



Technical FAQ

SnapLock

ONTAP 9.0, 9.1 and 9.2

Siddharth Agrawal, NetApp
Version 1.2
May 2017

Abstract

This FAQ is intended for internal NetApp field personnel and channel partners and answers frequently asked questions regarding NetApp® SnapLock® software, the NetApp WORM solution in NetApp ONTAP® 9.2. Do not distribute this FAQ directly to customers.

Data Classification

NetApp Confidential – Limited Use

Version History

Version	Date	Document Version History
Version 1.0	June 2016	Initial release
Version 1.1	Dec 2016	Updated the document for ONTAP 9.1
Version 1.2	May 2017	Updated the document for ONTAP 9.2

TABLE OF CONTENTS

Version History	1
1 Overview.....	7
1.1 SnapLock Release Summary.....	7
What was released for SnapLock in ONTAP 9.2?	7
What was released for SnapLock in ONTAP 9.1?	7
What was released for SnapLock in ONTAP 9.0?	7
Does the ONTAP 9 release affect 7-Mode SnapLock?	7
2 Requirements.....	8
Does SnapLock require any special hardware?	8
2.1 Licensing.....	8
Is a SnapLock license part of the Premium Bundle?	8
What is the SnapLock licensing scheme on ONTAP 9? Do I require different licenses for SLE and SLC?	8
What is the licensing scheme for existing SnapLock customers on 7-Mode?	8
3 Architecture.....	9
What is the difference between SnapLock Compliance (SLC) and SnapLock Enterprise (SLE)?	9
Which file access protocols are supported?	9
Can I copy a regular volume to a SnapLock aggregate or vice versa?.....	9
Can I use a WORM volume as a regular volume?.....	9
Does SnapLock automatically delete expired WORM data?	10
Can I use Snapshot copies on a SnapLock volume?	10
Can I convert an existing non SnapLock volume to a SnapLock volume?	10
Can I change any attributes on a WORM committed file before or after it expires?	10
Are there ways around SnapLock protection?	11
Has SnapLock been validated by a third party?	11
3.1 ComplianceClock.....	11
What prevents me from resetting the system clock to delete files early?	11
How does ComplianceClock work? Are there any changes in ComplianceClock from 7-Mode to ONTAP 9?	12

How is the VCC updated on SnapLock volumes? What happens to VCC when I offline a volume?	12
How does ComplianceClock behave when there is a SnapMirror relationship?	12
What happens to VCC in case of Storage Failover (SFO) or AggregateRelocate (ARL) of a SnapLock aggregate?	13
What happens to the ComplianceClock during a system shutdown or outage? I left my system powered off for a while, and now the ComplianceClock is behind the system clock. How do I fix this?	13
How do I find out more about ComplianceClock?	13
3.2 Transition from 7-Mode to ONTAP 9.....	13
Which tools are available to SnapLock customers to support transition from 7-Mode to ONTAP 9?.....	13
Is there any support for Copy-Free Transition for SnapLock aggregates?	13
Are there any requirements specific to SnapLock aggregates that must be met before carrying out transition from 7-Mode to ONTAP 9?	13
How can I transition the DR relationship for SnapLock volumes from 7-Mode to ONTAP 9?	14
Are there any best practices for transitioning 7-Mode audit log volumes to ONTAP 9?	14
4 Configuration	15
What happens when I remove the SnapLock license or mirror to a SnapLock volume on a system that does not have a SnapLock license?	15
What does an application need to do to integrate with SnapLock?	15
What granularity of data protection does SnapLock allow?	16
How do I manually commit a file to WORM state on a SnapLock volume?	16
How does autocommit work on ONTAP 9? What is the range for autocommit-period? Is it configurable per volume?	16
Since a Snapshot copy does not have a last access time field, how does the retention period get assigned to Snapshot copies?	17
How do I identify SnapLock volumes in a system?	17
How can I find the expiration time of the SnapLock volume?	17
Why can't I rename directories on a SnapLock volume?	17
Is it possible to resize a SnapLock FlexVol volume in ONTAP 9?	18
What is privileged-delete? How does it work in ONTAP 9?	18

What are the privileges for the <code>vsadmin-snaplock</code> role? Can the administrator with the <code>vsadmin</code> role perform privileged-delete?	18
What is the purpose of the SnapLock audit log?	18
What is a WORM appendable file?.....	19
How does SnapLock use fingerprints?	19
How can I get the list of expired files?	19
Can I use Advanced drive partitioning (ADP) with SnapLock aggregates?	20
Can I put SVM root volume on a SnapLock aggregate?.....	20
Is LUN creation allowed on SnapLock volumes?.....	20
5 Interoperability.....	21
Does System Manager support SnapLock?	21
Does Unified Manager support SnapLock?	21
Is it possible to enable compression and dedupe on a SnapLock volume?.....	21
Is SnapLock supported with NetApp Storage Encryption (NSE) drives?	21
Is SnapLock supported with NetApp Volume Encryption (NVE)?	21
How is SnapLock feature interoperability on ONTAP 9 compared to on 7-Mode?	21
Is there any new change in the programmatic interfaces (ZAPIs) for SnapLock on ONTAP 9?	26
5.1 Data Protection	30
What is LockVault? Is it still supported with SnapLock on ONTAP 9?	30
Is vaulting supported with SnapLock on ONTAP 9?	30
Does RAID-level SyncMirror work with SnapLock volumes?.....	30
How is the support for MetroCluster with SnapLock on ONTAP 9?.....	30
Do NDMP-initiated DUMP and restore preserve the WORM attributes of the files?	30
6 Miscellaneous Questions and Additional Resources.....	31
6.1 Reclaiming Disks	31
I accidentally locked a file for 30 years. How can I reuse the disk?	31
How do I sanitize SnapLock disks?	31
7 Appendix	32
Sample output for the entry in the audit log before and after a privileged-delete operation.....	32

8 **Contact Us** 37

LIST OF TABLES

Table 1) SnapMirror relationships between SnapLock volumes 9

Table 2) SnapLock feature interoperability on ONTAP 9 22

Table 3) List of ZAPIs for SnapLock on ONTAP 9 26

1 Overview

SnapLock is NetApp's high-performance compliance solution that provides data retention and WORM (write once read many) protection for retained data. With SnapLock, customers can create special-purpose volumes in which files can be stored and committed to a nonerasable, nonrewritable state either for a designated retention period or indefinitely.

