



Technical Report

NetApp SANtricity synchronous and asynchronous mirroring (11.62 and later)

Feature descriptions and deployment guide

Mitch Blackburn, NetApp
November 2024 | TR-4839

Abstract

NetApp® E-Series and EF-Series systems provide rock-solid local and remote data mirroring capabilities by using the NetApp SANtricity® synchronous and asynchronous mirroring features. This document describes the features and provides specific setup and operations guidance for customers who use next-generation E-Series or EF-Series systems with SANtricity OS 11.62 and later. With SANtricity OS 11.80 EF300 and EF600 now support asynchronous mirroring.

TABLE OF CONTENTS

Introduction	4
Document scope.....	4
Audience	4
Terminology.....	4
Overview of mirroring features	6
Comparison of asynchronous and synchronous mirroring.....	6
Typical use cases	7
Overview of mirroring operations	7
Asynchronous mirroring operations	7
Synchronous mirroring operations.....	11
Storage array requirements	13
Configuration requirements	13
Storage array limits.....	14
Application requirements.....	14
Unified Manager	14
System Manager	15
Mirroring and role-based access control.....	16
Connectivity requirements.....	16
FC interface (asynchronous or synchronous).....	16
iSCSI interface (asynchronous only)	18
Distances and speed	19
Host connections	19
Capacity and volume requirements	20
Mirror volume requirements.....	20
Free capacity	21
Reserved capacity	22
Connectivity and volume ownership	22
Asynchronous mirroring configuration and management.....	22
Determine synchronization times and performance impacts	23
Configure asynchronous mirroring.....	27
Test the asynchronous communications link	33
Manage asynchronous mirroring	35

Troubleshooting of common failures.....	39
Synchronous mirroring configuration and management	41
Balance link speeds and data traffic	41
Configure synchronous mirroring.....	42
Test the synchronous communications link	46
Manage synchronous mirroring	47
Where to find additional information	49
Version history	49

LIST OF TABLES

Table 1) Mirroring terms.	4
Table 2) Comparison between mirroring features.	6
Table 3) Resynchronization methods for asynchronous mirroring.....	9
Table 4) Resynchronization methods for synchronous mirroring.....	13
Table 5) Storage array limits for mirroring.	14
Table 6) Cable distances by cable type and speed.	19
Table 7) Transmission rates by carrier type.	19
Table 8) Candidates for mirror volumes.	20
Table 9) Approximate hours needed for initial or full synchronization by pipe size.....	25
Table 10) System configuration for performance example.	26
Table 11) Example performance results.....	26
Table 12) Troubleshooting of common failures.	40

LIST OF FIGURES

Figure 1) Periodic resynchronization in asynchronous mirroring.	10
Figure 2) Synchronous mirroring write process.	12
Figure 3) Unified Manager application interface.	15
Figure 4) System Manager application interface.	16
Figure 5) Single-fabric configuration.....	17
Figure 6) Fabric dedicated to mirroring.	18
Figure 7) Time from one protection point to the next.....	23

Introduction

Mirroring is a feature in NetApp® SANtricity® applications that replicates data between two storage arrays. NetApp systems offer two types of mirroring—asynchronous and synchronous:

- Asynchronous mirroring copies data volumes on demand or on a schedule, capturing only the data that has changed since the last image capture. This type of mirroring is optimal for periodic processes such as backup and archiving.
- Synchronous mirroring replicates data volumes in real time for continuous availability, capturing an identical copy of data at every moment. This type of mirroring is optimal for business continuity purposes such as disaster recovery.

Document scope

This document describes how to configure mirroring with the following NetApp SANtricity versions and controller models:

- SANtricity applications:
 - System Manager, OS version 11.62 or later
 - Unified Manager, version 4.2 or later
- Array models (dual controllers only), firmware versions 8.62 and later:
 - EF280 and E2800 storage arrays
 - E4000 storage arrays
 - EF570 and E5700 storage arrays
 - EF300 and EF600 storage arrays (firmware versions 8.80 and later)

Note: Synchronous and asynchronous mirroring are not supported with iSCSI on E4000 arrays.

Note: Synchronous mirroring is not supported on the EF600 or EF300 storage arrays.

Note: This document does not describe earlier SANtricity versions, earlier array models, or other types of SANtricity management applications, such as CLI and API. For detailed information about mirroring with earlier systems, see [TR-4656, SANtricity Synchronous and Asynchronous Mirroring \(11.61 and Earlier\), Feature Descriptions and Deployment Guide](#).

Audience

This document is for NetApp customers, partners, and OEMs with general experience in the following:

- NetApp E-Series systems and products, particularly in the area of data protection
- Disaster recovery solutions

Terminology

Table 1 lists terms that are used in this document.

Table 1) Mirroring terms.

Term	Definition
Array	An E-Series storage array, such as the E2800. Also called a “storage system.”
Initial (or full) synchronization	A process by which all the data from the primary site is copied to the secondary site.
Local storage array	The primary site, which holds the original data in a mirroring configuration.

