Technical Report

Configuring NetApp SnapLock with Symantec Enterprise Vault

Nathan Walker, Technical Marketing Engineering, NetApp
January 2010| TR-3501

ABSTRACT

Structured and nonstructured data is rapidly growing. Critical intellectual property includes enterprise messaging (primarily e-mail) and distributed file systems. As these data sets grow, it is increasingly clear that today's enterprises need a structured approach to protect the data. Increased acceptance of electronic communication requires an efficient method of storing, managing, and discovering the data. For legal and compliance purposes, enterprise customers require that their electronic data are archived and safely secured. Regulations require that requested data be able to be collected within a short period of time. An archive stack based on Symantec™ Enterprise Vault™ and NetApp® storage systems provides enterprises the capability to archive and protect both business-critical e-mail and unstructured file system data in a simple-to-manage unified storage environment. This technical report discusses the procedures required to configure archiving using SnapLock® with Symantec Enterprise Vault in a NetApp unified storage environment. Refer to the "Enterprise Vault Installation Guide for NetApp Storage Systems" technical report for additional details.

TABLE OF CONTENTS

1 INTRODUCTION

Archiving and compliance are becoming increasingly important in today's business environment. NetApp delivers a comprehensive suite of standards-based, simple-to-use, and cost-effective compliance solutions to hundreds of enterprises worldwide to enable them to meet the most stringent regulations today and in the future. Enterprise Vault on NetApp storage solutions provide a platform to enable companies to reduce costs and improve productivity and compliance. By taking advantage of the tight integration of Enterprise Vault on NetApp storage systems, you can create a scalable and easy-to-use archive-and-compliance platform to keep your enterprise data available and secure.

2 TECHNICAL BACKGROUND

A number of compliance regulations enacted globally require archives to be kept in an immutable state on write once read many (WORM) media. Archive applications and the data associated with them have unique performance, scalability, and availability needs that require flexibility in storage systems. NetApp FAS systems offer a unified architecture that supports both Fibre Channel and SATA disk drives simultaneously, allowing customers to balance archive application performance requirements with system cost. Our SnapLock technologies help you to create nonrewritable, nonerasable storage volumes to prevent files from being altered or deleted until a set retention date. You have the flexibility of mixing WORM and non-WORM disks within the same storage system, and the assurance that your Enterprise Vault archive data are protected.

3 PURPOSE AND SCOPE

The purpose of this paper is to demonstrate the configuration procedures for Enterprise Vault and NetApp storage systems. This paper assumes that the Enterprise Vault and NetApp storage systems are operating according to best practices and that the reader is ready to configure the Enterprise Vault partition. Because Enterprise Vault WORM partitions are typically found only in journaling environments, the focus of this paper is on this compliance use case. However, the concepts discussed herein can be applied to many other archiving scenarios. A detailed introduction to SnapLock can be found in TR-3263.

4 REFERENCE INFRASTRUCTURE

The following diagram illustrates how Enterprise Vault can be deployed on a NetApp storage system with support for the Enterprise Vault index, database, and WORM archives. These storage volumes can be contained within a single storage controller or spread across multiple devices. For many enterprises a remote disaster recovery facility typically mirrors the primary site.

