



Archive and compliance using SnapLock technology

ONTAP 9

NetApp

February 02, 2026

Table of Contents

Archive and compliance using SnapLock technology	1
Learn about ONTAP SnapLock	1
What you can do with SnapLock	1
SnapLock Compliance and Enterprise modes	2
Supported and unsupported features with SnapLock	3
FabricPool on SnapLock Enterprise aggregates	4
FlexGroup volumes	4
LUN support	4
MetroCluster support	4
Multi-admin verification (MAV) support	5
Storage efficiency	5
Encryption	5
7-Mode Transition	5
Configure SnapLock	6
Learn about configuring ONTAP SnapLock	6
Initialize the ONTAP Compliance Clock	6
Create an ONTAP SnapLock aggregate	8
Create and mount ONTAP SnapLock volumes	9
Set the ONTAP SnapLock retention time	12
Create an ONTAP SnapLock-protected audit log	17
Verify ONTAP SnapLock settings	18
Manage WORM files	20
Manage WORM files with ONTAP SnapLock	20
Commit files to WORM using ONTAP SnapLock	20
Commit snapshots to WORM on an ONTAP vault destination	24
Mirror WORM files with ONTAP SnapMirror for disaster recovery	28
Retain WORM files during litigation using ONTAP SnapLock Legal Hold	32
Delete WORM files with ONTAP SnapLock	33
Move an ONTAP SnapLock volume	34
Create a SnapLock security administrator account	35
Move a SnapLock volume	35
Lock an ONTAP snapshot for protection against ransomware attacks	36
Enable snapshot locking when creating a volume	38
Enable snapshot locking on an existing volume	38
Create a locked snapshot policy and apply retention	39

Archive and compliance using SnapLock technology

Learn about ONTAP SnapLock

SnapLock is a high-performance compliance solution for organizations that use WORM storage to retain files in unmodified form for regulatory and governance purposes.

SnapLock helps to prevent deletion, change, or renaming of data to meet regulations such as SEC 17a-4(f), HIPAA, FINRA, CFTC, and GDPR. With SnapLock, you can create special-purpose volumes in which files can be stored and committed to a non-erasable, non-writable state either for a designated retention period or indefinitely. SnapLock allows this retention to be performed at the file level through standard open file protocols such as CIFS and NFS. The supported open file protocols for SnapLock are NFS (versions 2, 3, and 4) and CIFS (SMB 1.0, 2.0, and 3.0).

Using SnapLock, you commit files and snapshots to WORM storage, and set retention periods for WORM-protected data. SnapLock WORM storage uses NetApp snapshot technology and can leverage SnapMirror replication, and SnapVault backups as the base technology for providing backup recovery protection for data. Learn more about WORM storage: [Compliant WORM storage using NetApp SnapLock - TR-4526](#).

You can use an application to commit files to WORM over NFS or CIFS, or use the SnapLock autocommit feature to commit files to WORM automatically. You can use a *WORM appendable file* to retain data that is written incrementally, like log information. For more information see [Use volume append mode to create WORM appendable files](#).

SnapLock supports data protection methods that should satisfy most compliance requirements:

- You can use SnapLock for SnapVault to WORM-protect snapshots on secondary storage. See [Commit snapshots to WORM](#).
- You can use SnapMirror to replicate WORM files to another geographic location for disaster recovery. See [Mirror WORM files](#).

SnapLock is a license-based feature of ONTAP. A single license entitles you to use SnapLock in strict Compliance mode, to satisfy external mandates like SEC Rule 17a-4(f), and a looser Enterprise mode, to meet internally mandated regulations for the protection of digital assets. SnapLock licenses are part of the [ONTAP One](#) software suite.

SnapLock is supported on all AFF and FAS systems as well as ONTAP Select. SnapLock is not a software-only solution; it is an integrated hardware and software solution. This distinction is important for strict WORM regulations such as SEC 17a-4(f), which requires an integrated hardware and software solution. For more information, refer to [SEC Guidance to Broker-Dealers on the Use of Electronic Storage Media](#).

What you can do with SnapLock

After you configure SnapLock, you can complete the following tasks:

- [Commit files to WORM](#)
- [Commit snapshots to WORM for secondary storage](#)
- [Mirror WORM files for disaster recovery](#)
- [Retain WORM files during litigation using Legal Hold](#)

- Delete WORM files using the privileged delete feature
- Set the file retention period
- Move a SnapLock volume
- Lock a snapshot for protection against ransomware attacks
- Review snapLock use with the Audit Log
- Use SnapLock APIs

SnapLock Compliance and Enterprise modes

SnapLock Compliance and Enterprise modes differ mainly in the level at which each mode protects WORM files:

SnapLock mode	Protection level	WORM file deleting during retention
Compliance mode	At the disk level	Cannot be deleted
Enterprise mode	At the file level	Can be deleted by the compliance administrator using an audited "privileged delete" procedure

After the retention period has elapsed, you are responsible for deleting any files you no longer need. Once a file has been committed to WORM, whether under Compliance or Enterprise mode, it cannot be modified, even after the retention period has expired.

You cannot move a WORM file during or after the retention period. You can copy a WORM file, but the copy will not retain its WORM characteristics.

The following table shows the differences in capabilities supported by SnapLock Compliance and Enterprise modes:

Capability	SnapLock Compliance	SnapLock Enterprise
Enable and delete files using privileged delete	No	Yes
Reinitialize disks	No	Yes
Destroy SnapLock aggregates and volumes during retention period	No	Yes, with the exception of the SnapLock audit log volume
Rename aggregates or volumes	No	Yes
Use non-NetApp disks	No	No
Use the SnapLock volume for audit logging	Yes	Yes, beginning with ONTAP 9.5

Supported and unsupported features with SnapLock

The following table shows the features that are supported with SnapLock Compliance mode, SnapLock Enterprise mode, or both:

Feature	Supported with SnapLock Compliance	Supported with SnapLock Enterprise
Consistency Groups	No	No
Encrypted volumes	Yes, learn more about Encryption and SnapLock .	Yes, learn more about Encryption and SnapLock .
FabricPools on SnapLock aggregates	No	Yes, beginning with ONTAP 9.8. Learn more about FabricPool on SnapLock Enterprise aggregates .
Flash Pool aggregates	Yes.	Yes.
FlexClone	You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume.	You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume.
FlexGroup volumes	Yes, beginning with ONTAP 9.11.1. Learn more about FlexGroup volumes .	Yes, beginning with ONTAP 9.11.1. Learn more about FlexGroup volumes .
LUNs	No. Learn more about LUN support with SnapLock .	No. Learn more about LUN support with SnapLock .
MetroCluster configurations	Yes, beginning with ONTAP 9.3. Learn more about MetroCluster support .	Yes, beginning with ONTAP 9.3. Learn more about MetroCluster support .
Multi-admin verification (MAV)	Yes, beginning with ONTAP 9.13.1. Learn more about MAV support .	Yes, beginning with ONTAP 9.13.1. Learn more about MAV support .
SAN	No	No
Single-file SnapRestore	No	Yes
SnapMirror active sync	No	No
SnapRestore	No	Yes
SMTape	No	No
SnapMirror Synchronous	No	No

SSDs	Yes.	Yes.
Storage efficiency features	Yes, beginning with ONTAP 9.9.1. Learn more about storage efficiency support .	Yes, beginning with ONTAP 9.9.1. Learn more about storage efficiency support .

FabricPool on SnapLock Enterprise aggregates

FabricPools are supported on SnapLock Enterprise aggregates beginning with ONTAP 9.8. However, your account team needs to open a product variance request documenting that you understand that FabricPool data tiered to a public or private cloud is no longer protected by SnapLock because a cloud admin can delete that data.



Any data that FabricPool tiers to a public or private cloud is no longer protected by SnapLock because that data can be deleted by a cloud administrator.

FlexGroup volumes

SnapLock supports FlexGroup volumes beginning with ONTAP 9.11.1; however, the following features are not supported:

- Legal-hold
- Event-based retention
- SnapLock for SnapVault (supported beginning with ONTAP 9.12.1)

You should also be aware of the following behaviors:

- The volume compliance clock (VCC) of a FlexGroup volume is determined by the VCC of the root constituent. All non-root constituents will have their VCC closely synced to the root VCC.
- SnapLock configuration properties are set only on the FlexGroup as a whole. Individual constituents cannot have different configuration properties, such as default retention time and autocommit period.

LUN support

LUNs are supported in SnapLock volumes only in scenarios where snapshots created on a non-SnapLock volume are transferred to a SnapLock volume for protection as part of SnapLock vault relationship. LUNs are not supported in read/write SnapLock volumes. Tamperproof snapshots however are supported on both SnapMirror source volumes and destination volumes that contain LUNs.

MetroCluster support

SnapLock support in MetroCluster configurations differs between SnapLock Compliance mode and SnapLock Enterprise mode.

SnapLock Compliance

- Beginning with ONTAP 9.3, SnapLock Compliance is supported on unmirrored MetroCluster aggregates.
- Beginning with ONTAP 9.3, SnapLock Compliance is supported on mirrored aggregates, but only if the aggregate is used to host SnapLock audit log volumes.
- SVM-specific SnapLock configurations can be replicated to primary and secondary sites using

MetroCluster.

