



Technical FAQ

## SnapLock FAQs

Manish M. Agarwal, Timothy Isaacs, Lokesh AP, Ajay Someshwar, NetApp  
May 2012

## TABLE OF CONTENTS

<b>1</b>	<b>General Questions</b>	<b>5</b>
	What is the difference between SnapLock Compliance and SnapLock Enterprise?	5
	Which file access protocols are supported?	5
	Which versions of Data ONTAP for SnapLock are currently supported?	5
<b>2</b>	<b>Licensing Questions</b>	<b>6</b>
	Can enterprise and compliance licenses coexist?	6
	What happens when I remove the SnapLock license or vol copy/mirror to a SnapLock volume on a system that does not have a SnapLock license?	6
<b>3</b>	<b>Questions About SnapLock Volumes</b>	<b>6</b>
	Can I convert an existing volume to a SnapLock volume?	6
	Can I use a WORM volume as a regular volume?	6
	Can I use Snapshot copies on a SnapLock volume?	6
	How do I identify SnapLock volumes in a system?	7
	How can I find the expiration date of the volume?	7
	Does SnapLock automatically delete expired WORM data?	7
	Is SnapLock supported with MultiStore <sup>®</sup> (vFiler units)?	7
<b>4</b>	<b>Questions About Hardware</b>	<b>7</b>
	Does SnapLock require any special hardware?	7
	Which hardware platform does SnapLock support?	7
	Does SnapLock support virtual storage appliances?	8
<b>5</b>	<b>Questions About Data ONTAP 8.1</b>	<b>8</b>
	With the inclusion of freebsd shell in Data ONTAP 8.x, are there any caveats for SnapLock?	8
	What are the new SnapLock features introduced in Data ONTAP operating in 8.1?	8
	Can I convert a 32-bit SnapLock aggregate to 64-bit aggregate?	8
	What are the commands that are deprecated in Data ONTAP 8.1? What are the equivalent commands?	8
	Can I enable compression on a SnapLock volume?	8
	Is it possible to enable compression and dedupe on a SnapLock volume?	8
	Is it possible to use SnapLock with 2TB drives and 64-bit aggregates?	8
	Are there any SnapLock APIs that have been deprecated in Data ONTAP 8.1?	9
<b>6</b>	<b>Backup of SnapLock Volumes</b>	<b>9</b>
	How do I protect/back up WORM data?	9
<b>7</b>	<b>ComplianceClock Behavior in Backup Relationships in Data ONTAP Releases Prior to 8.1</b>	<b>12</b>
	How does ComplianceClock behave with volume SnapMirror?	12
	How does ComplianceClock behave with qtree SnapMirror?	14
	How does ComplianceClock behave with SnapVault?	17
<b>8</b>	<b>ComplianceClock Behavior in Backup Relationships in Data ONTAP 8.1</b>	<b>17</b>
	How does ComplianceClock behave with volume SnapMirror?	17
	How does ComplianceClock behave with qtree SnapMirror?	20
	How does volume ComplianceClock behave in Data ONTAP 8.1?	22
<b>9</b>	<b>Application Integration</b>	<b>23</b>
	What does an application need to do to integrate with SnapLock?	23
	Can SnapLock be used without changing the application?	23

What granularity of data management does SnapLock allow? .....	24
What is WORM append? .....	24
How do I find out whether a particular application is certified with SnapLock? .....	24
<b>10 Reclaiming Disks .....</b>	<b>24</b>
I accidentally locked a file for 30 years. How can I reuse the disk? .....	24
How do I sanitize SnapLock disks?.....	24
<b>11 Miscellaneous Questions and Other Documents.....</b>	<b>24</b>
How do I find out more about ComplianceClock? .....	24
I have an application vendor who wants to integrate with SnapLock. How do I get that vendor started? ....	25
What regulations does SnapLock cover?.....	25
I do not see the regulation that I am looking for in the list. What do I do?.....	25
Are there Professional Services offered that are specific to SnapLock? .....	<b>Error! Bookmark not defined.</b>
Where can I get more information?.....	25
<b>Version History .....</b>	<b>25</b>

## LIST OF TABLES

Table 1) Supported Data ONTAP releases. ....	5
Table 2) Deprecated commands. ....	8
Table 3) VSM initialize.....	9
Table 4) VSM resync.....	10
Table 5) QSM resync.....	10
Table 6) Backup with SnapVault. ....	10
Table 7) Restore with SnapVault.....	11
Table 8) SnapMirror sync. ....	11
Table 9) SnapMirror resync.....	12
Table 10) VSM initialize command issued.....	13
Table 11) VSM initialize command completion time.....	13
Table 12) VSM break. ....	13
Table 13) VSM resync.....	14
Table 14) QSM update/break. ....	15
Table 15) QSM break. ....	15
Table 16) QSM resync.....	16
Table 17) ComplianceClock drift. ....	16
Table 18) VSM initialize command issued.....	18
Table 19) VSM initialize command completion time.....	18
Table 20) VSM update/break: example 1.....	19
Table 21) VSM update/break: example 2.....	19
Table 22) VSM resync.....	20
Table 23) QSM update/break.....	20

Table 24) QSM break.....	21
Table 25) QSM resync.....	21
Table 26) ComplianceClock drift.....	22

# 1 General Questions

## What is the difference between SnapLock Compliance and SnapLock Enterprise?

Answer: NetApp® SnapLock® Compliance (SLC) is designed to enable compliance in strictly regulated environments, such as those governed by SEC 17a-4. SnapLock Compliance provides an "untrusted storage administrator" model of operation in which the write once, read many (WORM) data, when committed to the storage system, is protected even from the storage administrator.

SnapLock Enterprise (SLE), in contrast, is designed for more flexible customer environments and operates under a "trusted storage administrator" model in which NetApp Data ONTAP® permits some administrative actions while still enforcing WORM protection.

**Note:** Even those with administrative access to the storage system are not permitted to modify individual files under WORM protection on a SnapLock Enterprise volume.