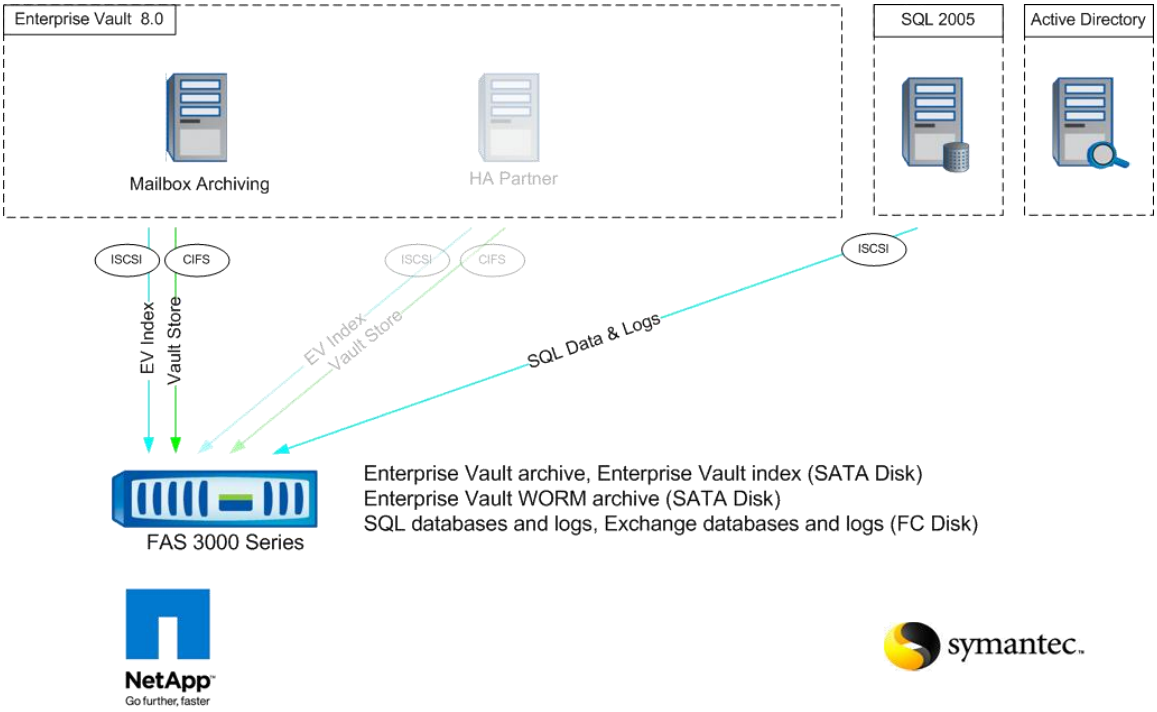

Figure 1) Enterprise Vault on NetApp storage solutions.

5 COMPLIANCE CLOCK

The initial step in setting up SnapLock is to initialize the compliance clock. This can be accomplished by executing the following command:

```
Storage controller> date -c initialize
```

This initializes the compliance clock to the current value of the system clock. Care should be taken to appropriately set the system clock before initializing the compliance clock. The compliance clock may be set only one time. There is no mechanism for resetting this clock. Using the compliance clock on a storage system with SnapLock volumes prevents retained data from being deleted prematurely by changing the regular clock of the storage system. Further details about the SnapLock compliance clock can be found in TR-3618.

6 SNAPLOCK AGGREGATE

Before committing data to a SnapLock Compliance volume in a production environment, you may have better project success by testing the application and completing use cases in a test environment using SnapLock Enterprise.

Make sure you have entered your license key before attempting to create the SnapLock volumes. The license keys for SnapLock Compliance and SnapLock Enterprise can be added through the GUI-based FilerView® or a console session. From FilerView, first select the Filer node, select Manage Licenses, scroll down to the appropriate SnapLock option, and enter your license. From the command line console or an SSH session, enter the command:

```
storage controller> license add your_license_code
```

In order to create the SnapLock aggregate you need to select the SnapLock checkbox when using the FilerView aggregate wizard.