SnapLock Enterprise

- SnapLock Enterprise aggregates are supported.
- Beginning with ONTAP 9.3, SnapLock Enterprise aggregates with privileged delete are supported.
- SVM-specific SnapLock configurations can be replicated to both sites using MetroCluster.

MetroCluster configurations and compliance clocks

MetroCluster configurations use two compliance clock mechanisms, the Volume Compliance Clock (VCC) and the System Compliance Clock (SCC). The VCC and SCC are available to all SnapLock configurations. When you create a new volume on a node, its VCC is initialized with the current value of the SCC on that node. After the volume is created, the volume and file retention time is always tracked with the VCC.

When a volume is replicated to another site, its VCC is also replicated. When a volume switchover occurs, from Site A to Site B, for example, the VCC continues to be updated on Site B while the SCC on Site A halts when Site A goes offline.

When Site A is brought back online and the volume switchback is performed, the Site A SCC clock restarts while the VCC of the volume continues to be updated. Because the VCC is continuously updated, regardless of switchover and switchback operations, the file retention times do not depend on SCC clocks and do not stretch.

Multi-admin verification (MAV) support

Beginning with ONTAP 9.13.1, a cluster administrator can explicitly enable multi-admin verification on a cluster to require quorum approval before some SnapLock operations are executed. When MAV is enabled, SnapLock volume properties such as default-retention-time, minimum-retention-time, maximum-retention-time, volume-append-mode, autocommit-period and privileged-delete will require quorum approval. Learn more about [MAV](#).

Storage efficiency

Beginning with ONTAP 9.9.1, SnapLock supports storage efficiency features, such as data compaction, cross-volume-deduplication, and adaptive compression for SnapLock volumes and aggregates. For more information about storage efficiency, see [ONTAP storage efficiency overview](#).

Encryption

ONTAP offers both software- and hardware-based encryption technologies for ensuring that data at rest cannot be read if the storage medium is repurposed, returned, misplaced, or stolen.

Disclaimer: NetApp cannot guarantee that SnapLock-protected WORM files on self-encrypting drives or volumes will be retrievable if the authentication key is lost or if the number of failed authentication attempts exceeds the specified limit and results in the drive being permanently locked. You are responsible for ensuring against authentication failures.



Encrypted volumes are supported on SnapLock aggregates.

7-Mode Transition

You can migrate SnapLock volumes from 7-Mode to ONTAP by using the Copy-Based Transition (CBT) feature of the 7-Mode Transition Tool. The SnapLock mode of the destination volume, Compliance or Enterprise, must match the SnapLock mode of the source volume. You cannot use Copy-Free Transition (CFT) to migrate

SnapLock volumes.

Configure SnapLock

Learn about configuring ONTAP SnapLock

Before you use SnapLock, you need to configure SnapLock by completing various tasks such as [install the SnapLock license](#) for each node that hosts an aggregate with a SnapLock volume, initialize the [Compliance Clock](#), create a SnapLock aggregate for clusters running ONTAP releases earlier than ONTAP 9.10.1, [create and mount a SnapLock volume](#), and more.

Initialize the ONTAP Compliance Clock

SnapLock uses the *volume Compliance Clock* to ensure against tampering that might alter the retention period for WORM files. You must first initialize the *system ComplianceClock* on each node that hosts a SnapLock aggregate.

Beginning with ONTAP 9.14.1, you can initialize or reinitialize the system Compliance Clock when there are no SnapLock volumes or no volumes with snapshot locking enabled. The ability to reinitialize enables system administrators to reset the system Compliance Clock in instances where it might have been incorrectly initialized or to correct clock drift on the system. In ONTAP 9.13.1 and earlier releases, once you initialize the Compliance Clock on a node, you cannot initialize it again.

Before you begin

To reinitialize the Compliance Clock:

- All nodes in the cluster must be in the healthy state.
- All volumes must be online.
- No volumes can be present in the recovery queue.
- No SnapLock volumes can be present.
- No volumes with snapshot locking enabled can be present.

General requirements for initializing the Compliance Clock:

- You must be a cluster administrator to perform this task.
- [The SnapLock license must be installed on the node.](#)

About this task

The time on the system Compliance Clock is inherited by the *volume Compliance Clock*, the latter of which controls the retention period for WORM files on the volume. The volume Compliance Clock is initialized automatically when you create a new SnapLock volume.

 The initial setting of the system Compliance Clock is based on the current hardware system clock. For that reason, you should verify that the system time and time zone are correct before initializing the system Compliance Clock on each node. Once you initialize the system Compliance Clock on a node, you cannot initialize it again when SnapLock volumes or volumes with locking enabled are present.

Steps

You can use the ONTAP CLI to initialize the Compliance Clock or, beginning with ONTAP 9.12.1, you can use System Manager to initialize the Compliance Clock.

System Manager

1. Navigate to **Cluster > Overview**.
2. In the **Nodes** section, click **Initialize SnapLock Compliance Clock**.
3. To display the **Compliance Clock** column and to verify that the Compliance Clock is initialized, in the **Cluster > Overview > Nodes** section, click **Show/Hide** and select **SnapLock Compliance Clock**.

CLI

1. Initialize the system Compliance Clock:

```
snaplock compliance-clock initialize -node node_name
```

The following command initializes the system Compliance Clock on node1:

```
cluster1::> snaplock compliance-clock initialize -node node1
```

Learn more about `snaplock compliance-clock initialize` in the [ONTAP command reference](#).

2. When prompted, confirm that the system clock is correct and that you want to initialize the Compliance Clock:

Warning: You are about to initialize the secure ComplianceClock of the node "node1" to the current value of the node's system clock. This procedure can be performed only once on a given node, so you should ensure that the system time is set correctly before proceeding.

The current node's system clock is: Mon Apr 25 06:04:10 GMT 2016

Do you want to continue? (y|n): y

3. Repeat this procedure for each node that hosts a SnapLock aggregate.

Enable Compliance Clock resynchronization for an NTP-configured system

You can enable the SnapLock Compliance Clock synchronization feature when an NTP server is configured.

Before you begin

- This feature is available only at the advanced privilege level.
- You must be a cluster administrator to perform this task.

- The SnapLock license must be installed on the node.
- This feature is available only for Cloud Volumes ONTAP, ONTAP Select, and VSIM platforms.

About this task

When the SnapLock secure clock daemon detects a skew beyond the threshold, ONTAP uses the system time to reset both the system and volume Compliance Clocks. A period of 24 hours is set as the skew threshold. This means that the system Compliance Clock is synchronized to the system clock only if the skew is more than a day old.

The SnapLock secure clock daemon detects a skew and changes the Compliance Clock to the system time. Any attempt at modifying the system time to force the Compliance Clock to synchronize to the system time fails, since the Compliance Clock synchronizes to the system time only if the system time is synchronized with the NTP time.

Steps

1. Enable the SnapLock Compliance Clock synchronization feature when an NTP server is configured:

```
snaplock compliance-clock ntp
```

The following command enables the system Compliance Clock synchronization feature:

```
cluster1::>*> snaplock compliance-clock ntp modify -is-sync-enabled true
```

Learn more about `snaplock compliance-clock ntp modify` in the [ONTAP command reference](#).

2. When prompted, confirm that the configured NTP servers are trusted and that the communications channel is secure to enable the feature:
3. Check that the feature is enabled:

```
snaplock compliance-clock ntp show
```

The following command checks that the system Compliance Clock synchronization feature is enabled:

```
cluster1::>*> snaplock compliance-clock ntp show
```

```
Enable clock sync to NTP system time: true
```

Learn more about `snaplock compliance-clock ntp show` in the [ONTAP command reference](#).

Create an ONTAP SnapLock aggregate

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.

Before you begin

- You must be a cluster administrator to perform this task.
- The SnapLock [license must be installed](#) on the node. This license is included in [ONTAP One](#).
- [The Compliance Clock on the node must be initialized](#).
- If you have partitioned the disks as “root”, “data1”, and “data2”, you must ensure that spare disks are available.

Upgrade considerations

When upgrading to ONTAP 9.10.1, existing SnapLock and non-SnapLock aggregates are upgraded to support the existence of both SnapLock and non-SnapLock volumes; however, the existing SnapLock volume attributes are not automatically updated. For example, data-compaction, cross-volume-dedupe, and cross-volume-background-dedupe fields remain unchanged. New SnapLock volumes created on existing aggregates have the same default values as non-SnapLock volumes, and the default values for new volumes and aggregates are platform dependent.

Revert considerations

If you need to revert to an ONTAP version earlier than 9.10.1, you must move all SnapLock Compliance, SnapLock Enterprise, and SnapLock volumes to their own SnapLock aggregates.

About this task

- You cannot create Compliance aggregates with the SyncMirror option.
- You can create mirrored Compliance aggregates in a MetroCluster configuration only if the aggregate is used to host SnapLock audit log volumes.



In a MetroCluster configuration, SnapLock Enterprise is supported on mirrored and unmirrored aggregates. SnapLock Compliance is supported only on unmirrored aggregates.

