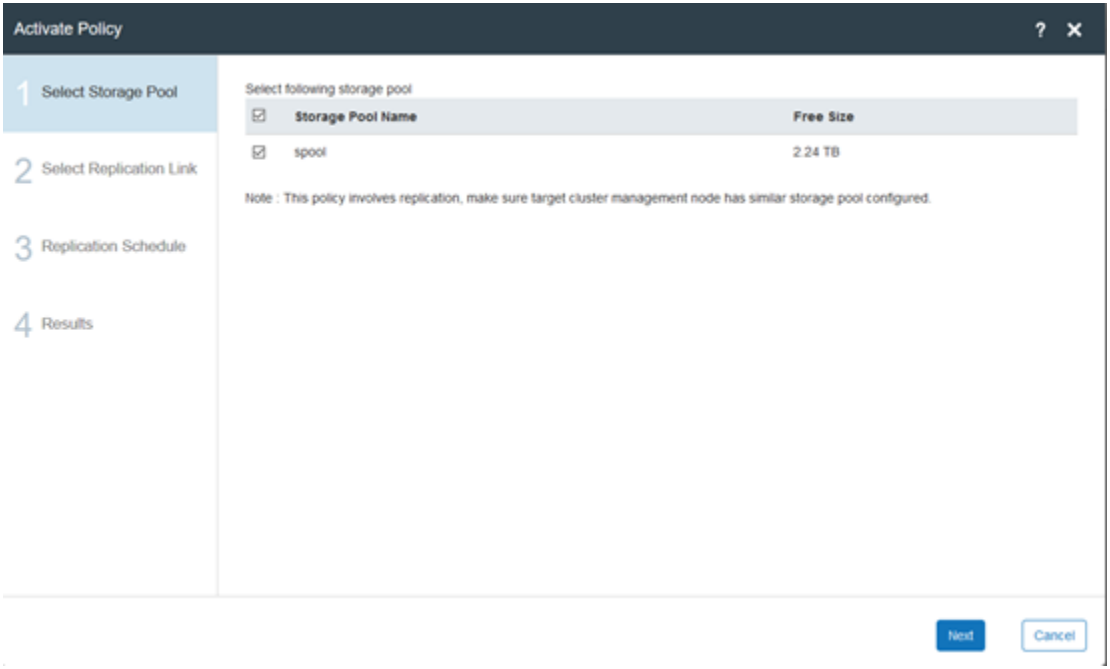

4 Select the replication link if the replication-related policy is activated. Select the replication link that has been configured.