## Which file access protocols are supported?

Answer: Currently, the supported open file protocols for SnapLock are NFS (Versions 2, 3, and 4), Common Internet File System protocol (CIFS), FTP, and HTTP.

## Which versions of Data ONTAP for SnapLock are currently supported?

Answer: SnapLock is supported in all releases following Data ONTAP 6.4.1 except for the following releases:

- Data ONTAP 7.2.2L1 and 7.2.4L1 releases that were for the FAS2000 platform only.
- SnapLock is disabled in Data ONTAP 7.3.0 and Data ONTAP 8.0. Data ONTAP 7.3.1 and Data ONTAP 8.1.0 support SnapLock.

**Note:** A mandatory upgrade is in effect for SnapLock Compliance. Table 1 shows the list of supported releases for each codeline.

Table 1) Supported Data ONTAP releases.

Current Data ONTAP Release	Minimum Required Upgrade Release
6.5.x	7.0.7, 7.2.5.1
7.0.x	7.0.7, 7.2.5.1
7.1.x	7.1.3, 7.2.5.1
7.2.x	7.2.5.1
7.3.x	7.3.1P2
8.1.x	NA

**Note:** Official Data ONTAP feature matrixes are available at <http://support.netapp.com/matrix/mtx/login.do>. In addition, the field technical leads (FTLs) have a more detailed "What Works with What" matrix. Check with the appropriate FTL for details. The list of FTLs, along with their areas of expertise, is available at [http://se-web.corp.netapp.com/fieldtech\\_teams/ftl/](http://se-web.corp.netapp.com/fieldtech_teams/ftl/).

## 2 Licensing Questions

### Can enterprise and compliance licenses coexist?

Answer: Yes. SnapLock Compliance (license “snaplock”) and SnapLock Enterprise (license “snaplock\_enterprise”) can coexist on the same system.

### What happens when I remove the SnapLock license or vol copy/mirror to a SnapLock volume on a system that does not have a SnapLock license?

Answer: SnapLock continues to enforce the WORM property of the data regardless of the state of licensing. The licenses are required only for the creation of new SnapLock volumes and for committing new files to WORM state on the SnapLock volumes.

## 3 Questions About SnapLock Volumes

### Can I convert an existing volume to a SnapLock volume?

Answer: It is not possible to convert volumes directly (in place). If the data must be locked, it can be copied to a SnapLock volume. The simplest way to accomplish this is to use the `vol copy` command to an equivalently sized or larger SnapLock Enterprise or SnapLock Compliance volume. The destination compliance volume should be empty when `vol copy` is issued. Note that `vol copy` does not cause the files to be committed to WORM. When the data resides on the SnapLock volume:

- A retention period must be set for the files.
- The files must be committed to WORM.

The retention period can be set individually on each file either by changing the `atime` (last access time) on the file or by using volume defaults. The commit to WORM can be accomplished either by committing individual files to WORM by removing the write permissions on the files or by using the `autocommit` feature of SnapLock, which automatically locks the files placed on the volume after the designated time period.

A script or an application can be used to accomplish these tasks.

### Can I use a WORM volume as a regular volume?

Answer: No. The unique properties and behavior of SnapLock volumes make them unsuitable for use as regular data storage volumes. In most ways, SnapLock volumes behave identically to regular volumes, but there are some very specific and critical differences in functionality and administration that make them unsuitable for use as regular volumes. For example:

- Renaming directories on SnapLock volumes is not allowed.
- Making the transition of file attributes from writable to read only commits a file to WORM state.
- SnapLock volumes have the `no_atime_update` volume option set, and explicit updates to `atime` are interpreted as changes to the retention period of the file.
- Administrative actions are restricted on SnapLock volumes.

### Can I use Snapshot copies on a SnapLock volume?

Answer: Yes and no. There are no restrictions on creation, deletion, or number of Snapshot™ copies on SnapLock volumes beyond the normal system limits because risk of data loss does not occur with these operations. However, SnapRestore® cannot be used to restore to a previous Snapshot copy on a SnapLock Compliance volume because it violates the immutability of compliance data. It is allowed on a SnapLock Enterprise volume because the storage administrator is trusted.

### How do I identify SnapLock volumes in a system?

Answer: The regular `vol status` command can be used to identify SnapLock volumes. SnapLock volumes have an entry in the options field that identifies a volume as either `snaplock_compliance` or `snaplock enterprise`.

### How can I find the expiration date of the volume?

Answer: The expiration date is the maximum retention time of WORM files on a SnapLock volume. SnapLock Compliance volumes cannot be deleted until the expiration date, as measured by ComplianceClock™, has passed. Data ONTAP 7.3.1 introduces a new `-w` flag to the `vol status` command to display the expiration date of SnapLock Compliance and enterprise volumes. A value of "-" is displayed for non SnapLock volumes.

### Does SnapLock automatically delete expired WORM data?

Answer: No. SnapLock never automatically deletes any files from the volume. It is the responsibility of the applications to delete expired records.

### Is SnapLock supported with MultiStore® (vFiler units)?

Answer: Yes, SnapLock is supported on vFiler® units.

NetApp advises that the root volume should not be created on a SnapLock volume. The root volume for the vFiler unit should be created on a regular volume, and then SnapLock volumes should be added to the vFiler unit. Note that the `snaplock.autocommit_period` option is not available on vFiler units. The option must be set in vFiler0 context and is applicable to all other vFiler units. Also, the privileged delete feature added in Data ONTAP 7.3.1 is not available on volumes owned by vFiler units.

Starting with Data ONTAP 8.1, the autocommit feature has been redesigned to be managed at the FlexVol® volume level. This feature can be configured for individual volumes using the option `snaplock_autocommit_period` setting. This provides the administrator the option of setting different criteria for automatically committing files to WORM within a NetApp storage controller.

## 4 Questions About Hardware

### Does SnapLock require any special hardware?

Answer: No. SnapLock is a licensed feature of Data ONTAP and is supported on almost all NetApp hardware platforms. (Refer to "Which versions of Data ONTAP for SnapLock are currently supported?" under "General Questions," General Questions earlier in this FAQ.) Note that SnapLock is *not* a software-only solution. It is an integrated hardware and software solution. This is important for strict WORM regulations such as SEC 17a-4, which requires an integrated hardware and software solution. For more information, refer to [SEC Interpretation: Electronic Storage of Broker-Dealer Records](#).