Term	Definition
Mirror consistency group	<p>A container for one or more volumes that participates in asynchronous mirroring. Each mirrored pair in a mirror consistency group shares the same synchronization settings. Thus, all mirrored pairs in the group synchronize simultaneously, which preserves a consistent recovery point at the secondary site.</p> <p>Note: Synchronous mirroring does not use mirror consistency groups.</p>
Mirror volume	Any volume that participates in mirroring. It can have either the primary or the secondary role.
Mirrored pair	<p>Two volumes that are linked together through a mirroring relationship. A mirrored pair consists of a primary volume and a secondary volume.</p> <p>In asynchronous mirroring, a mirrored pair always belongs to a mirror consistency group.</p>
Primary site	<p>The production site for any given mirrored pair or group. Mirroring occurs from the primary site to the secondary site.</p> <p>Sites can also mirror for each other; that is, one site can be the primary for one set of volumes and the secondary for another.</p>
Primary volume	The source volume in a mirrored pair.
Remote storage array	The secondary site, which holds a replica of the data in a mirroring configuration.
Reserved capacity	<p>The physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.</p> <p>Reserved capacity volumes are required so that the controller can persistently save the information that is necessary to maintain mirroring in an operational state. They contain information such as delta logs and copy-on-write data.</p>
Resynchronization	A process in which changed data is sent from the primary site to the secondary site after the initial or full synchronization has been completed.
Role	The volume's designation in the mirrored pair, either primary or secondary. You can reverse the volume's role, if necessary. Role changes are typically used for recovery when a failure has occurred at the primary site or when communication is lost between volume members.
Secondary site	The destination recovery site where volumes reside for any given mirrored pair or group. Sites can mirror for each other; that is, one site can be the primary for one set of volumes and the secondary for another.
Secondary volume	The destination volume in a mirrored pair, which contains the replica of the data. In synchronous mirroring, the secondary volume is an exact copy of the primary volume. In asynchronous mirroring, the secondary lags in time behind the primary.
Synchronization	Data transfer between the local and the remote storage arrays.

Overview of mirroring features

This section describes the differences between the two types of mirroring.

Comparison of asynchronous and synchronous mirroring

Asynchronous mirroring and synchronous mirroring offer several important differences and benefits, as follows:

- Asynchronous mirroring replicates changed data from the primary site to the secondary site at discrete points in time, providing a recovery point at the secondary site that is behind the current primary site. Asynchronous mirroring includes the following attributes:
 - Long distances are supported between arrays.
 - Block-level updates reduce bandwidth and time requirements by replicating only changed blocks.
 - Crash-consistent data is available at a recovery point at the secondary site.
 - A disaster recovery plan can be tested without affecting production and replication.
 - Data can be replicated between dissimilar NetApp E-Series storage arrays.
 - A standard IP or FC network can be used for replication.
- Synchronous mirroring continuously replicates data so that the secondary site is up to date. Synchronous mirroring includes the following attributes:
 - The supported distances between arrays are shorter than for asynchronous mirroring.
 - Block-level replication provides consistent data across two sites.
 - Up-to-date crash-consistent data is available at a disaster recovery site.
 - Data can be replicated between dissimilar NetApp E-Series storage arrays.
 - An FC network is used for replication.

Both types of mirroring have some effect on performance, but be aware that with synchronous mirroring, each write I/O is not complete until the write to the remote volume has been completed. Given the extra write latency that remote write confirmation messages incur, this consideration is important for high-performance environments.

Table 2 describes the main differences between the two types of mirroring.

Table 2) Comparison between mirroring features.

Attribute	Asynchronous mirroring	Synchronous mirroring
Replication method	Point in time: Copies are created on demand or at user-defined intervals. Data that has changed between the current and the previous copy is written to the secondary volume.	Continuous: Mirroring is automatically executed continuously, copying data from every host write. Writes are acknowledged to the host after the copy to the secondary is completed. Note: The remote write acknowledgment adds latency to all mirrored-volume write I/O.
Consistency groups	Each mirrored volume is a member of a mirror consistency group. All volumes in a group are synchronized at the same time so that a recovery point can include several volumes.	Not supported: Mirroring is performed on a volume basis only.
Reserved capacity (repository)	Reserved capacity (repository) volumes are required for each volume in a mirrored pair.	One reserved capacity (repository) volume is created for each controller with synchronous mirroring activated.

Attribute	Asynchronous mirroring	Synchronous mirroring
Allowed volume types	Standard or thin	Standard only
Communication between arrays	FC or iSCSI	FC only
Distance	Virtually unlimited: Distance is limited only by the capabilities of the network, network components, and required synchronization interval.	To meet latency and application performance requirements, the supported distance is limited to 6.2 miles (10km).

Typical use cases

Asynchronous mirroring is optimal for satisfying the demand for nonstop operations. In general, it is far more network efficient for periodic processes, such as backup and archiving.

Synchronous mirroring is optimal for continuous replication between a small number of systems over relatively short distances for business continuity purposes. It is **not** ideal for periodic processes, such as backup and archiving.