SnapLock is a license-based feature of ONTAP that works with application software to administer nonrewritable storage of data. There are two SnapLock versions: SnapLock Compliance (SLC) and SnapLock Enterprise (SLE). Both versions run on NetApp systems with either lower-cost SATA-based drives or higher-performance SAS or fiber-attached disk drives. This flexibility allows customers to buy the type and amount of storage that fits their business needs for SnapLock WORM storage.

This document provides answers to common questions about SnapLock in ONTAP 9. Unless otherwise stated, all references to SnapLock in this document mean SnapLock in ONTAP 9, which is a release for a clustered environment only. References to SnapLock products for Data ONTAP operating in 7-Mode are specifically called out.

1.1 SnapLock Release Summary

What was released for SnapLock in ONTAP 9.2?

Answer: The major addition for the ONTAP 9.2 release is the support for SnapLock with NetApp Volume Encryption (NVE). NVE is a software-based, data-at-rest encryption solution available starting with NetApp ONTAP® 9.1 management software. NVE allows ONTAP to encrypt data and to have that data be stored on disk without requiring self-encrypting drives. Customers can use any existing disk with NVE, which also includes NetApp Storage Encryption (NSE) drives for “double” or “layered” encryption.

What was released for SnapLock in ONTAP 9.1?

Answer: The major addition for the ONTAP 9.1 release is the support for SnapLock on All Flash FAS (AFF) systems and hybrid systems. SnapLock is not supported on hybrid aggregates or Solid State Drives (SSDs) in 9.0.

What was released for SnapLock in ONTAP 9.0?

Answer: ONTAP 9.0 is the first release in clustered mode that supports SnapLock. In addition to CLI and ZAPI support, this release includes System Manager and Unified Manager support for SnapLock.

Does the ONTAP 9 release affect 7-Mode SnapLock?

Answer: No. SnapLock in 7-Mode remains unaffected and will continue to be supported for that version.

2 Requirements

Does SnapLock require any special hardware?

Answer: No. SnapLock is a licensed feature of Data ONTAP and is supported on all the FAS platforms that support ONTAP 9. SnapLock is not a software-only solution. SnapLock is an integrated hardware and software solution. This solution is important for strict WORM regulations such as SEC 17a-4, which requires an integrated hardware and software solution. For more information, refer to the SEC Interpretation: [Electronic Storage of Broker-Dealer Records](#).

Note: SnapLock aggregates are not supported on hybrid aggregates or SSD drives in ONTAP 9.0, but they are supported in ONTAP 9.1.

2.1 Licensing

Is a SnapLock license part of the Premium Bundle?

Answer: No. A SnapLock license is not part of the Premium Bundle.

What is the SnapLock licensing scheme on ONTAP 9? Do I require different licenses for SLE and SLC?

Answer: Unlike in 7-Mode, in which we had a separate license for SnapLock Compliance (SLC) and one for SnapLock Enterprise (SLE), in ONTAP 9 a single SnapLock license enables both SLE and SLC functionality. 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 licensed functionality, all nodes must be licensed for the functionality.

What is the licensing scheme for existing SnapLock customers on 7-Mode?

Answer: In ONTAP 9, we have consolidated the SnapLock Compliance (SLC) and SnapLock Enterprise (SLE) features into a single entitlement enabled by a SnapLock Compliance license key.

- If customer is currently using 7-Mode SnapLock Compliance license and wants to migrate to ONTAP9, then they can use the same SLC key (SNAPLOCK_2) in ONTAP 9 as long as they are 28 character keys.
- Customers who own only SnapLock Enterprise will be provided a SnapLock Compliance license free of charge. The licenses have already been generated and applied to the serial numbers associated with them. These new license keys can be downloaded from the NetApp support Site. If licenses were not applied correctly due to system errors, please let the support team know.

3 Architecture

What is the difference between SnapLock Compliance (SLC) and SnapLock Enterprise (SLE)?

Answer: SnapLock Compliance (SLC) is designed to assist organizations in implementing a comprehensive archival solution for meeting strict regulatory requirements for data retention, such as SEC 17a-4. 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 but can be deleted after their retention periods expire. Moreover, a SnapLock Compliance volume cannot be deleted until all the data stored on it has passed its retention period.

SnapLock Enterprise (SLE), 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.

Which file access protocols are supported?

Answer: Currently, the supported open file protocols for SnapLock are NFS (versions 2, 3, and 4) and CIFS (SMB 1.0, 2.0, and 3.0).

Can I copy a regular volume to a SnapLock aggregate or vice versa?

Answer: No. You cannot copy a regular volume to a SnapLock aggregate or vice versa. In ONTAP 9, only like-to-like NetApp SnapMirror relationships are supported between regular, SLE, and SLC volumes. Table 1 shows the combinations that are supported.

Table 1) SnapMirror relationships between SnapLock volumes.

		SnapMirror Destination		
		SLC	SLE	Regular
SnapMirror Source	SLC	✓	✗	✗
	SLE	✗	✓	✗
	Regular	✗	✗	✓

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 `atime_update` volume option set to `false` and explicit updates to `atime` are interpreted as changes to the retention period of the file.
- Administrative actions are restricted on SnapLock volumes.

Does SnapLock automatically delete expired WORM data?

Answer: No. SnapLock never automatically deletes any files from the volume. Once the retention time expires, the files can be deleted either manually or using an application.

Can I use Snapshot copies on a SnapLock volume?

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

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

Answer: No. 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 step is to use `ndmptcopy` to an equivalently sized or larger SnapLock Enterprise or SnapLock Compliance volume. The destination compliance volume should be empty when `ndmptcopy` is initiated. `ndmptcopy` 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 by committing individual files to WORM either by removing the write permissions on the files or by using the autocommit feature of SnapLock. This feature automatically locks the files placed on the volume if there is no change to the file in the specified period. A script or an application can also be used to automate and accomplish these tasks.

Can I change any attributes on a WORM committed file before or after it expires?