### Which hardware platform does SnapLock support?

Answer: Support to hardware platform depends on the SnapLock edition.

SnapLock Compliance:

- Supports all current NetApp FAS and NearStore® platforms.
- Supports V-Series (formerly known as gFiler®) systems, available only on NetApp storage, and requires Data ONTAP 7.2.5.1 or a later supported release. SnapLock Compliance is not supported on heterogeneous (third-party) storage.
- SnapLock Compliance supports the FAS2000 series beginning with Data ONTAP 7.2.5.1.

SnapLock Enterprise:

- Is supported on all current NetApp FAS and NearStore platforms.
- Support for V-Series systems is available on both NetApp and heterogeneous storage. It requires Data ONTAP 7.2.5.1 or a later supported release.
- SnapLock Enterprise is supported on the FAS2000 series beginning with Data ONTAP 7.2.5.1.

**Does SnapLock support virtual storage appliances?**

Answer: No.

## 5 Questions About Data ONTAP 8.1

**With the inclusion of freebsd shell in Data ONTAP 8.x, are there any caveats for SnapLock?**

Answer: No. SnapLock is part of Data ONTAP, which runs in D-blade. FreeBSD shell users such as “root” or “diag” cannot unload D-blade. The FreeBSD shell users cannot tamper with SnapLock data or disks.

For additional information about FreeBSD shell and security, refer to [http://wikid.netapp.com/w/SnapLock/Questions\\_on\\_FreeBSD\\_security](http://wikid.netapp.com/w/SnapLock/Questions_on_FreeBSD_security).

**What are the new SnapLock features introduced in Data ONTAP operating in 8.1?**

Answer: The new SnapLock features are as follows:

- Volumes ComplianceClock
- Autocommit option for each volume
- Ability to enable compression on SnapLock volumes

**Can I convert a 32-bit SnapLock aggregate to 64-bit aggregate?**

Answer: Yes.

**What are the commands that are deprecated in Data ONTAP 8.1? What are the equivalent commands?**

Answer: Table 2 describes the commands.

Table 2) Deprecated commands.

Before Data ONTAP 8.1	Equivalent Command in Data ONTAP 8.1
date -c initialize	snaplock clock initialize
date -c	snaplock clock status

**Can I enable compression on a SnapLock volume?**

Answer: Yes. Compression is supported only on 64-bit SnapLock volumes.

**Is it possible to enable compression and dedupe on a SnapLock volume?**

Answer: Yes. Compression and dedupe can be enabled on a SnapLock volume. This is supported only on 64-bit volumes.

**Is it possible to use SnapLock with 2TB drives and 64-bit aggregates?**

Answer: Yes. Starting with Data ONTAP 8.1, 2TB drives are supported on SnapLock volumes. SnapLock is not limited to the size for a drive. If a particular platform supports a drive, SnapLock will support it as well.



## Are there any SnapLock APIs that have been deprecated in Data ONTAP 8.1?

Answer: Yes. The SnapLock API `snaplock-get-compliance-clock` is deprecated. You must use the APIs `snaplock-get-system-compliance-clock` to get the system `ComplianceClock` time and `snaplock-get-volume-compliance-clock` to get the volume `ComplianceClock` time.

## 6 Backup of SnapLock Volumes

### How do I protect/back up WORM data?

Answer: Several options are available. Common workflows associated with data protection are backup and restore or resync.

- Backup: Creating backup copies of data
- Restore or resync: Restoring the primary copy from a backup copy of data

In certain scenarios, the backup copy of data might be used to serve data when the primary copy is unavailable. Resync refers to the act of merging changes from the backup copy to the primary copy once the primary site has recovered. Both restore and resync require changing the primary copy based on the secondary copy. Note that the source and destination roles can be reversed during resync. In the following tables, these are identified as “source” and “destination.”

The remaining answer lists restrictions on the common workflows associated with protection of WORM data.

- Volume SnapMirror® (VSM): SnapLock supports VSM. The following details the restrictions on backup and restore scenarios using VSM.
  - Backup: It is not possible to mirror a regular volume or a SnapLock Enterprise volume to a SnapLock Compliance volume. However, it is possible to mirror a regular volume to an SLE volume. SnapMirror does not automatically do WORM commits for files transferred in such a configuration (regular volume to SLE volume). Also, there are no restrictions on mirroring a SnapLock volume (compliance or enterprise) to a non SnapLock volume. Table 3 summarizes the options.

Table 3) VSM initialize.

		VSM Destination		
		SLC	SLE	Regular Volume
VSM Source	SLC	✓	✓	✓
	SLE	✗	✓	✓
	Regular Volume	✗	✓	✓

**Note:** The previous text and table apply to VSM only. Also, although VSM relationships to mirror data from a regular volume to an SLC volume are not allowed, a `vol copy` command (which is like a “one-time VSM”) to do the same is allowed.

- Resync: Restores from a copy residing on SLC to a primary copy also residing on SLC are not enabled. Table 4 summarizes the combinations.

Table 4) VSM resync.

		VSM Destination		
		SLC	SLE	Regular Volume
VSM Resync Source	SLC	x	x	x
	SLE	✓	✓	✓
	Regular Volume	✓	✓	✓

Qtree SnapMirror (QSM): SnapLock is supported with QSM.

- **Backup:** All combinations of SnapLock Compliance, SnapLock Enterprise, and regular volumes can be used to establish a QSM relationship.
- **Resync:** Restore or resync is not allowed when the source is a non SnapLock or SLE volume and the destination is an SLC volume. Table 5 shows all the combinations that are allowed.

Table 5) QSM resync.

		QSM Resync Destination		
		SLC	SLE	Regular Volume
QSM Resync Source	SLC	✓	✓	✓
	SLE	x	✓	✓
	Regular Volume	x	✓	✓

**Note:** When the resync destination is a SnapLock Compliance volume, as part of the restore the data that is changed on that copy is archived into a dump file stored under an etc directory in the SnapLock Compliance volume itself.

