



Technical Report

Compliant WORM storage using NetApp SnapLock

ONTAP 9

Jeannine Walter, Dan Tulledge, NetApp
January 2023 | TR-4526

Abstract

Many businesses rely on some use of write once, read many (WORM) data storage to meet regulatory compliance requirements or simply to add another layer to their data protection roadmap. This document discusses the integration of NetApp® SnapLock® software, the NetApp WORM solution in NetApp ONTAP® 9, into the environments that require WORM data storage.

TABLE OF CONTENTS

Introduction	4
SnapLock fundamentals	4
What are SnapLock Compliance and SnapLock Enterprise?	5
What is Snapshot copy locking?	5
Using SnapLock.....	6
Licensing	6
Initializing SnapLock ComplianceClock	6
SnapLock volume and aggregate creation	7
SnapLock volume usage	8
SnapLock volume append mode	10
Application integration	11
Best practices for SnapLock	12
ComplianceClock best practices.....	12
SnapLock Compliance testing	13
Creating and growing production SnapLock volumes.....	13
SnapLock volume minimum, maximum, and default retention period values	13
Unspecified retention.....	14
Data protection	14
Event-based retention	16
Legal hold.....	17
SnapLock with SnapVault.....	18
Restore.....	19
Miscellaneous.....	19
Conclusion	19
Appendix A: Transition from 7-Mode to ONTAP 9	19
Data migration methods	20
Preparation.....	21
Data copy	21
Verification.....	26
Report.....	28
Where to find additional information	28
Contact us	28

Version history.....	28
-----------------------------	-----------

LIST OF TABLES

Table 1) SnapLock volume default values.....	14
Table 2) SnapLock volume transition combinations.	22

LIST OF FIGURES

Figure 1) Basic disaster recovery in 7-Mode.....	22
Figure 2) Basic disaster recovery in clustered mode.....	23
Figure 3) SnapLock with SnapVault in 7-Mode.	24
Figure 4) SnapLock with SnapVault in clustered mode.	24
Figure 5) SnapVault to disaster recovery cascade in 7-Mode.....	25
Figure 6) SnapVault to disaster recovery cascade in ONTAP 9.....	25

Introduction

Many businesses rely on some use of write once, read many (WORM) data storage to meet regulatory compliance or simply to add another layer of data protection to their critical files (or data). Why have so many companies implemented WORM data storage, given the myriad data storage options available? There are two primary reasons:

- Regulatory agencies recognize the ability of WORM data storage to help safeguard the permanence of archived data and therefore often stipulate that only nonerasable, nonrewritable WORM storage be used to meet their regulations.
- Businesses place a premium on protecting certain business records or critical data files from accidental or intentional alteration or deletion, and WORM functionality, such as nonerasable and nonrewritable data storage, can provide long-term data permanence.

To address issues faced by growing business requirements for WORM data storage and to alleviate issues inherent with traditional WORM storage solutions, NetApp introduced SnapLock software.

SnapLock allows companies to benefit from the data permanence functionality of traditional WORM storage using existing, easy-to-manage NetApp disk storage technologies. NetApp systems can now be configured with SnapLock Compliance and SnapLock Enterprise software for high levels of data integrity, performance, and retention and low TCO.

SnapLock helps customers meet internal and external requirements for retaining, protecting, and accessing regulated and reference data. With SnapLock, customers can create volumes in which files can be committed to the WORM state to prevent files from being altered or deleted until a specified retention date. You can back up this WORM data to disk or to tape for an additional level of data protection. With the NetApp fully integrated data protection portfolio, customers can simultaneously perform disk-to-disk backup and cross-platform replication to protect their most precious resource: their data. As regulatory rules change over time, the flexibility of SnapLock allows companies to implement these policy changes, making it scalable for the future of your business and the industry. Because SnapLock relies on industry-standard network storage protocols, you can achieve your data permanence objectives without sacrificing simplicity or performance.

NetApp unifies archival and compliance initiatives enterprise-wide on a single flexible platform that eliminates the need for separate storage silos. For companies experiencing unmitigated data growth and mounting compliance challenges, some data is archived purely for cost savings without the need for immutability; other data is retained for compliance or legal purposes for which immutability is desired or required. Because SnapLock is a volume-based solution, customers can mix and match both WORM volumes and non-WORM volumes on the same unit to reduce cost and complexity.

SnapLock fundamentals

SnapLock is the NetApp high-performance compliance solution that provides the capability of data retention and WORM protection for retained data. With SnapLock, customers can create nonmodifiable, nonerasable volumes to prevent files from being altered or deleted until a specified retention date.

SnapLock allows this retention to be performed at the file level through standard open file protocols such as CIFS and NFS.

SnapLock is a license-based feature of ONTAP that works with application software to administer nonrewritable storage of data. There are two types of SnapLock: SnapLock Compliance and SnapLock Enterprise. You can activate both types through a single add-on license in ONTAP.

Part of our proven NetApp ONTAP storage software, NetApp SnapLock software delivers high-performance, disk-based data permanence for HDD and SSD deployments.

SnapLock helps provide data integrity and retention, enabling electronic records to be both unalterable and rapidly accessible. Both SnapLock retention features are certified to meet strict records retention requirements as well as addressing an expanded set of retention requirements, including Legal Hold, Event-Based Retention, and Volume Append Mode.

SnapLock creates nonrewritable, nonerasable data on hard disk drives or flash media to prevent files from being altered or deleted until a predetermined or default retention date. SnapLock supports litigation hold and event-based retention, and incrementally appending files while remaining locked (for example audio or video surveillance and logging).

Both types run on NetApp systems with lower cost SATA-based drives, higher performance SAS or fiber-attached disk drives, and SSD/flash drives. This flexibility allows customers to buy the amount and type of storage that fit their business needs for SnapLock WORM storage.

What are SnapLock Compliance and SnapLock Enterprise?

Both SnapLock versions provide nonerasable, nonrewritable WORM data permanence using all types of disk and/or flash drives in a cost-efficient, highly available RAID configuration. From a data protection perspective, the process of committing data to an immutable WORM state on either SnapLock type can be thought of in the same manner as storing data on an optical platter. Similar to an optical platter burned with data, both SnapLock types protect data committed to WORM state from any possible alteration or deletion until their retention periods have been expired.

The SnapLock Compliance software feature is certified to meet strict records-retention requirements, such as SEC Rule 17a-4, FINRA, HIPAA, and CFTC—as well as national requirements for the German-speaking countries (DACH) and GDPR, which require the use of nonerasable storage. SnapLock Compliance provides an "untrusted storage administrator" model of operation in which the records and files committed to WORM storage on a SnapLock Compliance volume can never be altered or modified and can be deleted only after their retention periods expire. Moreover, a SnapLock Compliance volume cannot be deleted until all the records and files stored on it have passed their retention periods. In the case of SnapLock Compliance, any operation by the storage administrator that could compromise WORM data is not permitted, and there is protection not only at the file level but also at volume, aggregate, and disk levels.

SnapLock Enterprise, in contrast, operates under a "trusted storage administrator" model and is designed to help organizations meet self-regulated and best practice guidelines for protecting digital assets with WORM data storage. Data stored on a SnapLock Enterprise volume is protected from alteration or modification. SnapLock Enterprise has one main difference from SnapLock Compliance; because the stored data is not for strict regulatory compliance, SnapLock Enterprise volumes and the data they contain can be destroyed by an administrator with root privileges on the storage system that contains the SnapLock Enterprise volumes before the end of their retention periods.

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.

What is Snapshot copy locking?

Beginning with ONTAP 9.12.1, Snapshot copy locking is a SnapLock capability where Snapshot copies are rendered indelible manually or automatically with a retention period on the volume Snapshot copy policy. Snapshot copy locking is also referred to as tamper-proof Snapshot copy locking. Although it does require the SnapLock license and initialization of the compliance clock, Snapshot copy locking is unrelated to SnapLock Compliance or SnapLock Enterprise. There is no trusted storage administrator as with SnapLock Enterprise and it does not protect the underlying physical storage infrastructure as with SnapLock Compliance. The purpose of Snapshot copy locking is to prevent rogue or untrusted administrators from deleting Snapshot copies on the primary and secondary ONTAP systems. This capability is an improvement over SnapVaulting a Snapshot copy to a secondary system. Rapid recovery

of locked Snapshot copies on primary systems can be achieved in order to restore volumes corrupted by ransomware. For more information about Snapshot copy locking, see the [ONTAP 9.12.1 documentation](#).