Asynchronous mirroring is good for the following use cases:

- Remote backup consolidation
- Wide area content distribution
- Long-distance disaster recovery
- Backup of 24/7 application data
- Application development and testing on a point-in-time image of live data
- Application data replication to a secondary site

Synchronous mirroring is good for the following use cases:

- Data center–type environments to top-tier applications and data
- Protection of databases and other highly transactional applications
- Local and near-local disaster recovery; a secondary data center in the same metropolitan area
- Business continuity during scheduled maintenance of the primary storage array, with the secondary array acting as the primary

Overview of mirroring operations

This section describes the process for asynchronous and synchronous mirroring, including initial synchronization and resynchronization methods.

Asynchronous mirroring operations

Asynchronous mirroring copies data volumes on demand or on a schedule, which minimizes or avoids downtime that might result from data corruption or loss. This type of mirroring captures the state of the primary volume at a particular point in time and copies just the data that has changed since the last image capture. The primary site can be updated immediately, and the secondary site can be updated as bandwidth allows. The information is cached and sent later, as network resources become available.

After you configure it, asynchronous mirroring works as follows:

1. Initial synchronization begins. This step is a one-time transfer of the entire primary volume to the secondary volume. This step is required before incremental updates can occur. At the end of the initialization, the source system creates a point-in-time image and begins tracking changed blocks.

2. After the initialization is complete, scheduled or manually triggered updates can occur. Each update transfers only the new and changed blocks from the source to the remote system. This operation proceeds as follows:
 - a. The source storage array creates a new point-in-time image and begins tracking changes.
 - b. The changed blocks are sent to the destination system.
 - c. After the update is complete, both systems have the new point-in-time image, which becomes the baseline image for the next update.

Because asynchronous replication is periodic, the system can consolidate the changed blocks and conserve network bandwidth. There is minimal impact on write throughput and write latency.

The following sections provide more details on how asynchronous mirroring works.

Initial synchronization process for asynchronous mirroring

After you configure a mirrored pair, the data from the primary volume is copied in its entirety to the secondary volume. The owning controller on the primary side directs this process. During this process, the primary volume is fully accessible for I/O operations by all host systems that normally have access to it. Until the initial synchronization is complete, the secondary volume has minimal value as a recovery point.

The initial synchronization is completed in multiple phases. The first phase copies all the data from the primary to the secondary volume. For a thin volume, this approach includes only blocks that have been written on the primary volume and have not been unmapped. However, the first phase does not create a point-in-time image as the synchronization point. The first phase usually takes quite a long time (possibly many hours for large volumes), so the mirror is not considered to be synchronized from the first pass.

During the first phase, a delta log is used to track write requests to the primary volume. The second phase of the initial synchronization creates a point-in-time image on the primary volume to use as a synchronization point. Not using a point-in-time image for the first phase of the synchronization process minimizes the effect on performance during the initialization process.

During initial synchronization, the following steps are taken on the primary storage array:

1. The mirror state is set to initializing.
2. One of the delta logs (log A) in the primary side's reserved capacity (repository) is initialized to a clean state and then is activated to track updates to the primary volume from that point forward.
3. The other delta log (log B) is initialized in a completely set state, indicating that all the data regions of the primary and secondary are not synchronized.
4. A background synchronization process transfers the contents of the primary volume to the secondary volume. The process iterates through delta log B, copying any data that is flagged as unsynchronized in the delta log to the secondary volume. After copying an individual data segment, the delta log is updated to persistently save the progress of the synchronization process. Any interruption or reset of the controller causes the synchronization process to resume at the point of the most recent progress record. Such interruptions do not force the controller to restart the synchronization process from the beginning of the volume. Because the address range of an individual bit in the delta log is relatively small, a larger data chunk might be copied to the secondary. This action results in several bits in the delta log being marked complete in the synchronization progress.
5. A point-in-time image is created to capture the post-first-phase consistent image of the primary volume to be synchronized to the secondary array.
6. Data regions of the primary volume that had been written during step 4 as recorded in delta log A are written to the secondary volume. After copying an individual data segment, the delta log is updated to persistently save the progress of the synchronization process. Any interruption or reset of the controller causes the synchronization process to resume at the point of the most recent progress record. When all the data regions have been copied to the secondary, a message is sent to the

remote array so that an image copy can be created to capture the recovery point on the secondary volume.

7. The point-in-time image that was created to capture the primary synchronization image is deleted.
8. Steps 5 through 7 are repeated as long as the synchronization process is not completed within the configured synchronization interval.
9. When the background synchronization process is completed within the synchronization interval, the mirror state is set to Optimal.
10. The next synchronization process for normal mirror operation is scheduled based on the creation time of the point-in-time image that was used to complete initial synchronization.

The following steps are taken on the secondary array:

1. The secondary side receives the mirror state change to Initializing from the primary array and persists with the state change.
2. The secondary volume receives synchronization write commands from the primary site. Because this synchronization is the initial one, the updates are written directly to the secondary site's base volume, without first creating an image copy.
3. The secondary site receives a message from the primary site that the background synchronization process is complete. A point-in-time image is created to capture the synchronized image.

Resynchronization methods for asynchronous mirroring

After the initial synchronization is complete, the system must resynchronize every so often to create a new recovery point that meets your SLA. Two resynchronization methods, manual and automatic, are available for asynchronous mirroring. Table 3 describes the two methods.

Table 3) Resynchronization methods for asynchronous mirroring.