For example:

```
/vol/<vol_name>/etc/log/snapmirror_resync_archive/slcsec_1374c60e-44ba-11d9-9991-0050560669e7_qd
```

SnapVault® SnapLock is supported with SnapVault. SnapVault can be used with SnapLock for Snapshot copies to manage the retention of data on the SnapVault secondary at the Snapshot level as opposed to managing it at the volume or individual file level.

- **Backup:** SnapVault follows the restrictions of QSM, except that copies of SnapLock Compliance source volumes cannot be created on a SnapLock Compliance destination volume.

Table 6) Backup with SnapVault.

		SnapVault Source		
		SLC	SLE	Regular Volume
SnapVault Destination	SLC	x	✓	✓
	SLE	✓	✓	✓

		SnapVault Source		
	Regular Volume	✓	✓	✓

- **Restore:** Restore is not allowed when the source is a non SnapLock or SLE volume and the destination is an SLC volume. Table 7 shows all the allowed combinations.

Table 7) Restore with SnapVault.

		SnapVault Restore Destination		
		SLC	SLE	Regular Volume
SnapVault Restore Source	SLC	NA	✓	✓
	SLE		✓	✓
	Regular Volume	x	✓	✓

- **Dump and NDMP:**

- **Backup:** Dump and NDMP backups work exactly as they would on regular volumes. In addition, the data streams of these features have been augmented to preserve the WORM attributes of files on a SnapLock volume when backing up, restoring, or copying data.
- **Restore:** Restore and NDMP restores also work as on regular volumes, with the following two restrictions:

True incremental restores are not supported on SnapLock volumes. The restore command automatically converts the restore r to restore x if the destination volume is a SnapLock volume.

Restore and NDMP are bound by the same restrictions that apply to regular writes. This means that restore and NDMP cannot overwrite, modify, delete, or change attributes of any existing WORM files on a SnapLock volume. As a result, the restores must be done to a new SnapLock volume.

**Note:** For the WORM attributes to be meaningfully enforced, the restore must also be to a SnapLock volume. If a backup from a SnapLock volume is restored to a non SnapLock volume, the WORM attributes are preserved but are ignored and not enforced by Data ONTAP.

- **SnapMirror Sync:**

- **Backup:** SnapMirror Sync is only supported for SnapLock Enterprise volumes. Table 8 shows all the combinations that are supported.

Table 8) SnapMirror sync.

		SnapMirror Sync Destination		
		SLC	SLE	Regular Volume
SnapMirror Sync Source	SLC	x	x	x
	SLE	x	✓	✓
	Regular Volume	x	✓	✓

- **Resync:** Table 9 shows all the allowed combinations.

Table 9) SnapMirror resync.

		SnapMirror Sync: Resync Destination		
		SLC	SLE	Regular Volume
SnapMirror Sync Resync Source	SLC	x	x	x
	SLE	x	✓	✓
	Regular Volume	x	✓	✓

- **MetroCluster:**

- **Backup:** MetroCluster™ can be used to mirror both SnapLock Enterprise and SnapLock Compliance aggregates from one site to another. SnapLock volumes retain their characteristics at the mirror site.
- **Resync:** Resync from a copy residing on SnapLock Compliance to a primary copy also residing on a SnapLock Compliance volume is not enabled. The combinations are the same as for volume SnapMirror restores.

**Note:** RAID SyncMirror® behaves the same as MetroCluster.

## 7 ComplianceClock Behavior in Backup Relationships in Data ONTAP Releases Prior to 8.1

### How does ComplianceClock behave with volume SnapMirror?

Answer: Volume SnapMirror is block-for-block replication; it transfers the file system verbatim. The ComplianceClock value is transferred over from the source to the destination volume. If the source ComplianceClock value is behind the current ComplianceClock value at the destination system, then the destination system ComplianceClock value is drifted back to match the most recently transferred ComplianceClock value.

Typically, a volume SnapMirror relationship involves the following workflow:

- **Initialize:** Initializing the SnapMirror relationship between the source and destination volume. This causes the first data transfer from the source to the destination volume. The destination should have the source file system at the end of the operation.
- **Update:** Periodic SnapMirror updates between the source and destination volumes. This causes the file system changes to be propagated from the source to the destination volume.
- **Break:** Breaks the SnapMirror relationship between the source and destination volumes.
- **Resync:** Resynchronizes a broken-off destination to its former source.

ComplianceClock behavior during these operations includes:

- **Initialize**

For a SnapMirror initialize to take place, the destination volume needs to be in the restricted state. Restricted volumes have the effect of pulling back the system's current ComplianceClock value when they are brought online. That is because on completion of the initialize step the current ComplianceClock value on the destination system is drifted back to match the ComplianceClock value of the now-online destination volume. This drift is approximately equal to the time taken for the initialize step.

Table 10) VSM initialize command issued.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror initialize issued	01:00	01:00	01:00
SnapMirror initialize completion time	01:15	01:15	01:00

In Table 10, note that the `snapmirror initialize` took 15 minutes to complete. This had the effect of drifting the ComplianceClock value on the destination system back by 15 minutes.

We know that during the initialize step, the ComplianceClock value from the source is copied over to the destination volume. If this copied-over ComplianceClock value is behind the current ComplianceClock value on the destination system, then the ComplianceClock value on the destination system is drifted back to match the copied-over ComplianceClock value.

Table 11) VSM initialize command completion time.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror initialize issued	02:00	01:00	02:00
SnapMirror initialize completion time	02:15	01:15	01:00

In Table 11 the ComplianceClock value copied over was 01:00; therefore, when `snapmirror initialize` completed, the destination ComplianceClock value was drifted back to match 01:00.

- **Update**

A SnapMirror update propagates the file system changes between source and destination volumes, which include the ComplianceClock value. If this value is behind the current ComplianceClock value on the destination system and the relationship is subsequently broken, then the ComplianceClock value on the destination system is drifted back to match the propagated ComplianceClock value.