Using SnapLock

In ONTAP 9, apart from CLI and ONTAP REST API support, ONTAP System Manager and ONTAP System Manager support have also been introduced for SnapLock.

Licensing

In ONTAP 9, only a single SnapLock license is required to use both SnapLock Compliance and SnapLock Enterprise functionalities on a node. The SnapLock license is a node-locked license. Installing a node-locked license entitles a node to the licensed functionality. For the cluster to use the licensed functionality and to be compliant, all the nodes must be licensed for the functionality.

After the data has been committed to WORM, SnapLock continues to enforce the WORM property of the data regardless of the state of licensing. Moreover, after the SnapLock volumes have been created, you can commit files to WORM state even without a license. However, the state of licensing determines the configuration changes for SnapLock. The following operations require the SnapLock license to be enabled:

- Initializing ComplianceClock
- Creating a SnapLock aggregate
- Creating a SnapLock volume
- Turning on the autocommit scanner
- Creating a user with vsadmin-snaplock role
- Enabling privileged-delete
- Setting up the SnapLock audit log volume configuration

Modifying existing configurations is usually allowed even if the SnapLock license is not enabled. The following operations do not require the SnapLock license to be enabled:

- Deleting a SnapLock aggregate
- Deleting a SnapLock volume
- Turning off the autocommit scanner
- Deleting a user with vsadmin-snaplock role
- Disabling or permanently disabling privileged-delete
- Deleting the SnapLock audit log volume configuration

Initializing SnapLock ComplianceClock

In a data compliance environment, you cannot rely on a system clock because it can be arbitrarily modified by the administrator, thereby compromising the retention period of WORM files and NetApp Snapshot™ copies. Therefore, SnapLock relies on the ComplianceClock service in ONTAP, which is a software-based, tamper-resistant clock. The ComplianceClock can be initialized only once by the administrator on every node, after which it operates based on hardware ticks. Before initializing the ComplianceClock, the administrator must properly take note of the time zone and make sure that the storage system's time is as accurate as possible. After the ComplianceClock is initialized, the administrator cannot perform any action that causes any adjustment to the clock. This feature makes sure that the retention period of WORM files cannot be shortened by doing forward adjustments of the reference clock.

There are two types of ComplianceClock:

- **Volume ComplianceClock (VCC).** The volume ComplianceClock is a tamper-resistant reference time per volume. The VCC is stored only on SnapLock volumes and is used to determine the expiration of WORM files and WORM Snapshot copies in that volume. Because the VCC is per volume, VCC skewing of volume X does not affect the VCC of volume Y. The VCC is updated on SnapLock volumes lazily, only when they are in a consistency point, because of client writes or any other ONTAP operation. There is a forced update of the VCC after a large interval of time (24 hours).
- **System ComplianceClock (SCC).** A single system ComplianceClock is maintained per node. The SCC is used to update the VCC values and to provide the base VCC value for new volumes. The SCC is stored in the node root volume. It is updated in memory and on the node root volume every 15 seconds based on hardware ticks.

Each SnapLock volume maintains the following on-disk metadata for the VCC:

- VCC time: 64-bit VCC time stamp
- SCC time: 64-bit SCC time stamp (SCC time at last update)
- Node ID: unique identifier for the node (used for SCC-VCC association)
- SCC ID: unique identifier for the SCC (used for SCC-VCC association)

The VCC is initialized only once (either at creation or at upgrade) and cannot be reinitialized after that. The VCC obtains its starting value from the SCC at the time of volume creation. It uses the SCC as a reference time base for its updates. The VCC is updated as follows:

```
time elapsed since last update (delta) = current SCC time - SCC time at last VCC update
new VCC time = stored VCC time + delta
```

When a volume is brought online, the time elapsed since the last update is computed. If the node ID and SCC ID match, the on-disk VCC of the volume is updated. Therefore, if the volume is brought back online on the same system, there are no VCC skews, irrespective of the duration for which the volume was kept offline.

The ComplianceClock can be initialized exactly once on a node by running the following command:

```
snaplock compliance-clock initialize -node <nodename>
```

The ComplianceClock can be viewed by running the following command:

```
snaplock compliance-clock show
```

SnapLock volume and aggregate creation

After a SnapLock license has been installed and the ComplianceClock has been initialized on a storage system, the following steps show various aspects of a SnapLock volume, including creation and attempted destruction, operations such as WORM commits, attempted deletion or modification, and overall ease of use. Some thought and planning are required to optimize reliability and performance, and this information is contained in the best practice guidelines section later in this document.

You use the volume `-snaplock-type` option to specify a Compliance or Enterprise SnapLock volume type. For releases earlier than ONTAP 9.10.1, you must create a separate SnapLock aggregate. Beginning with ONTAP 9.10.1, SnapLock and non-SnapLock volumes can exist on the same aggregate; therefore, you are no longer required to create a separate SnapLock aggregate if you are using ONTAP 9.10.1. You create SnapLock volumes in the same way that you create non-SnapLock volumes.

Note: The comingled aggregates always reflect the strictest SnapLock type of all the volumes hosted on it. For example, an aggregate with one SnapLock Compliance volume, one SnapLock Enterprise volume and several non-SnapLock volumes will show as a SnapLock Compliance aggregate even though there are multiple types of volumes on the aggregate.

Beginning with ONTAP 9.11.1, FlexGroup volumes are supported with SnapLock. A FlexGroup volume is comprised of one or more FlexVol volumes, spread across multiple aggregates hosted across different

nodes of the cluster, tied together into a single scalable filesystem. FlexGroup volumes enable scale-out performance and storage. For more information on FlexGroup volumes, see [NetApp ONTAP FlexGroup volumes Best practices and implementation guide](#).

FlexGroup volume support for SnapLock includes core SnapLock capabilities, volume append mode, unspecified retention time, extended retention periods (beyond 2,038), audit log volume support, and autocommit scanner. FlexGroup SnapLock support for legal hold, event-based retention, and LockVault are not yet supported in ONTAP 9.11.1. Beginning with ONTAP 9.12.1, LockVault is supported with FlexGroup volumes.

Note: Reverting to an ONTAP version lower than 9.11.1 will be blocked if there are any FlexGroup SnapLock volumes present on the cluster.

SnapLock Compliance restrictions

There is a command that is modified to prevent it from carrying out its normal actions on SnapLock Compliance volumes:

- **Volume delete.** Allowing a SnapLock Compliance volume deletion before the expiration of the retention period of all the records and files on it violates the principle of WORM storage, especially in the regulated data archived space. The delete command succeeds on a SnapLock Compliance volume only if all the WORM records and files with specified retention periods have expired.

SnapLock volume usage

SnapLock provides retention granularity at the individual file level. There are several methods to commit a file to WORM on a SnapLock volume. Two are discussed below.

Note: Event Based Retention (EBR) is another way to commit files to WORM. See “Event-based retention”.

Manual commit

The first method is to copy or create a file on a SnapLock volume and then change the file attribute to read only using either the NFS or CIFS open protocol. Following the change to read only, the file is committed to an immutable state on the SnapLock volume. Committing a file to WORM state on a SnapLock volume includes two steps:

1. Change the last access date to the retention date on the file. If you want to use the default retention period for the volume, skip this step. The exact operation depends on the file protocol (CIFS, NFS, and so on) and the client operating system. Here is an example of how the operation can be done in a UNIX shell environment:

```
touch -a -t [retention date] [file]
```

2. Transition the file's attributes from a writable state to a read-only state. The exact operation depends on the file protocol (CIFS, NFS, and so on) and the client operating system, but the operation is always straightforward. This transition can easily be done manually with scripts or programmatically. Some examples for different environments are:

UNIX shell environment:

```
chmod -w [file]
```

Windows shell environment:

```
attrib +r [file]
```

When you commit a file to WORM state, volume ComplianceClock time is written to the ctime field of the file. The volume ComplianceClock is used for calculating a file's retention period.

Note: It is very important to note that for the committal to WORM state to occur, the file must experience a transition from a writable state to a read-only state. Files that are created read only do not

experience this transition and hence are not committed to WORM state. Applications should always make sure that the file is initially writable on the SnapLock volume before making it read only to commit to WORM state.

Autocommit

The second method of committing files to an immutable state on a SnapLock volume is accomplished using the autocommit option, which can be set on the SnapLock volume. In this case, the application is not required to set the read-only file attribute. The autocommit feature enables automated committing of a file to WORM state on a SnapLock volume if the file did not get changed for the autocommit-period duration. Each volume has its own autocommit-period volume option, giving flexibility to the administrator to configure different WORM policies on different volumes. By default, autocommit is disabled on the SnapLock volume. Minimum configurable value for autocommit-period is as little as 5 minutes, and maximum is up to 10 years.