Answer: **Before the retention date has been reached:**

Once a file on a SnapLock volume has been committed to the WORM state, only one attribute can be modified: The retention date of the file can be extended (pushed further out into the future). The retention period can never be shortened by pulling the retention date closer. This does not compromise the WORM protection on the file in any way. However, it does enable the retention period to be extended before the file expires, an important feature in cases such as legal holds on certain records.

After the retention date has been reached:

Once the file has expired (reached its retention time), two changes are possible:

- The retention date of the file can be extended (pushed further out into the future). This reenables

WORM protection on the file as if it were being committed for the first time.

- In case of CIFS, the write permissions on the file can be reenabled and it can be deleted. In NFS cases, changing of file permissions is not needed to delete the file.

File contents and any other file property still cannot be modified.

Are there ways around SnapLock protection?

Answer: There are no logical back doors around SnapLock through file protocol or administrative interfaces. SnapLock functionality has been tightly integrated with the NetApp Data ONTAP WAFL® (Write Anywhere File Layout) subsystem, which controls all logical access to the volume. Because every protocol must go through the protocol-independent WAFL layer to get access to the disks, SnapLock automatically provides protection for all current and future file protocols.

Significant effort was expended identifying and locking down all administrative functions that could cause the destruction or modification of SnapLock protected WORM data. This lock-down extends from the lowest-level administrative and maintenance functions to integration with other features, such as SnapMirror replication. However, SnapLock is not immune to disk errors and corruption. It relies on WAFL iron and underlying RAID to recover from RAID recoverable errors wherever possible.

Has SnapLock been validated by a third party?

Answer: SnapLock has been assessed by an independent third party (Cohasset) and was found to meet the SEC 17a-4 regulations. Validation report can be accessed on NetApp [fieldportal](#).

Additionally, KPMG in Germany has certified SnapLock for:

- German trade and tax law regulations (e.g. §§ 238 ff HGB , §§ 140 – 148 AO)
- Austrian trade and tax law regulations(UGB, especially §§ 190 and 216)
- Swiss trade and tax law regulations (OR §§ 957 ff.)

The certification is available in German language and can be directly accessed here:

- [Short version](#)
- [Long version](#)

Electronic distribution is permitted if the recipient signs the disclaimer only – therefore please **distribute only the URLs**.

3.1 ComplianceClock

What prevents me from resetting the system clock to delete files early?

Answer: 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 snapshots. 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 has to properly take note of the time zone and make sure that the storage system's time is as accurate as possible. Once the ComplianceClock is initialized, the administrator cannot perform any action that will cause a forward adjustment. This feature ensures that the retention period of WORM files can never be shortened by doing forward adjustments of the reference clock.

How does ComplianceClock work? Are there any changes in ComplianceClock from 7-Mode to ONTAP 9?

Answer: The ComplianceClock architecture was last changed in 7-Mode in Data ONTAP 8.1 to minimize the clock skews. There was no change in the architecture after that. 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 expiry of WORM files and snapshots in that volume. Since the VCC is per volume, VCC skewing of volume X does not affect the VCC of volume Y. The VCC will be updated on SnapLock volumes lazily, only when they are in Consistency Point (CP) 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 (similar to the Data ONTAP 7G ComplianceClock).

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

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

The ComplianceClock can be viewed using the following command:

```
snaplock compliance-clock show
```

Refer to [TR-3618](#) for more details on ComplianceClock.

How is the VCC updated on SnapLock volumes? What happens to VCC when I offline a volume?

