



Technical FAQ

SnapMirror FAQ for ONTAP

ONTAP 9.2

Siddharth Agrawal, NetApp
July 2017

Abstract

This FAQ, intended for internal NetApp field personnel and channel partners, answers frequently asked questions about NetApp® data replication in the NetApp ONTAP® architecture. Do not distribute this FAQ directly to customers. Use this FAQ with NetApp SnapMirror® product documents and the "Data ONTAP Data Protection Online Backup and Recovery Guide." In addition, detailed answers can be found in [TR-4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP](#) and the [SnapMirror in Clustered Data ONTAP SE technical presentation](#).

Internal

TABLE OF CONTENTS

1	Overview.....	5
	Why should customers care about secondary storage?	5
1.1	SnapMirror Release Summary	5
	What functionalities were introduced in ONTAP 9.1 for SnapMirror?	5
	What functionalities were introduced in ONTAP 9.0 for SnapMirror?	5
	What changed between clustered Data ONTAP 8.2 and 8.3?	6
	What changed between clustered Data ONTAP 8.1 and 8.2?	6
	What changed between clustered Data ONTAP 8.0 and 8.1?	6
1.2	SnapMirror Unified Replication.....	6
	What is SnapMirror Unified Replication?	6
	Why would I want to use SnapMirror Unified Replication?	7
2	Requirements.....	7
2.1	Licensing	7
	How is SnapMirror licensed?	7
	Are there any changes to the SnapMirror licensing process?	7
	Why did NetApp simplify the licensing process?	8
	How can I get a temporary SnapMirror license for customer demonstration?	8
	How do I get a license for SnapMirror Unified Replication?	8
	Do I need separate SnapRestore licenses for use with other third-party backup management software, such as Veeam?	8
3	Architecture	8
	What is a Snapshot copy?	8
	What is a Snapshot policy?.....	8
	What are the naming conventions for the scheduled Snapshot copies?	8
	What is a Snapshot prefix?	8
	What is a SnapMirror policy?	8
	What is a load-sharing mirror?	9
	Does SnapMirror support intercluster replication for DR?	9
	Does SnapMirror support intracluster replication?	9
	Does SnapMirror support cascaded relationships?	9
	Can DR and backup reside on the same destination?	9
	How do SnapMirror and SnapVault work in a unified way?	9
	Will Unified Replication work on any source volume?	9
	We are still running Data ONTAP operating in 7-Mode. Will Unified Replication work for us?	10
	We already use SnapMirror without Unified Replication, but we would like to upgrade to new releases more easily. How difficult is it to convert to Unified Replication?	10
	We currently have cascaded relationships. How do I implement Unified Replication in that situation? ...	10
	What is SnapMirror seeding? How do I perform it?	10
	To use SnapMirror, must the destination and source controllers have the same ONTAP version?	10

What is the deswizzling process?	10
How long does it take for the deswizzler to finish? Is there anything we can do to lessen the amount of work a deswizzler does?	10
What is secondary Snapshot management? How is it beneficial?	11
Can I restore data from a secondary Snapshot copy of a volume?	11
Can I restore part of a file from a Snapshot copy of a FlexVol volume?	11
What is SnapMirror resynchronization?	11
Can I resynchronize a SnapMirror relationship?	11
Can I resynchronize a SnapMirror relationship in reverse?	12
Does SnapMirror require that a common Snapshot copy be used for resynchronization after a break? ..	12
Is a Snapshot copy created for each update?	12
Can I mirror from a volume in a 32-bit aggregate to a volume in a 64-bit aggregate with SnapMirror? ...	12
Can I mirror from a volume in a 64-bit aggregate to a volume in a 32-bit aggregate with SnapMirror? ...	12
4 Configuration	12
What are the limits for concurrent SnapMirror transfers?	12
What is the limit for the number of Snapshot copies per volume?	12
What happens after I reach the limit of Snapshot copies per volume?	12
Can I configure the Snapshot expiry time for SnapVault relationships?	13
Can I preserve my Snapshot copies for SnapVault relationships after they reach their retention limit? ..	13
What is the fan-in peering limit in ONTAP 9?	13
What is the fan-out limit for SnapMirror in ONTAP 9?	13
Does SnapMirror support multipathing?	13
Does ONTAP record SnapMirror operations?	13
If I encounter a SnapMirror error—for example, a Snapshot copy not being mirrored—where is that error logged?	13
How is the status of a SnapMirror relationship displayed?	14
If a transfer fails, is the number of attempts configurable?	14
Are SnapMirror transfers automatically restarted if a storage failover occurs?	14
Does SnapMirror in clustered Data ONTAP set periodic transfer restart checkpoints?	14
If a SnapMirror relationship is deleted, is the destination volume also deleted?	14
Are any orchestration tools available for configuring and managing SnapMirror relationships?	14
5 Performance.....	14
How is space management handled in a SnapMirror relationship?	14
Does SnapMirror support storage efficiency?	14
Can I deduplicate or compress the destination volume if the source volume is not deduplicated or compressed?	15
Is SnapMirror network compression available in clustered Data ONTAP?	15
Are the efficiency ratios nearly the same when mirroring from a NetApp all-flash FAS system to a FAS system?	15
Can I throttle the intercluster SnapMirror transfer?	15
Why should I use SnapMirror over NDMP?	15
What is SnapMirror transfer priority?	15