Note: Autocommit might not be 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. The autocommit scanner runs a thread per volume to scan all the files in the volume. It can have scaling issues if the number of volumes in that node is high or if WAFL[®] is busy with client I/O.

Committing a file to WORM

A file can be committed to WORM in several ways:

- Create a file in a SnapLock volume:
 - File created by an application as read-write (RW) in a SnapLock volume that has no autocommit period specified. When ready, the application finally sets the file to read only, thereby committing the file to WORM.
 - File created by an application as RW in a SnapLock volume that has an autocommit period specified. At the end of the autocommit period, the file is automatically set to read only by the autocommit scanner. No need for the application to explicitly make the file read only.
- Copy an RW file into a SnapLock volume:
 - Whatever copy application is used to copy into the SnapLock volume, the end result will be an RW file that is not committed to WORM. Either the file will have to be manually committed to WORM by changing it to read only, or the autocommit functionality will have to be used to automatically commit the file.
- Copy a read-only file into a SnapLock volume:
 - Using NFSv3:
 - An NFS (UNIX-style) copy command will create a read-only file in the SnapLock volume and then try to copy data into it. This will fail. The file must be changed to RW before copying.
 - Using CIFS/SMB:
 - If the copy command results in an RW file, then the result is the same as when you copy an RW file into a SnapLock volume; the file must be manually committed to WORM, or autocommit must be enabled. (Using the `copy` command from the command line, a read-only file is written as an RW file.)
 - If the copy command results in a read-only file, then the file is automatically committed to WORM. For example, using Windows drag-and-drop or copy/paste, initially an RW file is created, the data is copied into it, and the file is then made read only (thereby committing it to WORM).

If in any doubt, the recommendation is to test the copy methods prior to implementation.

For more details about the procedure to commit WORM files, see the [ONTAP 9 Documentation Center](#).

SnapLock retention periods

Regardless of how files are committed to an immutable state on a SnapLock volume, it is important to understand the retention period settings. Every record committed to the WORM state on a SnapLock volume can have an individual retention period associated with it. ONTAP data management software enforces retention of these records until the retention period ends. After the retention period is over, the records can be deleted but not modified.

Each SnapLock volume has options that are set to control the minimum, maximum, and default retention periods. These values are `minimum-retention-period`, `maximum-retention-period`, and `default-retention-period`, respectively. `default-retention-period` can be set to any value between and including `minimum-retention-period` and `maximum-retention-period`. If an application does not specify any retention period when committing the file, the `default-retention-period` is used. If the application attempts to set a retention period that is less than the `minimum-retention-period`, then the `minimum-retention-period` is used instead. If the application attempts to set a retention period that is more than the `maximum-retention-period`, then the `maximum-retention-period` is used. These settings are beneficial in situations in which you are evaluating a new application and do not want to have files committed for an extended period of time. The `maximum-retention-period` check does not come into play when extending the retention period of a file.

SnapLock volume append mode

When a user commits a file in a SnapLock volume to WORM, the file cannot be deleted until the file retention time has expired. At no point in time can the file contents be modified before or even after expiration. A file's retention time can only be extended, not shortened. For logging purposes, a user might want to append to this WORM file.

With ONTAP 9, SnapLock allows creation of another type of file called a WORM-appendable file. The WORM append feature allows users to create a WORM file and append data to it. The data that is added to the file is committed automatically to WORM in chunks of 256K. Blocks are locked as they are written to specially defined appendable WORM files. The user can append logs to this file but cannot modify the existing contents of the file or delete the file until expiration. This is especially useful for log files that can be added to, but not altered or deleted. For example, use this approach when audio, video, or logging applications autonomously or automatically create files in an NFS or CIFS share.

Steps to create a WORM-appendable file on a SnapLock volume:

1. Create a zero-byte file.
2. Set the required retention time in the `atime` field of the file (optional).
3. Remove the write permissions on the file to make it WORM.
4. Add the write permissions to make the file writable (in this case, only appendable).
5. After logging is complete, the file can be made into a WORM read-only file by removing the write permissions of the file.

With ONTAP 9.3, a new volume option for SnapLock volumes was introduced to enable or disable SnapLock Volume Append Mode (VAM). When the VAM option is enabled, all newly created files with write permissions are made WORM-appendable files by default. This option can be toggled only on empty volumes (with no user data or Snapshot copies) to prevent disruption of applications already using the volume.

Command to enable VAM on an empty volume:

```
volume snaplock modify -vserver vs1 -volume vol1 -is-volume-append-mode-enabled true
```

On a VAM-enabled volume, the autocommit scanner code will look for WORM-appendable files with no writes in the last autocommit period amount of time in addition to regular non-WORM files. Those files will be converted to WORM-only status (that is, write access will be removed).

In summary: When a file is being written to in append mode, data is protected against overwriting or deletion in 256KB segments, for example, every time a 256KB segment is filled it is committed to WORM. Any data in an incomplete (and therefore uncommitted) 256KB segment remains writable and/or deleteable until a) the segment is filled, or b) the file is manually committed to WORM. In the special case of an append mode file in a VAM volume, the file can optionally be auto-committed to WORM using the volume defaults (between five minutes and 10 year in one-minute increments).

Application integration

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. Applications can integrate with SnapLock functionality by employing one of two basic methodologies:

- **Integration through NFS or CIFS.** This approach allows clients to perform the following operations needed to commit files to WORM:
 - a. Select the files that must be retained for a certain time period.
 - b. Select the retention period (this 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.
 - c. Commit the files to WORM state. 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.
 - d. When the retention period has expired (that is, the ComplianceClock value has surpassed the value of the atime), those files can be deleted.
- **Integration through the ONTAP REST APIs.** By implementing the functionality in the ONTAP REST APIs, applications can perform SnapLock functionality that is described in this document. For the detailed list of SnapLock APIs, see the [ONTAP 9 Documentation Center](#).

The SnapLock autocommit feature can be leveraged by applications to automatically transition files to the WORM state on SnapLock volumes. The autocommit feature is especially useful when applications simply copy files to SnapLock volumes and do not have the ability to programmatically transition them to the WORM state.

Setting file retention dates with SnapLock

In keeping with the SnapLock open protocol design, support for file retention periods was implemented without requiring the use of proprietary APIs or protocols. File retention dates can be set and queried programmatically through standard system call interfaces supplied by most operating systems or interactively through standard command-line tools. They can also be set through ONTAP REST APIs. As with SnapLock operations, operations for setting retention dates occur over standard network file system interfaces such as NFS and CIFS. This flexibility allows applications to utilize SnapLock from compiled code or scripts without requiring any libraries or software to be installed on the client systems.

The retention date for WORM files on a SnapLock volume is stored in the last access time stamp of the file metadata. To set a retention date for a WORM file, the application must explicitly set the file's last access time to the desired retention date before setting the file to read only and engaging the WORM commit operation. After being committed to WORM state, the access time of the file is immutable, with the only exception being the extension of the file retention period.

File deletion before the retention period

Privileged-delete functionality in SnapLock allows a privileged user with the vsadmin-snaplock role to delete an unexpired WORM file on a SnapLock Enterprise volume. However, the user cannot use the

privileged-delete functionality to delete a WORM file that has already expired. The deletion is logged in an audit file on a SnapLock Compliance audit log volume so that there is a nonerasable record of the file's previous existence and early deletion. In the SnapLock audit log file, you can find details related to the privileged deletion of a WORM file, such as whether a file was deleted, when it was deleted, who deleted the file, and the file fingerprint information. Therefore, the privileged-delete feature is also known as auditable delete.

File deletion after the retention period

It is important to note that ONTAP does not automatically delete files, including files that have reached or passed their retention dates. Instead, all deletions of such files must be handled by the application or some other process such as a script or batch job. After the retention date of a WORM file has been reached, ONTAP permits applications to change the file permissions back to writable from read only and then allow the file to be deleted. ONTAP does not allow any alteration or modification on the SnapLock file when it is back in a writable state. The only action allowed at this point is to delete the file or set a new retention date and change the file to read only to reengage SnapLock WORM protection.

Best practices for SnapLock

This section discusses best practices for SnapLock.