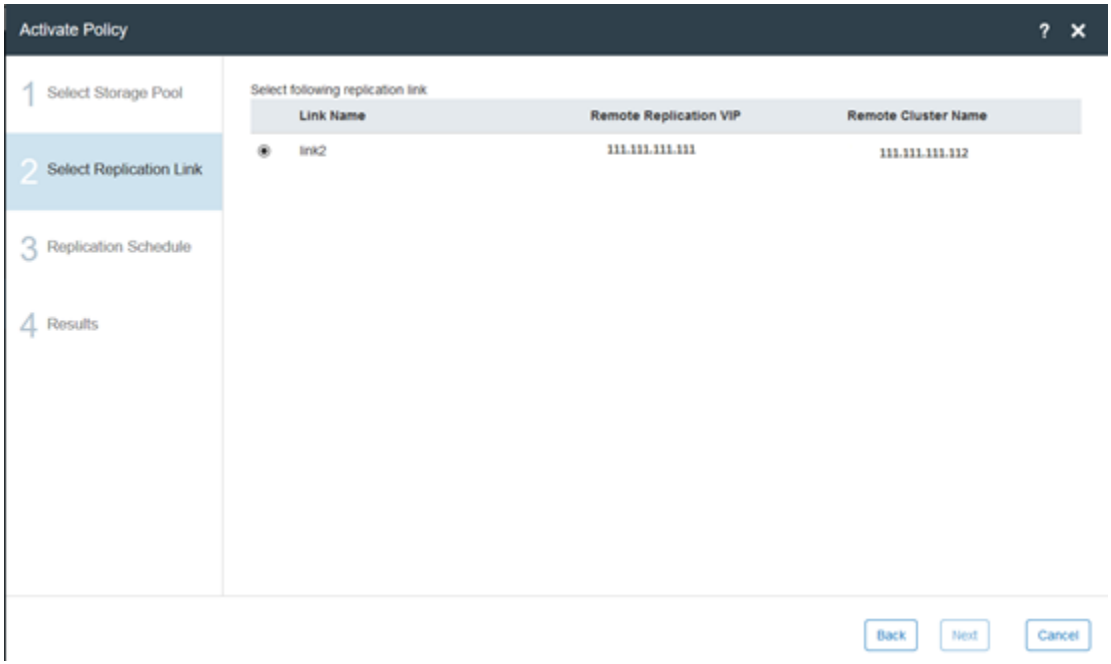

5 Provide replication job schedule information.

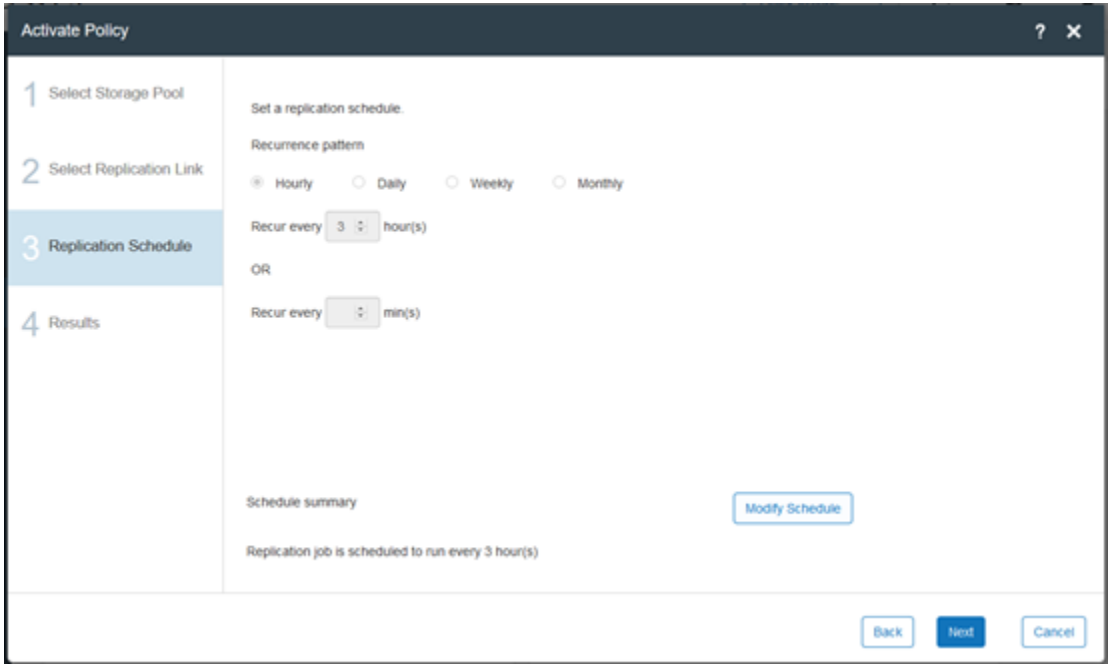

Once you activate the archival policy, storage for archival can be provisioned using the activated policy.

Note: Enterprise Vault partition rollovers show expected behavior only if there is a one-to-one mapping between the Enterprise Vault partitions and Veritas Access shares.

See “Storage provisioning using policies” on page 28.

Configuration of CIFS shares for archival using Veritas Access CLISH

Ensure that all the prerequisites are completed such as:

- Access appliance should be in enterprise or compliance mode.
- CIFS service should have been started.

- AD settings, configuration of storage pool and setting up the replication links should be complete.

To set up the archival folder for Enterprise Vault

1 Create a file system with a supported layout using the following command:

```
Storage> fs create striped testfs 10G 5 testpool stripeunit=128
blksize=1024 pdir_enable=no encrypt=off worm=yes
```

Note: To use WORM capabilities, make sure that you turn on the WORM flag during file system creation. You can also enable WORM flag for an already created file system using the `storage fs worm set minret maxret` command.

2 Create a CIFS share by adding a directory inside the file system, with share options as `rw`, `full_acl`, and `allow domain user`.

```
CIFS> share add testfs/testfs testfs
allow=accessra.local\evuser,noguest,rw,full_acl
```

Where the `testfs/testfs` is the CIFS share.

3 Go to the bash prompt by elevating to root. Create an empty directory inside the share, which is used as a root folder for archival data by Enterprise Vault.

```
cd /vx/testfs/testfs
mkdir -m 0777 ev_archival
```

Note: After the empty folder creation, full permission should be given to Enterprise Vault user for archival from Windows client (Enterprise Vault server). You can access the share from the Windows client (Enterprise Vault server) and configure partition.

The Veritas Access policies provide quick and easy configuration of archival storage for Enterprise Vault. You have an option to configure WORM or Non-WORM policies with or without replication.

See "Access Appliance GUI policies for archival storage" on page 18.

If you want to configure your archival policy with replication, perform the following steps.

To set up the episodic replication of archived files to target site

1. Create a replication unit for the root folder that was created for episodic replication.

   ```
   Replication> episodic repunit create source_rep1
   testfs/testfs/ev_archival
   ```

2. Create a replication schedule.

   ```
   Replication> episodic schedule create schedule1 */15
   ```

3. Create a replication job and enable the job.

   ```
   Replication> episodic job create replication job sourcerep targetrep
   src_to_destination schedule1 evpsn=yes
   ```

Storage provisioning using policies

You have to configure replication and set up the replication VIP and replication link.