Is there a way I can calculate the performance or the RPO of my Unified Replication relationships?	16
Is lag time displayed for the SnapMirror transfers?	16
What configuration requirements should I consider when positioning ONTAP SnapMirror replication for a customer?	16
6 Interoperability.....	17
Does SnapMirror work with the other solutions that NetApp sells, for example, NetApp AltaVault?	17
If I use SnapMirror with NetApp Volume Encryption (NVE), can my source volume be encrypted and my SnapMirror target be unencrypted (or conversely)?	17
If I use SnapMirror to mirror my encrypted volume to a different cluster, is the data encrypted “in flight” or “over the wire” by NVE?	17
Is replication between 7-Mode and clustered Data ONTAP systems supported?	17
What is the compatibility between different ONTAP systems for SnapMirror relationships using type <code>DP</code> and policy <code>async-mirror</code> ?.....	17
What is the compatibility between different ONTAP systems for SnapMirror relationships using type <code>XDP</code> and policy <code>async-mirror</code> or <code>mirror-vault</code> ?.....	18
What is the compatibility between different ONTAP systems for SnapVault relationships using type <code>XDP</code> and policy <code>vault</code> ?	18
Is SnapMirror to tape for supported for seeding?	19
Is SnapMirror to tape for backup application supported?	19
7 Contact Us.....	19
Who are the contacts for SnapMirror/SnapVault related questions?	19
References	19
Version History	19

LIST OF TABLES

Table 1) Considerations for positioning ONTAP SnapMirror to a customer.	16
Table 2) Permitted ONTAP releases for SnapMirror relationships using type <code>DP</code> and policy <code>async-mirror</code>	17
Table 3) Permitted ONTAP releases for SnapMirror relationships using type <code>XDP</code> and policy <code>async-mirror</code> or <code>mirror-vault</code>	18
Table 4) Permitted ONTAP releases for SnapVault relationships using type <code>XDP</code> and policy <code>vault</code>	18

1 Overview

For data-driven enterprises, safeguarding customer data and business-critical applications has never been so important. Determining the best strategy to protect that data is crucial for the business's future. There are several approaches to increasing data availability in the face of hardware, software, or even site failures. Backups provide a way to recover lost data from an archival medium (tape or disk). Redundant hardware technologies also help mitigate the damage caused by hardware issues or failures. Mirroring provides a third mechanism to facilitate data availability and minimize downtime. NetApp® SnapMirror® technology offers a fast and flexible enterprise solution for mirroring or replicating data over LANs and WANs. SnapMirror is a key component in enterprise data protection (DP) strategies.

Why should customers care about secondary storage?

Answer: The secondary storage is not merely a backup repository. There are many data protection use cases that can be run on a secondary system. Customers can use it for creating replica copies for disaster recovery (DR), to offload tape backup processes from the primary system, to distribute datasets to other locations, and to create read/write clones for test and development environments using NetApp FlexClone® technology. Customers can also mirror their primary systems to the secondary system and then create NetApp Snapshot™ copies on secondary systems. This approach saves considerable space on primary systems while maintaining a secondary system that can be a reliable DR destination.

1.1 SnapMirror Release Summary

What functionalities were introduced in ONTAP 9.1 for SnapMirror?

Answer: Beginning with NetApp ONTAP® 9.1 software, we introduced the following functionalities for SnapMirror:

- You can use SnapMirror to back up NAS file services to the cloud from primary or secondary ONTAP storage. You must run a combination of NetApp AltaVault™ 4.3, SnapCenter® 2.0, and ONTAP 9.1 to use this functionality.
- Data Protection (DP) Bundle systems can be purchased as target storage for data protection such as backup consolidation. The bundle includes NetApp FAS8200 series hybrid storage systems configured as two-node switchless clusters with Premium Bundle software. The DP Bundle allows customers to use SnapMirror without purchasing a SnapMirror license for selected primary systems on NetApp Data ONTAP 8.3.2 and later.
- You can use a SnapMirror license to enable NetApp SnapVault® capabilities in NetApp SnapManager® or NetApp SnapDrive® software. Those products did not allow you to use a SnapMirror license to enable SnapVault in ONTAP 9.0.

What functionalities were introduced in ONTAP 9.0 for SnapMirror?

Answer: Beginning with ONTAP 9.0, we introduced the following functionalities for SnapMirror:

- If there are volume-level SnapMirror relationships between two storage virtual machines (SVMs), you can convert these relationships to an SVM disaster recovery relationship.
- You can use either a SnapMirror license or a SnapVault license to enable SnapVault. In previous releases, you could use only a SnapVault license to enable SnapVault.
- You can use timestamp-based naming conventions and ordinal naming conventions for configuring the scheduled Snapshot copies.
- You can configure the expiry time for Snapshot copies to protect them from deletion until the expiry time elapses. You cannot configure the expiry time for Snapshot copies of infinite volumes.
- Global network throttling is available for all SnapMirror and SnapVault transfers at a per-node level.

- You can create Snapshot copies on a SnapMirror SnapVault destination for longer-term retention. You don't have to create and maintain a corresponding longer-term Snapshot copy at the primary destination.
- SnapVault enables you to transfer a Snapshot copy from the SnapMirror primary volume even if the secondary volume already has a Snapshot copy with the same name.
- If secondary SnapVault Snapshot copies are *not* used with SnapVault relationships that are part of a cascade, all cascading configurations for asynchronous disaster recovery (ADR) and the vault and logical replication engine with storage efficiency disaster recovery (LRSE-DR) are supported.

What changed between clustered Data ONTAP 8.2 and 8.3?

Answer: Clustered Data ONTAP 8.3 introduced SnapMirror Unified Replication. SnapMirror Unified Replication refers to the use of SnapMirror with the same (unified) logical replication engine as in SnapVault. This unified relationship type is designated Extended Data Protection (XDP) and provides single baseline functionality at the volume level. This permits DR and backup to occur on the same destination, drastically reducing storage and network bandwidth, which translates immediately into cost savings.

In addition, Unified Replication also supports version-flexible replication, which simplifies upgrading of replication relationships to newer ONTAP releases. The destination controller no longer needs to have a clustered Data ONTAP major version number equal to or higher than the major version of the source controller, allowing customers to have nondisrupted upgrades. The destination must have at least Data ONTAP 8.3, but the source can be ONTAP 8.3 or later. A source system with ONTAP 8.3 will not support version-flexible replication with a destination system on ONTAP 8.2 or older releases. Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site for more details.