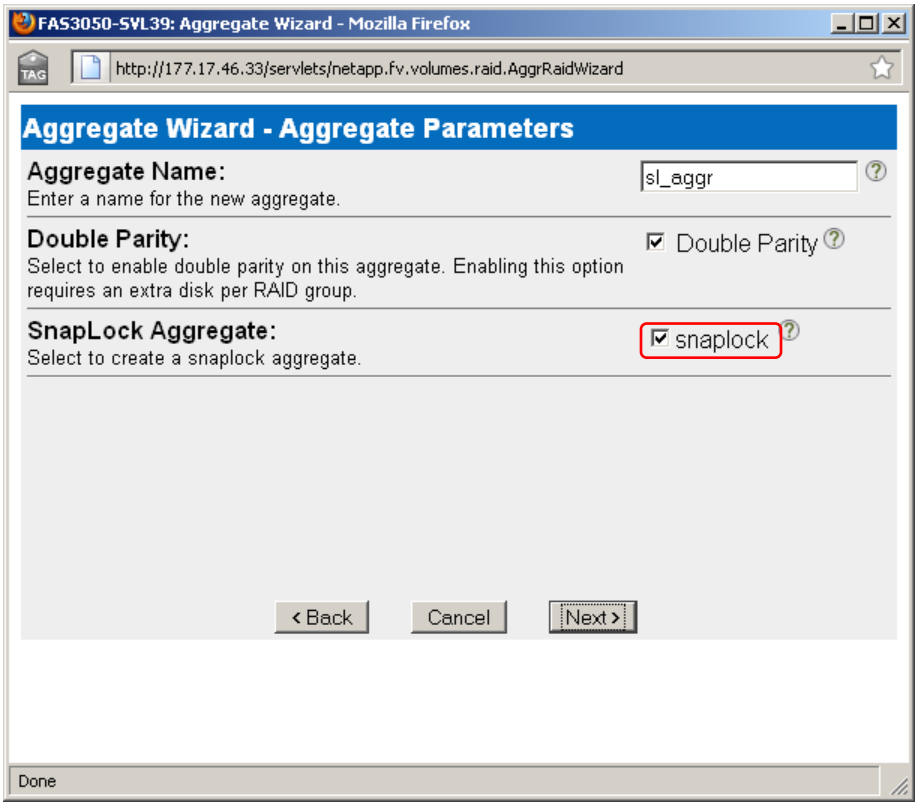

Figure 2) Aggregate Wizard SnapLock enable screen.

By selecting the SnapLock checkbox you are prompted to create either a compliance or an enterprise aggregate, according to your needs in the following screen.

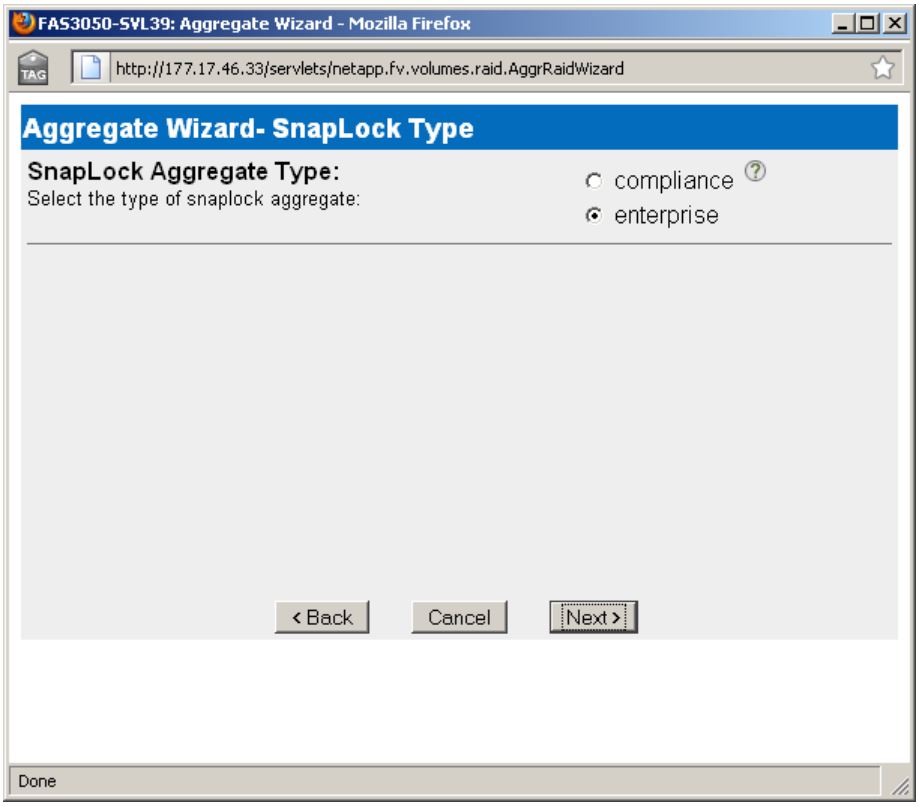

Figure 3) Aggregate wizard SnapLock type selection screen.