To configure replication

1. On the **Quick actions** tab on top right corner of the Access Appliance GUI Dashboard, select **Provision Storage**.
2. Under **Storage type**, select **Storage for Enterprise Vault**. Select any one of the options, **Without Replication** or **With Replication** as required. Click **Next**.

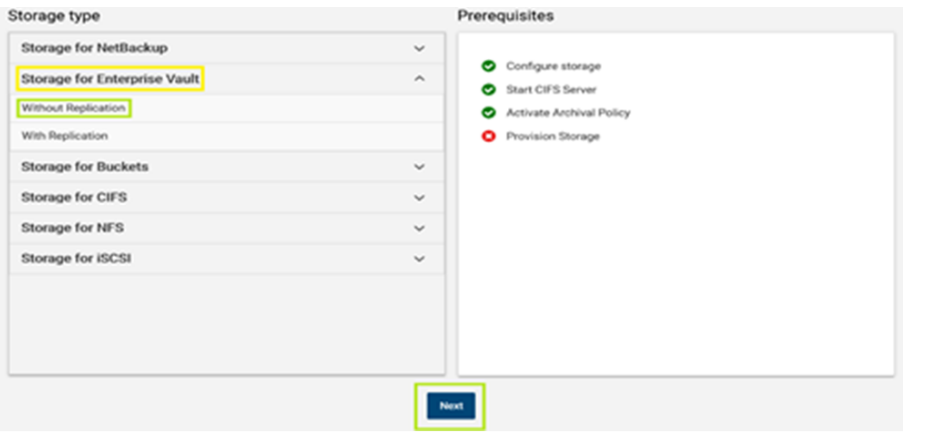

3 In the **EV setup** screen, select **Create CIFS share**.

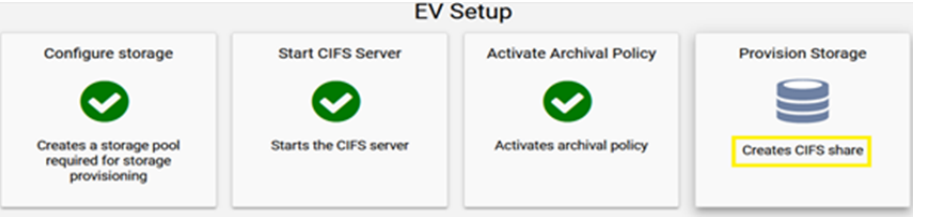

4 Expand **Basic options** to configure the basic CIFS options.

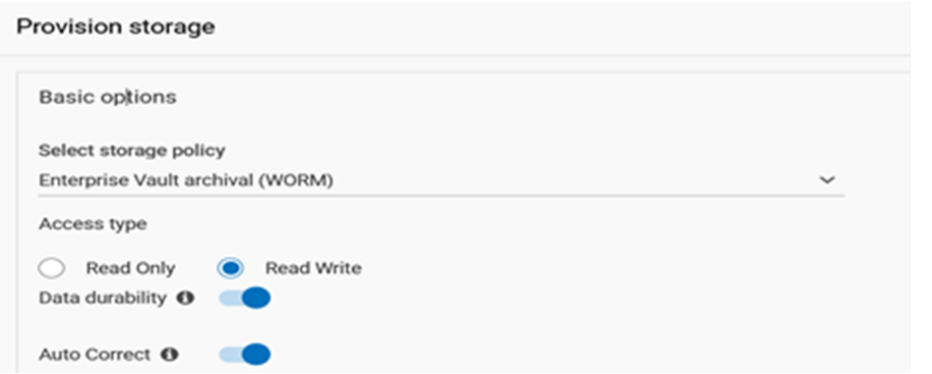

Under **Advanced CIFS export options**, check **Allow user and user group** to give access to the AD user.

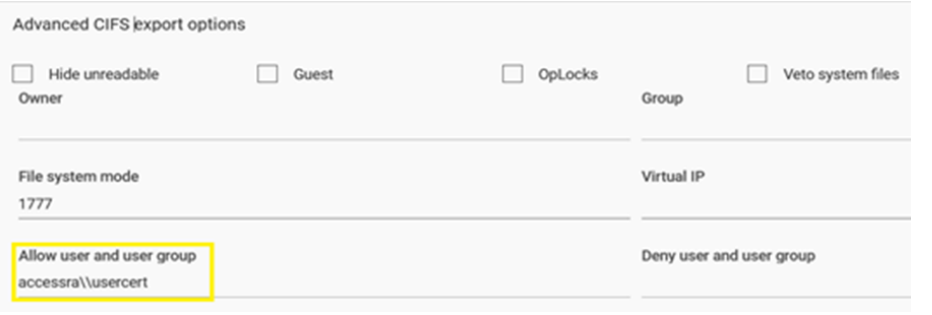

Add the share details and click **Provision Storage**.