Clustered Data ONTAP 8.3 also supports network compression and FlexClone mirroring. It lets you resume SnapMirror transfers after the system instantiated nondisruptive operations (NDO). It also lets you receive the history of SnapMirror operations.

What changed between clustered Data ONTAP 8.1 and 8.2?

Answer: Clustered Data ONTAP 8.2 introduced SnapVault to clustered Data ONTAP, began supporting SnapMirror cascading, and began supporting SnapMirror to tape for seeding only. It is multitenancy-ready, and replication can be managed by SVM administrators. Clustered Data ONTAP 8.2 also offers improved scalability (number of concurrent transfers) and increased data transfer speeds over version 8.1.

What changed between clustered Data ONTAP 8.0 and 8.1?

Answer: Clustered Data ONTAP 8.1 allows intercluster asynchronous volume replication (replication between volumes hosted on different clusters), enabling disaster recovery replication to a cluster in a remote site. Also, in clustered Data ONTAP 8.1, the SnapMirror function allows savings gained with NetApp storage efficiency technologies, such as FAS deduplication and compression, to be maintained in the replica.

1.2 SnapMirror Unified Replication

What is SnapMirror Unified Replication?

Answer: SnapMirror Unified Replication refers to the use of SnapMirror with the same (unified) logical replication engine as in SnapVault. This unified relationship type is designated Extended Data Protection (XDP) and provides single baseline functionality at the volume level, drastically reducing storage and network bandwidth, which translates immediately into cost savings.

SnapMirror provides support for applications to fail over to a secondary volume and continue operating, as well as the capability to fail back to the primary location later. This capability is sometimes referred to as disaster recovery (DR).

SnapVault supports long-term retention of data for compliance, backup, archiving, and other reasons. This support is sometimes referred to as vaulting or backup. SnapMirror Unified Replication brings together the powerful capabilities of SnapMirror and SnapVault at the volume level.

Why would I want to use SnapMirror Unified Replication?

Answer: SnapMirror Unified Replication has four major benefits for end users:

- Only one baseline copy of a volume is needed on the secondary storage. (Without Unified Replication, both SnapMirror and SnapVault need their own baseline copies.)
- Less network traffic is required between the primary and secondary (a single baseline plus fewer Snapshot copies over time).
- Upgrading of replication relationships to newer ONTAP releases is simpler. (Without Unified Replication, you can replicate only from later to earlier releases, and doing so can cause serious operational issues with complex replication topologies. With Unified Replication, you can replicate from earlier to later *and* from later to earlier, as long as both sides are ONTAP 8.3 or later.)
- If a primary volume that is the source of a SnapMirror relationship becomes corrupted, the secondary system inevitably becomes corrupted as well. With Unified Replication, you can now recover the primary volume from any arbitrary SnapVault or SnapMirror Snapshot copy.

Note: BRE will be deprecated after the ONTAP 9.5 release.

2 Requirements

2.1 Licensing

How is SnapMirror licensed?

Answer: Data protection mirroring, both within the same cluster and to another cluster, requires a SnapMirror license on each cluster.

Are there any changes to the SnapMirror licensing process?

Answer: Yes. To simplify our licensing process, we have made the following changes:

- For ONTAP 9.0 and later releases, if you purchase a SnapMirror license, you can use SnapVault at either the ONTAP 9 source or the ONTAP 9 destination. A SnapVault license is no longer necessary if you have a SnapMirror license at both source and destination.
- For ONTAP 8.3.x, if you purchase a SnapMirror license, you receive a free SnapVault license.
- For pre-8.3.x systems, you still need a separate SnapVault license (bundled or a la carte).

These changes apply to relationships between pre-ONTAP 9 and ONTAP 9 systems, and to relationships between two ONTAP 9 systems. The new process reduces the number of licenses a customer has to manage, and it makes SnapVault easier to access and use.

Note: Beginning with ONTAP 9.1, DP Bundle systems can be purchased as target storage for data protection. The bundle includes NetApp FAS8200 series hybrid storage systems configured as two-node switchless clusters with Premium Bundle software. The DP Bundle allows customers to use SnapMirror without purchasing a SnapMirror license for selected primary systems on ONTAP 8.3.2 and later. For more details on qualified primary NetApp systems, refer to the [DP Bundle technical FAQ](#) on the Field Portal.

Why did NetApp simplify the licensing process?

Answer: We expect the simplification to increase the adoption and use of SnapVault and secondary storage.

How can I get a temporary SnapMirror license for customer demonstration?

Answer: You can get demo evaluation SnapMirror licenses from the [licensing portal](#). To generate the license keys, you will need the serial number of the controller. These license keys are valid for 90 days from the generation date.

How do I get a license for SnapMirror Unified Replication?

Answer: There is no separate license for SnapMirror Unified Replication. If you do not already run ONTAP 8.3 or later releases, you need to upgrade both primary and secondary systems to it. If you already have licenses for SnapMirror or SnapVault, you can start using Unified Replication. You enable Unified Replication by specifying the relationship type as XDP in the profile.

Do I need separate SnapRestore licenses for use with other third-party backup management software, such as Veeam?

Answer: Yes. However, a NetApp SnapRestore® license is needed only on the source cluster.

3 Architecture

What is a Snapshot copy?

Answer: A Snapshot copy is a read-only image of a volume that captures the state of the file system at a point in time. NetApp ONTAP allows you to create and delete Snapshot copies, and manage Snapshot schedules according to your requirements.

What is a Snapshot policy?

Answer: A Snapshot policy controls creation and retention of backup copies for a volume. Cluster administrators and SVM administrators can use the `volume snapshot policy` commands to create and manage all Snapshot copy policies for a volume. A Snapshot policy can have up to five schedules.