Answer: The VCC is initialized only once (either at creation or upgrade) and cannot be reinitialized after that. The VCC obtains its starting value from the System ComplianceClock (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
```

Each volume stores the VCC and reference SCC time at which the volume was last updated. 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 will be no VCC skews irrespective of the duration for which the volume was kept offline.

How does ComplianceClock behave when there is a SnapMirror relationship?

Answer: If there is a SnapMirror relationship between SnapLock volumes, the SnapLock source type and the source and destination volume 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 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.

What happens to VCC in case of Storage Failover (SFO) or AggregateRelocate (ARL) of a SnapLock aggregate?

Answer: Once the aggregate moves to the destination node, VCC of all the SnapLock volumes hosted on that aggregate is reassociated with the new host node's SCC, resulting in the loss of some clock ticks. This does not skew VCC too much as compared to SCC. VCC will lose some ticks until the volume moves to the partner node.

What happens to the ComplianceClock during a system shutdown or outage? I left my system powered off for a while, and now the ComplianceClock is behind the system clock. How do I fix this?

Answer: You cannot fix this situation. Once the ComplianceClock has been initialized, it cannot be modified under any circumstances. This is to prevent any tampering with the retention date. The ComplianceClock is a software clock, so it will not run during the outage, but its last state will be persistently stored before the shutdown. When the system is brought back up, the ComplianceClock will be running behind real time. 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. This is safe to do from a compliance perspective, assuring that files under SnapLock protection can never be prematurely deleted.

How do I find out more about ComplianceClock?

Answer: For more details on ComplianceClock, refer to "[TR-3618: Guide to Understanding SnapLock ComplianceClock](#)."

3.2 Transition from 7-Mode to ONTAP 9

Which tools are available to SnapLock customers to support transition from 7-Mode to ONTAP 9?

Answer: The 7MTT v3.1 (Copy-Based Transition) supports transitions involving SnapLock volumes in ONTAP 9. The tool supports both SnapLock Compliance (SLC) and SnapLock Enterprise (SLE) volumes.

Is there any support for Copy-Free Transition for SnapLock aggregates?

Answer: No. Copy-Free Transition (CFT) is not supported on SnapLock aggregates.

Are there any requirements specific to SnapLock aggregates that must be met before carrying out transition from 7-Mode to ONTAP 9?

Answer: A major design change from 7-Mode is the like-to-like transition between regular, SLE, and SLC volumes. This change means that SLC can transition only to SLC and that SLE can transition only to SLE. Non SnapLock to SnapLock transitions or vice versa fail. In the case of SLC volumes, `snapmirror resync` is not allowed if it results in the loss of data at the resync destination.

How can I transition the DR relationship for SnapLock volumes from 7-Mode to ONTAP 9?

Answer: DR relationships involving SnapLock volumes can be transitioned from 7-Mode to ONTAP 9 using Copy-Based Transition (CBT) and following these steps:

- Transition the destination volume separately.
- Make the transitioned volume the destination volume for the DR relationship by creating a relationship between the 7-Mode source volume and the ONTAP 9 transitioned destination volume. Skip this step for SLC because `snapmirror resync` cannot be done for SLC because it might result in data loss.
- Transition the source volume.
- For SLE, break the DR relationship between the 7-Mode source and the ONTAP 9 destination. For SLC, a `snapmirror break` will be performed followed by reestablishing `snapmirror` between the source and destination volumes.
- Do a SnapMirror resync between the transition source volume and the destination volume.

In summary, only parallel transitions are supported for transitioning SLC DR relationships, but both staggered and parallel transitions are supported for SLE DR relationships.

Are there any best practices for transitioning 7-Mode audit log volumes to ONTAP 9?

Answer: Audit log volumes in 7-Mode are node specific, whereas in ONTAP 9 they are Vserver specific. During transition, users must decide where to place this volume in the ONTAP 9 destination. This behavior should be acceptable because log volumes are generated only by operations on SLE volumes for which the administrator is trusted.

4 Configuration

What happens when I remove the SnapLock license or 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. Moreover, once the SnapLock volumes have been created, you can commit files to WORM state even without 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
- 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
- Disabling or permanently disabling privileged-delete
- Deleting the SnapLock audit log volume configuration

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. SnapLock 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 state. This can be done at an individual file level (by removing the write permissions on the file). It can also be done 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.

What granularity of data protection does SnapLock allow?

Answer: SnapLock allows data protection at the file level by enabling users to set the retention period at an individual file level. If the retention period is not set explicitly at the individual file level, then SnapLock uses the volume's default retention period. 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.

How do I manually commit a file to WORM state on a SnapLock volume?

Answer: 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 ensure that the file is initially writable on the SnapLock volume before making it read-only to commit to WORM state. For more details on the procedure, refer to the [Archive and Management Compliance Guide](#).

How does autocommit work on ONTAP 9? What is the range for autocommit-period? Is it configurable per volume?

Answer: 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 will have 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. The minimum configurable value for autocommit-period is 2 hours and the maximum is 24 hours.

Note: Autocommit is not suited to 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 an unmodifiable file that 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.

Since a Snapshot copy does not have a last access time field, how does the retention period get assigned to Snapshot copies?

Answer: Each SnapLock volume has a default retention period assigned to it. The value for this default retention period is initially set to a minimum retention period of 0 years for SLE volumes and a maximum retention period of 30 years for SLC volumes. Each Snapshot copy is committed with this default retention period at first. The retention period can be extended later, if needed.

How do I identify SnapLock volumes in a system?

Answer: You can use the following command to identify SnapLock volumes.

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

A SnapLock type entry in the Options field identifies a SnapLock volume as either `compliance` or `enterprise`.

How can I find the expiration time of the SnapLock volume?

Answer: The expiration time 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 the ComplianceClock, has been reached. You can use the following command to check the retention periods and the expiry time of a volume. The Expiry Time field in the output displays the expiration time for the specified volume.

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

Why can't I rename directories on a SnapLock volume?

Answer: WORM directories behave in the same way as they do on regular volumes with one important exception: They cannot be renamed or moved once created. Regulatory compliance requires that WORM data not only be nonerasable and nonrewritable but also must be locked down in the same location at which it was created. For ONTAP WORM implementation, this means that the directory path to WORM files must be locked down and never change. If the path were allowed to change, then multiple possibilities for problems would present themselves, which might cause serious problems with the SEC.

The first and most obvious problem is that the front-end databases, which our application partners use to keep track of the location of WORM data on the volume, would lose track of any content contained underneath the moved or renamed directory. This issue might cause problems with the SEC because customers are required by SEC 240.17a-4 to be able to generate an index of all archived content and verify the integrity of all the content on demand. If customers can't do so, they face substantial penalties and fines.

The second and more serious problem is that a security hole would be introduced into our WORM model.