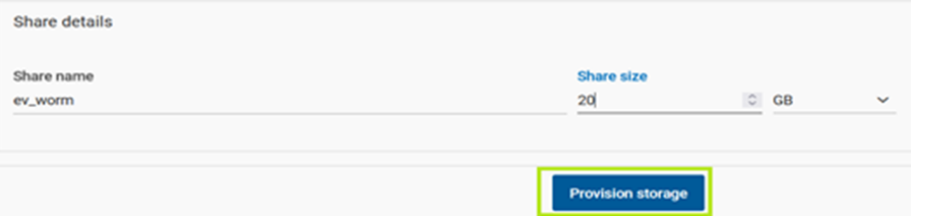

The share configuration creates an empty directory called **ev_archival** inside the share, which is used for archival data by Enterprise Vault.

In the example above, CIFS network share path used for Enterprise Vault is `/vx/demoshare/ev_archival`

Note: After the share is created, full permissions should be given to the `EV user` for the `ev_archival` directory in a CIFS share from the Windows client (Enterprise Vault server).

To give ownership of `ev_archival` directory to `evuser` in `EVSC` domain using the `chown` command:

```
# chown "EVSC\evuser" ev_archival
```

Configuring Access Appliance storage with Enterprise Vault store partition

Once a CIFS share is created for archival storage provisioning, you can configure the Enterprise Vault store partition. Access Appliance is now listed as a storage type in the **New Partition** dialog of Enterprise Vault. You can create a new partition under the vault store as shown in Figure 5-1.

Figure 5-1 Creating a new partition

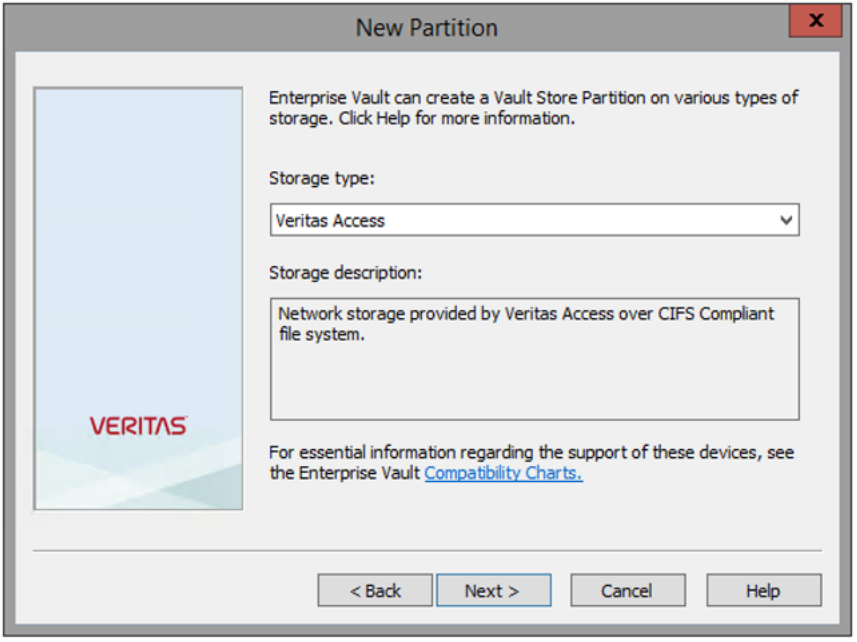

Select **Access Appliance** as the Storage Type.