What are the naming conventions for the scheduled Snapshot copies?

Answer: ONTAP provides two naming conventions for scheduled Snapshot copies: timestamp-based and ordinal. You can configure Snapshot copies with either of them.

What is a Snapshot prefix?

Answer: A Snapshot prefix is an optional string that you can specify for naming automatic Snapshot copies. It provides more flexibility than using schedule names. If a prefix is specified in the Snapshot schedule, then the schedule name is not used to name Snapshot copies.

What is a SnapMirror policy?

Answer: Cluster administrators can use the `snapmirror policy` commands to create and manage all data protection mirror and vault policies. SnapMirror policy rules determine whether the policy is applied to a SnapVault relationship or to a data protection mirror copy. A SnapMirror policy in which the `vserver` parameter contains the cluster name is a clusterwide policy. A SnapMirror policy in which the `vserver` parameter contains an SVM name is an SVM-wide policy.

What is a load-sharing mirror?

Answer: A load-sharing (LS) mirror is a read-only, full copy of a volume created on cluster nodes to help distribute read workload. Copies are synchronized according to a configured schedule. NetApp recommends LS mirrors for the root volume of each SVM across nodes, because an LS mirror can be promoted if the original source volume fails. An LS mirror can be written to only by accessing the source volume admin path. For a CIFS client to have read/write access, create a share by prepending `/.admin` to the beginning of the share's path (the path that is shared, not the share name). For an NFS client to have read/write access, prepend `/.admin` to the beginning of the path used in the client's `mount` command. The best practice is to create LS mirrors on all nodes, including the node hosting the source volume. LS mirrors are used within an intracluster deployment. The two use cases are read acceleration within a cluster and namespace resiliency.

Note: LS mirroring for nonroot volumes is not supported in ONTAP 9.3 and later. In ONTAP 9.3, an FPVR can temporarily reenable LS mirroring. In ONTAP 9.5, it will be fully deprecated.

Does SnapMirror support intercluster replication for DR?

Answer: Yes. Starting with clustered Data ONTAP 8.1, you can create a data protection mirror with a source volume on one cluster replicated to a destination volume on a different cluster. You cannot create LS mirrors between different clusters.

Does SnapMirror support intracluster replication?

Answer: Yes. Clustered Data ONTAP 8.1 supports replication between volumes on the same cluster. You can use data protection mirrors to create a local copy of a volume or to copy a volume from one SVM to another within the same cluster. You can create LS mirrors only within the same SVM.

Does SnapMirror support cascaded relationships?

Answer: Yes. SnapMirror cascading is supported in clustered Data ONTAP 8.2 and later.

Can DR and backup reside on the same destination?

Answer: Yes. SnapMirror Unified Replication provides an option to do both vaulting and replication through a single engine (logical replication with storage efficiency) and on the same destination, where you can choose to fail over to the recent state of business with a recovery point objective (RPO) of two hours. As an added benefit, a failback to a backup state of business is also possible if the most recent state of business is corrupted by malware.

How do SnapMirror and SnapVault work in a unified way?

Answer: With Unified Replication, SnapMirror and SnapVault use the same logical replication engine. On the secondary system, there is a single baseline copy of the volume being replicated, and all Snapshot copies are associated with this volume regardless of how they were created or replicated. Therefore, you can recover data to the source system from any Snapshot copy on the secondary system.

Will Unified Replication work on any source volume?

Answer: Unified Replication supports all volume types that SnapMirror or SnapVault previously supported on ONTAP 8.3 and later. There are currently no exceptions. However, in some extreme edge cases in which the rate of data change in the source volume is very high *and* the customer needs a very low RPO, it might be better not to use Unified Replication. For RPO-related conversations, refer to the [Unified Replication RPO Calculator](#) on the Field Portal. If an RPO is lower than the calculator can compute, contact the appropriate product manager to discuss FPVR options.

We are still running Data ONTAP operating in 7-Mode. Will Unified Replication work for us?

Answer: No. Unified Replication works only on clustered Data ONTAP 8.3 and later. There are no plans to support earlier clustered Data ONTAP releases or any 7-Mode releases.

We already use SnapMirror without Unified Replication, but we would like to upgrade to new releases more easily. How difficult is it to convert to Unified Replication?

Answer: For each SnapMirror relationship that you currently have configured, you simply break that relationship and re-create it using the `-type XDP` parameter. There is no need to do a rebaseline. From that point onward, you are using Unified Replication. However, keep expectations in line with the environment and the configuration. There will be overhead and time involved with the metadata conversion. For example, a data warehouse rebuild can take anywhere from 10 to 90 minutes for 1TB of source volume data depending on the file sizes residing in the volume.

We currently have cascaded relationships. How do I implement Unified Replication in that situation?

Answer: SnapMirror Unified Replication can be used in a cascade relationship. However, because the relationship type is XDP, you can have only one XDP relationship type in the cascade in ONTAP 8.3 because of the ONTAP programmatic enforcement. There is no such restriction in ONTAP 9.0 and later.

What is SnapMirror seeding? How do I perform it?

Answer: *Seeding* refers to the initial transfer of data for a newly created SnapMirror relationship. When you create a SnapMirror relationship by using the `snapmirror create` command, an initial transfer is not automatically performed. The `create` command simply establishes the relationship and the metadata that defines it. To perform the initial transfer, follow the `snapmirror create` command with the `snapmirror initialize` command. Alternatively, use the `snapmirror initialize` command alone to perform the initial transfer as soon as the relationship is created. If the SnapMirror relationship does not exist, the `initialize` command creates the relationship and performs the initial transfer.

To use SnapMirror, must the destination and source controllers have the same ONTAP version?