If you do not select the SnapLock checkbox, none of the aggregates created for SnapLock will display.

To create a SnapLock aggregate from the command line, include the `-L` flag with the `aggr create` command. The following example shows how to create a SnapLock enterprise aggregate called sl_aggr, with a default RAID-DP® group and RAID size, using 14 disk drives.

```
Storage controller> aggr create sl_aggr -L enterprise 14
```

Next, make sure the aggregate is available and online using the `aggr status` command:

```
Storage controller> aggr status
           Aggr State           Status            Options
       sle_aggr online          raid_dp, aggr     snaplock_enterprise
```

Certain settings for all SnapLock aggregates within a storage controller can be managed with the `options` command. If specified, the autocommit_period setting defines the amount of time after being created or modified that a file will remain in a non-WORM state before being committed. The autocommit_period setting will likely be defined by your legal department.

```
Storage controller> options snaplock
snaplock.autocommit_period   2h
snaplock.compliance.write_verify off
snaplock.log.default_retention 6m
snaplock.log.maximum_size    10m
```

7 SNAPLOCK FLEXVOL

After creating the SnapLock aggregate, the next step in the process is to create a FlexVol® volume or a traditional volume on top of the new aggregate. Aggregates are capable of storing numerous flexible volumes; the amount of storage allocated to each volume can vary. Every FlexVol volume and traditional volume created on a SnapLock aggregate takes on the SnapLock properties of the underlying aggregate, including either the SnapLock Compliance or SnapLock Enterprise characteristics of the aggregate.

The steps for creating the volume on a SnapLock aggregate are exactly the same as volume creation on a regular read-write aggregate. From FilerView click **add** under the Volumes node to start the wizard. Select the SnapLock checkbox on the volume parameters screen. If you do not select this setting, only traditional aggregates will appear in the drop-down box in the subsequent screen.