ComplianceClock best practices

The ComplianceClock is a software-based clock that is independent from the system clock and is updated based on the hardware ticks. Make sure that all SnapLock volumes and volumes other than SnapLock on a NetApp system with ComplianceClock enabled are taken offline (or restricted) for only brief periods. After the ComplianceClock has been initialized, it cannot be modified under any circumstances. This is to prevent any tampering with the retention date. Because the ComplianceClock is a software clock, it does not run during the outage, but its last state is persistently stored before the shutdown. When the system is brought back up, the ComplianceClock is running behind real time. As a result of this, the customer might need to wait a few extra hours or days, depending on how long the system was powered off, to delete expired files. However, this is a safe implementation from a compliance perspective, assuring that files under SnapLock protection cannot be prematurely deleted.

Additionally, the volume ComplianceClock (VCC) is updated only if the system ComplianceClock (SCC) association of the volume matches the system. This is determined using the SCC ID and node ID. Both the node ID and SCC ID of a SnapLock volume must match the corresponding values for the system to establish an association. It is required to establish the association between VCC and SCC before updating the VCC to make sure that the current SCC and last updated SCC are from the same time base.

The node ID is needed to detect change of node due to volume copy or physical movement of disks. The SCC ID is needed to detect change to SCC association due to SCC reinitialization. A change in SCC association implies that the SCC time stored on the volume does not correspond to the SCC time of the system. Therefore, SCC time of the volume gets discarded, and the VCC delta is assumed to be zero.

Upon change of SCC association, the volume's ComplianceClock metadata is updated to establish association with the new SCC. This might cause a VCC skew. In order to minimize such skews, all operations that can cause changes to the SCC association (SCC/node ID) should update the VCC before proceeding. Following is a list of such scenarios (this list might not be exhaustive):

- Volume restrict
- Volume offline
- Aggregate offline
- Aggregate relocate

- Halt/reboot

SnapLock Compliance testing

IT organizations implementing new, comprehensive archival solutions that include application software and storage on a SnapLock Compliance volume often require testing from the proof-of-concept stage through final acceptance. Even after the acceptance milestone has been reached, additional testing might arise as a natural part of upgrade efforts on any piece of the archival infrastructure. Testing an application that uses SnapLock Compliance volumes can potentially have hazards. The SnapLock Compliance traditional volume or aggregate, by design, cannot be destroyed until the retention period of all the files residing on it has expired. If the retention period is set for a long period by accident, the disks that make up the SnapLock Compliance aggregate are not available for reuse until all immutable files have reached the end of their respective retention periods.

Using physical volumes

For both initial and ongoing testing, storage administrators are advised to create a permanent dedicated test volume consisting of the minimum possible number of drives. When testing archiving on a SnapLock Compliance volume, be sure that retention dates are set for each file or record. Files committed to SnapLock without having a retention date set are by default set to the maximum retention period (30 years) and cannot be removed before then, unless the SnapLock volume's default or maximum retention date options specify otherwise. The SnapLock default retention period for the volume should be set to some value other than the default when the volume was created. After files have reached the end of their retention periods, the containing test SnapLock Compliance volume can be destroyed to reclaim the space.

Using the NetApp appliance simulator

Another method of testing SnapLock Compliance process is to use the ONTAP simulator that is available on the NetApp Support site. The simulator runs in a VMware virtual machine (VM) and has all of the functionality of ONTAP found on NetApp storage systems. The ComplianceClock value can be set in the simulator. Both types of SnapLock volumes can be created for testing and integration activities. After the testing is complete, the simulator can be deleted, and all disk space is reclaimed, even in the case of SnapLock Compliance testing. For more information about the use of the ONTAP simulator, refer to the accompanying documentation.

Creating and growing production SnapLock volumes

For SnapLock storage, it is important to consider how directories are treated. After directories are created on a SnapLock volume, they cannot be renamed regardless of their access permissions. This is important to remember when using Microsoft Windows Explorer to create a new folder. Windows Explorer first creates a directory called New Folder. Attempting to rename this directory to something more useful is not possible on a SnapLock volume. You can create and rename folders in a volume other than SnapLock until they are correctly named and then copy them to the SnapLock volume. Manually creating directories on SnapLock volumes in either the Microsoft or the UNIX environment is better handled using the `mkdir` command in a CLI. Although directories cannot be renamed, they can be deleted as long as no files committed to WORM state are contained within their hierarchy.

SnapLock volume minimum, maximum, and default retention period values

When a SnapLock volume is created, default values are set for the volume minimum, maximum, and default retention periods for files residing on the volume. Table 1 contains the default values.

Table 1) SnapLock volume default values.

Option	SnapLock Enterprise	SnapLock Compliance
minimum-retention-period (min.)	0	0
maximum-retention-period (max.)	30 years	30 years
default-retention-period	Minimum	Maximum
autocommit_period	None	None

These values are conservative values and probably do not reflect your company's standards. NetApp recommends that these values be reviewed and reset to values that correlate to your company's business and legal requirements.

The minimum-retention-period value prevents a retention period for a file residing on the volume to be set to a value less than the minimum period. The minimum-retention-period is used if the requested retention period is less than this value. The maximum-retention-period represents the furthest time in the future that a file can be immutable. If the requested file retention period is beyond the maximum value, the maximum-retention-period is used. If no value is specified in the retention period field, the default-retention-period is used.

Unspecified retention

Customers might not always know what retention period to apply to newly created data, but they need it to be protected immediately. Unspecified retention time (URT) enables users to lock the data after creation and to specify a retention period at a later date.

Files with URT are retained until the file is set with an absolute retention time.

You can set URT on the file by doing the following:

- Use the file retention CLI, NetApp Manageability SDK, and REST APIs
- Use EBR with the retention period unspecified
- Enable autocommit with the default retention period of the SnapLock volume unspecified
- Remove write permission from NFS and CIFS clients when the default retention period of the SnapLock volume is unspecified

The conversion of absolute retention time (ART) to URT to ART is allowed only if the new retention time is ahead of the earlier retention time.

When using URT, the best practice is to set an explicit retention period to establish baseline (that is, minimum) retention control. After the initial retention is applied, you can change the ART to URT. This will make sure that when an ART is later set, it cannot be set to a value lower than what was originally set.

Data protection

ONTAP has numerous capabilities built-in or available as add-on options to promote data protection and high availability. However, attaining the levels of data protection mandated by regulatory agencies requires a more comprehensive enterprise strategy. At a minimum, NetApp recommends the following data protection strategies for consideration in a robust archival solution. NetApp professional services or

a qualified technology partner can work with you to identify the most advantageous data protection strategy for your specific business and technology needs.

Replication to remote site

To comply with data retention rules, regulatory agencies might require that a second copy of archived data be kept at a remote site. The most straightforward and natural way to comply with this requirement is to replicate data from a primary NetApp system to a secondary NetApp system in a separate location.

There are three integrated NetApp solutions available to seamlessly perform data replication:

- The easiest and most robust solution is to use the NetApp SnapMirror® feature in asynchronous mode to replicate data to a remote location. Asynchronous SnapMirror replicates SnapLock data to a remote NetApp SnapLock volume while maintaining all the WORM attributes. SnapMirror is an add-on license product available with ONTAP.
- The second solution, ndmpcopy, is a free utility and is already bundled with ONTAP. Like SnapMirror, ndmpcopy maintains WORM aspects of the original files in the replicated copy.
- The third solution is NetApp MetroCluster. Only SnapLock Enterprise aggregates are supported in MetroCluster in ONTAP 9.0. With ONTAP 9.3, SnapLock Enterprise aggregates with privileged delete are supported on MetroCluster. SnapLock Compliance is also supported on unmirrored aggregates of MetroCluster starting with ONTAP 9.3. SnapLock Compliance on MetroCluster mirrored aggregates is supported only if the aggregate is only used to host SnapLock audit log volumes. SVM-specific SnapLock configurations can be replicated to both sites using MetroCluster.

Note: In all three replication cases, WORM attributes such as retention times of the SnapLock files and volumes are preserved and mirrored from the source to the destination.

ComplianceClock behavior with replication

In the case of a SnapMirror relationship between SnapLock volumes, the types of SnapLock source and destination volumes must match. Volume SnapMirror does a block-level copy from the source to the destination. The source sends its computed in-core volume ComplianceClock (VCC) time to the destination. As a result, the destination VCC time ends up being the same as the source VCC time. The destination VCC time is updated with every SnapMirror update. If the mirroring relationship is broken, the destination volume is mounted read-write, and its VCC software starts operating using the destination system ComplianceClock (SCC) time as the reference. Therefore, no skews are introduced as a result of the break.