Answer: No. Starting with 8.3, version-flexible replication simplifies upgrading of replication relationships to newer ONTAP releases. The destination controller no longer needs to have a clustered Data ONTAP major version number equal to or higher than the major version of the source controller, allowing customers to have nondisrupted upgrades. The destination must have at least Data ONTAP 8.3, but the source can be ONTAP 8.3 or later. A source system with ONTAP 8.3 will not support version-flexible replication with a destination system on ONTAP 8.2 or older releases. Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site for more details.

What is the deswizzling process?

Answer: When we replicate a NetApp FlexVol® volume using the block replication engine (type `DP`), we need to make the internal representation of that volume compatible with the new block layout on the new system. In more technical terms, all the indirect blocks inside the FlexVol volume need to be cleaned up to point to the correct physical block numbers on the new home aggregate. This process occurs in the background with a scanner, the "deswizzler." The deswizzler starts deswizzling from the oldest Snapshot copy to the active file system.

How long does it take for the deswizzler to finish? Is there anything we can do to lessen the amount of work a deswizzler does?

Answer: Deswizzling time depends on the size of the volume, the number of Snapshot copies in the volume, differences between copies, the biggest file size on the volume, and the value of `maxfiles`. Deswizzling is a low-priority background process that is interrupted by other conflicting I/O operations.

To lessen the amount of work a deswizzler does:

- Delete unnecessary Snapshot copies while replicating from source to destination.
- Disable scheduled Snapshot copies until deswizzling is finished.
- Perform less frequent SnapMirror updates from secondary system.

What is secondary Snapshot management? How is it beneficial?

Answer: There are two parts to using SnapVault to back up data (Snapshot copies) from a primary system to a secondary system: Snapshot policy and Snapshot schedule. The Snapshot policy controls creation and retention of backup copies on the **primary system**. The Snapshot schedule controls the transfer of backup copies from the primary system to the secondary system. The schedule also defines the retention time for the backup copies on the secondary system. In a typical scenario, customers keep Snapshot copies on the primary system only for faster recovery of mission-critical data and only for a shorter duration. The backup Snapshot copies are eventually transferred to the secondary system for longer duration and deleted from the primary system when they reach the end of their retention time on the primary system. Because the minimum retention count for any Snapshot copy is 1, there is always one backup Snapshot copy of each type on the primary system. Prior to ONTAP 9.0, SnapVault does not delete any Snapshot copies from the primary system after the transfer. Thus, a Snapshot copy might remain on the primary system for months or even years after it has been transferred to the secondary system, locking up expensive space on the primary system.

In ONTAP 9.0, we introduced secondary Snapshot management. With secondary Snapshot management, you can create Snapshot copies on the secondary system itself instead of creating them on the primary system and then transferring them over the network. Thus, backup Snapshot copies that need to be maintained longer (weekly or monthly) can be created directly on the secondary system while daily ones are being transferred from the primary system. Prior to ONTAP 9.0, each Snapshot policy rule identified a type of backup Snapshot copy to be transferred from the primary system. Now each rule specifies whether the associated Snapshot copy is to be transferred from the primary system or is to be created locally on the secondary system.

Can I restore data from a secondary Snapshot copy of a volume?

Answer: Depending on the extent of the corruption, you can recover a single file, incremental changes since a previous Snapshot copy, or the entire volume. You can do this from any Snapshot copy that is created through Unified Replication.

Can I restore part of a file from a Snapshot copy of a FlexVol volume?

Answer: Yes. You can restore a range of data from a file in a Snapshot copy to an existing file in the active file system by using the `volume snapshot partial-restore-file` command. Partial file restores can be used only to restore specific pieces of a LUN, and NFS or CIFS container files. To avoid inconsistent data, write operations are not allowed on the object that you are restoring.

What is SnapMirror resynchronization?

Answer: When you set up data protection mirror relationships between the source and destination volumes, and the source becomes unavailable, the destination can be used to provide data access, with minimum disruption. When the source volume becomes available again, the administrator of the source cluster can use the `snapmirror resync` command to resynchronize the data from the destination volume to serve data from the source volume.

Can I resynchronize a SnapMirror relationship?

Answer: Yes. After breaking a SnapMirror relationship, you can perform a SnapMirror `resync` command to re-establish replication in the same direction from the last common SnapMirror Snapshot copy.

Can I resynchronize a SnapMirror relationship in reverse?

Answer: Yes. In NetApp OnCommand® System Manager, if you right-click a broken (failed-over) SnapMirror relationship, one of the options in the menu is Reverse/Resync. Selecting this option performs all the steps necessary to reverse the replication. To do so in the CLI, you must first delete the existing SnapMirror relationship on the current destination, then create a new one in the reverse direction from the new destination, and then resynchronize in SnapMirror. In either the OnCommand System Manager or CLI case, both volumes must have a common SnapMirror Snapshot copy to resynchronize.

Does SnapMirror require that a common Snapshot copy be used for resynchronization after a break?

Answer: If a common SnapMirror Snapshot copy still exists on both volumes, it is automatically be used for resynchronization. If the SnapMirror Snapshot copies have been removed, any common Snapshot copy that was replicated by SnapMirror can be used for resynchronization. Any Snapshot copies on the destination that are newer than the copy selected for resynchronizing and that no longer exist on the source are removed from the destination. Any Snapshot copies on the source newer than the one selected for resynchronizing is rereplicated to the destination.

Is a Snapshot copy created for each update?

Answer: Yes. An update or initialization will force a Snapshot copy creation on the source volume. In addition, these Snapshot copies will have SnapMirror owner information associated with them, preventing their deletion.

Can I mirror from a volume in a 32-bit aggregate to a volume in a 64-bit aggregate with SnapMirror?

Answer: Yes. For additional reference, refer to [TR-3978: 64-Bit Aggregates: Overview and Best Practices](#).

Can I mirror from a volume in a 64-bit aggregate to a volume in a 32-bit aggregate with SnapMirror?