- **Break**

If the last propagated ComplianceClock value is behind the current ComplianceClock value on the destination system, then breaking the SnapMirror relationship causes the ComplianceClock value on the destination system to drift back to match the last propagated ComplianceClock value from the source.

Table 12) VSM break.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror initialize issued	02:00	01:00	02:00
SnapMirror initialize completion time	02:05	01:05	01:00
SnapMirror update issued	03:05	01:05 (assume ComplianceClock time not updated for some reason such as a volume being offline)	02:00

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror update completion time	03:06	01:05	02:01
SnapMirror break issued	03:30	01:05	01:05

In Table 12, we assume that the ComplianceClock value on the source system does not get updated starting at step 2. Notice that the last ComplianceClock value propagated to the destination system is 01:05 (step 3); therefore, when the SnapMirror break is issued, the ComplianceClock value on the destination system is drifted back to match 01:05.

- **Resync**

A SnapMirror resync involves transmitting data from source to destination to reestablish the SnapMirror relationship. This includes the ComplianceClock value; thus if the ComplianceClock value transferred over is behind the current ComplianceClock value on the destination system and the relationship is subsequently broken, then the ComplianceClock value on the destination will be drifted back to match the transferred-over ComplianceClock value.

Table 13) VSM resync.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror resync issued	02:00	01:00	02:00
SnapMirror resync completion time	02:05	01:05	02:05
SnapMirror break issued	02:30	01:30	01:00

In Table 13 the latest ComplianceClock value transmitted from source to destination is 01:00. The ComplianceClock value on the destination drifts from 02:30 to 01:00 when the SnapMirror break is issued.

#### Caveat

Prior to Data ONTAP 7.0.1, all volume combinations "non SnapLock to non SnapLock," "SnapLock to non SnapLock," and "non SnapLock to SnapLock" could result in a ComplianceClock skew on the destination systems if the source ComplianceClock value was behind the current ComplianceClock value on the destination system. The ComplianceClock skews caused by the combination of non SnapLock to non SnapLock and SnapLock to non SnapLock had no compliance reason and hence were fixed (to not cause a skew) in Data ONTAP 7.0.1. Note that the combination of "non SnapLock to SnapLock Enterprise" will cause a ComplianceClock skew (if source ComplianceClock value is behind destination ComplianceClock value) and that "non SnapLock to SnapLock Compliance" is not a supported configuration.

#### How does ComplianceClock behave with qtree SnapMirror?

Answer: SnapMirror qtree replication is logical replication; all of the files and directories in the source file system are created on the destination file system. The destination qtree is read only, but the volume on which it is located must be online and writable.

In qtree SnapMirror the ComplianceClock time is sent from the source to destination only if the replication is from a SnapLock volume to a SnapLock volume. When the source ComplianceClock value is behind the destination, unlike volume SnapMirror, which drifts the destination ComplianceClock value to match the last propagated source ComplianceClock value, qtree SnapMirror aims to drift the destination ComplianceClock value to match the current ComplianceClock value on the source system.

Following are ComplianceClock behaviors during initialize, update, break, and resync, in the context of qtree SnapMirror.

- **Initialize**

Unlike volume SnapMirror, in qtree SnapMirror the destination volume does not need to be in the restricted state. This means that lengthy SnapMirror initialize operations will not cause a ComplianceClock drift on the destination system.

However, if the ComplianceClock value in the source qtree is behind the current ComplianceClock value on the destination system, then the idea is to drift the ComplianceClock value on the destination system to match the current ComplianceClock value on the source system when the relationship is broken. Note that volume SnapMirror in such a situation drifts the ComplianceClock value on the destination system to match the ComplianceClock value that was transferred over; however, in qtree SnapMirror the idea is to match the current ComplianceClock value on the source system.

This drift does not reflect upon completion of the SnapMirror initialize operation but rather on breaking the SnapMirror relationship.

- **Update**

A SnapMirror update propagates the file system changes between source and destination qtrees, which include the ComplianceClock value. If this value is behind the current ComplianceClock value on the destination system and the relationship is subsequently broken, then the ComplianceClock value on the destination system is drifted back to match the current ComplianceClock value on the source system.

- **Break**

On breaking the SnapMirror relationship, ComplianceClock changes, if any, reflect on the destination system.

Table 14) QSM update/break.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror initialize issued	02:00	01:00	02:00 (01:00 gets copied to destination)
SnapMirror initialize completion time	02:05	01:05	02:05
SnapMirror update issued	02:30	01:30	02:30 (01:30 gets copied to destination)
SnapMirror update completion time	02:31	01:31	02:31
SnapMirror break issued (quiesced first)	02:45	01:45	01:45

In Table 14, during initialize the ComplianceClock value transferred over from the source was 01:00. Thirty minutes later a SnapMirror update is issued; the previous value of 01:00 is overwritten by 01:30. Fifteen minutes later a SnapMirror break is issued; this period is added to the most recent ComplianceClock value from the source (01:30), and then the destination ComplianceClock value is drifted back to match it (01:45).

Table 15) QSM break.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
-----------	--	---------------------------------------	--

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror initialize issued	02:00	01:00	02:00
SnapMirror initialize completion time	02:05	01:05	02:05
SnapMirror break issued (quiesced first)	02:30	01:30	01:30

The example in Table 15 is similar to Table 14. No update is issued only at this point. The most recent and only ComplianceClock value from the source is 01:00, and the total wait period before the SnapMirror break is 30 minutes, setting the destination ComplianceClock value to 01:30.

- **Resync**

The underlying idea is the same. If the ComplianceClock value transferred over from the source is behind the current ComplianceClock value on the destination (and the relationship is subsequently broken), then the ComplianceClock value on the destination system is drifted back to match the current ComplianceClock value on the source system.

Table 16) QSM resync.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror resync issued	02:00	01:00	02:00
SnapMirror resync completion time	02:05	01:05	02:05
SnapMirror break issued (quiesced first)	02:30	01:30	01:30

**Note:** The destination ComplianceClock value is *Not* in sync with source ComplianceClock value.

It is possible that the destination system ComplianceClock value will *not* be in sync with the source at the time of the break, as demonstrated in Table 17.