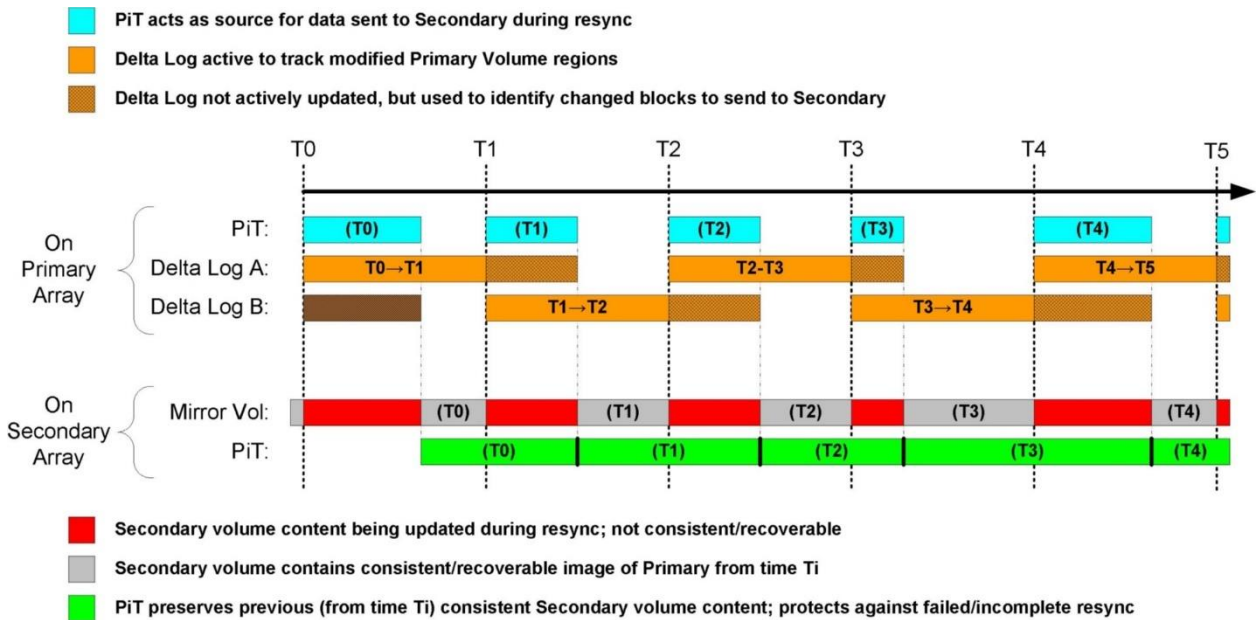
Method	Description
Manual	To immediately resynchronize data on all mirrored pairs in the group, use manual resynchronization.
Automatic (periodic)	To schedule resynchronization times, use the automatic method. You can specify the time (in minutes) from the beginning of the previous update to the beginning of the next update. For example, if the synchronization interval is set at 30 minutes and the synchronization process starts at 4 p.m., the next process starts at 4:30 p.m.

Periodic resynchronization for asynchronous mirroring

Each mirror consistency group has a configurable attribute that specifies the resynchronization interval for all mirrored pairs in the group. This interval is the amount of time between automatically sending updates of modified primary-side data to the secondary side. For example, if the interval is 30 minutes and the first resynchronization interval is encountered at 3:15 p.m., then the ensuing resynchronization intervals start at 3:45 p.m., 4:15 p.m., and so on. Because of differing quantities of modified data on the primary volume and because of performance variations in the intercontroller communications link, the actual amount of time to complete a resynchronization varies. Thus, the first resynchronization cycle might be from 3:15 to 3:39 p.m., the second from 3:45 to 4:06, the third from 4:15 to 4:42, and so on. Figure 1 shows an example.

Note: In Figure 1, "PiT" refers to a point-in-time image.

Figure 1) Periodic resynchronization in asynchronous mirroring.



When a resynchronization interval arrives (for example, at time T_1 as shown in Figure 1), the controller on the primary side of the mirrored pair manages the synchronization process with the following steps:

1. The mirror synchronization activity state changes to Active.
2. A point-in-time image is created to capture the T_1 image of the primary volume.
3. Delta log B is cleared and is activated to track primary volume updates that occur after T_1 .
4. Delta log A is deactivated to no longer track writes. A background synchronization process transfers the regions of the primary T_1 point-in-time image that were modified between T_0 and T_1 , as recorded in delta log A. The process iterates through the delta log, copying any data that is flagged as unsynchronized in the delta log to the secondary volume. After copying an individual data segment, the delta log is updated to persistently save the progress of the synchronization process. Any interruption or reset of the controller causes the synchronization process to resume at the point of the most recent progress record. Such interruptions do not force the controller to restart the synchronization process from the beginning of the volume.
5. Copy-on-write operations that result from write requests to the primary volume are minimized by performing the copy-on-write only for regions that are flagged in delta log A. The background synchronization process accesses only data regions that are flagged in the delta log. If a write request to the primary site affects a region that is not flagged, saving the original data is unnecessary because the synchronization process does not access that region.
6. When the background synchronization process is complete (all flagged regions from delta log A have been copied to the secondary site), the mirror synchronization activity state is set to Idle.
7. The point-in-time image that was created to capture the initial synchronization image is deleted, leaving no active point-in-time images on the primary mirror repository.