If someone wanted to modify a WORM file in a covert way, that person could simply move the directory containing the target file to an obscure place (or rename it as "..."). Then the person could recreate the original directory, copy all original files except the file he/she wants to tamper with into this new directory, and then copy that desired file to a non-WORM volume to edit it. Finally, the person could move the tampered file back into the new directory and commit the tampered file as WORM.

The applications are free to create and delete temporary directories on WORM volumes as long as they don't commit any WORM files inside those directories. Regular directory semantics dictate that you cannot delete a non-empty directory. Therefore, if you create WORM files in a directory, the directory will be permanently non-empty and, as a result, become impossible to delete until the specified retention period expires.

Is it possible to resize a SnapLock FlexVol volume in ONTAP 9?

Answer: Yes. It is possible to both shrink and grow NetApp SnapLock FlexVol® volumes in ONTAP 9. However, you cannot shrink SnapLock FlexVol volumes below the actual used capacity.

What is privileged-delete? How does it work in ONTAP 9?

Answer: 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 file that exceeded its retention period. The deletion is logged in an audit file on a SnapLock Compliance log volume so that there is a nonerasable record of the file's previous existence and early deletion. Therefore, the privileged-delete feature is also known as auditable delete.

What are the privileges for the `vsadmin-snaplock` role? Can the administrator with the `vsadmin` role perform privileged-delete?

Answer: These are the capabilities the `vsadmin-snaplock` role has:

- Managing his/her own user account local password and key information
- Managing volumes, quotas, Snapshot copies, and files
- Configuring NFS and CIFS
- Performing SnapLock related management including privileged-delete
- Monitoring jobs
- Configuring LDAP, DNS, and NIS
- Monitoring the network interface

The `vsadmin` can also perform SnapLock related management excluding privileged-delete.

What is the purpose of the SnapLock audit log?

Answer: SnapLock audit logging logs operations on SnapLock volumes and WORM files. One of the primary requirements of log files is to have the details of the WORM files that were removed before their expiry date using the privileged-delete operation. Hence, the log files have to be retained at least until the retention time of the files that were removed. So, the log files themselves will be WORM files and only SLC volumes can be made into audit log volumes. The following SnapLock operations are logged in the SnapLock log volume:

- Privileged-delete of a retained WORM file
- Changing the SnapLock log volume
- Adding or deleting the `vsadmin-snaplock` role for a user
- Turning privileged-delete to `enabled` on an SLE volume
- Turning privileged-delete to `disabled` from the `enabled` state on an SLE volume
- Turning privileged-delete to `permanently-disabled` from the `enabled` state on an SLE volume
- Turning privileged-delete to `permanently-disabled` from the `disabled` state on an SLE volume

In the SnapLock audit log file, you can find details related to the privileged-deleted WORM file, such as whether a file has been deleted, when it was deleted, and who deleted it. For details about the entry in the audit log before and after a privileged-delete operation, refer to section 7 of the Appendix of this document.

What is a WORM appendable file?

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

How does SnapLock use fingerprints?

Answer: A fingerprint is useful for checking data integrity. The fingerprint operation allows you to generate a fingerprint on a per-file basis by using one of the hashing algorithms—MD5 or SHA-256. Doing so enables the user to verify the integrity of the file at any given time. A file fingerprint is exported outside ONTAP through a CLI and ZAPI for user and partner applications. SnapLock does not store any file fingerprint on disk anywhere in the system. A file fingerprint calculates hash digest on the fly for a file as requested by the user over a CLI or ZAPI. You can issue the file fingerprint using this command:

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

Once the file fingerprint is issued with that command, a session id is generated. You can use this session id to get the status with the following command:

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

Once the status is 'Completed', use the same session id from the above command and issue the following command to get the fingerprint output:

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

The file type will be "worm" in the case of a SnapLock file and "worm_appendable" in the case of a WORM appendable file. The file type will be "worm_active_log" for an active WORM log file, "worm_log" for a closed WORM log file, and "regular" for regular or non SnapLock files.

How can I get the list of expired files?

Answer: There is no CLI to get the list of expired files. However, you can use the file fingerprint to learn

whether a file has expired or not.

Can I use Advanced drive partitioning (ADP) with SnapLock aggregates?

Answer: Yes. ADP is supported with SnapLock aggregates on FAS systems on ONTAP 9.0, and on AFF systems on ONTAP 9.1. Low end platforms with small numbers of drives may need to be carefully considered to make sure that sufficient redundancy is available for the root partitions, and sufficient capacity is available for the SnapLock aggregate.

Can I put SVM root volume on a SnapLock aggregate?

Answer: No. You cannot put an SVM root volume on a snaplock aggregate. You need to create a separate non-SnapLock aggregate to put the SVM root volume on it.

Is LUN creation allowed on SnapLock volumes?

Answer: No. ONTAP disallows the creation of LUNs on SnapLock volumes.

5 Interoperability

Does System Manager support SnapLock?

Answer: Yes. NetApp OnCommand® System Manager support was introduced for SnapLock in ONTAP 9. You can use System Manager to perform SnapLock related operations. Following is a summary of the supported System Manager actions:

- Installing the SnapLock license on the nodes
- Setting the SnapLock ComplianceClock on each node that ticks independently of the system clock
- Creating and managing SnapLock Compliance and SnapLock Enterprise aggregates
- Creating and managing SnapLock Compliance and SnapLock Enterprise volumes
- Setting different SnapLock options on SnapLock volumes, namely autocommit-period, default-retention-period, minimum-retention-period, and maximum-retention-period
- Creating a DR relationship for SLC or SLE volumes
- Creating a vault relationship from a non SnapLock volume to a SnapLock volume

Does Unified Manager support SnapLock?

Answer: Yes. Unified Manager 7.0 supports SnapLock in ONTAP 9.

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

Answer: Yes. Compression and dedupe can be enabled on a SnapLock volume, bringing a storage efficiency advantage along with the WORM data.

Is SnapLock supported with NetApp Storage Encryption (NSE) drives?

Answer: Yes. Starting with ONTAP 9, both SnapLock Compliance and SnapLock Enterprise are supported in combination with NSE drives.

Is SnapLock supported with NetApp Volume Encryption (NVE)?

Yes. Starting with ONTAP 9.2, both SnapLock Compliance and SnapLock Enterprise are supported with NVE.

How is SnapLock feature interoperability on ONTAP 9 compared to on 7-Mode?

Answer: The following table lists all of the 7-Mode features and whether SnapLock interoperates with them on ONTAP 9. This table also includes the ONTAP 9 equivalents of those features and features exclusive to ONTAP 9.

Table 2) SnapLock feature interoperability on ONTAP 9

7-Mode			ONTAP 9			Comments
Feature	SLE	SLC	Feature	SLE	SLC	
Data Protection						
MetroCluster™	Yes	Yes*	MetroCluster	Yes	No**	*User needs to back up data that will be deleted by resync process before initiating it in 7-Mode with SLC aggregate. **Privilege-delete is not supported because creation of audit log volume, which is of type SLC to support privilege-delete, is not allowed in MCC.
NDMP/Dump	Yes	Yes, with some restrictions	NDMP/Dump	Yes	Yes, with same restrictions as 7-Mode	
QSM	Yes	Yes	No equivalent	N/A	N/A	
No equivalent	N/A	N/A	LRSE DR	No	No	
Single file SnapRestore	Yes	No	Single file SnapRestore	Yes	No	
SnapRestore (aggr and volume)	Yes	No	SnapRestore (volume)	Yes	No	
Snapshot	Yes	Yes	Snapshot	Yes	Yes	

7-Mode			ONTAP 9			Comments
RAID SyncMirror®	Yes	Yes (same restrictions as 7-Mode)	SyncMirror	Yes	No	
VSM	Yes	Yes*	BRE	Yes	Yes**	*No resync support with VSM. **Resync process identifies if there is data that will be deleted as a result of resync. If there is, then user must create a clone of the volume and then proceed.
Open Systems SnapVault®	No	No	A2N	No	No	
SMTape	No	No	SMTape	No	No	
vFiler® unit DR	Yes	No	SVM DR	No	No	vFiler DR is not supported with SLC because of lack of resync support.
Protocol						
FCP and iSCSI	No	No	FCP and iSCSI	No	No	
NFS v2, v3, v4	Yes	Yes	NFS v2, v3, v4	Yes	Yes	
SMB 1.0, 2.0, 3.0	Yes	Yes	SMB 1.0, 2.0, 3.0	Yes	Yes	

7-Mode			ONTAP 9			Comments
Performance						
Flash Cache™	Yes	Yes	Flash Cache	Yes	Yes	Starting with ONTAP 9.1