Table 17) ComplianceClock drift.

Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror initialize issued	02:00	01:00	02:00
SnapMirror initialize completion time	02:05	01:05	02:05
SnapMirror update issued	02:30	01:30	02:30 (01:30 gets copied to destination)
SnapMirror update completion time	02:31	01:31	02:31
<b>Note:</b> ComplianceClock time on the source drifts back to 00:00 after 15 minutes of the update being issued, due to some reason.			



Operation	Actual System Clock (Source and Destination)	ComplianceClock Time on Source System	ComplianceClock Time on Destination System
SnapMirror break issued	03:00	00:15	02:00 (01:30 [copied during last update] + 30 minutes)

When the update is issued, the ComplianceClock value transferred over from the source to the destination is 01:30. The relationship is broken 30 minutes later; therefore, the destination gets a ComplianceClock time of 02:00 (01:30 + 30). The source, in contrast, drifted back to 00:00 due to some reason (say, a volume previously offline coming online) 15 minutes after the update was issued. Therefore, at 02:45 (actual time) the ComplianceClock time on the source drifted back to 00:00. The relationship was broken 15 minutes after the drift took place (03:00 actual time); therefore, the ComplianceClock time on the source advanced 15 minutes to 00:15.

#### Caveat

- To allow a qtree SnapMirror resync between the SnapLock Compliance qtrees after upgrading Data ONTAP, you need to make at least one update from the qtree source to the destination volume before attempting the resync.
- In qtree SnapMirror the ComplianceClock value is sent from the source to the destination only if the replication is from a SnapLock volume to a SnapLock volume. Therefore, a volume combination of “non SnapLock to SnapLock” will *not* cause a ComplianceClock skew (if the source ComplianceClock value is behind the destination ComplianceClock value), which is unlike volume SnapMirror behavior (see the volume SnapMirror caveat).

#### How does ComplianceClock behave with SnapVault?

Answer: Even though SnapVault uses qtree SnapMirror under the covers, ComplianceClock does not behave the same way. In qtree SnapMirror, in order to do clock adjustments at “SnapMirror break” time, the source ComplianceClock value is stored as metadata on the destination qtree. In SnapVault the source ComplianceClock value is *not* stored as metadata on the destination qtree, making ComplianceClock drifts impossible. This effectively means that when using SnapVault, if the source ComplianceClock value is behind the current ComplianceClock value on the destination, there will be no drift in the destination ComplianceClock value. In Data ONTAP 8.1, the later references to ComplianceClock mean volume ComplianceClock.

## 8 ComplianceClock Behavior in Backup Relationships in Data ONTAP 8.1

#### How does ComplianceClock behave with volume SnapMirror?

Answer: Starting with Data ONTAP 8.1, we have introduced the new volume ComplianceClock software apart from the system ComplianceClock software. During SnapLock volume creation, volume ComplianceClock will use system ComplianceClock as a reference clock. Volume SnapMirror does a block-level copy from the source to the destination. The following cases are possible with VSM:

- **Both source and destination are SnapLock volumes.** The source would send its computed in-core volume ComplianceClock time to the destination. As a result, the destination volume ComplianceClock time would end up to be the same as the source volume ComplianceClock time. The destination volume ComplianceClock time would be updated with every SnapMirror update. If the mirroring relationship is broken, the destination volume will be mounted read-write, and its volume ComplianceClock software will start operating using the destination system ComplianceClock time as the reference. Therefore, no skews would be introduced as a result of the break.

- **Source is SnapLock, but destination is non SnapLock.** The destination volume's ComplianceClock metadata will not change due to VSM.
- **Source is non SnapLock, and destination is SnapLock.** The volume ComplianceClock value would not change in the destination.
- **Initialize:** For a SnapMirror initialize to take place, the destination volume needs to be in the restricted state.

Table 18) VSM initialize command issued.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror initialize issued	01:00	01:00	01:00	01:00
SnapMirror initialize completion time	01:15	01:15	01:15	01:00

During the initialize step, the volume ComplianceClock value from the source is copied over to the destination volume. In Data ONTAP 8.1, if the mirror relationship is broken, the destination volume will be mounted read-write, and its VCC will start operating using the destination SCC time as the reference. Therefore, no skews would be introduced as a result of break.

Table 19) VSM initialize command completion time.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror initialize issued	02:00	01:00	01:00	02:00
SnapMirror initialize completion time	02:15	01:15	01:15	01:00

- **Update**

A SnapMirror update propagates the file system changes between source and destination volumes, which include the volume ComplianceClock value.

- **Break**

If the SnapMirror relationship is broken, the destination volume will be mounted read-write, and its volume ComplianceClock software will start operating using the destination system ComplianceClock time as the reference. Therefore, no skews would be introduced as a result of break.

Table 20) VSM update/break: example 1.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror initialize issued	02:00	01:00	01:00	02:00
SnapMirror initialize completion time	02:05	01:05	01:05	01:00
SnapMirror update issued	03:05	01:05 (assume ComplianceClock not updated for some reason such as a volume being offline)	02:05	01:00
SnapMirror update completion time	03:06	01:05	02:06	02:06
SnapMirror break issued	03:30	01:05	01:29	01:29

Table 21) VSM update/break: example 2.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror initialize issued	02:00	01:00	01:00	02:00
SnapMirror initialize completion time	02:05	01:05	01:05	01:00
SnapMirror update issued	03:05	01:05 (assume ComplianceClock not updated for some reason such as a volume being offline)	02:05	01:00
SnapMirror update completion time	03:06	01:05	02:06	02:06
SnapMirror break issued	03:30	01:05	01:29	01:29

- **Resync**

A VSM resync will be treated similar to a VSM update. In the case of a VSM flip-resync, the roles of the source and destination would get reversed.