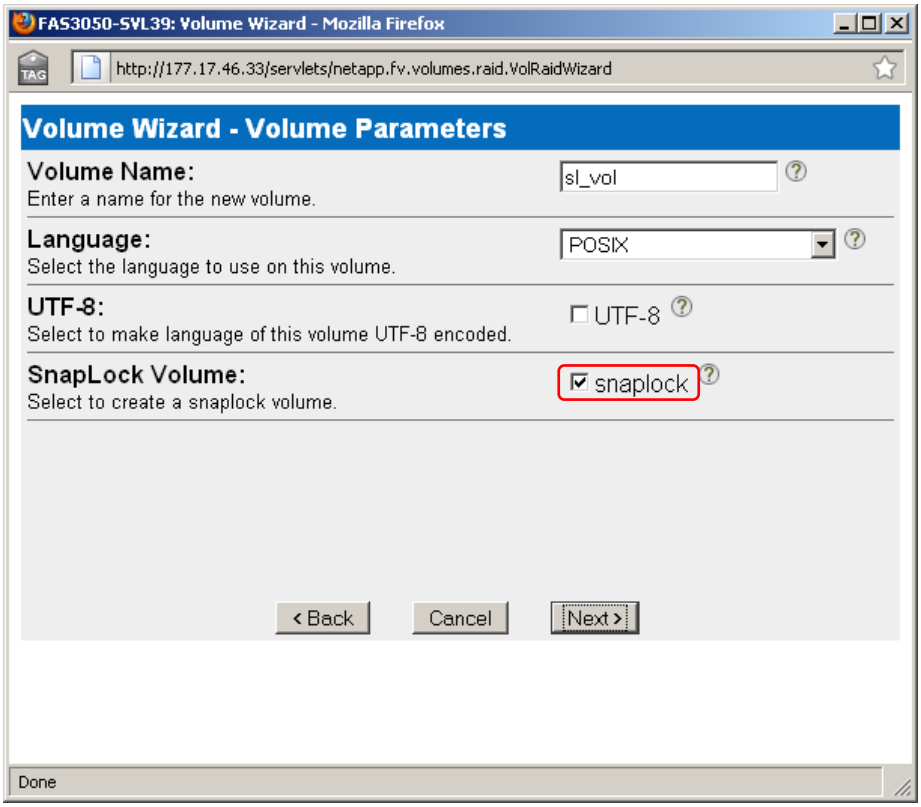

Figure 4) New FlexVol SnapLock selection.

The FlexVol create command line syntax to create the same volume is as follows:

```
Storage controller> vol create sl_vol sl_aggr 100g

Creation of volume 'sl_mvol' with size 100g on containing aggregate

'sl_aggr' has completed.
```

Where sl_vol is the name of the FlexVol volume and sl_aggr is the underlying aggregate from the previous section. The trailing 100g denotes that the size of the FlexVol volume will be 100GB.

Finally, verify that the FlexVol volume is available and online with the `vol status` command:

```
Storage controller> vol status

         Volume State           Status            Options

         sl_vol online          raid_dp, flex     no_atime_update=on,

                                                  snaplock_enterprise
```

Three extra volume options are available with a SnapLock FlexVol volume. On the command line, type `vol options your_volume_name` to display these values and settings. The `minimum_retention_period` applies to the smallest amount of time the WORM file or snapshot must be kept in a SnapLock volume. You set this retention period so that applications or users do not assign noncompliant retention periods to retained records in regulatory environments. Until you explicitly reconfigure it, the minimum retention period is 0.

The `maximum_retention_period` specifies the longest period a WORM file or snapshot must be kept in a SnapLock volume. You normally set this retention period so that applications or users do not assign excessive retention periods to retained records in regulatory environments. Until you explicitly reconfigure it, the maximum retention period is 30 years. Set this value to the shortest retention period defined in Enterprise Vault.

The `default_retention_period` specifies the retention period assigned to any WORM file or snapshot on the SnapLock volume that was not explicitly assigned a retention period. The default maximum retention period is 30 years. You set this retention period so that a default retention period is assigned to WORM files in the case in which users or applications failed to assign a retention period. Set this value to the longest retention period defined in Enterprise Vault. To define an infinite retention period use the following syntax:

```
Storage controller> vol options sl_vol infinite
```

8 CIFS SHARE

The final step on the storage system in preparing a SnapLock volume for Enterprise Vault is to create a CIFS share. There are no steps that pay particular attention to SnapLock when creating and configuring the share. In FilerView this can be performed by navigating to CIFS, Shares, Add. From the command line it is a single command:

```
Storage controller> cifs shares -add sl_vol /vol/sl_vol
```

Typically only the Enterprise Vault service account has access to the partition:

```
Storage controller> cifs access sl_vol domain\ev_account "Full control"

Storage controller> cifs access -delete sl_vol everyone
```

9 ENTERPRISE VAULT SNAPLOCK TASKS

The majority of tasks and activities are unchanged when working with a WORM archive partition. There are a few differences when setting up the vault store partition as well as backups.

First start by creating the archive folder on the CIFS share. This should be done from the command line using the `mkdir` command. When new folders are created using the Windows® explorer, the initial

name is New Folder. By design, SnapLock does not allow folders to be renamed, so you should create the folder for the SnapLock archive from the command line.

NetApp advises having a vault store dedicated to SnapLock archives. This creates better separation between compliance and noncompliance archives. When you create the vault store, configure the safety copies to be removed immediately.