The following steps occur on the secondary site:

1. The secondary site receives the mirror synchronization activity state change to Active from the primary array and persists with the state change.
2. The secondary volume receives synchronization write commands from the primary side. Because a point-in-time image was created at the end of the previous synchronization cycle, the consistent secondary image is preserved during the synchronization process.

3. The secondary site receives a mirror synchronization activity state change to Idle from the primary side. A point-in-time image is created to protect the newly established consistent image; any old point-in-time image for a previous consistent image is deleted, and the mirror state is persisted.

When the next resynchronization interval arrives (time T2), the primary side's actions are repeated, except delta log A and delta log B are reversed. On the secondary side, the processing is the same as previously described for the first resynchronization interval. This pattern repeats for all subsequent resynchronization intervals.

Use of thin-provisioned volumes in asynchronous mirroring

Thin-provisioned volumes are allowed to participate in a mirror consistency group only when they are paired with another thin-provisioned volume. Thin-provisioned volumes that are in a group follow the same rules as for normal volumes (for example, security settings, data assurance, and capacity). They also have the following operational differences from fully provisioned volumes:

- If an existing thin volume is selected for the secondary volume of a mirrored pair, that thin volume is reinitialized with a new reserved capacity volume.
- The secondary thin volume is initialized before the initial (or any full) synchronization process begins.
- Only provisioned blocks in the primary thin volume are transferred during the initial synchronization process.
- Automatic expansion must be enabled. Thin volumes with automatic expansion disabled are not shown as candidates for mirroring in the NetApp SANtricity applications.
- The secondary thin volume parameters that control that volume's growth are set to match the primary thin volume parameters. When you select an existing thin volume for the secondary volume, a warning appears and allows you to cancel the operation.
- The alert thresholds can be changed only on the primary side of the mirrored pair. Any changes to those parameters on the primary side are automatically propagated to the secondary side.
- Removing a thin-provisioned volume from a mirror consistency group does not cause any changes to the parameters that control that volume's growth.

Synchronous mirroring operations

Synchronous mirroring replicates data volumes in real time to promote continuous availability. The copy is identical to production data at every moment because each time a write to the primary volume occurs, a write to the secondary volume occurs. The host does not receive an acknowledgment that the write was successful until the secondary volume is successfully updated with the changes.

After you configure it, synchronous mirroring works as follows:

1. Initial synchronization begins. The first step is a one-time baseline transfer of the entire dataset. This step is required before incremental updates can occur.
2. After the initial synchronization is complete, the two sides attempt to stay synchronized through normal operations. Each update transfers the new and changed blocks from the primary volume to the secondary volume. This operation proceeds as follows:
 - a. Changed data blocks are written to the primary volume.
 - b. The primary storage array sends the changed blocks to the secondary array.
 - c. The secondary array sends an acknowledgment after the changed blocks have been written to the cache (or to the disk for write-through caching).
 - d. The primary storage array sends an acknowledgment to the host.
 - e. After the operation is complete, both systems are in a synchronized state. From the application point of view, both systems have the same consistent dataset.

Because synchronous replication is continuous, the replication link between the two sites must provide sufficient bandwidth capabilities. For more information, see section 6, “Connectivity requirements.”

The following sections provide more in-depth information about how synchronous mirroring works.

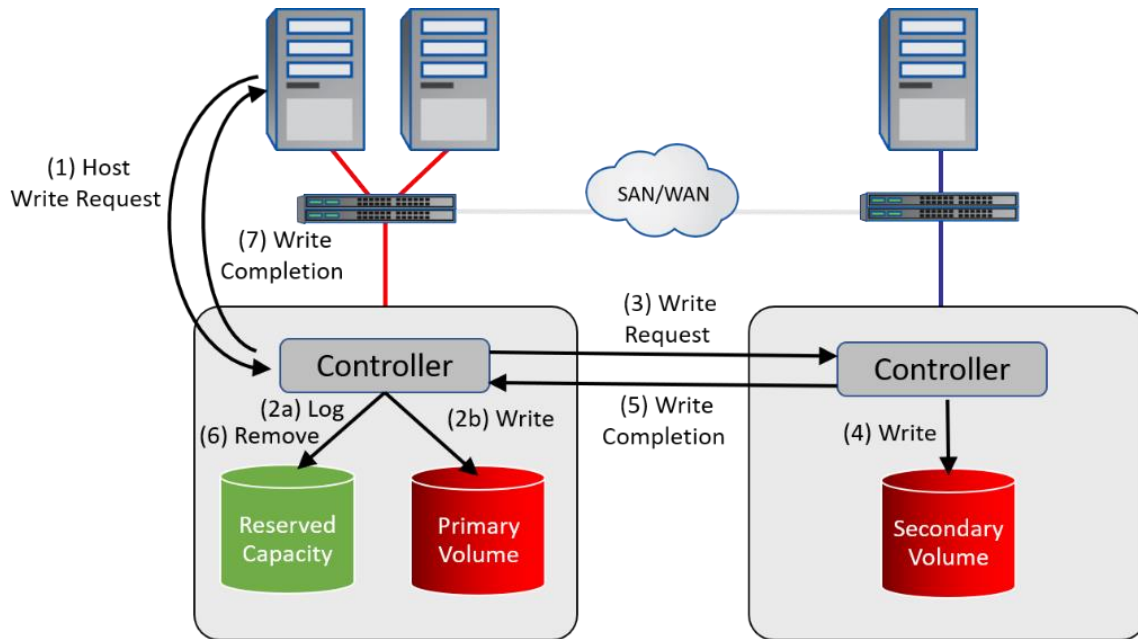
Initial synchronization process for synchronous mirroring