Steps

1. Create a SnapLock aggregate:

```
storage aggregate create -aggregate <aggregate_name> -node <node_name>
-diskcount <number_of_disks> -snaplock-type <compliance|enterprise>
```

The following command creates a SnapLock Compliance aggregate named aggr1 with three disks on node1:

```
cluster1::> storage aggregate create -aggregate aggr1 -node node1
-diskcount 3 -snaplock-type compliance
```

Learn more about `storage aggregate create` in the [ONTAP command reference](#).

Create and mount ONTAP SnapLock volumes

You must create a SnapLock volume for the files or snapshots that you want to commit to the WORM state. Beginning with ONTAP 9.10.1, any volume you create, regardless of

the aggregate type, is created by default as a non-SnapLock volume. You must use the `-snaplock-type` option to explicitly create a SnapLock volume by specifying either Compliance or Enterprise as the SnapLock type. By default, the SnapLock type is set to non-snaplock.

Before you begin

- The SnapLock aggregate must be online.
- You should [verify that a SnapLock license is installed](#). If a SnapLock license is not installed on the node, you must [install](#) it. This license is included with [ONTAP One](#). Prior to ONTAP One, the SnapLock license was included in the Security and Compliance bundle. The Security and Compliance bundle is no longer offered but is still valid. Although not currently required, existing customers can choose to [upgrade to ONTAP One](#).
- [The Compliance Clock on the node must be initialized](#).

About this task

With the proper SnapLock permissions, you can destroy or rename an Enterprise volume at any time. You cannot destroy a Compliance volume until the retention period has elapsed. You can never rename a Compliance volume.

You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume. The clone volume will be of the same SnapLock type as the parent volume.

 LUNs are not supported in SnapLock volumes. LUNs are supported in SnapLock volumes only in scenarios where snapshots created on a non-SnapLock volume are transferred to a SnapLock volume for protection as part of SnapLock vault relationship. LUNs are not supported in read/write SnapLock volumes. Tamperproof snapshots however are supported on both SnapMirror source volumes and destination volumes that contain LUNs.

Perform this task using ONTAP System Manager or the ONTAP CLI.

System Manager

Beginning with ONTAP 9.12.1, you can use System Manager to create a SnapLock volume.

Steps

1. Navigate to **Storage > Volumes** and click **Add**.
2. In the **Add Volume** window, click **More Options**.
3. Enter the new volume information, including the name and size of the volume.
4. Select **Enable SnapLock** and choose the SnapLock type, either Compliance or Enterprise.
5. In the **Auto-Commit Files** section, select **Modified** and enter the amount of time a file should remain unchanged before it is automatically committed. The minimum value is 5 minutes and the maximum value is 10 years.
6. In the **Data Retention** section, select the minimum and maximum retention period.
7. Select the default retention period.
8. Click **Save**.
9. Select the new volume in the **Volumes** page to verify the SnapLock settings.

CLI

1. Create a SnapLock volume:

```
volume create -vserver <SVM_name> -volume <volume_name> -aggregate
<aggregate_name> -snaplock-type <compliance|enterprise>
```

Learn more about `volume create` in the [ONTAP command reference](#). The following options are not available for SnapLock volumes: `-nvfail`, `-atime-update`, `-is-autobalance-eligible`, `-space-mgmt-try-first`, and `vmalign`.

The following command creates a SnapLock Compliance volume named `vol1` on `aggr1` on `vs1`:

```
cluster1::> volume create -vserver vs1 -volume vol1 -aggregate aggr1
-snaplock-type compliance
```

Mount a SnapLock volume

You can mount a SnapLock volume to a junction path in the SVM namespace for NAS client access.

Before you begin

The SnapLock volume must be online.

About this task

- You can mount a SnapLock volume only under the root of the SVM.
- You cannot mount a regular volume under a SnapLock volume.

Steps

1. Mount a SnapLock volume:

```
volume mount -vserver SVM_name -volume volume_name -junction-path path
```

Learn more about `volume mount` in the [ONTAP command reference](#).

The following command mounts a SnapLock volume named `vol1` to the junction path `/sales` in the `vs1` namespace:


```
cluster1::> volume mount -vserver vs1 -volume vol1 -junction-path /sales
```

Set the ONTAP SnapLock retention time

You can set the retention time for a file explicitly, or you can use the default retention period for the volume to derive the retention time. Unless you set the retention time explicitly, SnapLock uses the default retention period to calculate the retention time. You can also set file retention after an event.

About retention period and retention time

The *retention period* for a WORM file specifies the length of time the file must be retained after it is committed to the WORM state. The *retention time* for a WORM file is the time after which the file no longer needs to be retained. A retention period of 20 years for a file committed to the WORM state on 10 November 2020 6:00 a.m., for example, would yield a retention time of 10 November 2040 6:00 a.m.



i Beginning with ONTAP 9.10.1, you can set a retention time up to October 26, 3058 and a retention period up to 100 years. When you extend retention dates, older policies are converted automatically. In ONTAP 9.9.1 and earlier releases, unless you set the default retention period to infinite, the maximum supported retention time is January 19 2071 (GMT).

Important replication considerations

When establishing a SnapMirror relationship with a SnapLock source volume using a retention date later than January 19th 2071 (GMT), the destination cluster must be running ONTAP 9.10.1 or later or the SnapMirror transfer will fail.

Important revert considerations

ONTAP prevents you from reverting a cluster from ONTAP 9.10.1 to an earlier ONTAP version when there are any files with a retention period later than "January 19, 2071 8:44:07 AM".

Understanding the retention periods

A SnapLock Compliance or Enterprise volume has four retention periods:

- Minimum retention period (min), with a default of 0

- Maximum retention period (`max`), with a default of 30 years
- Default retention period, with a default equal to `min` for both Compliance mode and Enterprise mode beginning with ONTAP 9.10.1. In ONTAP releases earlier than ONTAP 9.10.1, the default retention period depends on the mode:
 - For Compliance mode, the default is equal to `max`.
 - For Enterprise mode, the default is equal to `min`.
- Unspecified retention period.



In releases prior to ONTAP 9.10.1, if you do not set the retention time explicitly before committing a Compliance-mode file to the WORM state, and you do not modify the defaults, the file will be retained for 30-years. This change *cannot* be undone. Similarly, in ONTAP 9.10.1 and later, if you do not set the retention time explicitly before committing an Enterprise-mode file to the WORM state, and you do not modify the defaults, the file will be retained for 0 years, or, effectively, not at all.

Beginning with ONTAP 9.8, you can set the retention period on files in a volume to `unspecified`, to enable the file to be retained until you set an absolute retention time. You can set a file with absolute retention time to unspecified retention and back to absolute retention as long as the new absolute retention time is later than the absolute time you previously set.

Beginning with ONTAP 9.12.1, WORM files with the retention period set to `unspecified` are guaranteed to have a retention period set to the minimum retention period configured for the SnapLock volume. When you change the file retention period from `unspecified` to an absolute retention time, the new retention time specified must be greater than the minimum retention time already set on the file.

Set the default retention period

You can use the `volume snaplock modify` command to set the default retention period for files on a SnapLock volume.

Before you begin

The SnapLock volume must be online.

About this task

The following table shows the possible values for the default retention period option:



The default retention period must be greater than or equal to (\geq) the minimum retention period and less than or equal to (\leq) the maximum retention period.

Value	Unit	Notes
0 - 65535	seconds	
0 - 24	hours	
0 - 365	days	
0 - 12	months	

Value	Unit	Notes
0 - 100	years	Beginning with ONTAP 9.10.1. For earlier ONTAP releases, the value is 0 - 70.
max	-	Use the maximum retention period.
min	-	Use the minimum retention period.
infinite	-	Retain the files forever.
unspecified	-	Retain the files until an absolute retention period is set.

The values and ranges for the maximum and minimum retention periods are identical, except for `max` and `min`, which are not applicable. For more information about this task, see [Set the retention time overview](#).

You can use the `volume snaplock show` command to view the retention period settings for the volume. Learn more about `volume snaplock show` in the [ONTAP command reference](#).



After a file has been committed to the WORM state, you can extend but not shorten the retention period.

Steps

1. Set the default retention period for files on a SnapLock volume:

```
volume snaplock modify -vserver SVM_name -volume volume_name -default
-retention-period default_retention_period -minimum-retention-period
min_retention_period -maximum-retention-period max_retention_period
```

Learn more about `volume snaplock modify` in the [ONTAP command reference](#).



The following examples assume that the minimum and maximum retention periods have not been modified previously.

The following command sets the default retention period for a Compliance or Enterprise volume to 20 days:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retention-period 20days
```

The following command sets the default retention period for a Compliance volume to 70 years:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -maximum
-retention-period 70years
```

The following command sets the default retention period for an Enterprise volume to 10 years:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retenction-period max -maximum-retention-period 10years
```

The following commands set the default retention period for an Enterprise volume to 10 days:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -minimum
-retenction-period 10days
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retenction-period min
```

The following command sets the default retention period for a Compliance volume to infinite:

```
cluster1::> volume snaplock modify -vserver vs1 -volume vol1 -default
-retenction-period infinite -maximum-retention-period infinite
```

Set the retention time for a file explicitly

You can set the retention time for a file explicitly by modifying its last access time. You can use any suitable command or program over NFS or CIFS to modify the last access time.

About this task

After a file has been committed to WORM, you can extend but not shorten the retention time. The retention time is stored in the `atime` field for the file.