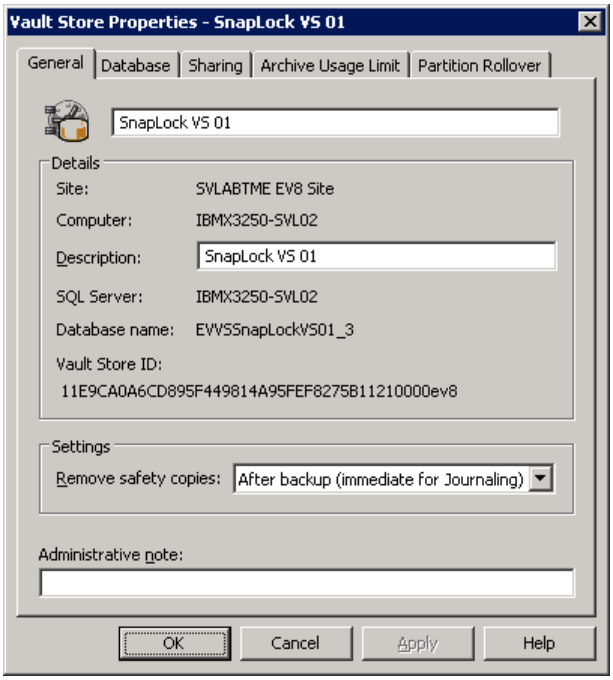

Figure 5) Enterprise Vault store partition properties.

Now open the Enterprise Vault administration console and start the vault store partition wizard. Make sure you select “NetApp device” as the storage type.

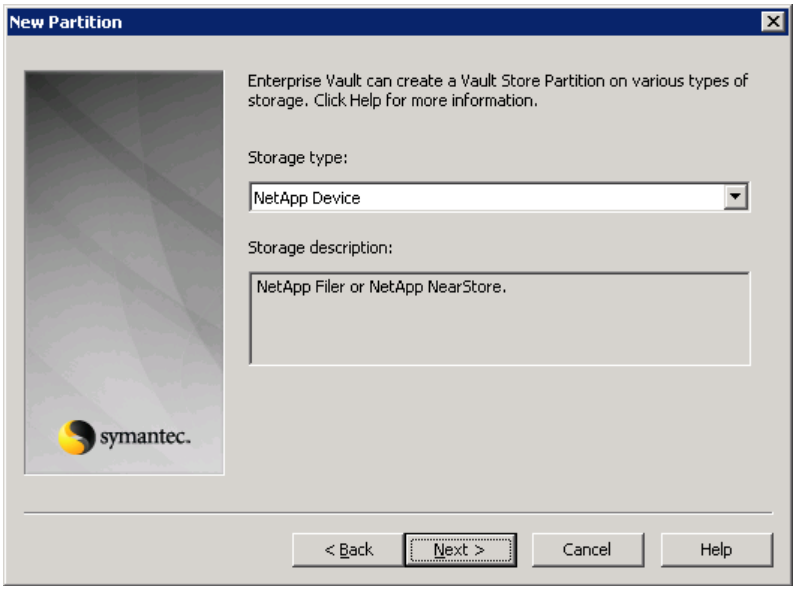

Figure 6) Vault store properties.

If you select NetApp storage, the next screen in the wizard will display some options specific to NetApp. Click the top checkbox to signal Enterprise Vault that it is writing to a SnapLock partition.