Disk-to-disk backup

NetApp offers an efficient disk-based backup solution called NetApp SnapVault® that leverages block-level incrementals for reliable, low-overhead backup and recovery that are suitable for any environment. Storage-efficient (block incremental) daily (or more frequent) Snapshot copies are backed up to secondary storage (using SnapVault technology) and protected against modification or deletion until a specified retention date (using SnapLock technology). Note that vaulting of SnapLock volumes is not supported. A SnapVault transfer fails if the source of the SnapVault relationship is a SnapLock volume. Retention periods for these WORM Snapshot copies can be specified through the volume's default retention period. The retention period for WORM Snapshot copies can be extended, but not reduced. In ONTAP 9, this feature is known as SnapLock with SnapVault.

With non-WORM Snapshot copies, after the maximum count of Snapshot copies to be retained is reached, the oldest retained Snapshot copies are deleted when new Snapshot copies are added. However, older WORM Snapshot copies cannot be deleted until their retention period has expired. In the event that more WORM Snapshot copies need to be retained than the maximum allowed, volume clones must be used to overcome this limit.

Tape backup

Although NetApp offers substantial performance improvement and storage capabilities for near-line data storage over optical or tape-based storage, tape backups are still a valuable part of an overall data protection strategy for many enterprises. If the SnapLock volume is not mirrored to another site, NetApp recommends that regulated data archived on a SnapLock volume also be backed up to another medium, whether tape or disk, using NDMP-initiated DUMP and restore to preserve the WORM characteristics of the source files. This is prudent for making multiple copies of regulated data available for redundant recovery scenarios. 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. However, 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 volume other than SnapLock, the WORM attributes are preserved but are ignored and not enforced by ONTAP data management software.

Physical security

SnapLock is designed to completely preserve data in an immutable state. SnapLock is unable to prevent data loss in the event of physical destruction of the disks. In the same sense that an optical media platter or paper document can be physically destroyed, the disks in a SnapLock aggregate can be removed and destroyed. In any scenario, the storage medium is only as resilient as the physical security of its location. NetApp storage systems with SnapLock volumes should be housed in locked cabinets in a restricted area to minimize the risk of physical tampering.

Security hardening

It is very important to ensure that your data is able to meet your organization's security objectives. ONTAP includes many security capabilities and features along with SnapLock that can make your installation more resilient and productive. [TR-4569: Security Hardening Guide for NetApp ONTAP 9](#) This technical report provides guidance and configuration settings for ONTAP 9 to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability. NetApp encryption and SnapLock

Customers might want to encrypt data to comply with overlapping regulations, such as when regulatory compliance retention requirements conflict with privacy regulations. Customers might also want an additional layer of protection for expired and deleted compliant data. Starting with ONTAP 9, both SnapLock Compliance and SnapLock Enterprise are supported in combination with NetApp Storage Encryption (NSE) drives. Both SnapLock Compliance and SnapLock Enterprise are also supported with NetApp Volume Encryption (NVE).

Note: If you are using NVE, you can enable encryption only on new, empty SnapLock volumes. You cannot enable encryption on an existing SnapLock volume.

Encryption also offers the ability to cryptographically shred data by deleting the encryption key, thereby rendering the encrypted data unreadable.

Note: Electronically shredding (either intentionally or accidentally) compliant data before its expiration might open a customer to litigation. It is the customer's responsibility to make sure that the encryption keys are protected and can be recovered in the case of a disaster. Failure to do so could result in SnapLock data being permanently destroyed, which might in turn be a compliance violation.

Event-based retention

Event-based retention is defined as an instruction specifying that a file shall be disposed of a fixed period of time after a specified event. In other words, the retention policy starts at the time of such an event, some undetermined period of time after the file was committed to the WORM state.

Event-based retention (EBR) was introduced in ONTAP 9.3 and can be applied to any SnapLock volume (SnapLock Compliance or SnapLock Enterprise). EBR can be achieved on SnapLock Enterprise volumes by using a combination of infinite retention on files and privileged delete (a feature of SnapLock Enterprise that allows a special user to delete WORM files even before their retention date has passed), alongside application-level bookkeeping. The idea is to configure a SnapLock Enterprise volume with a default-retention-period value of infinite, set up an application to copy files to the volume and track every file copied, and then transition the files to the WORM state by giving them an indefinite retention period (use autocommit to skip this step).

When an event occurs (usually one event per file), the retention policy for a file kicks in. Because the application keeps track of individual files, after the event occurs, the application simply makes a note to delete the file at a certain point in the future as dictated by the retention policy. After that time is reached, the application uses privileged delete to perform this deletion.

The user needs to create a retention policy and then apply it to the SnapLock volume:

```
vserver::> snaplock event-retention policy create -name employee_exit -retention-period "10 years"
```

The user can create, modify, delete, and list SVM-wide event retention policies. The policy can be applied to a single file or the entire directory. When an EBR policy is applied to a file that is not in a WORM state, it will get committed to WORM. When an EBR policy is applied to a WORM file, the retention time of that file can only be extended (not reduced) as a result of application of the EBR policy.

One simple, real-world example is in the context of a healthcare organization that is using an enterprise content management (ECM) solution such as IBM FileNet to manage patient records on a SnapLock Enterprise volume. This company is required by HIPAA to retain immutable copies of patient records for seven years after the death of a patient. Although FileNet has been integrated with SnapLock, this type of business process requires extended integration through ONTAP REST APIs. Because the date of death is not known in advance, all patient records are initially set with infinite retention. Seven years after the event has occurred, FileNet executes a delete preprocess that contains the code to remove the patient data from the SnapLock volume. Within FileNet this is registered as an auditable event. Similarly, within the NetApp storage system, the privileged delete operation is logged to a SnapLock compliance volume that cannot be modified or removed by even the storage administrator. The integration of an ECM application such as FileNet with SnapLock and privileged delete helps companies to satisfy regulatory record retention requirements on a flexible and cost-efficient storage platform.

For more information about the privileged delete operation and the ONTAP REST APIs that exposes this functionality, refer to the [ONTAP 9 Documentation Center](#) on the NetApp Support site.

Legal hold

ONTAP 9.3 introduced a feature of legal hold. Legal hold is a feature by which files, folders, a volume, or list of volumes can be held in a tamper-proof state for an indefinite time period for litigation purposes. This hold prevents deletion of the specified objects until the legal hold is removed. This legal hold can be released at any time. If a previous hold of any sort or the original retention period has not expired when the legal hold is removed, the original retention period or previous hold remains in effect.

A legal hold is allowed only on SnapLock Compliance volumes. Up to 255 legal holds per file and 65,535 litigations per volume can be applied. There is no restriction on the number of files per litigation. It solely depends on space available in the volume. A volume under legal hold has indefinite retention. After all the legal holds are removed from the volume, it returns back to its previous retention. All litigation-related metadata is stored in the public inode space of the volume. Users are not allowed to modify any of this data. Legal-hold begin and end operations are audit logged under the path `/snaplock_log/legal_hold_logs/`.

The litigation names can be from a non-ASCII character set (UTF8) if the volume language settings allow a UTF character set. The litigation names must not begin with "." or "%" and must not contain "/" or spaces. Also, the maximum length of a litigation name is restricted to 80 characters.

SnapMirror with block-level replication replicates legal hold metadata. Backup and restore of legal hold information are only supported for full-volume backups and restores. Subvolume granular (for example, a qtree, dir, or file) backup or restore does not retain legal hold information.

The command interface to apply or remove a legal hold would be like this:

```
vserver::> snaplock legal-hold begin -litigation-name litigation1 -volume voll -path / vserver::>
snaplock legal-hold end -litigation-name litigation1 -volume voll -path /
```

Different CLI options are also available to view the status:

```
show' command displays the holds on a particular volume. vserver::> snaplock legal-hold show -
volume voll
Operation

hold    16842755      vs1    voll Completed
hold    16842757      vs1    voll Completed
'dump-litigations' command displays the litigation within a given SVM.
vserver::> snaplock legal-hold dump-litigations -output-volume out -output-directory-path /dl
```

SnapLock with SnapVault

ONTAP allows backing up a flexible volume other than a SnapLock volume to a SnapLock Enterprise or a SnapLock Compliance volume. That is, you can create a SnapVault relationship between a FlexVol volume as source and SnapLock volume as destination. Snapshot copies are backed up to secondary storage (by using SnapVault technology) and protected against modification or deletion until a specified retention date (by using SnapLock technology).