Table 22) VSM resync.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror resync issued	02:00	01:00	01:00	02:00
SnapMirror resync completion time	02:05	01:05	01:05	01:00
SnapMirror break issued	02:30	01:30	01:30	01:30

**How does ComplianceClock behave with qtree SnapMirror?**

Answer: SnapMirror qtree replication is logical replication; all of the files and directories in the source file system are created on the destination file system. The destination qtree is read only, but the volume on which it is located must be online and writable. In qtree SnapMirror the ComplianceClock time is sent from the source to destination only if the replication is from a SnapLock volume to a SnapLock volume.

- **Initialize**

Unlike volume SnapMirror, in qtree SnapMirror the destination volume does not need to be in the restricted state.

- **Update**

A QSM SnapMirror update propagates the file system changes between source and destination qtrees, which include the volume ComplianceClock value. With every qtree SnapMirror update, the source will send its volume ComplianceClock value to the destination. The source volume ComplianceClock time and the destination in-core volume ComplianceClock time will be stored in the qtree metafile.

- **Break**

At the time of QSM break, the destination will compute the updated source volume ComplianceClock time based on the stored timestamps in the qtree metafile and the current destination in-core volume ComplianceClock time. If the computed source volume ComplianceClock time is less than the current destination volume ComplianceClock time, the destination volume ComplianceClock time will be set to the computed source volume ComplianceClock time.

Table 23) QSM update/break.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror initialize issued	02:00	01:00	01:00	02:00 (01:00 gets copied to destination qtree metafile)
SnapMirror initialize completion time	02:05	01:05	01:05	02:05

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror update issued	02:30	01:30	01:30	02:30 (01:30 gets copied to destination)
SnapMirror update completion time	02:31	01:31	01:31	02:31
SnapMirror break issued (quiesced first)	02:45	01:45	01:45	01:45

The example in Table 24 is similar to Table 23; no update is issued only at this point.

Table 24) QSM break.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror initialize issued	02:00	01:00	01:00	02:00 (01:00 gets copied to destination qtree metafile)
SnapMirror initialize completion time	02:05	01:05	01:05	02:05
SnapMirror break issued (quiesced first)	02:30	01:30	01:30	01:30

- **Resync**

A QSM resync will be treated similarly to a QSM update. In the case of a VSM flip-resync, the roles of the source and destination would get reversed.

Table 25) QSM resync.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror resync issued	02:00	01:00	01:00	02:00 (01:00 gets copied to destination qtree metafile)
SnapMirror resync completion time	02:05	01:05	01:05	02:05

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror break issued (quiesced first)	02:30	01:30	01:30	01:30

**Note:** The destination ComplianceClock value *Not* in sync with source ComplianceClock value.

It is possible that the destination volume ComplianceClock value will *not* be in sync with the source at the time of the break, as demonstrated in Table 25.

Table 26) ComplianceClock drift.

Operation	Actual System Clock (Source and Destination)	ComplianceClock on Source System (Before Data ONTAP 8.1)	ComplianceClock on Source Volume (Starting with Data ONTAP 8.1)	ComplianceClock on Destination Volume (Starting with Data ONTAP 8.1)
SnapMirror initialize issued	02:00	01:00	01:00	02:00
SnapMirror initialize completion time	02:05	01:05	01:05	02:05
SnapMirror update issued	02:30	01:30	01:30	02:30 (01:30 gets copied to destination)
SnapMirror update completion time	02:31	01:31	01:31	02:31
<b>Note:</b> The volume ComplianceClock value on the source drifts back to 00:00 after 15 minutes of the update being issued.				
SnapMirror break issued	03:00	00:15	00:15	02:00 (01:30 [copied during last update] + 30 minutes)

### How does volume ComplianceClock behave in Data ONTAP 8.1?

Answer: The volume ComplianceClock behavior for volume SnapMirror, aggregate copy, or volume copy changes when the source and the destination are running the same or different Data ONTAP versions:

- Source and destination storage system running different Data ONTAP versions:
  - If the source and destination are SnapLock volumes, the destination volume ComplianceClock time is updated as per the source ComplianceClock time.
  - If the source is a non SnapLock volume or aggregate and destination is a SnapLock volume or aggregate, the destination volume ComplianceClock time is updated as per the source ComplianceClock value when the SnapMirror relationship is broken.

- Source and destination storage system running the same Data ONTAP version:
  - If the source and the destination are SnapLock volumes, the volume ComplianceClock time of the destination volume is updated as per the source volume.
  - If the source is a non SnapLock volume and the destination is a SnapLock volume, the volume ComplianceClock time of the destination volume is reinitialized according to the destination system ComplianceClock time.

## 9 Application Integration

SnapLock enables WORM capability; it is typically used with either custom scripts or some other application that manages the WORM data.

### What does an application need to do to integrate with SnapLock?

Answer: SnapLock is very easy to integrate with other applications because it allows the use of standard open protocols (NFS and CIFS) to set and manage the WORM data. It does this by using the `atime` (last access time stamp) file attribute to represent the retention period for the file. It also uses the removal of write access on the file to trigger the commit to WORM. Typically, applications do the following:

1. Select the files that must be retained for a certain time period (this is typically dictated by the governing regulations).
2. Select the retention period (this too is typically dictated by regulations). The retention period can be set on a file basis (allowing file-level granularity), or volume-level defaults can be used to set the retention period on files that do not specify a retention period and that reside on the volume.
3. Commit the files to WORM status. This can be done either at an individual file level (by removing the write permissions on the file) or by using the `autocommit` feature to automatically commit to WORM files that have not changed for a specified period of time.
4. When the retention period has expired (that is, the ComplianceClock value has surpassed the value of the `atime`), those files can be deleted.

**Note:** Many of these tasks can be automated.

### Can SnapLock be used without changing the application?

Answer: Starting with Data ONTAP 7.2, a time-delay commit to WORM with an adjustable timer is available. The administrator can set the global option `snaplock.autocommit_period` to specify the delay. This starts a low-priority scanner that scans the WORM volumes on the system for candidate files that have not changed for the specified period. When it finds any file that matches the criteria, it automatically commits the file to WORM. The minimum delay that can be specified is two hours. This feature, also called *autocommit*, permits existing applications to be used with SnapLock without any application changes.