Figure 5-2

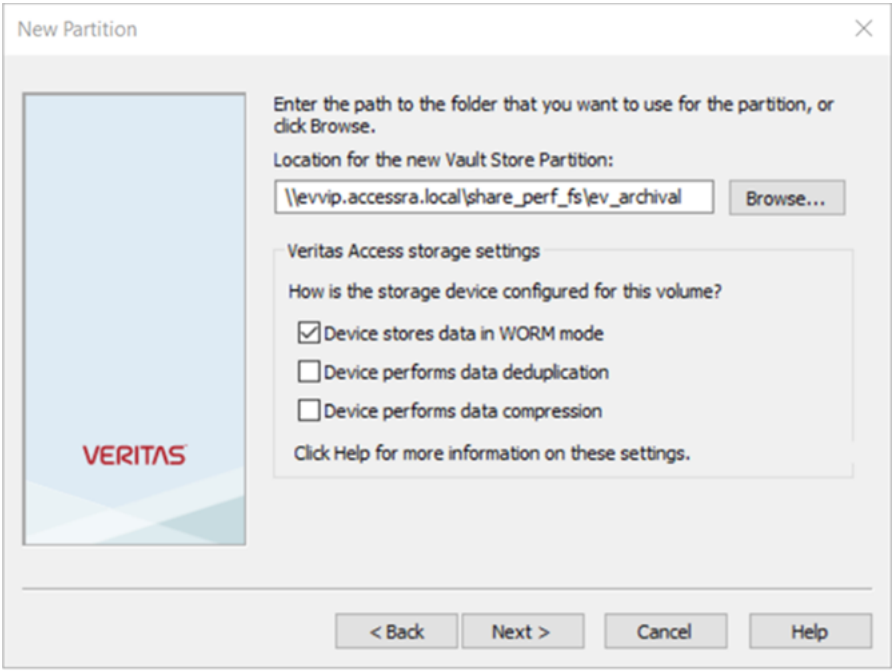

When you configure the Enterprise Vault store partition, specify the full path of the directory created in the CIFS share configuration step as a vault store partition location as shown in Figure 5-2.

`demoshare1` is the Access Appliance CIFS share and `ev_archival` is the empty directory in this share that is created through the GUI archival policy.

Note: Veritas supports one-to-one-to-one mapping between the file system, Access Appliance share, and Enterprise Vault partition. It is recommended to use one Access Appliance file system (after creating CIFS share for the file system) as one Enterprise Vault partition instead of creating multiple directories on one Access Appliance file system and use it as a separate partition for Enterprise Vault.

Episodic replication job failover and failback

Typically, the source cluster drives an episodic replication session. However, in some situations, it may be useful for the destination cluster to drive the episodic replication session. Access Appliance supports a failover and a failback feature for

episodic replication jobs. This feature enables control of episodic replication jobs to be temporarily relocated from the source cluster to the destination (target) cluster.

Job failover and failback is useful for:

- Planned failover
  In cases where the source cluster is taken down for routine maintenance or for moving applications to another cluster, a planned failover procedure is available for moving episodic replication jobs from the source cluster to the destination cluster.
- Unplanned failover or Disaster recovery
  In cases where the source cluster fails unexpectedly, an unplanned failover procedure is available for moving episodic replication jobs to the destination cluster.

Note: In the event of a planned or unplanned failover from the source cluster to the destination cluster, there should be at least one successful sync attempt. The successful sync ensures that a consistent point in time image is present on the destination cluster that can be used for the failover.

With job failover and failback, you use the `Replication> episodic job failover` command to move control from the source cluster to the destination cluster. You use the `Replication> episodic job failback` to restore control to the source cluster. The *link_name* is the link of one of the destination clusters. The *link_name* argument can be empty when the source cluster is not available, in which case the job failover can be run from one of the destination clusters.

Essentially, job failover takes job and episodic replication unit definitions from the episodic replication database on the source cluster and copies them to the episodic replication database on the destination cluster.

Warning: Job failover assumes that all episodic replication job names and episodic replication unit names are unique across all Access Appliance clusters on your network. Before you use the episodic replication failover feature, make sure that these names are unique.

After a job failover or failback, you must manually start or enable the episodic replication job to start preconfigured schedules. Link throttle information should be reconfigured after the job failover or failback.

Job failover does not automatically move the NFS or the CIFS share information that is associated with job failover episodic replication units from the source cluster to the destination cluster. Share information has to be done manually.

Table 5-1 Job failover and failback commands

Command	Definition
`episodic job failover`	Transfer control of an episodic replication job from the source cluster to the destination cluster.
`episodic job failback`	Return control of an episodic replication job from the destination cluster to the source cluster.

Process summary

The steps you take for job failover and failback vary depending on the type of failover or failback you perform. Failover and failback types include:

- Planned failover: This is performed to switch the job roles of source and destination clusters.
- Failback after a planned failover: This is performed after a planned failover to switch the job roles of the clusters back to the original state.
- Unplanned failover: This is performed when the primary cluster is down and disaster recovery has to be performed.
- Failback after an uplanned failover: This is performed after an unplanned failover or disaster recovery.

Each process is summarized in the following sections. Typically, you would use the planned failover and planned failback processes in most situations.

Overview of the planned failover process

Consider two Access Appliance clusters, cluster A and cluster B. Here, cluster A is the original source cluster that replicates data to cluster B which is the original destination cluster. If the source cluster A is taken down for routine maintenance or for moving applications to cluster B, you can use a planned failover operation so that the replication direction switches from cluster B (new source) to cluster A (new destination).

For planned failovers, most of the failover steps are run from the source cluster.

- From the source cluster (cluster A):
  - Stop all applications that access the replicated files. This step is recommended, but not required.
  - Use the `replication episodic job sync` *`job_name`* command to run the job and make sure files on the source cluster and destination cluster are synchronized.

- Use the `replication episodic job failover force=yes/no job_name current_cluster_link` command to move control of the job from the source cluster to the destination cluster.