You cannot explicitly set the retention time of a file to `infinite`. That value is only available when you use the default retention period to calculate the retention time.

Steps

1. Use a suitable command or program to modify the last access time for the file whose retention time you want to set.

In a UNIX shell, use the following command to set a retention time of 21 November 2020 6:00 a.m. on a file named `document.txt`:

```
touch -a -t 202011210600 document.txt
```



You can use any suitable command or program to modify the last access time in Windows.

Set the file retention period after an event

Beginning with ONTAP 9.3, you can define how long a file is retained after an event occurs by using the SnapLock *Event Based Retention (EBR)* feature.

Before you begin

- You must be a SnapLock administrator to perform this task.

Create a SnapLock administrator account

- You must have logged in on a secure connection (SSH, console, or ZAPI).

About this task

The *event retention policy* defines the retention period for the file after the event occurs. The policy can be applied to a single file or all the files in a directory.

- If a file is not a WORM file, it will be committed to the WORM state for the retention period defined in the policy.
- If a file is a WORM file or a WORM appendable file, its retention period will be extended by the retention period defined in the policy.

You can use a Compliance-mode or Enterprise-mode volume.



EBR policies cannot be applied to files under a Legal Hold.

For advanced usage, see [Compliant WORM Storage Using NetApp SnapLock](#).

Using EBR to extend the retention period of already existing WORM files

EBR is convenient when you want to extend the retention period of already existing WORM files. For example, it might be your firm's policy to retain employee W-4 records in unmodified form for three years after the employee changes a withholding election. Another company policy might require that W-4 records be retained for five years after the employee is terminated.

In this situation, you could create an EBR policy with a five-year retention period. After the employee is terminated (the "event"), you would apply the EBR policy to the employee's W-4 record, causing its retention period to be extended. That will usually be easier than extending the retention period manually, particularly when a large number of files is involved.

Steps

1. Create an EBR policy:

```
snaplock event-retention policy create -vserver SVM_name -name policy_name  
-retention-period retention_period
```

The following command creates the EBR policy `employee_exit` on `vs1` with a retention period of ten years:

```
cluster1::>snaplock event-retention policy create -vserver vs1 -name  
employee_exit -retention-period 10years
```

2. Apply an EBR policy:

```
snaplock event-retention apply -vserver SVM_name -name policy_name -volume  
volume_name -path path_name
```

The following command applies the EBR policy `employee_exit` on `vs1` to all the files in the directory `d1`:

```
cluster1::>snaplock event-retention apply -vserver vs1 -name
employee_exit -volume vol1 -path /d1
```

Related information

- [snaplock event-retention policy create](#)
- [snaplock event-retention apply](#)

Create an ONTAP SnapLock-protected audit log

If you are using ONTAP 9.9.1 or earlier, you must first create a SnapLock aggregate and then you must create a SnapLock-protected audit log before performing a privileged delete or SnapLock volume move. The audit log records the creation and deletion of SnapLock administrator accounts, modifications to the log volume, whether privileged delete is enabled, privileged delete operations, and SnapLock volume move operations.

Beginning with ONTAP 9.10.1, you no longer create a SnapLock aggregate. You must use the `-snaplock-type` option to [explicitly create a SnapLock volume](#) by specifying either Compliance or Enterprise as the SnapLock type.

Before you begin

If you are using ONTAP 9.9.1 or earlier, you must be a cluster administrator to create a SnapLock aggregate.

About this task

You cannot delete an audit log until the log file retention period has elapsed. You cannot modify an audit log even after the retention period has elapsed. This is true for both SnapLock Compliance and Enterprise modes.

 In ONTAP 9.4 and earlier, you cannot use a SnapLock Enterprise volume for audit logging. You must use a SnapLock Compliance volume. In ONTAP 9.5 and later, you can use either a SnapLock Enterprise volume or a SnapLock Compliance volume for audit logging. In all cases, the audit log volume must be mounted at the junction path `/snaplock_audit_log`. No other volume can use this junction path.

You can find the SnapLock audit logs in the `/snaplock_log` directory under the root of the audit log volume, in subdirectories named `privdel_log` (privileged delete operations) and `system_log` (everything else). Audit log file names contain the timestamp of the first logged operation, making it easy to search for records by the approximate time that operations were executed.

- You can use the `snaplock log file show` command to view the log files on the audit log volume.
- You can use the `snaplock log file archive` command to archive the current log file and create a new one, which is useful in cases where you need to record audit log information in a separate file.

Learn more about `snaplock log file show` and `snaplock log file archive` in the [ONTAP command reference](#).



A data protection volume cannot be used as a SnapLock audit log volume.

Steps

1. Create a SnapLock aggregate.

[Create a SnapLock aggregate](#)

2. On the SVM that you want to configure for audit logging, create a SnapLock volume.

[Create a SnapLock volume](#)

3. Configure the SVM for audit logging:

```
snaplock log create -vserver SVM_name -volume snaplock_volume_name -max-log-size size -retention-period default_retention_period
```



The minimum default retention period for audit log files is six months. If the retention period of an affected file is longer than the retention period of the audit log, the retention period of the log inherits the retention period of the file. So, if the retention period for a file deleted using privileged delete is 10 months, and the retention period of the audit log is 8 months, the retention period of the log is extended to 10 months. For more information about retention time and default retention period, see [Set the retention time](#).

The following command configures SVM1 for audit logging using the SnapLock volume logVol. The audit log has a maximum size of 20 GB and is retained for eight months.

```
SVM1::> snaplock log create -vserver SVM1 -volume logVol -max-log-size 20GB -retention-period 8months
```

Learn more about `snaplock log create` in the [ONTAP command reference](#).

4. On the SVM that you configured for audit logging, mount the SnapLock volume at the junction path `/snaplock_audit_log`.

[Mount a SnapLock volume](#)

Verify ONTAP SnapLock settings

You can use the `volume file fingerprint start` and `volume file fingerprint dump` commands to view key information about files and volumes, including the file type (regular, WORM, or WORM appendable), the volume expiration date, and so forth.

Steps

1. Generate a file fingerprint:

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

```
svm1::> volume file fingerprint start -vserver svm1 -file  
/vol/sle/vol/f1  
File fingerprint operation is queued. Run "volume file fingerprint show  
-session-id 16842791" to view the fingerprint session status.
```

The command generates a session ID that you can use as input to the `volume file fingerprint dump` command.



You can use the `volume file fingerprint show` command with the session ID to monitor the progress of the fingerprint operation. Make sure that the operation has completed before attempting to display the fingerprint.

2. Display the fingerprint for the file:

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

```
svm1::> volume file fingerprint dump -session-id 33619976  
Vserver:svm1  
Session-ID:33619976  
Volume:slc_vol  
Path:/vol/slcs_vol/f1  
Data  
Fingerprint:MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=Metadata  
Fingerprint:8imjqJXiNcqgXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=Fingerprint  
Algorithm:SHA256  
Fingerprint Scope:data-and-metadata  
Fingerprint Start Time:1460612586  
Formatted Fingerprint Start Time:Thu Apr 14 05:43:06 GMT 2016  
Fingerprint Version:3  
**SnapLock License:available**  
Vserver UUID:acf7ae64-00d6-11e6-a027-0050569c55ae  
Volume MSID:2152884007  
Volume DSID:1028  
Hostname:my_host  
Filer ID:5f18eda2-00b0-11e6-914e-6fb45e537b8d  
Volume Containing Aggregate:slc_aggr1  
Aggregate ID:c84634aa-c757-4b98-8f07-eefe32565f67  
**SnapLock System ComplianceClock:1460610635  
Formatted SnapLock System ComplianceClock:Thu Apr 14 05:10:35  
GMT 2016  
Volume SnapLock Type:compliance  
Volume ComplianceClock:1460610635  
Formatted Volume ComplianceClock:Thu Apr 14 05:10:35 GMT 2016  
Volume Expiry Date:1465880998**
```

```
Is Volume Expiry Date Wraparound:false
Formatted Volume Expiry Date:Tue Jun 14 05:09:58 GMT 2016
Filesystem ID:1028
File ID:96
File Type:worm
File Size:1048576
Creation Time:1460612515
Formatted Creation Time:Thu Apr 14 05:41:55 GMT 2016
Modification Time:1460612515
Formatted Modification Time:Thu Apr 14 05:41:55 GMT 2016
Changed Time:1460610598
Is Changed Time Wraparound:false
Formatted Changed Time:Thu Apr 14 05:09:58 GMT 2016
Retention Time:1465880998
Is Retention Time Wraparound:false
Formatted Retention Time:Tue Jun 14 05:09:58 GMT 2016
Access Time:-
Formatted Access Time:-
Owner ID:0
Group ID:0
Owner SID:-
Fingerprint End Time:1460612586
Formatted Fingerprint End Time:Thu Apr 14 05:43:06 GMT 2016
```

Manage WORM files

Manage WORM files with ONTAP SnapLock

You can manage WORM files in the following ways:

- Commit files to WORM
- Commit snapshots to WORM on a vault destination
- Mirror WORM files for disaster recovery
- Retain WORM files during litigation
- Delete WORM files

Commit files to WORM using ONTAP SnapLock

You can commit files to WORM (write once, read many) either manually or by committing them automatically. You can also create WORM appendable files.