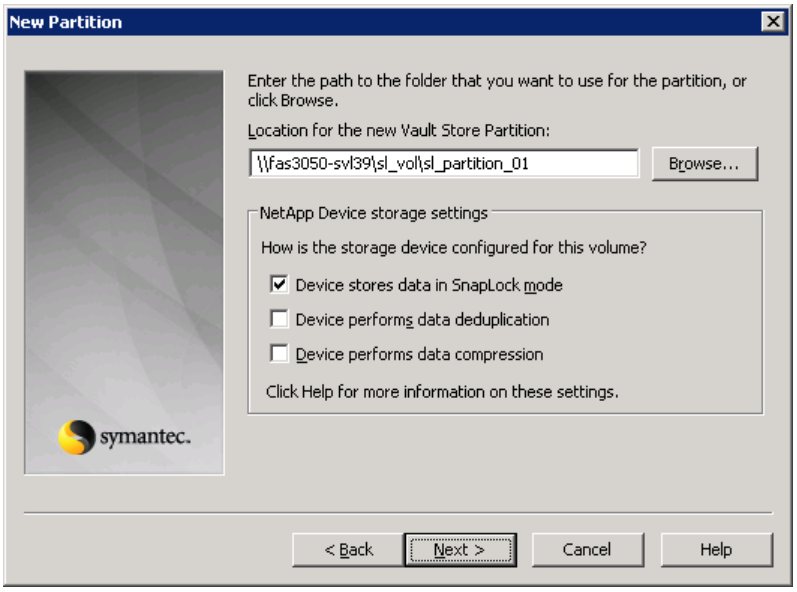

Figure 7) New partition SnapLock select.

If your vault store is configured to remove safety copies immediately, then the next screen is not relevant.

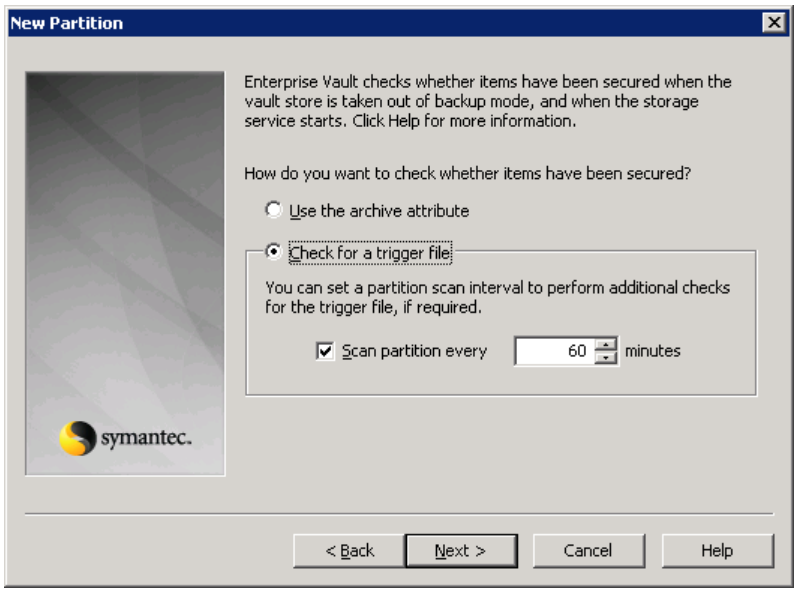

Figure 8) New partition trigger file check.

When safety copies remain in the source mailboxes, Enterprise Vault needs to receive a signal that the backups are complete. There are two ways this happens: either through setting the archive attribute or by creating a trigger file at the root of the partition. Enterprise Vault writes archive files to the SnapLock partition with the archive and read-only bits set. So only the trigger file allows you to remove safety copies, if used. Because both bits are set, the Enterprise Vault archive files are immediately committed to WORM and set to an immutable state. After they are committed you cannot change any file attributes utill the retention period has passed.

To view the retention period of a particular file on a SnapLock volume, open the file properties and take note of the accessed date. In this case it is set to the furthest date recognized by Windows.