The SnapMirror policy associated with the relationship defines the number of Snapshot copies of a particular snapmirror-label that are retained on the destination SnapLock volume. The default retention period of the SnapLock volume defines the retention period for the Snapshot copies backed (transferred) up to this volume (destination). This results in setting the snaplock-expiry-time for the Snapshot copies. You can also extend a Snapshot copy's snaplock-expiry-time beyond the default set expiry time. On every scheduled transfer (or manual update) operation, an attempt to delete old Snapshot copies corresponding to the snapmirror-label is made to maintain the retention count. However, if these Snapshot copies have an expiry-time that is in the future, then the Snapshot copies are not deleted. The SnapLock destination volume continues to accumulate (retain/backup) more Snapshot copies even if it means exceeding the number specified in the SnapMirror policy. For example, if a customer wants to retain 15 daily Snapshot copies on an SnapLock destination volume whose default-retention-period is one month (30 days). The SnapVault transfer schedule has been set as daily. Because the expiration time will be set for 30 days in the future, on transfer of the 16th Snapshot copy, the oldest Snapshot copy is not deleted. Only on the 31st day, transfer of the 31st Snapshot copy results in deletion of the oldest Snapshot copy (because its retention period would have expired).

CLI commands to set retention-period or retention-count:

```
snapmirror policy add-rule -vserver vserver -policy test_lv -snapmirror-label sle
-keep 15

volume snaplock modify -volume test_dst -default-retention-period "30days"
```

Note: If the destination is a SnapLock Compliance volume, be aware that the default retention period is 30 years. It is a best practice to set the default retention time on the destination volume PRIOR to initiating SnapVault Snapshot copy transfers to the SnapLock Compliance volume.

Restore

To restore data in a SnapLock volume, the procedure is the same as restoring any other NetApp Snapshot copy except in the case of SnapLock Compliance. For SnapLock Enterprise data, NetApp FlexClone®, the `snapmirror restore` operation, and NetApp SnapRestore® all for SnapLock Enterprise data in the same that they work for NetApp Snapshot data. SnapRestore operations are extremely valuable for file and data recovery or for reverting to a previous known good state. But in the case of SnapLock Compliance volumes, allowing a SnapRestore or FlexClone recovery to a previous state might result in a loss of all the data written since the Snapshot copy was created. Allowing this violates the WORM integrity of the data and regulatory compliance.

To restore data from a SnapLock Compliance volume, the options include `ndmp copy`, `ndmp restore`, and `snapmirror restore` operations. For Snapshot copies that contain LUN data where the LUN is locked in a Snapshot copy on a SnapLock volume, the LUN must be restored to a non-SnapLock volume through the `snapmirror restore` operation.

You can create a FlexClone volume of SnapLock Compliance data but there is a caveat: The clone inherits the properties of the original. The created FlexClone is read only, and it inherits the retention of the original Snapshot copy. You cannot delete the FlexClone volume until the retention time of the original Snapshot copy is met. If the retention time is short, this might not be a problem, but if the retention time is long, it can create issues with FlexClone volumes that cannot be deleted.

The best method for restoring SnapLock Compliance data remains `snapmirror restore`.

Miscellaneous

The behavior of hard links on a SnapLock volume is no different from that on a flexible volume. A hard link to a file can be created in the same directory or across directories. The destination file can be RW, WORM, or WORM_APPEND (VAM or non-VAM). A hard link to a WORM file makes the hard link a WORM too. If so, the hard link cannot be removed until the underlying inode expires.

Conclusion

SnapLock Compliance and SnapLock Enterprise are designed to be critical pieces of a comprehensive data archival solution for businesses that require higher performance and lower TCO alternatives for WORM storage functionality. SnapLock benefits over traditional WORM storage include performance improvements, and advanced data protection while significantly reducing storage management costs. These SnapLock benefits address the needs of corporations with immutability storage requirements for compliance and regulatory purposes. SnapLock helps to enforce strict data protection and immutability standards across your organization with less complexity than competing product offerings.

The powerful data permanence and data integrity features of SnapLock combine with the low TCO by leveraging the existing ONTAP data management software and storage product line; ongoing low operational costs through the use of storage efficient replication technologies; and the use of open, industry-standard protocols for simplified data access and application integration. Together, these provide an unrivaled solution in the WORM data storage space.

For more information about solutions-based products from NetApp, see www.netapp.com/products.

Appendix A: Transition from 7-Mode to ONTAP 9

ONTAP 9.0 is the first release to introduce SnapLock in the clustered environment. If you are an existing NetApp customer with 7-Mode storage systems, taking advantage of the clustered environment capabilities of ONTAP 9 means that your existing 7-Mode environment needs to transition. Transitioning to the clustered environment involves identifying your current environment, defining the transition scope,

designing the optimal configuration of the destination systems, planning how to migrate the data and configurations, and making necessary environmental updates.

This section addresses the key knowledge that is necessary for migrating SnapLock data from 7-Mode systems to ONTAP 9 systems. Note that this section covers the recommendations only specific to SnapLock. For a thorough understanding of the fundamentals of transition from 7-Mode to a clustered environment, see [TR-4052: Successfully Transitioning to Clustered Data ONTAP](#). The technical report provides guidance for scoping, designing, and transitioning your 7-Mode storage environment to a clustered environment. In addition, it discusses the key ONTAP considerations for migrating, such as NetApp SnapMirror®, qtrees, and NetApp FlexClone® volumes. Additionally, a number of ONTAP classes are available through NetApp University that cover the fundamentals of transition:

- [NetApp Transition Fundamentals](#) (web-based)
- [Planning and Implementing Transition Using the 7-Mode Transition Tool](#) (web-based)
- [Transitioning to Clustered Data ONTAP](#) (web-based)

After you've gained a fundamental understanding, you are ready to start planning your transition.

Note: Transition of 7-Mode SnapLock volumes is not supported if the SnapLock volumes contain LUNs.

Best practice

Following the procedure described here does not automatically mean legal compliance. Rather, it is intended as a starting point for planning so that the important steps are included in the overall migration plan. Consultation with your legal department is highly recommended to make sure that your unique specific compliance requirements are met.

Data migration methods

Prior to starting data migration to the clustered environment, identify the recommended migration method based on the application type, the application environment, and other factors. There are several migration tools available today, each with its own benefits and considerations. These tools can be classified into the following two categories.

Replication-based migration

This migration method uses NetApp SnapMirror technology and is available with both the 7-Mode Transition Tool (7MTT) and Transition Data Protection (TDP) SnapMirror. The TDP SnapMirror relationship is reported by ONTAP and refers to the type of SnapMirror relationship where the source is 7-Mode and the destination ONTAP. The key benefits of replication-based migrations are that Snapshot copies and storage efficient replication savings are retained through the migration activity.

Migration of the SnapLock volumes can be accomplished either by using a manual TDP SnapMirror or by using the 7MTT v3.3.3. However, NetApp recommends that you use the 7MTT to transition 7-Mode volumes because the tool provides prechecks to verify both 7-Mode and the cluster environment in every step of the migration process, which helps you to avoid many potential issues. The tool significantly simplifies the migration of all protocols, network, and services configurations along with the data migration.

Note: 7MTT supports only copy-based transition (CBT) and does not have support for copy-free transition (CFT) for SnapLock volumes.

Copy-based migration

Host-based and application-based migration methods use tools that are not directly provided or supported by NetApp (because they are not NetApp products).

Host-based tools that are commonly used for data migration are as follows:

- Logical volume managers (LVMs) from various vendors
- ScriptLogic Secure Copy
- Rsync
- Robocopy/Richcopy
- PEER Software PeerSync
- Data Dynamics StorageX

Some of the aforementioned tools are general data migration tools, and others are offerings from NetApp partners that built specific capabilities into their products to address transition from 7-Mode to a clustered environment. Note that both application-based and host-based migration methods are copy based and not replication based. As a result, Snapshot copies and storage efficient replication savings are not retained through the data migration activity.

Note: NetApp does not directly support third-party tools. If customers use third-party tools for data migration and encounter issues with the tools themselves unrelated to ONTAP or other NetApp products, they need to contact the vendor's customer support department.

You should be aware of the versions of ONTAP operating in 7-Mode that are supported for transitioning to ONTAP 9. If the source 7-Mode system has only 64-bit aggregates and volumes, you can transition them to ONTAP 9. However, if the source 7-Mode system has 32-bit aggregates or volumes with 32-bit Snapshot copies, you must first upgrade to ONTAP 8.1.4 P4 or 8.2.1. After upgrading, you must expand the 32-bit aggregates to 64-bit and then find and remove any 32-bit data including Snapshots.

Copy-based migration methods can be used to migrate data regardless of the source and destination aggregate types. However, replication-based migration methods cannot migrate a 32-bit aggregate from a source 7-Mode storage system to a 64-bit aggregate in ONTAP 9. If you are unsure whether you have 32-bit Snapshot copies present in your 64-bit aggregate, contact NetApp Support for assistance.