**Note:** Autocommit is not suited for all cases because it is very difficult to know when an application has completed writing to the file. If the commit to WORM occurs prematurely, it leaves behind a nonmodifiable file, which the application might not accept.

Autocommit can take significantly longer to commit files to WORM than the time specified in the `autocommit` period. This is because files are committed to WORM using a low-priority scanner.

It can take longer to commit files to WORM under either or both of the following conditions:

- The volume has a large number of files.
- The system has a large amount of traffic on it.

Starting with Data ONTAP 8.1, this option has been moved to the volume level and needs to be configured separately for all individual SnapLock volumes.

### **What granularity of data management does SnapLock allow?**

Answer: SnapLock can use the retention period set at an individual file level. It can also use the volume default retention period to automatically set the retention period on individual files that do not have an explicit retention period set on them. The default retention period allows the application to use the same retention period on all files residing in the same volume by not explicitly setting a retention period value.

### **What is WORM append?**

Answer: SnapLock has a special feature for applications that need to store logs of events on tamperproof WORM storage. Applications can create a WORM append file, which receives special treatment by SnapLock. On a WORM append file, the data is committed to WORM state in 256K chunks. After the data is written to byte  $n*256K+1$  of the file, the previous 256K segment is committed to WORM and cannot be altered.

### **How do I find out whether a particular application is certified with SnapLock?**

Answer: Application vendors can choose to certify an application with SnapLock. Refer to <http://tme-webtools1/CerTrac/CerTracMenu.aspx> for a list of applications that have been certified. For an alternate recent link, visit [http://my.sharepoint.com/personal/netapp\\_itimothy/Lists/SnapLock%20Partners%20Page/AllItems.aspx](http://my.sharepoint.com/personal/netapp_itimothy/Lists/SnapLock%20Partners%20Page/AllItems.aspx).

**Note:** This list contains applications that are certified by NetApp. There are many more applications that are integrated with SnapLock but have not gone through the official certification process.

## **10 Reclaiming Disks**

### **I accidentally locked a file for 30 years. How can I reuse the disk?**

Answer: If the file is on a SnapLock Compliance volume, then the only way to get the disks back is by waiting for the file to expire. There is no way to bypass the retention period; having said that, NetApp might have a partner that can degauss the disks. For more information, contact the SnapLock product manager.

SnapLock Enterprise volumes can be destroyed at any time by the system administrator. By using the privileged delete feature (available in Data ONTAP 7.3.1), individual files can be deleted and operation logged into a WORM log by the system administrator on a SnapLock Enterprise volume.

### **How do I sanitize SnapLock disks?**

Answer: For SnapLock Compliance volumes, when all the WORM data has expired, the volume can be destroyed. For SnapLock Enterprise, the system administrator can destroy the volumes at any time. In the case of FlexVol volumes, when all SnapLock flexible volumes on an aggregate are destroyed, the aggregates can also be destroyed. After the destruction of the traditional volume or the aggregate, the disks can be sanitized by using the `disk sanitize` command.

## **11 Miscellaneous Questions and Other Documents**

The following section provides references for additional information.

### **How do I find out more about ComplianceClock?**

Answer: Refer to "TR-3618: Understanding SnapLock ComplianceClock" by at <http://media.netapp.com/documents/tr-3618.pdf>.



## I have an application vendor who wants to integrate with SnapLock. How do I get that vendor started?

Answer: Contact the SnapLock technical marketing engineer (TME), who can explain the integration process to the vendor.

## What regulations does SnapLock cover?

Answer: Many regulations do not have clear requirements for the storage layer. For an assessment of the regulations that apply to SnapLock, visit

[http://my.sharepoint.com/personal/netapp\\_amanish/Lists/Regulations/AllItems.aspx](http://my.sharepoint.com/personal/netapp_amanish/Lists/Regulations/AllItems.aspx). Most regulations can be viewed as a subset of the SEC 17a-4 regulation in the United States. SnapLock has been assessed by an independent third party (Cohasset), and it has been found to meet the SEC 17a-4 regulations.

**Note:** The marketing team for SnapLock is collecting reference material and customers for each regulation. This information will help with new sales. If you have sold SnapLock to comply with any regulation, inform the SnapLock TME.

## I do not see the regulation that I am looking for in the list. What do I do?

Answer: The guidelines are as follows:

- If the regulation has an explicit certification process, the SnapLock team can evaluate the business case to fund a certification. Very few regulations have an associated certification process. Of all the places where SnapLock has been used, only Germany and Korea have explicit certifications.
- Alternatively, if the regulation has very detailed requirements and the target customer segment requires additional documentation, NetApp can hire an external expert to do an assessment, as Cohasset did for SEC 17a-4.
- For the rest, SnapLock engineering, product management, marketing, and technical marketing can review the regulation language and assess whether SnapLock meets the requirements.

## Where can I get more information?

Answer: For more information, contact [dl-compliance](#). Some of the additional references are as follows:

- [TR-3618: Understanding SnapLock ComplianceClock](#)
- [TR-3263: WORM Storage on Magnetic Disks Using SnapLock Compliance and SnapLock Enterprise](#)
- [TR-3446: SnapMirror Async Overview and Best Practices Guide](#)
- [TR-3752: Hardware Upgrade of WORM Data](#)
- Detailed FAQ: <http://wikid.netapp.com/w/SnapLock/FAQs>

**Note:** Unless stated otherwise, the answers can be found here: <http://wikid.netapp.com/w/SnapLock/FAQs>

## Version History

Version	Date	Document Version History	Author Information
Version 1.0	February 2009	Initial release	Manish Agarwal
Version 1.1	September 2009	Updated to include section on ComplianceClock behavior In backup relationships	Timothy Isaacs
Version 1.2	January 2012	Updated to include sections on questions about Data ONTAP 8.1	Lokesh A P and Ajay T. Someshwar

Refer to the [Interoperability Matrix Tool](#) (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

[Go further, faster®](#)