- From the destination cluster (cluster B):
  - Use the `replication episodic job enable job_name` command to enable the job or run a sync on the destination cluster.
  - Use the `replication episodic job sync job_name` command to ensure that the episodic replication job is in a well-defined state and incremental episodic replication can be resumed.

Once the job is failed over, job control remains on the destination cluster until a planned failback is activated.

Note: While doing a planned replication episodic job failover from source to target, ensure that no shares are configured with the source file system as the source is marked as a target post failover.

Overview of the planned failback process

After a job failover has been accomplished and the source cluster is ready to take back control of the episodic replication task, you can use the job failback feature to release control from the original destination cluster (cluster B) and return it to the original source cluster (cluster A).

- From the original destination cluster (cluster B):
  - Stop all applications that access the replicated files. This step is recommended, but not required.
  - Use the `replication episodic job sync job_name` command to run the job and make sure files on the source cluster and destination cluster are synchronized.
  - Use the `Replication> episodic job disable job_name` command to disable the job.

- From the original source cluster (cluster A):
  - Use the `replication episodic job failback force=yes/no job_name current_cluster_link` command to move control of the job from the destination cluster back to the original source cluster.
  - Use the `replication episodic job enable job_name` command to enable the job or run a sync on the source cluster.

- Use the `replication episodic job sync job_name` command to ensure that the episodic replication job is in a well-defined state and incremental episodic replication can be resumed.

Once the planned failback process is completed, cluster A becomes the source cluster again and cluster B becomes the destination cluster again.

Overview of the unplanned failover process

In some cases (for example, unexpected equipment failure), you may need to run an unplanned failover for episodic replication jobs. The unplanned failover process differs from the planned failover process.

Consider two Access Appliance clusters, cluster A and cluster B. Here, cluster A is the original source cluster that replicates data to cluster B which is the original destination cluster. If the source cluster A fails unexpectedly, you have to perform an unplanned failover operation or a disaster recovery operation. After disaster recovery is performed, cluster B is marked as the new source cluster.

This section shows an overview of the steps you take to perform an unplanned failover.

For unplanned failovers, all the commands are run from the destination cluster (cluster B).

- Make sure that you are logged into the destination cluster.
- Use the `replication episodic job failover force=yes/no job_name` command to failover the job.

Overview of the unplanned failback process

After an unplanned failover, when the source cluster (cluster A) comes up, you can use the following unplanned failback process to return control to the original source cluster (cluster A):

- Make sure that you are logged into the source cluster.

  Note: Before starting the failback process, verify that the episodic replication service is running on the primary node. If the service is not running on the primary node, stop the service using the `replication episodic service stop` command and start it again using the `Replication> episodic service start` command.

- Use the `replication episodic job failover force=yes/no job_name current_cluster_link` command to configure the current source cluster as a

valid target to the new source cluster. This command should be run from the original source cluster (cluster A).

- Use the `replication episodic job sync job_name` command from the new source cluster to synchronize file system data with the newly added destination cluster.
- Use the `replication episodic job failback force=yes/no job_name current_cluster_link` command to move control of the episodic replication job from the destination cluster back to the source cluster. This command should be run from the original source cluster (cluster A).
- Use the `replication episodic job sync job_name` command to ensure that the episodic replication job is in a well-defined state and incremental episodic replication can be resumed.

Once the unplanned failback is completed, cluster A becomes the source cluster again and cluster B becomes the destination cluster again.

Note: An administrator can use the `Replication> episodic job destroy force` command to clean up local job configuration. Configuration of the other clusters, which are part of the job, are not modified and any episodic replication units are disassociated from job. The `Replication> episodic job destroy force` and `Replication> episodic repunit destroy force` commands should be used in the event of an unrecoverable configuration or episodic replication direction mismatch.

Continuous replication failover and failback

Typically, the source cluster drives a replication session. However, in some situations, it may be useful for the destination cluster to drive the replication session. Access Appliance supports a failover and a failback feature for continuous replication. This feature enables control of replication to be temporarily relocated from the source cluster to the destination (target) cluster.

Continuous replication failover and failback is useful for:

- Planned failover
  In cases where the source cluster is taken down for routine maintenance or for moving applications to another cluster, a planned failover procedure is available for moving replication from the source cluster to the destination cluster.
- Unplanned failover or Disaster recovery
  In cases where the source cluster fails unexpectedly, an unplanned failover procedure is availble for moving replication to the destination cluster.

With failover and failback, you can use the `replication continuous failover` command to move control from the source cluster to the destination cluster. You use the `replication continuous failback` to restore control to the source cluster.

Continuous replication failover does not automatically move the NFS or the CIFS share information that is associated with file system from the source cluster to the destination cluster. Share information has to be done manually.