When the mirror relationship is first established, the data from the primary volume is copied in its entirety to the secondary volume. The owning controller on the primary side directs this process. During the copy or mirror initialization, the primary volume is fully accessible for I/O operations by all host systems that normally have access to it. Until the initial synchronization is complete, the secondary volume has minimal value as a recovery point.

Following are the steps in the write process, indicating the corresponding numbers with the illustration in Figure 2:

1. When a primary controller (the owning controller of the primary volume) receives a write request from a host (1), the controller first logs information about the write request on the reserved capacity volume (the information is placed in a queue), (2a). In parallel, it writes the data to the primary volume (2b).
2. The controller then initiates a remote write operation to copy the affected data blocks to the secondary volume at the remote site (3).
3. When the remote write operation is complete (4 and 5), the primary controller removes the log record from the mirror repository volume (deletes it from the queue) (6).
4. Finally, the controller sends an I/O completion confirmation back to the host system (7).

Figure 2) Synchronous mirroring write process.



Resynchronization methods for synchronous mirroring

If a link interruption or a volume error prevents communication with the secondary storage array, the current owner of the primary volume changes the mirrored pair to an Unsynchronized status. The host can continue to issue write requests to the primary volume, but remote writes to the secondary volume do not take place. In this case, the mirrored volume pairs are running out of synchronization. Events that might cause a mirrored pair to become unsynchronized include the following:

- Primary controller failure or reset
- Remote volume error
- Secondary controller failure
- Link failures due to switch errors

When connectivity is restored between the controller owner of the primary volume and the controller owner of the secondary volume, resynchronization occurs. Only the blocks of data that have changed on the primary volume during the link interruption are copied to the secondary volume. The changed blocks are written to a delta bitmap table, which is part of the reserved capacity (repository) volume.

Note: During resynchronization, because of the process that the controller uses, the secondary volume is not a suitable candidate for disaster recovery. The controller processes changed blocks sequentially from the lowest to the highest logical block address of the primary volume.

Table 4 describes the two resynchronization methods, manual and automatic, that are available with synchronous mirroring.

Table 4) Resynchronization methods for synchronous mirroring.

Method	Description
Manual	<p>When the manual option is selected and communication fails between the primary volume and the secondary volume, the remote mirrored pair changes to Unsynchronized status. Any write requests to the primary volume are logged, and a Needs Attention status appears for the storage array.</p> <p>After the controller owner of the primary volume detects that communication has been restored, the remote mirrored pair stays in Unsynchronized status until you issue a <code>resume</code> command. The <code>resume</code> starts the resynchronization process.</p> <p>Because you can manage the resynchronization process in a way that provides the best opportunity to recover data, NetApp recommends that you use the manual resynchronization option.</p>
Automatic	<p>The automatic option enables automatic resynchronization after a communication failure between the primary storage array and the secondary storage array. Immediately after the controller owner of the primary volume detects that communication has been restored, that controller owner starts resynchronizing the primary volume and the secondary volume.</p> <p>Any communication disruptions between the primary storage array and the secondary storage array while resynchronization is in progress could result in a mix of new data and old data on the secondary volume. This situation renders the data unusable in a disaster recovery situation. For this reason, NetApp recommends that you use manual resynchronization and create a point-in-time image of the secondary volume before you begin the resynchronization. Then, if a failure occurs during the resynchronization, you have a crash-consistent image at the secondary site at the point at which synchronization was lost. (Asynchronous mirroring is not subject to this resynchronization issue.)</p>

Storage array requirements

Mirroring features are supported on NetApp EF280, E2800, E4000, EF570, E5700, EF300, and EF600 storage arrays, as well as on some earlier E-Series models.

Note: Synchronous and asynchronous mirroring are not supported with iSCSI on E4000 arrays.

Note: Synchronous mirroring is not supported on the EF600 or EF300 storage arrays.

Configuration requirements

For mirroring operations, you must have two storage arrays that meet the following requirements:

- Each storage array must have two controllers (duplex configuration).

- Each controller in both the primary and secondary storage arrays must have an Ethernet management port that is configured, and it must be connected to your network.
- For the instructions in this document, the storage arrays must have a minimum firmware version of 8.62. They can each run different OS versions.
- The remote storage array must have enough free capacity to create a secondary volume that is equal to or greater than the primary volume that you want to mirror.
- For asynchronous mirroring, the controllers must have FC or iSCSI host ports. For synchronous mirroring, the controllers must have FC host ports. However, controllers can still receive host I/O on a different interface, such as SAS.