Commit files to WORM manually

You commit a file to WORM manually by making the file read-only. You can use any suitable command or program over NFS or CIFS to change the read-write attribute of a file to read-only. You might choose to

manually commit files if you want to ensure an application has finished writing to a file so that the file isn't committed prematurely or if there are scaling issues for the autocommit scanner because of a high number of volumes.

Before you begin

- The file you want to commit must reside on a SnapLock volume.
- The file must be writable.

About this task

The volume ComplianceClock time is written to the `ctime` field of the file when the command or program is executed. The ComplianceClock time determines when the retention time for the file has been reached.

Steps

1. Use a suitable command or program to change the read-write attribute of a file to read-only.

In a UNIX shell, use the following command to make a file named `document.txt` read-only:

```
chmod -w document.txt
```

In a Windows shell, use the following command to make a file named `document.txt` read-only:

```
attrib +r document.txt
```

Commit files to WORM automatically

The SnapLock autocommit feature enables you to commit files to WORM automatically. The autocommit feature commits a file to WORM state on a SnapLock volume if the file did not change for the autocommit-period duration. The autocommit feature is disabled by default.

Before you begin

- The files you want to autocommit must reside on a SnapLock volume.
- The SnapLock volume must be online.
- The SnapLock volume must be a read-write volume.

 The SnapLock autocommit feature scans through all of the files in the volume and commits a file if it meets the autocommit requirement. There might be a time interval between when the file is ready for autocommit and when it is actually committed by the SnapLock autocommit scanner. However, the file is still protected from modifications and deletion by the file system as soon as it is eligible for autocommit.

About this task

The *autocommit period* specifies the amount of time that files must remain unchanged before they are autocommitted. Changing a file before the autocommit period has elapsed restarts the autocommit period for the file.

The following table shows the possible values for the autocommit period:

Value	Unit	Notes
none	-	The default.
5 - 5256000	minutes	-
1 - 87600	hours	-
1 - 3650	days	-
1 - 120	months	-
1 - 10	years	-



The minimum value is 5 minutes and the maximum value is 10 years.

Steps

- Autocommit files on a SnapLock volume to WORM:

```
volume snaplock modify -vserver SVM_name -volume volume_name -autocommit
-period autocommit_period
```

Learn more about `volume snaplock modify` in the [ONTAP command reference](#).

The following command autocommits the files on volume `vol1` of SVM `vs1`, as long as the files remain unchanged for 5 hours:

```
cluster1::>volume snaplock modify -vserver vs1 -volume vol1 -autocommit
-period 5hours
```

Create a WORM appendable file

A WORM appendable file retains data written incrementally, like log entries. You can use any suitable command or program to create a WORM appendable file, or you can use the SnapLock *volume append mode* feature to create WORM appendable files by default.

Use a command or program to create a WORM appendable file

You can use any suitable command or program over NFS or CIFS to create a WORM appendable file. A WORM appendable file retains data written incrementally, like log entries. Data is appended to the file in 256 KB chunks. As each chunk is written, the previous chunk becomes WORM-protected. You cannot delete the file until the retention period has elapsed.

Before you begin

The WORM appendable file must reside on a SnapLock volume.

About this task

Data does not have to be written sequentially to the active 256 KB chunk. When data is written to byte

$n \times 256\text{KB} + 1$ of the file, the previous 256 KB segment becomes WORM-protected.

Any unordered writes beyond the current active 256 KB chunk will result in the active 256KB chunk being reset to the latest offset and will cause writes to older offsets to fail with a 'Read Only File System (ROFS)' error. The write offsets are dependent on the client application. A client that does not conform to the WORM append file write semantics can cause incorrect termination of the write contents. Therefore, it is recommended to either ensure that the client follows the offset restrictions for unordered writes, or to ensure synchronous writes by mounting the file system in synchronous mode.

Steps

1. Use a suitable command or program to create a zero-length file with the desired retention time.

In a UNIX shell, use the following command to set a retention time of 21 November 2020 6:00 a.m. on a zero-length file named `document.txt`:

```
touch -a -t 202011210600 document.txt
```

2. Use a suitable command or program to change the read-write attribute of the file to read-only.

In a UNIX shell, use the following command to make a file named `document.txt` read-only:

```
chmod 444 document.txt
```

3. Use a suitable command or program to change the read-write attribute of the file back to writable.



This step is not deemed a compliance risk because there is no data in the file.

In a UNIX shell, use the following command to make a file named `document.txt` writable:

```
chmod 777 document.txt
```

4. Use a suitable command or program to start writing data to the file.

In a UNIX shell, use the following command to write data to `document.txt`:

```
echo test data >> document.txt
```



Change the file permissions back to read-only when you no longer need to append data to the file.

Use volume append mode to create WORM appendable files

Beginning with ONTAP 9.3, you can use the SnapLock *volume append mode* (VAM) feature to create WORM appendable files by default. A WORM appendable file retains data written incrementally, like log entries. Data is appended to the file in 256 KB chunks. As each chunk is written, the previous chunk becomes WORM-protected. You cannot delete the file until the retention period has elapsed.

Before you begin

- The WORM appendable file must reside on a SnapLock volume.
- The SnapLock volume must be unmounted and empty of snapshots and user-created files.

About this task

Data does not have to be written sequentially to the active 256 KB chunk. When data is written to byte $n \times 256\text{KB} + 1$ of the file, the previous 256 KB segment becomes WORM-protected.

If you specify an autocommit period for the volume, WORM appendable files that are not modified for a period greater than the autocommit period are committed to WORM.



VAM is not supported on SnapLock audit log volumes.

Steps

1. Enable VAM:

```
volume snaplock modify -vserver SVM_name -volume volume_name -is-volume-append-mode-enabled true|false
```

Learn more about `volume snaplock modify` in the [ONTAP command reference](#).

The following command enables VAM on volume `vol1` of `SVMvs1`:

```
cluster1::>volume snaplock modify -vserver vs1 -volume vol1 -is-volume-append-mode-enabled true
```

2. Use a suitable command or program to create files with write permissions.

The files are WORM-appendable by default.

Commit snapshots to WORM on an ONTAP vault destination

You can use SnapLock for SnapVault to WORM-protect snapshots on secondary storage. You perform all of the basic SnapLock tasks on the vault destination. The destination volume is automatically mounted read-only, so there is no need to explicitly commit the snapshots to WORM.

Before you begin

- If you want to use System Manager to configure the relationship, both the source and the destination clusters must be running ONTAP 9.15.1 or later.
- On the destination cluster:
 - [Install the SnapLock license](#).
 - [Initialize the Compliance Clock](#).
 - If you are using the CLI with an ONTAP release earlier than 9.10.1, [create a SnapLock aggregate](#).
- The protection policy must be of type "vault".
- The source and destination aggregates must be 64-bit.

- The source volume cannot be a SnapLock volume.
- If you are using the ONTAP CLI, the source and destination volumes must be created in [peered clusters and SVMs](#).

About this task

The source volume can use NetApp or non-NetApp storage.



You cannot rename a snapshot that is committed to the WORM state.

You can clone SnapLock volumes, but you cannot clone files on a SnapLock volume.



LUNs are not supported in SnapLock volumes. LUNs are supported in SnapLock volumes only in scenarios where snapshots created on a non-SnapLock volume are transferred to a SnapLock volume for protection as part of SnapLock vault relationship. LUNs are not supported in read/write SnapLock volumes. Tamperproof snapshots, however, are supported on both SnapMirror source volumes and destination volumes that contain LUNs.

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 use the volume '-snaplock-type' option to specify a Compliance or Enterprise SnapLock volume type. In ONTAP releases earlier than ONTAP 9.10.1, the SnapLock mode, Compliance or Enterprise, is inherited from the aggregate. Version-flexible destination volumes are not supported. The language setting of the destination volume must match the language setting of the source volume.

A SnapLock volume that is a vault destination has a default retention period assigned to it. The value for this period is initially set to a minimum of 0 years for SnapLock Enterprise volumes and a maximum of 30 years for SnapLock Compliance volumes. Each NetApp snapshot is committed with this default retention period at first. The retention period can be extended later, if needed. For more information, see [Set retention time overview](#).

Beginning with ONTAP 9.14.1, you can specify retention periods for specific SnapMirror labels in the SnapMirror policy of the SnapMirror relationship so that the replicated snapshots from the source to the destination volume are retained for the retention-period specified in the rule. If no retention period is specified, the default-retention-period of the destination volume is used.

Beginning with ONTAP 9.13.1, you can instantaneously restore a locked snapshot on the destination SnapLock volume of a SnapLock vault relationship by creating a FlexClone with the `snaplock-type` option set to `non-snaplock` and specifying the snapshot as the "parent-snapshot" when executing the volume clone creation operation. Learn more about [creating a FlexClone volume with a SnapLock type](#).

For MetroCluster configurations, you should be aware of the following:

- You can create a SnapVault relationship only between sync-source SVMs, not between a sync-source SVM and a sync-destination SVM.
- You can create a SnapVault relationship from a volume on a sync-source SVM to a data-serving SVM.
- You can create a SnapVault relationship from a volume on a data-serving SVM to a DP volume on a sync-source SVM.