CIFS shares should be removed before failover and failback operations. After failover and failback operation are complete, add the CIFS shares again.

Recommendations while configuring Access Appliance for continuous replication:

- The number of RVGs configured under continuous replication on Access appliance should be restricted to 4.
- The number of file systems under one RVG should be limited to 12.
- Veritas recommends that you to add one file system at a time to an RVG. Once the file system is added and the status of continuous replication becomes consistent (up-to-date), add another file system to the RVGs.
- Ensure that continuous replication IPs configured on Access Appliance (both primary and secondary sites) are properly resolved by the configured DNS server or by `/etc/hosts` file on both nodes of the primary and secondary sites.
- Addition of multiple file systems to an existing Replicated Volume Group is supported only through CLISH.
- The failover and failback happens at the RVG level. So, it is recommended to keep all partitions (corresponding file systems on Access Appliance) of one Enterprise Vault store in one RVG. This reduces the need of failing over or failing back multiple RVGs.
  For example:
  If *ev_vault1* uses *share1* (closed partition) and *share2*(open partition) on Access Appliance which maps to file systems, *fs1 (share1)*, and *fs2 (share2)*.
  If *ev_vault2* uses *share3* and *share4* on Access Appliance which maps to file systems, *fs3 (share3)*, and *fs4 (share4)*.
  In this scenario, you can create *rvg1* for *fs1* and add *fs2* to *rvg1*. Similarly, you can create *rvg2* for *fs3* and *fs4*. In case, you have to failover *ev_vault1*, you can failover *rvg1* and Enterprise Vault will be functional during disaster recovery. You are not required to do anything to *rvg2* unless you want to failover *ev_vault2* also.
- If there are more than 4 Enterprise Vault stores in an environment, you are required to put the file systems (partitions) of different Enterprise Vault stores in one RVG as the maximum number of recommended RVGs on Access Appliance is 4.

- The bandwidth available for replication between primary and secondary sites should be able to handle the rate at which data is being written on the primary site. With synchronous continuous replication, the writes on file systems on primary site are completed after it is replicated to the secondary site. Hence, Veritas recommends that you have a high bandwidth for replication with synchronous continuous replication.

Process summary

The steps you take for failover and failback vary depending on the type of failover or failback you perform. Failover and failback types include:

- Planned failover: This is performed to switch the job roles of source and destination clusters.
- Failback after a planned failover: This is performed after a planned failover to switch the roles of the clusters back to the original state.
- Unplanned failover: This is performed when the primary cluster is down and disaster recovery has to be performed.
- Failback after an uplanned failover: : This is performed after an unplanned failover or disaster recovery.

Each process is summarized in the following sections. Typically, you would use the planned failover and planned failback processes in most situations.

Overview of the planned failover process

Consider two Access Appliance clusters, cluster A and cluster B. Here, cluster A is the original source cluster that replicates data to cluster B which is the original destination cluster. If the source cluster A is taken down for routine maintenance or for moving applications to cluster B, you can use a planned failover operation so that the replication direction switches from cluster B (new source) to cluster A (new destination). You should perform the planned failover before cluster A is taken down.

For planned failover, run the following command on the source cluster (cluster A):

```
Replication> continuous failover fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once a planned failover happens, the roles of primary and secondary are switched. It will online the file system at new primary site and offline the file system at new secondary site. Now, cluster B is the primary site and cluster A is the secondary

site. The file systems come online at the new primary site (cluster B) and the file systems at new secondary site (cluster A) are offline.

Note: Planned failover command should be run when both the source and the destination clusters are reachable from each other. It should be run from the source cluster and replication should be in progress. If you have NFS/CIFS shares on the source cluster, it is recommended that you should stop the NFS/CIFS server before planned failover.

Overview of the planned failback process

After a planned failover has been accomplished and the original source cluster (cluster A) is ready to take back control of the replication task, you can use the failback feature to release control from the original destination cluster (cluster B) and return it to the original source cluster (cluster A).

For planned failback, run the following command:

```
Replication> continuous failback fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once a planned failback happens, the roles of primary and secondary are switched back to the original state (the pre-failover state). Now, cluster A is the primary site again and cluster B is the secondary site again.

Note: Planned failback command should be run when both the source and the destination clusters are reachable from each other. It should be run from the destination cluster (which was the original source cluster) and replication should be in progress. If you have NFS/CIFS shares on the source cluster, it is recommended that you should stop the NFS/CIFS server before planned failback.

Overview of the unplanned failover process

In some cases (for example, unexpected equipment failure), you may need to run an unplanned failover for replication. The unplanned failover process differs from the planned failover process.

If cluster A is the original source cluster that replicates data to cluster B which is the original destination cluster and if cluster A fails unexpectedly, you can perform an unplanned failover or a disaster recovery. This marks the cluster B as the new source cluster and you can access the data/applications replicated to it.