Flash Pool™	No	No	Flash Pool	Yes	Yes	Starting with ONTAP 9.1
All flash	No	No	All flash	Yes	Yes	Starting with ONTAP 9.1
FlexCache®	No	No	Not supported	N/A	N/A	
FlexShare®	No	No	FlexShare	Yes	Yes	
No equivalent	N/A	N/A	Loadshare mirror	No	No	
Block reallocation	No	No	Block reallocation	No	No	
Storage Efficiency						
Compression	Yes	Yes	Compression	Yes	Yes	Inline compression does not work with WORM append.
Dedupe	Yes	Yes	Dedupe	Yes	Yes	
Thin provisioning	Yes	Yes	Thin provisioning	Yes	Yes	
FlexClone®	Yes	Yes	FlexClone	Yes	Yes	Clone inherits the SnapLock properties of the source.
LUN cloning	No	No	LUN cloning	No	No	
File cloning	No	No	File cloning	No	No	

7-Mode			ONTAP 9			Comments
Fractional reserve	Yes	Yes	Fractional reserve	Yes	Yes	
Data Management						
FlexVol	Yes	Yes	FlexVol	Yes	Yes	
No equivalent	N/A	N/A	Flex groups	No	No	
No equivalent	N/A	N/A	Infinite Volume	No	No	
Shared storage with MultiStore®/SnapMover®	Yes	Yes	No equivalent	N/A	N/A	
No equivalent	N/A	N/A	ARL	Yes	Yes	Required to support nondisruptive tech refresh.
DataMotion™ for Volumes	No	No	No equivalent	N/A	N/A	
DataMotion for vFiler	Yes	No	No equivalent	N/A	N/A	
Volume copy	Yes	Yes	No equivalent	N/A	N/A	In ONTAP 9, there is no volume copy. NetApp SnapMirror® is used.
No equivalent	N/A	N/A	Vol move	No	No	
No equivalent	N/A	N/A	SFxOD	No	No	
Other Core Features						
V-Series	Yes	No	FlexArray®	Yes	No	
No equivalent	N/A	N/A	VSA	Yes	No	

7-Mode			ONTAP 9			Comments
Large aggregates	Yes	Yes	Large aggregates	Yes	Yes	
IPv6	Yes	Yes	IPv6	Yes	Yes	
NSE	Yes	No	NSE	Yes	Yes	
No equivalent	N/A	N/A	NVE	Yes	Yes	
vStorage API	No	No	vStorage API	No	No	
No equivalent	N/A	N/A	Antivirus	Yes	Yes	
HA	Yes	Yes	HA	Yes	Yes	

Is there any new change in the programmatic interfaces (ZAPIs) for SnapLock on ONTAP 9?

Answer: Yes. In addition to the 12 7-Mode ZAPIs, 7 new ZAPIs were introduced in ONTAP 9. Following is the list of ZAPIs that are available for SnapLock on ONTAP 9.

Table 3) List of ZAPIs for SnapLock on ONTAP 9

7-Mode ZAPI	ONTAP 9 ZAPI	Description
snaplock-get-system-compliance-clock	snaplock-get-node-compliance-clock	Gets SnapLock Node ComplianceClock date and time.
snaplock-set-system-compliance-clock	snaplock-set-node-compliance-clock	Sets the ComplianceClock of the node equal to the current system clock.
snaplock-get-volume-compliance-clock	snaplock-get-volume-compliance-clock	Gets the SnapLock Volume ComplianceClock date and time.
snaplock-get-log-volume	snaplock-get-vserver-log-iter	Displays audit log configuration.

7-Mode ZAPI	ONTAP 9 ZAPI	Description
snaplock-set-log-volume	snaplock-create-vserver-log	- Creates audit log configuration for a Vserver. <code>is-force</code> option from 7-Mode is removed. - Extra inputs from 7-Mode: <code>Max-log-size</code> (optional) <code>Default-retention-time</code> (optional)
snaplock-log-archive	snaplock-archive-vserver-log	- Archives the active WORM log file. - In ONTAP 9, only <code>base-name</code> is supplied as input, unlike in 7-Mode, in which the <code>log-volume</code> name can also be provided.
snaplock-log-status-list-info	snaplock-get-vserver-active-log-files-iter	Displays status of active SnapLock log files.
snaplock-get-snaplock-retention-time	snaplock-get-snaplock-retention-time	Gets retention period of a WORM file.
snaplock-set-snaplock-retention-time	snaplock-set-snaplock-retention-time	Sets retention period on a file.
file-snaplock-retention-time-list-info	No equivalent ZAPI exists in ONTAP 9	
file-get-snaplock-retention-time-list-info-max	No equivalent ZAPI exists in ONTAP 9	
snaplock-privileged-delete-file	snaplock-file-privileged-delete	Executes a privileged-delete on a SnapLock file.

7-Mode ZAPI	ONTAP 9 ZAPI	Description
snaplock-get-options	volume-get-snaplock-attrs	- Returns SnapLock attributes of a given SnapLock volume. - Unlike 7-Mode, which returned only the <code>privdel</code> SnapLock option, it returns all SnapLock attributes.
snaplock-set-options	volume-set-snaplock-attrs	- Sets the given SnapLock attributes of a SnapLock volume. - Unlike 7-Mode, which only allows setting the <code>privdel</code> option, it can set any of the modifiable SnapLock attributes: - Minimum-retention-period, default-retention-period, maximum-retention-period. - Autocommit period
No equivalent ZAPI exists in 7-Mode	snaplock-volume-set-privileged-delete	Sets the SnapLock privileged-delete option on a SnapLock Enterprise volume (on/off/disallowed).
No equivalent ZAPI exists in 7-Mode	snaplock-destroy-vserver-log	Deletes audit log configuration for a Vserver.
No equivalent ZAPI exists in 7-Mode	snaplock-modify-vserver-log	Modifies audit log configuration for a Vserver.
No equivalent ZAPI exists in 7-Mode	file-fingerprint-get-iter	Returns information for one or several fingerprint operations.
No equivalent ZAPI exists in 7-Mode	file-fingerprint-start	Starts the fingerprint computation of a file. It returns session id and dump return fingerprint output for a given session id.
No equivalent ZAPI exists in 7-Mode	file-fingerprint-abort	Aborts an ongoing fingerprint operation.

7-Mode ZAPI	ONTAP 9 ZAPI	Description
No equivalent ZAPI exists in 7-Mode	<code>file-fingerprint-dump</code>	Returns the fingerprint information for a file for a given session id.

5.1 Data Protection

What is LockVault? Is it still supported with SnapLock on ONTAP 9?

Answer: NetApp LockVault™ technology is a disk-based regulatory compliance solution for unstructured data in 7-Mode. In ONTAP 9, this feature is known as SnapLock with SnapVault. The feature 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) daily 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) that tracks changes between Snapshot copies and stores them on a WORM volume so the log cannot be modified. In ONTAP 9, creation of a Compliance Journal that tracks the file changes per transfer is not supported.

Is vaulting supported with SnapLock on ONTAP 9?

Answer: Vaulting of SnapLock volumes is not supported. A SnapMirror transfer fails if the source of the vault relationship is a SnapLock volume.

Does RAID-level SyncMirror work with SnapLock volumes?

Answer: SyncMirror works with SLE aggregates as long as the hardware platform supports SyncMirror. However, SyncMirror is not allowed on SLC aggregates.

How is the support for MetroCluster with SnapLock on ONTAP 9?

Answer: Only SLE aggregates are supported in MCC configuration in ONTAP 9. Vserver-specific SnapLock configurations are replicated to both sites. There are few scenarios in which compliant data can be lost; therefore, there is no support for SLC with MCC.

- Resync after an unplanned switchover (USO): As part of USO, plex1 of the aggregate owned by the node in the disaster site is switched over. An example is the case in which plex1 does not have up-to-date data after NVLog replay, as part of the USO processing, is complete. The user starts committing new WORM data to plex1. When the disaster site is back up, plex1 resyncs with plex0 in the disaster site and in the process deletes any data on plex0 that is not on plex1.
- Alternate loss of plexus: One plex of the mirrored aggregate, say, plex0, goes offline and new WORM data gets committed to plex1 only. A little later, plex1 goes offline. If at this point plex0 becomes available, the user can forcefully bring plex0 online and new WORM data is committed to plex0 only. Whenever plex1 comes online, the resync operation results in the loss of data.
- NVsave file loss: During switchover, since the DR partner root won't ever be mounted, we will never find any NVsave to replay from the DR partner. If the NVsave file existed before USO for an aggregate, then data loss occurs because of the missing NVsave file after USO.

Do NDMP-initiated DUMP and restore preserve the WORM attributes of the files?

Answer: Yes. The data streams of these features were 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 non SnapLock volume, the WORM attributes are preserved but ignored and not enforced by Data ONTAP.

6 Miscellaneous Questions and Additional Resources

The following section provides references for additional information.

6.1 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 disk back is by waiting for the file to expire. There is no way to bypass the retention period; however, NetApp might have a partner that can degauss the disk. For more information, contact the NetApp Global Services team.

SnapLock Enterprise volumes can be destroyed even before the expiry time by the SnapLock administrator. By using the privileged-delete feature, individual files on a SnapLock Enterprise volume can be deleted by a privileged user with the `vsadmin-snaplock` role and if the operation is logged in the SnapLock audit log.

How do I sanitize SnapLock disks?

Answer: When all the WORM data has expired in a SnapLock Compliance volume, the volume can be deleted. For SnapLock Enterprise, the system administrator can delete the volumes at any time. When all SnapLock flexible volumes on an aggregate are deleted, the aggregate can also be deleted. After the deletion of the SnapLock aggregate, the disks can be sanitized by using the `disk sanitize` command.

7 Appendix

Sample output for the entry in the audit log before and after a privileged-delete operation