Storage array limits

E-Series systems have different limits on the number of volumes and mirrors that they support, as is shown in Table 5.

Table 5) Storage array limits for mirroring.

Maximum	E2800/EF280	E4000	E5700/EF570	EF300	EF600
Volumes per system (including reserved capacity volumes)	512	512	2,048	1024	2048
Mirrors per system (both asynchronous and synchronous)	32	32	128	64	256
Asynchronous mirror consistency groups per system	4	4	4	4	4
Volumes per asynchronous mirror consistency group	32	32	64	64	64
Asynchronous mirrors per system	32	32	128	64	256
Synchronous mirrors per system	16	16	128	NA	NA

Application requirements

Use the NetApp SANtricity Unified Manager application (version 4.2 and later) for mirroring configuration and use the SANtricity System Manager application (version 11.62 and later) for synchronization management.

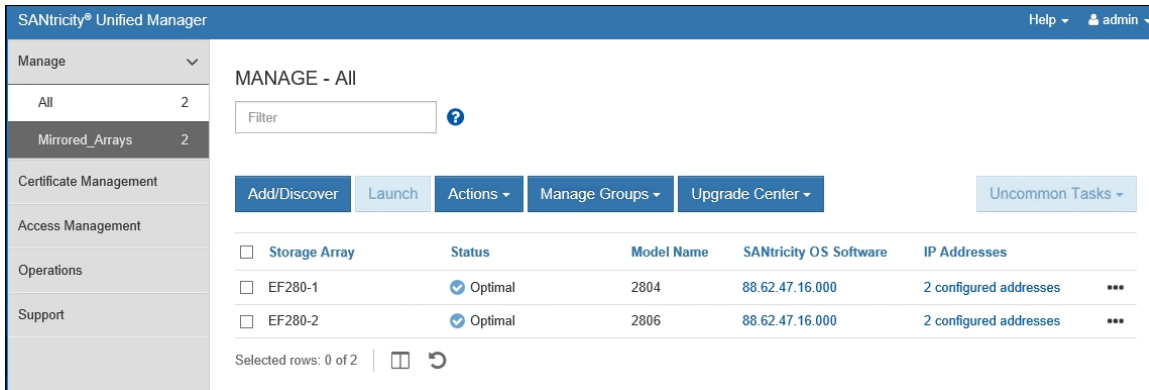
Unified Manager

Unified Manager is part of the SANtricity Web Services Proxy that is installed separately on a networked Windows or Linux host. To access Unified Manager, open a browser from the host and then enter the proxy's fully qualified domain name (FQDN) and port number, followed by /um. For example:

```
https://<WSP Server FQDN>:<port>/um
```

To then access the Unified Manager web interface, enter the administrator user name and password. In Unified Manager, before you can create mirrored pairs, you must discover the storage arrays in the network that you want to use for mirroring. See Figure 3.

Figure 3) Unified Manager application interface.



Before you configure mirrored pairs in Unified Manager, review these requirements:

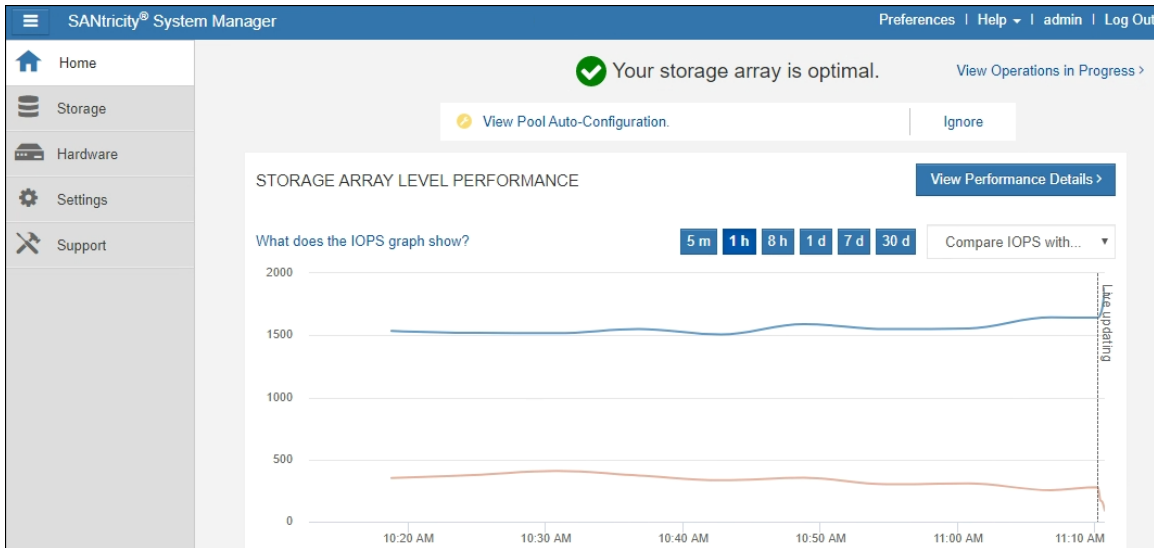
- The SANtricity Web Services Proxy and Unified Manager must be version 4.2 or later. If you have an earlier version, you must upgrade to access the mirroring configuration tasks in Unified Manager. You can download the Web Services Proxy and Unified Manager from the [NetApp Support site](#) or from [Docker Hub](#).
- The Web Services Proxy service must be running.
- Unified Manager must show valid Secure Sockets Layer (SSL) certificates for the storage array. You can accept a self-signed certificate or install your own security certificate by navigating to Certificate Management. For more information, see [TR 4813, Managing Certificates for NetApp E-Series Storage Systems](#).
- Both the local and remote storage arrays that you want to use for mirroring must be discovered in Unified Manager. You can discover a single array, or you can scan a range of IP addresses to discover multiple arrays at once by using the Add/Discover arrays wizard. For instructions, see the online help in the Unified Manager interface.