For unplanned failover, run the following command from the original destination cluster (cluster B):

```
Replication> continuous failover fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once an unplanned failover happens, the destination cluster (cluster B) becomes the new source cluster. It onlines the file system at the new source cluster (cluster B).

Note: Though the commands used for planned and unplanned failover are the same, the intention and pre-requisites are different. For unplanned failover, the source cluster should be unreachable from the destination cluster.

Overview of the unplanned failback process

After an unplanned failover, when the original source cluster (cluster A) comes up, you can use the following unplanned failback process to retain the replication between source and destination cluster.

When the original source cluster (cluster A) comes up, it still acts as the primary cluster. At this time, both the source and the destination clusters show the same status.

For unplanned failback, run the following command from the original source cluster (cluster A):

```
Replication> continuous failback fs_name
```

Where *fs_name* is the name of the file system which is configured under continuous replication.

Once, an unplanned failback happens, the original source cluster (cluster A) becomes the new destination cluster. The file system is still online at new source cluster (cluster B).

Note: Unplanned failback command should be run from the original source cluster. If you have NFS/CIFS shares on the original source cluster, it is recommended that you stop the NFS/CIFS server before unplanned failback.

Configuring replication failover and exporting CIFS share from target cluster

Perform replication failover. After the replication job to failover to the target cluster is complete and the target cluster becomes the primary cluster, you can allow Enterprise Vault to continue further archival on the appliance by performing the following steps:

To set up the archival folder for Enterprise Vault from target cluster

1. Create a CIFS share with the same name as the share present on the source cluster on the replicated file system.

   See "Configuration of CIFS shares for archival using Veritas Access CLISH" on page 26.

2. Add a directory inside the file system, with share options as `rw`, `full_acl`, and `allow domain user`.

   ```
   CIFS> share add testfs testfs/test_directory
   allow=accessra.local\evuser,noguest,rw,full_acl
   ```

   Where the `testfs/test_directory` is the CIFS share.

3. Go to the bash prompt by elevating to root. Create an empty directory inside the share, which is used as a root folder for archival data by Enterprise Vault.

   ```
   cd /vx/testfs/testfs
   mkdir -m 0777 ev_archival
   ```

Note: After the empty folder creation, full permission should be given to Enterprise Vault user for archival from Windows client (Enterprise Vault server). You can access the share from the Windows client (Enterprise Vault server) and configure partition.

---

4. Check the IP address to which the CIFS share is mapped using the `cifs share show <share_name>` command.
5. Update the DNS, modify the IP address for the FQDN which was used for the source cluster.

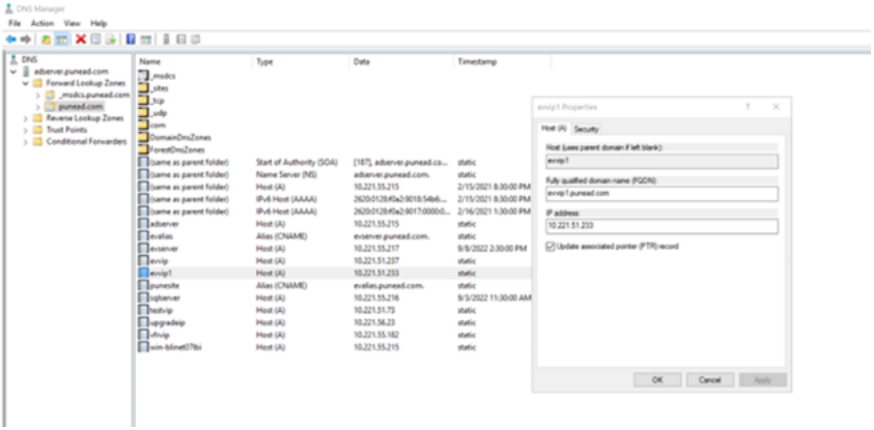

Chapter 6

Troubleshooting

This chapter includes the following topics:

Log locations for troubleshooting

- For issues related to accessing CIFS shares, see Access Appliance CIFS-related logs in the `/log/VRTSnas/` directory. The CIFS log files are present in the `cifs_<operation>_<sub operation>.log` folder. For example, `cifs share add` command log file is named as `cifs_share_add.log`.
- See Windows Event Log for errors related to the Enterprise Vault server.
  - Go to **Run**. Type **eventvwr**.
  - The Windows Event Viewer opens up.
    Select **Applications and Services logs**-> **Microsoft** -> **Veritas Enterprise Vault**.
- See **Recent Activity** at the top right corner of Access Appliance GUI for GUI-related status information

Additional resources

See the following documentation on SORT for more information on Access Appliance and Enterprise Vault.

- *Access Appliance Initial Configuration Guide*
- *Access Appliance Administration Guide*
- *Access Appliance Online Help*

- Enterprise Vault product documentation

Index