Preparation

Before you transition a SnapLock volume from 7-Mode to ONTAP 9, you must prepare the 7-Mode storage system and cluster and create a transition peer relationship between the 7-Mode system and the storage virtual machine (SVM).

You must also make sure that SnapMirror is licensed on the 7-Mode storage system and SnapLock is licensed on the destination cluster. If you are transitioning a 7-Mode VSM relationship between SnapLock volumes, SnapMirror licenses are also required on the destination clusters along with a SnapLock license.

Data copy

Following are the recommended migration approaches for the most common migration scenarios involving SnapLock systems.

Scenario 1: Standalone volume

Transitioning a standalone volume is easily accomplished using the 7MTT (recommended) or by using manual TDP SnapMirror. This process involves creating a SnapMirror relationship between the 7-Mode source and ONTAP 9 destination, performing a baseline transfer, performing incremental updates, monitoring the data copy operation, breaking the SnapMirror relationship, and moving client access from the 7-Mode volume to the ONTAP 9 volume.

You can transition 7-Mode SnapLock Compliance volumes only to SnapLock Compliance volumes and SnapLock Enterprise volumes only to SnapLock Enterprise volumes in ONTAP 9. Table 2 shows the combinations that are supported for SnapLock volume transition.

Table 2) SnapLock volume transition combinations.

SnapMirror destination			
	SnapLock Compliance	SnapLock Enterprise	Regular FlexVol volume
SnapLock Compliance	✓	✗	✗
SnapLock Enterprise	✗	✓	✗
Regular FlexVol volume	✗	✗	✓

Note: Audit log volumes in 7-Mode are node specific, whereas in ONTAP 9 they are SVM specific.

During transition, users must decide where to place SnapLock audit log volumes in the ONTAP 9 destination. This behavior should be acceptable because log volumes are generated only by operations on SnapLock Enterprise volumes for which the administrator is trusted.

Scenario 2: Basic disaster recovery by using SnapMirror

The basic disaster recovery scenario addresses the most common case of a single source and destination volume that are in a volume SnapMirror relationship. Migration of the volumes, associated Snapshot copies, and SnapMirror relationship is easily accomplished by using the 7MTT (recommended) or by using manual TDP SnapMirror.

Parallel transition for SnapLock Compliance volumes


In the case of SnapLock Compliance volumes, you can transition the primary and secondary volumes of a 7-Mode SnapMirror relationship in parallel and in the same cutover window. You must then manually set up the volume SnapMirror relationship in ONTAP after transition.

A 7-Mode SnapMirror relationship between SnapLock Compliance volumes must be transitioned in parallel. SnapMirror resynchronization of a TDP relationship with SnapLock Compliance volumes is not supported because it might result in data loss. Therefore, you cannot establish a SnapMirror disaster recovery relationship between 7-Mode primary volumes and ONTAP secondary volumes with SnapLock Compliance volumes.

Staggered transition for SnapLock Enterprise volumes

When transitioning a 7-Mode volume SnapMirror relationship, you can use staggered transition (transition secondary first and then primary) only for SnapLock Enterprise volumes. A SnapMirror disaster recovery relationship between 7-Mode primary volumes and ONTAP secondary volumes is supported only for SnapLock Enterprise volumes, not for SnapLock Compliance volumes.

Figure 1) Basic disaster recovery in 7-Mode.




Transition steps

Disaster recovery relationships involving SnapLock volumes can be transitioned from 7-Mode to ONTAP 9 using CBT by following these steps:

1. Transition the destination volume separately.
2. Make the transitioned volume the destination volume for the disaster recovery relationship by creating a relationship between the 7-Mode source volume and the ONTAP 9 transitioned destination volume. Skip this step for SnapLock Compliance because SnapMirror resync cannot be done for SnapLock Compliance as it might result in data loss.
3. Transition the source volume.
4. For SnapLock Enterprise, break the disaster recovery relationship between the 7-Mode source and the ONTAP 9 destination. For SnapLock Compliance, a SnapMirror break is performed followed by reestablishing SnapMirror between source and destination volumes.
5. Do a SnapMirror resync between transition source volume and destination volume.

In summary, only parallel transitions are supported for transitioning SnapLock Compliance disaster recovery relationships, but both staggered and parallel transitions are supported for SnapLock Enterprise disaster recovery relationships.

Figure 2) Basic disaster recovery in clustered mode.




Scenario 3: SnapLock with SnapVault (LockVault)

SnapLock with SnapVault is a disk-based regulatory compliance solution for unstructured data in 7-Mode. In ONTAP 9, this feature is known as SnapLock with SnapVault®. It delivers a capacity-efficient regulatory solution by making backups compliant (one copy of data serves two purposes) and by saving only block-level incremental changes. Storage-efficient (block incremental) Snapshot copies are backed up to secondary storage (using SnapVault technology) and protected against modification or deletion until a specified retention date (using SnapLock technology).

There is one notable difference in this feature compared to 7-Mode. In 7-Mode, there is support for a compliance journal (file log), which tracks changes between Snapshot copies and stores them on a WORM volume, so the log cannot be modified either. In ONTAP 9, creation of a compliance journal that tracks the file changes per transfer is not supported.


Figure 3) SnapLock with SnapVault in 7-Mode.



A LockVault relationship from 7-Mode cannot be transitioned to ONTAP 9. The vault destination can be transitioned as a standalone volume. As part of volume transition, if there are 32-bit WORM Snapshot copies in a LockVault destination, then it cannot be transitioned. SnapVault in 7-Mode is qtree based, whereas in clustered environments, SnapVault is volume based. As a result, it is necessary to create a new SnapVault relationship in ONTAP and determine the best course of action for the 7- Mode SnapVault repository. The primary volume can be migrated normally using the 7MTT or a manual TDP SnapMirror relationship. Movement of the secondary volume depends on the retention period for the repository assuming daily backups:

- If the retention period is greater than three months, the repository should be migrated to an ONTAP volume for archiving (not as a secondary volume for a new SnapVault relationship in ONTAP).
- If the retention period is less than or equal to three months, maintain the 7-Mode repository for the retention period, then retire the 7-Mode repository.

Figure 4) SnapLock with SnapVault in clustered mode.




Movement of the primary volume to ONTAP and the establishment of the new SnapVault relationship carry no dependency on the 7-Mode secondary volume (because the ONTAP SnapVault relationship is new). The new SnapVault relationship in ONTAP requires a baseline transfer to be completed.

Note: The `snap restore` command is not allowed on the SnapVault destination.

Scenario 4: SnapVault to disaster recovery cascade

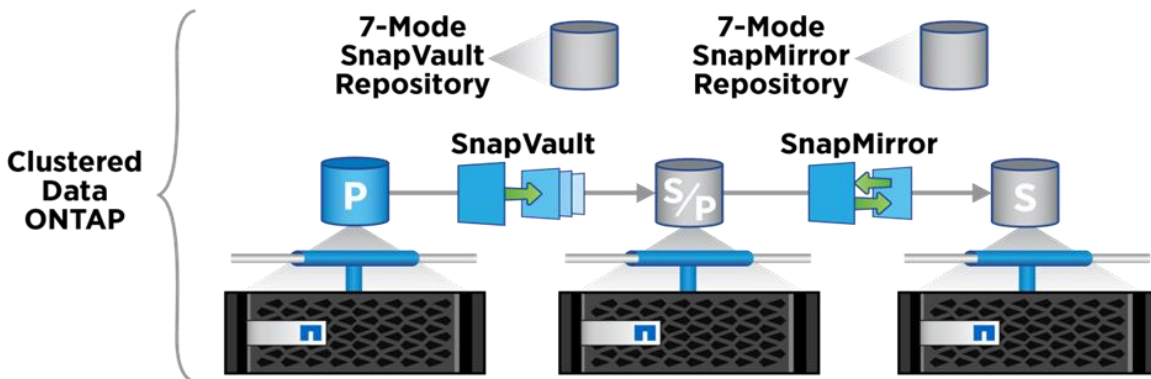
The SnapVault to disaster recovery cascade scenario addresses the case in which there are a SnapVault primary and secondary, and the SnapVault secondary is the primary volume for the SnapMirror relationship (which uses a separate secondary volume).

Figure 5) SnapVault to disaster recovery cascade in 7-Mode.