Answer: Yes. For additional reference, refer to [TR-3978: 64-Bit Aggregates: Overview and Best Practices](#).

4 Configuration

What are the limits for concurrent SnapMirror transfers?

Answer: There is no limit on the number of SnapMirror transfers that can be scheduled at one time. NetApp SnapMirror optimizes the concurrent active transfers depending on the system utilization and the platform type. If there are more transfers scheduled than the system can handle, they are queued and started when one of the current transfers completes.

What is the limit for the number of Snapshot copies per volume?

Answer: NetApp ONTAP supports 255 Snapshot copies per volume. However, in a SnapMirror relationship, the number of usable Snapshot copies is limited to 251 on the destination volume. ONTAP reserves 4 Snapshot copies.

What happens after I reach the limit of Snapshot copies per volume?

Answer: After you reach the Snapshot copy retention limit defined by a SnapMirror policy, the oldest Snapshot copy is automatically deleted to create space before a new Snapshot copy is transferred. You can also manually delete the older Snapshot copies.

Can I configure the Snapshot expiry time for SnapVault relationships?

Answer: Yes. As an administrator, you can set and modify the expiry time for Snapshot copies. This setting prevents the Snapshot copies from deletion until the expiry time elapses. However, if you want to delete Snapshot copies that have not expired, you can set the expiry time to 0 before deleting them.

In some cases, customer might want to avoid rogue administrators from deleting the Snapshot copies that have not expired. You can use NetApp SnapLock® software with SnapVault to prevent deletion of Snapshot copies until a defined retention period expires. This feature delivers a capacity-efficient regulatory compliant solution by making backups nondeletable 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). For more information, refer to [SnapLock Technical FAQ for ONTAP 9](#).

Can I preserve my Snapshot copies for SnapVault relationships after they reach their retention limit?

Answer: Yes. You can configure the policy rule by using the `snapmirror policy add-rule` command, or modify the rule by using the `snapmirror policy modify-rule` command. Use the `-preserve` parameter to preserve all the Snapshot copies.

What is the fan-in peering limit in ONTAP 9?

Answer: The fan-in limit is 255:1.

What is the fan-out limit for SnapMirror in ONTAP 9?

Answer: The fan-out limit for SnapMirror in ONTAP 9 is 1:8.

Does SnapMirror support multipathing?

Answer: Yes. SnapMirror supports multipathing with different types and speeds of networks. Only active-active multipathing is available; SnapMirror doesn't allow active-passive failover mode.

Does ONTAP record SnapMirror operations?

Answer: The `snapmirror show-history` command displays the history of SnapMirror operations. A maximum of one day's worth or 100 entries of history is stored for each SnapMirror relationship. SnapMirror history has two modes of operations:

- **Bulk mode.** In this mode, no destination endpoint needs to be specified, but this mode can generate a lot of output. History is displayed as follows:
 - When an SVM context exists, history of operations is displayed for all relationships in the destination SVM.
 - When no SVM context exists, history of operations is displayed for all relationships in the destination cluster.
- **Pointed mode.** In this mode, a destination volume endpoint, a relationship ID, or an operation ID needs to be specified. Only matching history entries are displayed.

If I encounter a SnapMirror error—for example, a Snapshot copy not being mirrored—where is that error logged?

Answer: All the SnapMirror events are logged in the `snapmirror_audit.log` and `snapmirror_error.log` files under the `/etc/log` directory on the node where the destination volume resides. This node might be different from the one where the command was issued. You can determine the node running the operation by running the `snapmirror show -fields destination-volume-node` command. OnCommand System Manager 3.0 allows viewing of the SnapMirror log files.

How is the status of a SnapMirror relationship displayed?

Answer: To display the status of a mirror by command line, use the `snapmirror show` command. If you need more detailed information, add the destination path and instance to the command.

If a transfer fails, is the number of attempts configurable?

Answer: Yes, the number of tries (not retries—the first update counts as a try) is configurable. The default value is 8 tries, with 1 try every 60 seconds. Note that a value of 0 for `tries` does not mean no retries; it means unlimited retries. You can use `snapmirror create` or `snapmirror modify` to change the retry count. When all the retries are exhausted, ONTAP aborts the transfer and logs the error.

Are SnapMirror transfers automatically restarted if a storage failover occurs?

Answer: Yes. Data ONTAP 8.3 and later provide NDO restartability. NDO events, such as takeover, giveback, ARL, and node reboots, do not affect SnapMirror operations, because SnapMirror automatically restarts the interrupted or missed scheduled update on volume-to-volume SnapMirror relationships after the reboot completes. If a schedule update was aborted, it might be restarted before the next schedule point if an NDO event occurs before the next schedule point.

Be aware of the following caveats about NDO restartability:

- Manual operations are not restarted.
- If a storage efficiency scan starts after the completion of a SnapVault transfer, and the scan is interrupted by an NDO event, the scan not automatically restarted.

Does SnapMirror in clustered Data ONTAP set periodic transfer restart checkpoints?

Answer: Yes.

If a SnapMirror relationship is deleted, is the destination volume also deleted?

Answer: No. The volume remains on the destination system. This is true for both DP and LS mirror relationships and volumes.

Are any orchestration tools available for configuring and managing SnapMirror relationships?

Answer: Yes. You can configure and manage SnapMirror relationships by using the following tools:

- Use **NetApp SnapCenter** software to simplify application consistent backup, restore, and clone management with integrated workflows. It provides role-based access control, monitoring, notification, logging, reporting, and scheduling for application and database plug-ins.
- Use **NetApp OnCommand System Manager** to set up and manage the clusters. It provides easy-to-use workflows for NetApp SnapVault and SnapMirror management.
- Use **NetApp OnCommand Unified Manager and OnCommand Insight** to discover and manage the data protection cluster.