```
<LogEntry LogVersion="3" Date="20160420 085519 GMT" SCCdate="20160420 084714 GMT" VCCdate="20160420
084714 GMT" Node="vsim7" Vserver="vs1">
<snaplock_pre_privileged_delete_1>
  <sequence_number>9bae8bba-06d5-11e6-b1b3-005056bb6335</sequence_number>
  <file_pathname>/vol/sle_vol1/file.1</file_pathname>
  <file_expires>infinite</file_expires>
  <record_expires>infinite</record_expires>
  <file_fingerprint><fingerprint>
    <fingerprint-info>
      <fingerprint-algorithm>SHA256</fingerprint-algorithm>
      <fingerprint-start-time>1461142518</fingerprint-start-time>
      <formatted-fingerprint-start-time>Wed Apr 20 08:55:18 GMT 2016</formatted-
fingerprint-start-time>
      <fingerprint-scope>DATA_AND_METADATA</fingerprint-scope>
      <fingerprint-version>3</fingerprint-version>
      <fingerprint-input-path>/vol/sle_vol1/file.1</fingerprint-input-path>
      <vserver-name>vs1</vserver-name>
      <vserver-uuid>37eacca1-06bb-11e6-9edf-005056bb6335</vserver-uuid>
      <volume-name>sle_vol1</volume-name>
      <volume-msid>2160232325</volume-msid>
      <volume-dsid>1028</volume-dsid>
      <volume-containing-aggregate>sle_aggr</volume-containing-aggregate>
      <aggregate-id>fb20fab3-27fa-4e88-a468-1245c36e7d14</aggregate-id>
      <snaplock-license>available</snaplock-license>
      <snaplock-system-compliance-clock>1461142033</snaplock-system-compliance-clock>
      <formatted-snaplock-system-compliance-clock>Wed Apr 20 08:47:13 GMT
2016</formatted-snaplock-system-compliance-clock>
      <volume-snaplock-type>enterprise</volume-snaplock-type>
      <snaplock-volume-compliance-clock>1461142033</snaplock-volume-compliance-clock>
      <formatted-snaplock-volume-compliance-clock>Wed Apr 20 08:47:13 GMT
2016</formatted-snaplock-volume-compliance-clock>
```