This approach deals directly with the fact that the SnapVault secondary as well as the SnapMirror primary and secondary volumes contain data in Snapshot copies that is not able to be directly restored in a clustered environment. The first step is to use the 7MTT or manual TDP SnapMirror to migrate the primary SnapVault volume to a clustered environment. Prior to cutover of the primary SnapVault volume, you must break the SnapVault relationship. After the primary SnapVault volume is on a clustered environment, create a new SnapVault relationship (using a new destination volume). After the SnapVault volume is established, a new SnapMirror relationship can be created between the SnapVault secondary volume and a new SnapMirror destination volume. Both the SnapVault and SnapMirror secondary volumes in 7-Mode are subject to the retention approach outlined in “Scenario 3: SnapLock with SnapVault (LockVault)” (being retained on 7-Mode or moved to a separate clustered environment volume based on the retention period). If Snapshot copy retention periods are short (weeks or at most a few months), then it is likely easier to allow the SnapMirror secondary volume to stay on 7-Mode until the Snapshot copies expire. Both the new SnapVault relationship and the new SnapMirror relationship in a clustered environment require a baseline transfer to complete in order to establish new relationships.

Figure 6) SnapVault to disaster recovery cascade in ONTAP 9.



Recovering from a disaster at the 7-Mode site during transition

SnapMirror disaster recovery (DR) relationship between 7-Mode primary volumes and ONTAP secondary volumes is supported only for SnapLock Enterprise volumes

If you have established a SnapMirror disaster recovery relationship between the 7-Mode primary volume and the clustered environment secondary volume and if a disaster occurs at the 7-Mode primary site, you can direct client access to the clustered environment secondary volume. After the 7-Mode primary volume comes back online after the disaster, you must transition the 7-Mode primary volume. Because all SnapMirror relationships to the 7-Mode primary volume are broken and deleted at this stage, you can transition a standalone volume for this type of transition. After transition to the ONTAP primary volume is

complete, you can resynchronize the ONTAP primary volume for the data written on the ONTAP secondary volume. You can then redirect the clients to the clustered environment primary volume.

Warning

SnapMirror resynchronization from clustered environment volumes to the 7-Mode volumes is not supported. Therefore, if you reestablish the disaster recovery relationship between the 7-Mode primary volume and the clustered environment secondary volume after the disaster, any data written on the secondary clustered environment volumes is lost.

Verification

Data copies created in the previous step are identical, and all the SnapLock metadata during the operation is preserved. However, this verification step is required to generate a persistent record of all the files and their contents in both the source and the destination copies of the SnapLock data. This record is useful to verify at a later time that the contents and other properties such as the retention period of the WORM data did not change in process of the copy. The verification results are especially useful if the source is discarded or destroyed (so it is not around to verify the contents at a later date) after the copy.

7MTT v3.3.3 offers a SnapLock Chain of Custody feature which performs post-transition data verification for files in transitioned SnapLock volumes. You can trigger the Chain of Custody operation for the SnapLock volumes in the 7MTT project after the transition is complete. You can perform this operation for all SnapLock volumes in the project or for a subset of SnapLock volumes in the project. The Chain of Custody verification is supported for both compliance and enterprise SnapLock volumes. The Chain of Custody verification is supported only for read-write SnapLock volumes and is not supported for read-only SnapLock volumes.

For details on 7MTT Chain of Custody, see the [ONTAP 9 Documentation Center](#) and specific chain of custody documentation [here](#).

If you have not used 7MTT's Chain-of-Custody feature, you can perform verification tasks manually. After the WORM data has been copied over to the destination, a check can be run to test the following conditions:

- The relevant metadata and contents of the files that were copied over are the same as the source.
- The retention periods in effect on the source and the destination are the same.
- Options related to SnapLock (both volume level and systemwide) are the same on both sides.

To avoid dealing with constantly changing data, it is advisable to do the comparison based on some recent Snapshot copy of the data.

Test 1: Relevant metadata and contents of files that were copied over are the same as the source

This can be done by generating and comparing “fingerprints” of the files on the source and the destination or by doing a byte-by-byte comparison of the contents and the metadata. The fingerprint operation allows you to generate a fingerprint on a per-file basis by using either one of the hashing algorithms: MD5 or SHA-256. NetApp recommends the use of SHA-256. This enables the user to verify integrity of the file at any given time. The file fingerprint is exported externally outside ONTAP using a CLI and NetApp Manageability SDK for user and partner applications. SnapLock does not store any file fingerprints on disk anywhere in the system. The file fingerprint calculates a hash digest on the fly for a file as requested by the user over a CLI or NetApp Manageability SDK. You can issue the file fingerprint using the following command:

```
volume file fingerprint start -file <file_path>
```

After the file fingerprint is issued using the preceding command, a session-id is generated. You can use this session-id to get the status using the following command:

```
volume file fingerprint show -session-id <session-id>
```

After the status is Completed, use the same session-id from the preceding command and issue the following command to get the fingerprint output:

```
volume file fingerprint dump -session-id <session_id>
```

The file type is `worm` in the case of a SnapLock file, `worm_appendable` in the case of a WORM-appendable file, `worm_active_log` in the case of an active WORM log file, `worm_log` in the case of a closed WORM log file, and `regular` in the case of regular files or files other than SnapLock.

When comparing the file metadata, NetApp recommends using the following file attributes:

- File type
- File size
- User ID of the file owner
- Group ID of the file owner
- Security ID (SID) for the owner (visible only from CIFS)
- Time of last modification (mtime)
- Time of last access (atime): with SnapLock this represents the file retention period for the file
- File creation time (this is only visible from a CIFS client, not visible to an NFS client)
- Time of last status change (ctime: only visible from NFS clients, not visible from CIFS clients)
- File permissions and other security attributes

Note: If instead of doing the verification immediately after the copy, the verification is done after the new system has been in use for a while, then the file metadata might not match completely. For example, the retention times on the WORM files might have been extended (they cannot be shortened). Moreover, new WORM or non-WORM content might have been created. Expired WORM files might also have been deleted. In these cases, depending on the situation, the verification could be relaxed to take this into account.

Test 2: Retention periods in effect on the source and destination are similar

It should be ascertained that the retention periods in effect on either end are the same. Matching the time of last access (done earlier) makes sure that the retention time stamp is the same on either end.

However, for the absolute time stamps to make sense, the value of ComplianceClock also needs to be compared on either end. ComplianceClock on the new system should either be in sync or be behind the source. If ComplianceClock on the new system is behind, it results in the retention period being that much longer.

Test 3: SnapLock options are the same on source and destination

Finally, the volume and system options related to SnapLock should be compared to make sure that they are the same on both source and destination. Following are the relevant SnapLock options:

- Volume name
- SnapLock type: enterprise or compliance
- Minimum retention period
- Default retention period
- Maximum retention period
- Autocommit period

- Expiration time
- ComplianceClock time
- Privileged delete option: enabled, disabled, or permanently disabled

These details can be obtained using the following command:

```
volume snaplock show -vserver <vserver name> -volume <volume name>
```

Report

This is a summary of actions taken with details that are pertinent from an audit perspective. The report should include information from the verification phase. The report should be stored as a WORM record with retention period equal to the maximum retention period of any record in the volume.

The following is a checklist of items that might be relevant for a report for the migration:

- Reason for carrying out data migration.
- Details of individuals carrying out the migration.
- System information for the source and destination nodes. It can be obtained using the following command: `system node show -node <nodename>`.
- Volume information for the migrated SnapLock volumes on source and destination.
- ComplianceClock value on the source and destination at both the start of the operation and the end of it.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites.

- TR-4052: Successfully Transitioning to Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4052.pdf>
- ONTAP 9 documentation
<https://docs.netapp.com/us-en/ontap/index.html>

Contact us

Let us know how we can improve this technical report. Contact us at doccomments@netapp.com.

Include TECHNICAL REPORT 4526 in the subject line.

Version history

Version	Date	Document version history
Version 1.0	July 2016	Siddharth Agrawal: This is the first publicly available version of this technical report.
Version 2.0	March 2018	Arpan Merchant: Updates for ONTAP 9.3.
Version 3.0	May 2021	Jeannine Walter: Updates for ONTAP 9.9.1
Version 3.1	September 2021	Additional updates added: Unspecified retention and Restore.
Version 3.2	December 2021	Jeannine Walter: Updates for ONTAP 9.10.1.

Version	Date	Document version history
Version 3.3	February 2022	Dan Tulledge: Additional updates for ONTAP 9.10.1
Version 3.4	May 2022	Dan Tulledge: Updates for ONTAP 9.11.1
Version 3.5	January 2023	Dan Tulledge: Updates for ONTAP 9.12.1

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.

Copyright Information

Copyright © 2023 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

TR-4526-0123