5 Performance

How is space management handled in a SnapMirror relationship?

Answer: Starting in Data ONTAP 8.2, when the size of the source volume that is in a SnapMirror relationship increases, the destination volume size also increases automatically.

Does SnapMirror support storage efficiency?

Answer: Yes. If the source volume is deduplicated or compressed, that storage efficiency is preserved from source to destination. This is true for all types of mirrors and volume move and volume copy

operations. With SnapMirror Unified Replication, you can enable additional deduplication and compression on the destination volume. However, if you enable additional compression on the destination volume, storage efficiencies are not preserved from source to destination.

Can I deduplicate or compress the destination volume if the source volume is not deduplicated or compressed?

Answer: Yes. You can configure additional deduplication and compression on the destination volume for SnapMirror Unified Replication. However, if you enable additional compression on the destination volume, the transfer does not preserve storage efficiency present on the source system. For all the other cases, storage efficiencies are preserved from source to destination.

Note: For more details on NetApp storage efficiencies, refer to [TR-4476](#).

Is SnapMirror network compression available in clustered Data ONTAP?

Answer: Yes. Beginning with clustered Data ONTAP 8.3, it is possible to compress SnapMirror relationships over the network. This feature is disabled by default and is controlled through the `network-compression-enabled` option in the SnapMirror policy. Network compression is enabled and disabled on a policy basis, so all relationships associated with a particular policy would have network compression enabled or disabled.

Are the efficiency ratios nearly the same when mirroring from a NetApp all-flash FAS system to a FAS system?

Answer: It depends. If logical replication (XDP SnapMirror or SnapVault) is used, preserving storage efficiencies on the FAS destination volume requires that all the storage efficiencies are enabled: adaptive compression, volume deduplication, and data compaction. You don't need to enable inline compression and deduplication, because no compression or deduplication is done inline; enabling the features preserves the storage efficiencies. Inline data compaction is required to compact blocks as they are rewritten to the destination volume.

If physical replication (DP SnapMirror) is used, adaptive compression and volume deduplication savings are preserved regardless of the destination volume storage efficiency settings, because that's how physical replication is done. Inline data compaction must still be enabled to achieve those savings.

Can I throttle the intercluster SnapMirror transfer?

Answer: To limit the amount of bandwidth used by intercluster SnapMirror, you can apply a throttle to intercluster SnapMirror relationships. You can set a throttle through the CLI when creating a relationship, by adding the `-throttle` option and a value in kilobytes, or by modifying an existing relationship with the `snapmirror modify` command. You can also specify the throttle value by using OnCommand System Manager to update a SnapMirror relationship.

Why should I use SnapMirror over NDMP?

Answer: SnapMirror transfers can happen quickly and with less impact on the source system than that of NDMP backups. As a best practice, use NetApp Snapshot copies and perform SnapMirror replication from a primary system as a first stage of backup to significantly shorten or eliminate backup windows. Then you can perform NDMP backup to tape from the secondary system.

What is SnapMirror transfer priority?

Answer: Starting with Data ONTAP 8.2, you can set a priority on SnapMirror relationships. Currently two levels are supported: normal and low. Normal transfers are given preference over low-priority transfers.

Is there a way I can calculate the performance or the RPO of my Unified Replication relationships?

Answer: Yes. You can use the [Unified Replication RPO Calculator](#) on the Field Portal to help with timing-related conversations.

Is lag time displayed for the SnapMirror transfers?

Answer: Yes. Starting with ONTAP 8.2, lag time is displayed in the output of the `snapmirror show` command.

What configuration requirements should I consider when positioning ONTAP SnapMirror replication for a customer?

Answer: See the following table.

Table 1) Considerations for positioning ONTAP SnapMirror to a customer.

Configuration Requirement	Customer Implication
Intercluster replication requires at least one intercluster logical interface (intercluster LIF) per node.	A LIF requires an IP address; the customer must dedicate at least one IP address per node for replication. The intercluster LIF can be assigned to a data port or a dedicated intercluster port.
The customer must have adequate bandwidth over the WAN for the initial transfer.	Tape seeding is available, and beginning with version 8.3, there is network compression functionality in clustered Data ONTAP. Depending on the bandwidth available and the amount of data to transfer, initialization could take an extended amount of time. Transfers that are interrupted restart from a periodically created checkpoint.
The customer business requirement is for disaster recovery.	Use SnapMirror functionality. The source and destination volumes maintain the same Snapshot history, and the destination can be made writable for conversion to primary storage.
The customer business requirement is for older Snapshot copies at the destination for backup or archiving.	Use SnapVault functionality. Certain Snapshot copies from the source are selected to be replicated, and the destination cannot be made writable.
The customer business requirement is for synchronous replication.	There is currently no synchronous SnapMirror replication capability in ONTAP. Check to see whether NetApp MetroCluster™ satisfies the customer business requirement.
The customer has Ethernet infrastructure for replication.	SnapMirror over Fibre Channel is not provided in clustered Data ONTAP.
The customer needs to create a destination SVM with the proper language type.	Both the source and destination SVM can have different language types, but the source and destination volumes <i>must</i> have the same language type. The SVM language type can be set <i>only</i> at the time of SVM creation.
The customer needs matching implementation of storage efficiency on source and destination volumes.	The customer cannot implement deduplication or compression on the destination volume without enabling deduplication or compression at the source volume.

6 Interoperability

Does SnapMirror work with the other solutions that NetApp sells, for example, NetApp AltaVault?

Answer: Yes. You can perform protection operations of NAS file services data by using the ONTAP CLI. Using SnapMirror, you can back up NAS file services data to AltaVault for replication to the cloud and restore it from the cloud back to an ONTAP volume.

If I use SnapMirror with NetApp Volume Encryption (NVE), can my source volume be encrypted and my SnapMirror target be unencrypted (or conversely)?