```

        <formatted-volume-expiry-date>infinite</formatted-volume-expiry-date>
    <file-metadata>
        <file-fingerprint-info>
            <path>/vol/sle_vol1/file.1</path>
            <fsid>1028</fsid>
            <fileid>97</fileid>
            <file-type>worm</file-type>
            <file-size>1048576</file-size>
            <creation-time>1461142021</creation-time>
            <formatted-creation-time>Wed Apr 20 08:47:01 GMT 2016</formatted-
creation-time>
            <modified-time>1461142021</modified-time>
            <formatted-modified-time>Wed Apr 20 08:47:01 GMT 2016</formatted-
modified-time>
            <changed-time>1461141542</changed-time>
            <is-changed-time-wraparound>false</is-changed-time-wraparound>
            <formatted-changed-time>Wed Apr 20 08:39:02 GMT 2016</formatted-
changed-time>
            <formatted-retention-time>infinite</formatted-retention-time>
            <owner-id>0</owner-id>
            <group-id>0</group-id>
            <data-
fingerprint>MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=</data-fingerprint>
            <metadata-
fingerprint>JgY8a6p+dAvTDSC1g19b4PXF3jqDpJUTXnex/g20A3Q=</metadata-fingerprint>
        </file-fingerprint-info>
    </file-metadata>
    <fingerprint-end-time>1461142518</fingerprint-end-time>
    <formatted-fingerprint-end-time>Wed Apr 20 08:55:18 GMT 2016</formatted-
fingerprint-end-time>
</fingerprint-info>
</fingerprint>
</file_fingerprint>
    <client_user>compliance_admin</client_user>
    <client_ip_address>localhost</client_ip_address>
    <client_transport>ssh</client_transport>

```

```

</snaplock_pre_privileged_delete_1>
</LogEntry>
<LogEntry LogVersion="3" Date="20160420 085521 GMT" SCCdate="20160420 084716 GMT" VCCdate="20160420
084716 GMT" Node="vsim7" Vserver="vs1">
<snaplock_post_privileged_delete_1>
  <sequence_number>9bae8bba-06d5-11e6-b1b3-005056bb6335</sequence_number>
  <file_pathname>/vol/sle_vol1/file.1</file_pathname>
  <file_expires>infinite</file_expires>
  <record_expires>infinite</record_expires>
  <file_fingerprint><fingerprint>
  <fingerprint-info>
    <fingerprint-algorithm>SHA256</fingerprint-algorithm>
    <fingerprint-start-time>1461142518</fingerprint-start-time>
    <formatted-fingerprint-start-time>Wed Apr 20 08:55:18 GMT 2016</formatted-
fingerprint-start-time>
    <fingerprint-scope>DATA_AND_METADATA</fingerprint-scope>
    <fingerprint-version>3</fingerprint-version>
    <fingerprint-input-path>/vol/sle_vol1/file.1</fingerprint-input-path>
    <vserver-name>vs1</vserver-name>
    <vserver-uuid>37eaccal-06bb-11e6-9edf-005056bb6335</vserver-uuid>
    <volume-name>sle_vol1</volume-name>
    <volume-msid>2160232325</volume-msid>
    <volume-dsid>1028</volume-dsid>
    <volume-containing-aggregate>sle_aggr</volume-containing-aggregate>
    <aggregate-id>fb20fab-27fa-4e88-a468-1245c36e7d14</aggregate-id>
    <snaplock-license>available</snaplock-license>
    <snaplock-system-compliance-clock>1461142033</snaplock-system-compliance-clock>
    <formatted-snaplock-system-compliance-clock>Wed Apr 20 08:47:13 GMT
2016</formatted-snaplock-system-compliance-clock>
    <volume-snaplock-type>enterprise</volume-snaplock-type>
    <snaplock-volume-compliance-clock>1461142033</snaplock-volume-compliance-clock>
    <formatted-snaplock-volume-compliance-clock>Wed Apr 20 08:47:13 GMT
2016</formatted-snaplock-volume-compliance-clock>
    <formatted-volume-expiry-date>infinite</formatted-volume-expiry-date>
    <file-metadata>
      <file-fingerprint-info>

```

```

        <path>/vol/sle_vol1/file.1</path>
        <fsid>1028</fsid>
        <fileid>97</fileid>
        <file-type>worm</file-type>
        <file-size>1048576</file-size>
        <creation-time>1461142021</creation-time>
        <formatted-creation-time>Wed Apr 20 08:47:01 GMT 2016</formatted-
creation-time>

        <modified-time>1461142021</modified-time>
        <formatted-modified-time>Wed Apr 20 08:47:01 GMT 2016</formatted-
modified-time>

        <changed-time>1461141542</changed-time>
        <is-changed-time-wraparound>false</is-changed-time-wraparound>
        <formatted-changed-time>Wed Apr 20 08:39:02 GMT 2016</formatted-
changed-time>

        <formatted-retention-time>infinite</formatted-retention-time>
        <owner-id>0</owner-id>
        <group-id>0</group-id>
        <data-
fingerprint>MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfylg=</data-fingerprint>

        <metadata-
fingerprint>JgY8a6p+dAvTDSC1g19b4PXF3jqDpJUTXnex/g20A3Q=</metadata-fingerprint>

        </file-fingerprint-info>
    </file-metadata>
    <fingerprint-end-time>1461142518</fingerprint-end-time>
    <formatted-fingerprint-end-time>Wed Apr 20 08:55:18 GMT 2016</formatted-
fingerprint-end-time>
</fingerprint-info>
</fingerprint>
</file_fingerprint>

    <client_user>compliance_admin</client_user>
    <client_ip_address>localhost</client_ip_address>
    <client_transport>ssh</client_transport>
    <result>Success</result>
</snaplock_post_privileged_delete_1>

```

```
</LogEntry>
```

8 Contact Us

Do you have a question that you think we should include in this FAQ?

Send it to docfeedback@netapp.com and include SnapLock ONTAP 9 TECHNICAL FAQ in the subject line.

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 © 1994–2016 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp

Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-XXXX-XXXX