The following illustration shows the procedure for initializing a SnapLock vault relationship:

Steps

You can use the ONTAP CLI to create a SnapLock vault relationship or, beginning with ONTAP 9.15.1, you can use System Manager to create a SnapLock vault relationship.

System Manager

1. If the volume doesn't already exist, on the source cluster, navigate to **Storage > Volumes** and select **Add**.
2. In the **Add Volume** window, choose **More Options**.
3. Enter the volume name, size, export policy and share name.
4. Save your changes.
5. On the destination cluster, navigate to **Protection > Relationships**.
6. Above the **Source** column, select **Protect** and choose **Volumes** from the menu.
7. In the **Protect volumes** window, choose **Vault** as the protection policy.
8. In the **Source** section, select the cluster, storage VM, and volume you want to protect.
9. In the **Destination** section, under **Configuration details**, select **Lock destination snapshots**, and then choose **SnapLock for SnapVault** as the locking method. **Locking method** is not displayed if the policy type selected is not of type `vault`, if the SnapLock license is not installed, or if the Compliance Clock is not initialized.
10. If it is not already enabled, select **Initialize SnapLock Compliance Clock**.
11. Save your changes.

CLI

1. On the destination cluster, create a SnapLock destination volume of type `DP` that is either the same or greater in size than the source volume:

```
volume create -vserver <SVM_name> -volume <volume_name> -aggregate
<aggregate_name> -snaplock-type <compliance|enterprise> -type DP
-size <size>
```

The following command creates a 2GB SnapLock Compliance volume named `dstvolB` in `SVM2` on the aggregate `node01_aggr`:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate
node01_aggr -snaplock-type compliance -type DP -size 2GB
```

2. On the destination cluster, [set the default retention period](#).
3. [Create a new replication relationship](#) between the non-SnapLock source and the new SnapLock destination you created.

This example creates a new SnapMirror relationship with destination SnapLock volume `dstvolB` using a policy of `XDPDefault` to vault snapshots labeled daily and weekly on an hourly schedule:

```
cluster2::> snapmirror create -source-path SVM1:srcvolA -destination
-path SVM2:dstvolB -vserver SVM2 -policy XDPDefault -schedule hourly
```



Create a custom replication policy or a [custom schedule](#) if the available defaults are not suitable.

4. On the destination SVM, initialize the SnapVault relationship created:

```
snapmirror initialize -destination-path <destination_path>
```

The following command initializes the relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2:

```
cluster2::> snapmirror initialize -destination-path SVM2:dstvolB
```

5. After the relationship is initialized and idle, use the `snapshot show` command on the destination to verify the SnapLock expiry time applied to the replicated snapshots.

This example lists the snapshots on volume `dstvolB` that have the SnapMirror label and the SnapLock expiration date:

```
cluster2::> snapshot show -vserver SVM2 -volume dstvolB -fields snapmirror-label, snaplock-expiry-time
```

Related information

- [Cluster and SVM peering](#)
- [Volume backup using SnapVault](#)
- [snapmirror initialize](#)

Mirror WORM files with ONTAP SnapMirror for disaster recovery

You can use SnapMirror to replicate WORM files to another geographic location for disaster recovery and other purposes. Both the source volume and destination volume must be configured for SnapLock, and both volumes must have the same SnapLock mode, Compliance or Enterprise. All key SnapLock properties of the volume and files are replicated.

Prerequisites

The source and destination volumes must be created in peered clusters with peered SVMs. For more information, see [Cluster and SVM peering](#).

About this task

- Beginning with ONTAP 9.5, you can replicate WORM files with the XDP (extended data protection) type SnapMirror relationship rather than the DP (data protection) type relationship. XDP mode is ONTAP version-independent, and is able to differentiate files stored in the same block, making it much easier to resync replicated Compliance-mode volumes. For information on how to convert an existing DP-type relationship to an XDP-type relationship, see [Data Protection](#).

- A resync operation on a DP type SnapMirror relationship fails for a Compliance-mode volume if SnapLock determines that it will result in a loss of data. If a resync operation fails, you can use the `volume clone` command to make a clone of the destination volume. You can then resync the source volume with the clone.
- A SnapMirror relationship of a SnapLock volume only supports the `MirrorAllSnapshots` policy of type `async-mirror`. The retention period of a SnapLock volume is determined by the maximum retention period among all the WORM files that it holds. Because the destination is a DR copy of the source, the retention period of the destination SnapLock volume will be same as the source.
- A SnapMirror relationship of type XDP between SnapLock compliant volumes supports a resync after a break even if data on the destination has diverged from the source post the break.

On a resync, when data divergence is detected between the source and the destination beyond the common snapshot, a new snapshot is cut on the destination to capture this divergence. The new snapshot and the common snapshot are both locked with a retention time as follows:

- The volume expiry time of the destination
- If the volume expiry time is in the past or has not been set, then the snapshot is locked for a period of 30 days
- If the destination has legal-holds, the actual volume expiry period is masked and shows up as ‘indefinite’; however, the snapshot is locked for the duration of the actual volume expiry period.

If the destination volume has an expiry period that is later than the source, the destination expiry period is retained and will not be overwritten by the expiry period of the source volume post the resync.

If the destination has legal-holds placed on it that differ from the source, a resync is not allowed. The source and destination must have identical legal-holds or all legal-holds on the destination must be released before a resync is attempted.

A locked snapshot on the destination volume created to capture the divergent data can be copied to the source using the CLI by running the `snapmirror update -s snapshot` command. The snapshot once copied will continue to be locked at the source as well.

- SVM data protection relationships are not supported.
- Load-sharing data protection relationships are not supported.

The following illustration shows the procedure for initializing a SnapMirror relationship:

System Manager

Beginning with ONTAP 9.12.1, you can use System Manager to set up SnapMirror replication of WORM files.

Steps

1. Navigate to **Storage > Volumes**.
2. Click **Show/Hide** and select **SnapLock Type** to display the column in the **Volumes** window.
3. Locate a SnapLock volume.
4. Click  and select **Protect**.
5. Choose the destination cluster and the destination storage VM.
6. Click **More Options**.
7. Select **Show legacy policies** and select **DPDefault (legacy)**.
8. In the **Destination Configuration details** section, select **Override transfer schedule** and select **hourly**.
9. Click **Save**.
10. To the left of the source volume name, click the arrow to expand the volume details, and on the right side of the page, review the remote SnapMirror protection details.
11. On the remote cluster, navigate to **Protection Relationships**.
12. Locate the relationship and click the destination volume name to view the relationship details.
13. Verify that the destination volume SnapLock type and other SnapLock information.

CLI

1. Identify the destination cluster.
2. On the destination cluster, [install the SnapLock license](#), [initialize the Compliance Clock](#), and, if you are using an ONTAP release earlier than 9.10.1, [create a SnapLock aggregate](#).
3. On the destination cluster, create a SnapLock destination volume of type **DP** that is either the same size as or greater in size than the source volume:

```
volume create -vserver SVM_name -volume volume_name -aggregate
aggregate_name -snaplock-type compliance|enterprise -type DP -size size
```



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 use the volume -snaplock-type option to specify a Compliance or Enterprise SnapLock volume type. In ONTAP releases earlier than ONTAP 9.10.1, the SnapLock mode—Compliance or Enterprise—is inherited from the aggregate. The language setting of the destination volume must match the language setting of the source volume.

The following command creates a 2 GB SnapLock Compliance volume named `dstvolB` in `SVM2` on the aggregate `node01_aggr`:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate node01_aggr -snaplock-type compliance -type DP -size 2GB
```

4. On the destination SVM, create a SnapMirror policy:

```
snapmirror policy create -vserver SVM_name -policy policy_name
```

The following command creates the SVM-wide policy SVM1-mirror:

```
SVM2::> snapmirror policy create -vserver SVM2 -policy SVM1-mirror
```

5. On the destination SVM, create a SnapMirror schedule:

```
job schedule cron create -name schedule_name -dayofweek day_of_week -hour hour -minute minute
```

The following command creates a SnapMirror schedule named weekendcron:

```
SVM2::> job schedule cron create -name weekendcron -dayofweek "Saturday, Sunday" -hour 3 -minute 0
```

6. On the destination SVM, create a SnapMirror relationship:

```
snapmirror create -source-path source_path -destination-path destination_path -type XDP|DP -policy policy_name -schedule schedule_name
```

The following command creates a SnapMirror relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2, and assigns the policy `SVM1-mirror` and the schedule `weekendcron`:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path SVM2:dstvolB -type XDP -policy SVM1-mirror -schedule weekendcron
```



The XDP type is available in ONTAP 9.5 and later. You must use the DP type in ONTAP 9.4 and earlier.

7. On the destination SVM, initialize the SnapMirror relationship:

```
snapmirror initialize -destination-path destination_path
```

The initialization process performs a *baseline transfer* to the destination volume. SnapMirror makes a snapshot of the source volume, then transfers the copy and all the data blocks that it references to the destination volume. It also transfers any other snapshots on the source volume to the destination volume.

The following command initializes the relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

Related information

- [Cluster and SVM peering](#)
- [Volume disaster recovery preparation](#)
- [Data protection](#)
- [snapmirror create](#)
- [snapmirror initialize](#)
- [snapmirror policy create](#)

Retain WORM files during litigation using ONTAP SnapLock Legal Hold

Beginning with ONTAP 9.3, you can retain Compliance-mode WORM files for the duration of a litigation by using the *Legal Hold* feature.

Before you begin

- You must be a SnapLock administrator to perform this task.

Create a SnapLock administrator account

- You must have logged in on a secure connection (SSH, console, or ZAPI).

About this task

A file under a Legal Hold behaves like a WORM file with an indefinite retention period. It is your responsibility to specify when the Legal Hold period ends.

The number of files you can place under a Legal Hold depends on the space available on the volume.

Steps

1. Start a Legal Hold:

```
snaplock legal-hold begin -litigation-name <litigation_name> -volume <volume_name> -path <path_name>
```

The following command starts a Legal Hold for all the files in `vol1`:

```
cluster1::>snaplock legal-hold begin -litigation-name litigation1 -volume vol1 -path /
```

2. End a Legal Hold:

```
snaplock legal-hold end -litigation-name <litigation_name> -volume
```

```
<volume_name> -path <path_name>
```

The following command ends a Legal Hold for all the files in vol1:

```
cluster1::>snaplock legal-hold end -litigation-name litigation1 -volume
vol1 -path /
```

Related information

- [snaplock legal-hold begin](#)
- [snaplock legal-hold end](#)

Delete WORM files with ONTAP SnapLock

You can delete Enterprise-mode WORM files during the retention period using the privileged delete feature. Before you can use this feature, you must create a SnapLock administrator account and then using the account, enable the feature.

Create a SnapLock administrator account

You must have SnapLock administrator privileges to perform a privileged delete. These privileges are defined in the vsadmin-snaplock role. If you have not already been assigned that role, you can ask your cluster administrator to create an SVM administrator account with the SnapLock administrator role.

Before you begin

- You must be a cluster administrator to perform this task.
- You must have logged in on a secure connection (SSH, console, or ZAPI).

Steps

1. Create an SVM administrator account with the SnapLock administrator role:

```
security login create -vserver SVM_name -user-or-group-name user_or_group_name
-application application -authmethod authentication_method -role role -comment
comment
```

The following command enables the SVM administrator account SnapLockAdmin with the predefined vsadmin-snaplock role to access SVM1 using a password:

```
cluster1::> security login create -vserver SVM1 -user-or-group-name
SnapLockAdmin -application ssh -authmethod password -role vsadmin-
snaplock
```

Learn more about `security login create` in the [ONTAP command reference](#).

Enable the privileged delete feature

You must explicitly enable the privileged delete feature on the Enterprise volume that contains the WORM files

you want to delete.

About this task

The value of the `-privileged-delete` option determines whether privileged delete is enabled. Possible values are `enabled`, `disabled`, and `permanently-disabled`.



`permanently-disabled` is the terminal state. You cannot enable privileged delete on the volume after you set the state to `permanently-disabled`.

Steps

1. Enable privileged delete for a SnapLock Enterprise volume:

```
volume snaplock modify -vserver SVM_name -volume volume_name -privileged  
-delete disabled|enabled|permanently-disabled
```

The following command enables the privileged delete feature for the Enterprise volume `dataVol` on `SVM1`:

```
SVM1::> volume snaplock modify -vserver SVM1 -volume dataVol -privileged  
-delete enabled
```

Delete Enterprise-mode WORM files

You can use the privileged delete feature to delete Enterprise-mode WORM files during the retention period.

Before you begin

- You must be a SnapLock administrator to perform this task.
- You must have created a SnapLock audit log and enabled the privileged delete feature on the Enterprise volume.

About this task

You cannot use a privileged delete operation to delete an expired WORM file. You can use the `volume file retention show` command to view the retention time of the WORM file that you want to delete. Learn more about `volume file retention show` in the [ONTAP command reference](#).

Step

1. Delete a WORM file on an Enterprise volume:

```
volume file privileged-delete -vserver SVM_name -file file_path
```

The following command deletes the file `/vol/dataVol/f1` on the SVM `SVM1`:

```
SVM1::> volume file privileged-delete -file /vol/dataVol/f1
```

Move an ONTAP SnapLock volume

Beginning with ONTAP 9.8, you can move a SnapLock volume to a destination aggregate

of the same type, either Enterprise to Enterprise, or Compliance to Compliance. You must be assigned the SnapLock security role to move a SnapLock volume.

Create a SnapLock security administrator account

You must have SnapLock security administrator privileges to perform a SnapLock volume move. This privilege is granted to you with the *snaplock* role, introduced in ONTAP 9.8. If you have not already been assigned that role, you can ask your cluster administrator to create a SnapLock security user with this SnapLock security role.

Before you begin

- You must be a cluster administrator to perform this task.
- You must have logged in on a secure connection (SSH, console, or ZAPI).

About this task

The *snaplock* role is associated with the admin SVM, unlike the *vsadmin-snaplock* role, which is associated with the data SVM.

Step

1. Create an SVM administrator account with the SnapLock administrator role:

```
security login create -vserver SVM_name -user-or-group-name user_or_group_name
-application application -authmethod authentication_method -role role -comment
comment
```

The following command enables the SVM administrator account *SnapLockAdmin* with the predefined *snaplock* role to access admin SVM *cluster1* using a password:

```
cluster1::> security login create -vserver cluster1 -user-or-group-name
SnapLockAdmin -application ssh -authmethod password -role snaplock
```

Learn more about `security login create` in the [ONTAP command reference](#).

Move a SnapLock volume

You can use the `volume move` command to move a SnapLock volume to a destination aggregate.

Before you begin

- You must have created a SnapLock-protected audit log before performing SnapLock volume move.

[Create an audit log](#).

- If you are using a version of ONTAP earlier than ONTAP 9.10.1, the destination aggregate must be the same SnapLock type as the SnapLock volume you want to move; either Compliance to Compliance or Enterprise to Enterprise. Beginning with ONTAP 9.10.1, this restriction is removed and an aggregate can include both Compliance and Enterprise SnapLock volumes, as well as non-SnapLock volumes.
- You must be a user with the SnapLock security role.

Steps

1. Using a secure connection, log in to the ONTAP cluster management LIF:

```
ssh snaplock_user@cluster_mgmt_ip
```

2. Move a SnapLock volume:

```
volume move start -vserver SVM_name -volume SnapLock_volume_name -destination -aggregate destination_aggregate_name
```

3. Check the status of the volume move operation:

```
volume move show -volume SnapLock_volume_name -vserver SVM_name -fields volume,phase,vserver
```

Lock an ONTAP snapshot for protection against ransomware attacks

Beginning with ONTAP 9.12.1, you can lock a snapshot on a non-SnapLock volume to provide protection from ransomware attacks. Locking snapshots ensures that they can't be deleted accidentally or maliciously.

You use the SnapLock compliance clock feature to lock snapshots for a specified period so that they cannot be deleted until the expiration time is reached. Locking snapshots makes them tamperproof, protecting them from ransomware threats. You can use locked snapshots to recover data if a volume is compromised by a ransomware attack.

Beginning with ONTAP 9.14.1, snapshot locking supports long-term retention snapshots on SnapLock vault destinations and on non-SnapLock SnapMirror destination volumes. Snapshot locking is enabled by setting the retention period using SnapMirror policy rules associated with an [existing policy label](#). The rule overrides the default retention period set on the volume. If there is no retention period associated with the SnapMirror label, the default retention period of the volume is used.

Tamperproof snapshot requirements and considerations

- If you are using the ONTAP CLI, all nodes in the cluster must be running ONTAP 9.12.1 or later. If you are using System Manager, all nodes must be running ONTAP 9.13.1 or later.
- [The SnapLock license must be installed on the cluster](#). This license is included in [ONTAP One](#).
- [The compliance clock on the cluster must be initialized](#).
- When snapshot locking is enabled on a volume, you can upgrade the clusters to a version of ONTAP later than ONTAP 9.12.1; however, you cannot revert to an earlier version of ONTAP until all locked snapshots have reached their expiration date and are deleted and snapshot locking is disabled.
- When a snapshot is locked, the volume expiry time is set to the expiry time of the snapshot. If more than one snapshot is locked, the volume expiry time reflects the largest expiry time among all snapshots.
- The retention period for locked snapshots takes precedence over the snapshot keep count, which means the keep count limit is not honored if the snapshot retention period for locked snapshots has not expired.
- In a SnapMirror relationship, you can set a retention period on a mirror-vault policy rule, and the retention period is applied for snapshots replicated to the destination if the destination volume has snapshot locking enabled. The retention period takes precedence over keep count; for example, snapshots that have not passed their expiry will be retained even if the keep count is exceeded.

- You can rename a snapshot on a non-SnapLock volume. Snapshot rename operations on the primary volume of a SnapMirror relationship are reflected on the secondary volume only if the policy is MirrorAllSnapshots. For other policy types, the renamed snapshot is not propagated during updates.
- If you are using the ONTAP CLI, you can restore a locked snapshot with the `volume snapshot restore` command only if the locked snapshot is the most recent. If there are any unexpired snapshots later than the one being restored, the snapshot restore operation fails.

Features supported with tamperproof snapshots

- [Cloud Volumes ONTAP](#)
- FlexGroup volumes

Snapshot locking is supported on FlexGroup volumes. Snapshot locking occurs only on the root constituent snapshot. Deleting the FlexGroup volume is allowed only if the root constituent expiration time has passed.

- FlexVol to FlexGroup conversion

You can convert a FlexVol volume with locked snapshots to a FlexGroup volume. Snapshots remain locked after the conversion.

- SnapMirror asynchronous

The compliance clock must be initialized on both the source and destination.

- SVM data mobility (used for migrating or relocating an SVM from a source cluster to a destination cluster)

Supported beginning with ONTAP 9.14.1.

- SnapMirror policy rules using the `-schedule` parameter
- SVM DR

The compliance clock must be initialized on both the source and destination.

- Volume clone and file clone

You can create volume clones and file clones from a locked snapshot.

- FlexCache volumes

Supported beginning with ONTAP 9.16.1.

Unsupported features

The following features currently are not supported with tamperproof snapshots:

- Consistency groups
- [FabricPool](#)

Tamperproof snapshots provide immutable protections that cannot be deleted. Because FabricPool requires the ability to delete data, FabricPool and snapshot locks cannot be enabled on the same volume.

- SMtape
- SnapMirror active sync

- SnapMirror synchronous

Enable snapshot locking when creating a volume

Beginning with ONTAP 9.12.1, you can enable snapshot locking when you create a new volume or when you modify an existing volume by using the `-snapshot-locking-enabled` option with the `volume create` and `volume modify` commands in the CLI. Beginning with ONTAP 9.13.1, you can use System Manager to enable snapshot locking.

System Manager

1. Navigate to **Storage > Volumes** and select **Add**.
2. In the **Add Volume** window, choose **More Options**.
3. Enter the volume name, size, export policy and share name.
4. Select **Enable Snapshot locking**. This selection is not displayed if the SnapLock license is not installed.
5. If it is not already enabled, select **Initialize SnapLock Compliance Clock**.
6. Save your changes.
7. In the **Volumes** window, select the volume you updated and choose **Overview**.
8. Verify that **SnapLock Snapshot Locking** displays as **Enabled**.

CLI

1. To create a new volume and enable snapshot locking, enter the following command:

```
volume create -vserver <vserver_name> -volume <volume_name> -snapshot  
-locking-enabled true
```

The following command enables snapshot locking on a new volume named vol1:

```
> volume create -volume vol1 -aggregate aggr1 -size 100m -snapshot  
-locking-enabled true  
Warning: snapshot locking is being enabled on volume "vol1" in  
Vserver "vs1". It cannot be disabled until all locked snapshots are  
past their expiry time. A volume with unexpired locked snapshots  
cannot be deleted.  
Do you want to continue: {yes|no}: y  
[Job 32] Job succeeded: Successful
```

Enable snapshot locking on an existing volume

Beginning with ONTAP 9.12.1, you can enable snapshot locking on an existing volume using the ONTAP CLI. Beginning with ONTAP 9.13.1, you can use System Manager to enable snapshot locking on an existing volume.

System Manager

1. Navigate to **Storage > Volumes**.
2. Select  and choose **Edit > Volume**.
3. In the **Edit Volume** window, locate the **Snapshots (Local) Settings** section and select **Enable snapshot locking**.

This selection is not displayed if the SnapLock license is not installed.

4. If it is not already enabled, select **Initialize SnapLock Compliance Clock**.
5. Save your changes.
6. In the **Volumes** window, select the volume you updated and choose **Overview**.
7. Verify that **SnapLock Snapshot locking** displays as **Enabled**.

CLI

1. To modify an existing volume to enable snapshot locking, enter the following command:

```
volume modify -vserver <vserver_name> -volume <volume_name> -snapshot  
-locking-enabled true
```

Create a locked snapshot policy and apply retention

Beginning with ONTAP 9.12.1, you can create snapshot policies to apply a snapshot retention period and apply the policy to a volume to lock snapshots for the specified period. You can also lock a snapshot by manually setting a retention period. Beginning with ONTAP 9.13.1, you can use System Manager to create snapshot locking policies and apply them to a volume.

Create a snapshot locking policy

System Manager

1. Navigate to **Storage > Storage VMs** and select a storage VM.
2. Select **Settings**.
3. Locate **Snapshot Policies** and select .
4. In the **Add Snapshot Policy** window, enter the policy name.
5. Select  **Add**.
6. Provide the snapshot schedule details, including the schedule name, maximum snapshots to keep, and SnapLock retention period.
7. In the **SnapLock Retention Period** column, enter the number of hours, days, months or years to retain the snapshots. For example, a snapshot policy with a retention period of 5 days locks a snapshot for 5 days from the time it is created, and it cannot be deleted during that time. The following retention period ranges are supported:
 - Years: 0 - 100
 - Months: 0 - 1200
 - Days: 0 - 36500
 - Hours: 0 - 24
8. Save your changes.

CLI

1. To create a snapshot policy, enter the following command:

```
volume snapshot policy create -policy <policy_name> -enabled true  
-schedule1 <schedule1_name> -count1 <maximum snapshots> -retention-period1  
<retention_period>
```

The following command creates a snapshot locking policy:

```
cluster1> volume snapshot policy create -policy lock_policy -enabled  
true -schedule1 hourly -count1 24 -retention-period1 "1 days"
```

A snapshot is not replaced if it is under active retention; that is, the retention count will not be honored if there are locked snapshots that have not yet expired.

Apply a locking policy to a volume

System Manager

1. Navigate to **Storage > Volumes**.
2. Select  and choose **Edit > Volume**.
3. In the **Edit Volume** window, select **Schedule snapshots**.
4. Select the locking snapshot policy from the list.
5. If snapshot locking is not already enabled, select **Enable snapshot locking**.
6. Save your changes.

CLI

1. To apply a snapshot locking policy to an existing volume, enter the following command:

```
volume modify -volume <volume_name> -vserver <vserver_name> -snapshot  
-policy <policy_name>
```

Apply retention period during manual snapshot creation

You can apply a snapshot retention period when you manually create a snapshot. Snapshot locking must be enabled on the volume; otherwise, the retention period setting is ignored.

System Manager

1. Navigate to **Storage > Volumes** and select a volume.
2. In the volume details page, select the **Snapshots** tab.
3. Select  **Add**.
4. Enter the snapshot name and the SnapLock expiration time. You can select the calendar to choose the retention expiration date and time.
5. Save your changes.
6. In the **Volumes > Snapshots** page, select **Show/Hide** and choose **SnapLock Expiration Time** to display the **SnapLock Expiration Time** column and verify that the retention time is set.

CLI

1. To create a snapshot manually and apply a locking retention period, enter the following command:

```
volume snapshot create -volume <volume_name> -snapshot <snapshot_name>  
-snaplock-expiry-time <expiration_date_time>
```

The following command creates a new snapshot and sets the retention period:

```
cluster1> volume snapshot create -vserver vs1 -volume vol1 -snapshot  
snap1 -snaplock-expiry-time "11/10/2022 09:00:00"
```

Apply retention period to an existing snapshot

System Manager

1. Navigate to **Storage > Volumes** and select a volume.
2. In the volume details page, select the **Snapshots** tab.
3. Select the snapshot, select , and choose **Modify SnapLock Expiration Time**. You can select the calendar to choose the retention expiration date and time.
4. Save your changes.
5. In the **Volumes > Snapshots** page, select **Show/Hide** and choose **SnapLock Expiration Time** to display the **SnapLock Expiration Time** column and verify that the retention time is set.

CLI

1. To manually apply a retention period to an existing snapshot, enter the following command:

```
volume snapshot modify-snaplock-expiry-time -volume <volume_name> -snapshot <snapshot name> -snaplock-expiry-time <expiration_date_time>
```

The following example applies a retention period to an existing snapshot:

```
cluster1> volume snapshot modify-snaplock-expiry-time -volume vol1 -snapshot snap2 -snaplock-expiry-time "11/10/2022 09:00:00"
```

Modify an existing policy to apply long-term retention

In a SnapMirror relationship, you can set a retention period on a mirror-vault policy rule, and the retention period is applied for snapshots replicated to the destination if the destination volume has snapshot locking enabled. The retention period takes precedence over keep count; for example, snapshots that have not passed their expiry will be retained even if the keep count is exceeded.

Beginning with ONTAP 9.14.1, you can modify an existing SnapMirror policy by adding a rule to set long-term retention of snapshots. The rule is used to override the default volume retention period on SnapLock vault destinations and on non-SnapLock SnapMirror destination volumes.

1. Add a rule to an existing SnapMirror policy:

```
snapmirror policy add-rule -vserver <SVM name> -policy <policy name> -snapmirror-label <label name> -keep <number of snapshots> -retention-period [<integer> days|months|years]
```

The following example creates a rule that applies a retention period of 6 months to the existing policy called "lockvault":

```
snapmirror policy add-rule -vserver vs1 -policy lockvault -snapmirror -label test1 -keep 10 -retention-period "6 months"
```

Learn more about `snapmirror policy add-rule` in the [ONTAP command reference](#).

Copyright information

Copyright © 2026 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, nonsublicensable, 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).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.