Answer: Yes. The source volume and destination volume in a SnapMirror relationship can have different encryption settings.

If I use SnapMirror to mirror my encrypted volume to a different cluster, is the data encrypted “in flight” or “over the wire” by NVE?

Answer: No. SnapMirror sits above the NetApp WAFL[®] layer, so the data sent by SnapMirror is not encrypted by NVE. For more information, refer to [NVE Technical FAQ](#).

Is replication between 7-Mode and clustered Data ONTAP systems supported?

Answer: Yes. Typically, SnapMirror is supported only between systems operating in the same mode. But, during a project involving transition from a 7-Mode system to a clustered Data ONTAP system, SnapMirror is used for migration of NAS data. However, you cannot create regular SnapMirror relationships for DR and backup that are reversible and can be resynchronized.

What is the compatibility between different ONTAP systems for SnapMirror relationships using type `DP` and policy `async-mirror`?

Answer: For SnapMirror relationships using type `DP` and policy `async-mirror`, you can build a relationship from an earlier ONTAP release to one that is no more than two releases later. The following relationships are permitted:

Table 2) Permitted ONTAP releases for SnapMirror relationships using type `DP` and policy `async-mirror`.

Source ONTAP Release	Permitted Destination ONTAP Release					
	8.1.x	8.2.x	8.3.x	9.0	9.1	9.2
8.1.x	Yes	Yes	Yes	No	No	No
8.2.x	No	Yes	Yes	Yes	No	No
8.3.x	No	No	Yes	Yes	Yes	No
9.0	No	No	No	Yes	Yes	Yes
9.1	No	No	No	No	Yes	Yes
9.2	No	No	No	No	No	Yes

What is the compatibility between different ONTAP systems for SnapMirror relationships using type `xdp` and policy `async-mirror` or `mirror-vault`?

Answer: SnapMirror relationships using type `xdp` and policy `async-mirror` or `mirror-vault`, also known as version-flexible replication, are available only with Data ONTAP 8.3 and later releases. The following relationships are permitted:

Table 3) Permitted ONTAP releases for SnapMirror relationships using type `xdp` and policy `async-mirror` or `mirror-vault`.

Source ONTAP Release	Permitted Destination ONTAP Release						
	8.1.x	8.2.x	8.3.x	8.3.2P4	9.0	9.1	9.2
8.1.x	No	No	No	No	No	No	No
8.2.x	No	No	No	No	No	No	No
8.3.x	No	No	Yes	Yes	Yes	Yes	No
8.3.2P4	No	No	Yes	Yes	Yes	Yes	Yes
9.0	No	No	Yes	Yes	Yes	Yes	Yes
9.1	No	No	Yes	Yes	Yes	Yes	Yes
9.2	No	No	No	Yes	Yes	Yes	Yes

What is the compatibility between different ONTAP systems for SnapVault relationships using type `xdp` and policy `vault`?

Answer: SnapVault relationships using type `xdp` and policy `vault` are available only with Data ONTAP 8.2 and later releases. The version of ONTAP running on the secondary volume can be older or newer than the version running on the primary volume. When the primary and secondary volumes run different versions of ONTAP, they should not be more than two major releases apart. The following relationships are permitted:

Table 4) Permitted ONTAP releases for SnapVault relationships using type `xdp` and policy `vault`.

Source ONTAP Release	Permitted Destination ONTAP Release					
	8.1.x	8.2.x	8.3.x	9.0	9.1	9.2
8.1.x	No	No	No	No	No	No
8.2.x	No	Yes	Yes	Yes	No	No
8.3.x	No	Yes	Yes	Yes	Yes	No
9.0	No	Yes	Yes	Yes	Yes	Yes
9.1	No	No	Yes	Yes	Yes	Yes
9.2	No	No	No	Yes	Yes	Yes

Is SnapMirror to tape for supported for seeding?

Answer: Yes. Starting with Data ONTAP 8.2, SnapMirror to tape is supported for seeding.

Is SnapMirror to tape for backup application supported?

Answer: Yes, SnapMirror to tape is supported.

7 Contact Us

To let us know how we can improve this FAQ, contact us at docfeedback@netapp.com. Include "SnapMirror FAQ" in the subject line.

Who are the contacts for SnapMirror/SnapVault related questions?

Answer: Use the ng-snapmirror mailing list. This list is monitored by the product management and technical marketing engineer groups.

References

The following references were used in this FAQ:

- ONTAP 9 Product Documentation
<http://library-clnt.dmz.netapp.com/documentation/docweb/index.html?productID=62320>
- Unified Replication SE Training Presentation
<https://fieldportal.netapp.com/content/272825>
- Unified Replication RPO Calculator
<https://fieldportal.netapp.com/content/262135>
- DP Bundle Technical FAQ
<https://fieldportal.netapp.com/content/508265>
- TR-3978: 64-Bit Aggregates: Overview and Best Practices
<http://www.netapp.com/us/media/tr-3978.pdf>
- SnapLock Technical FAQ for ONTAP 9
<https://fieldportal.netapp.com/content/417102>
- TR-4476: NetApp Data Compression, Deduplication, and Data Compaction
<http://www.netapp.com/us/media/tr-4476.pdf>
- NVE Technical FAQ
<https://fieldportal.netapp.com/content/467500>

Version History

Version	Date	Document Version History
Version 4.0	July 2017	Siddharth Agrawal: Updated the document for ONTAP 9.2
Version 3.2	January 2016	Mike Worthen
Version 3.1	August 2015	Mike Worthen
Version 3.0	April 2015	Vincent Goveas
Version 2.2	November 2013	Amit Prakash Sawant

Version	Date	Document Version History
Version 2.1	July 2013	Amit Prakash Sawant
Version 2.0	April 2013	Amit Prakash Sawant
Version 1.0	February 2012	Larry Touchette

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 © 1994–2017 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, 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.