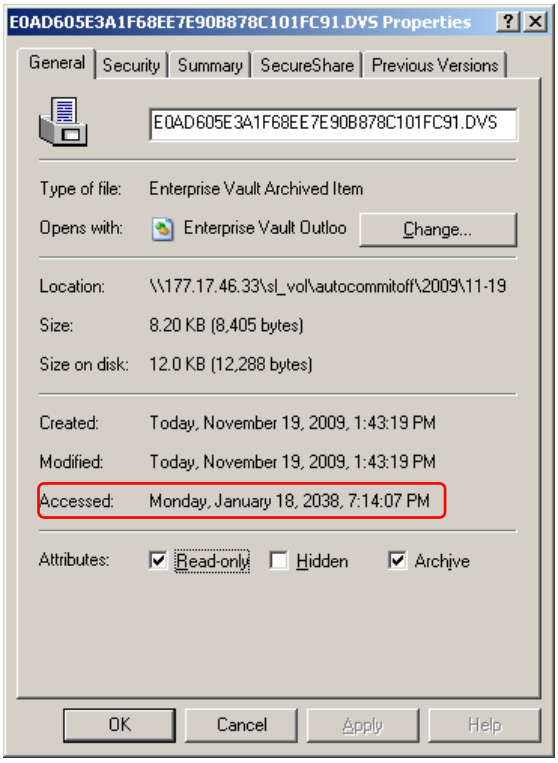

Figure 9) Viewing a SnapLock retention period.

Care should be taken if your organization has a legal requirement to turn on autocommit_period and you are using the ignore archive bit trigger file. The autocommit_period applies to all volumes on a storage controller. As discussed above, there are two mechanisms for committing content to the SnapLock volume: either by setting the read-only bit on the file to be committed or by defining the snaplock.autocommit_period option on the storage controller. If a file has not been modified in a time that exceeds the autocommit_period, it will be committed to SnapLock. Enterprise Vault looks at the create time of the archive bit trigger file to determine whether or not to remove safety copies from the archive source and replace them with shortcuts.

The following diagram shows the backup states that are sensitive to the autocommit period.

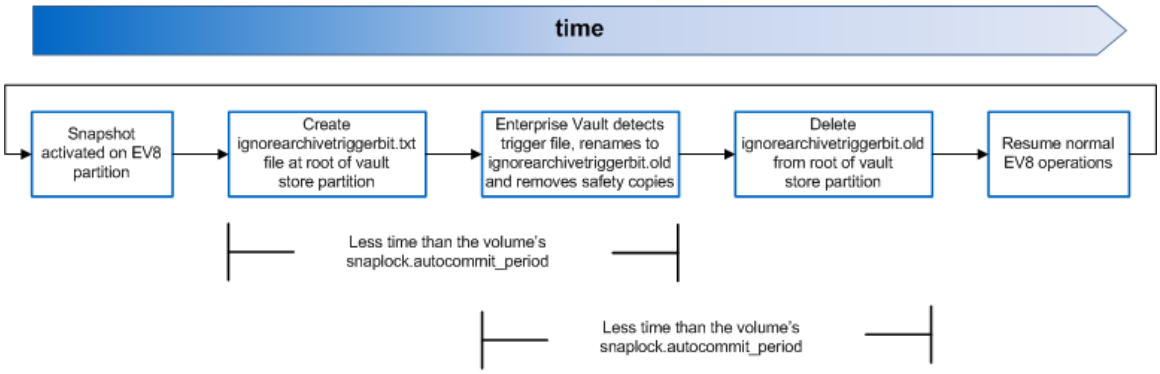

Figure 10) Backup states with autocommit_period.

Many backup programs such as Symantec NetBackup™ have pre- and post-scripting capabilities that are used to execute the Enterprise Vault PowerShell cmdlets and to create the trigger file. If the backups fail to activate or terminate properly, it is possible that the trigger file would be committed to disk. Therefore you would want to create a scheduled task to remove the ignorearchivebittrigger.txt and ignorearchivebittrigger.old files if they were not recently modified. The task should be scheduled to run more frequently than the length of the autocommit_period. Further details on the trigger file for Enterprise Vault backups can be found in Symantec KB 273151. An excellent Symantec article on the Enterprise Vault PowerShell cmdlets can be found here.

10 CONCLUSION

NetApp's unified system architecture provides the manageability, performance, flexibility, and availability necessary to easily run an e-mail archive environment facing today's increasingly complex business demands. We have worked with Symantec to integrate our storage technologies with Enterprise Vault to offer you an archiving platform that can meet your needs today and tomorrow. You can rest easy knowing that solutions built upon NetApp storage solutions and Enterprise Vault can meet your requirements to store, manage, and discover archived content in a simple-to-manage solution. With NetApp you get value and performance features that maximize your data protection and provide you with outstanding efficiencies.

www.netapp.com