System Manager

System Manager is a storage-provisioning application that is included with the controller's operating system. To access System Manager, open a browser from a host that is connected to the controller's management port and then enter the controller's IP address or domain name. To then access the System Manager web interface, enter the administrator user name and password. See Figure 4.

In System Manager, you can manage synchronization schedules and other mirroring tasks.

Figure 4) System Manager application interface.



Mirroring and role-based access control

If you use role-based access control (RBAC) in the SANtricity applications, be aware of the following requirements:

- You must have storage administrator rights on the primary storage array.
- You must know the administrator password for the secondary system.
- The legacy management interface must be on. With the advent of RBAC support, NetApp has provided a means to turn off the native SYMbol API for enhanced security. To set up mirroring relationships, this interface must be on temporarily until the configuration is complete. You can change this setting in System Manager at Settings > System > Change Management Interface.

Connectivity requirements

Asynchronous mirroring and synchronous mirroring have different connectivity requirements:

- Asynchronous can use either FC, iSCSI, or both connection types for communication between storage arrays. However, there is no failover from one channel type to the other.
- Synchronous mirroring supports only FC for communication between storage arrays.

Note: Although controllers are limited to FC or iSCSI for mirroring, controllers can receive host I/O by using a different protocol, such as SAS.

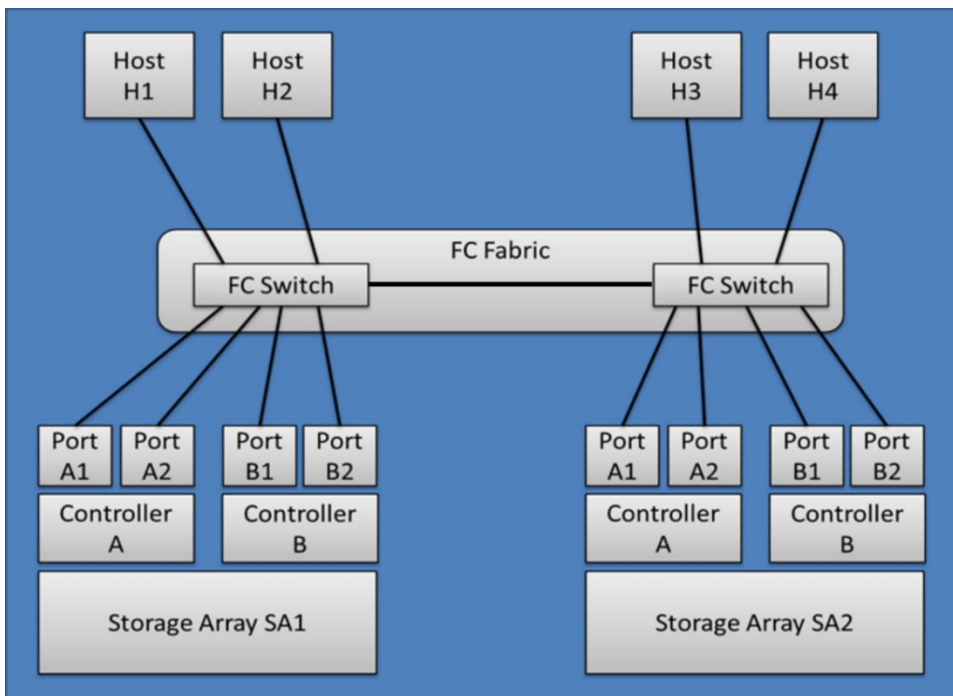
FC interface (asynchronous or synchronous)

Mirroring through an FC interface requires the following:

- Each controller of the storage array dedicates its highest-numbered FC host port to mirroring operations.
- If the controller has both base FC ports and host interface card (HIC) FC ports, the highest-numbered port is on a HIC. Any host that is logged in to the dedicated port is logged out, and no host login requests are accepted. I/O requests on this port are accepted only from controllers that are participating in mirroring operations.
- The dedicated mirroring ports must be attached to an FC fabric environment that supports the directory service and name service interfaces.

In particular, FC-AL and point-to-point are not supported as connectivity options between the controllers that are participating in mirror relationships. In the configuration that Figure 5 shows, a single fabric is used to provide total connectivity among all participating devices, including host systems and storage arrays.

Figure 5) Single-fabric configuration.



Note that certain I/O paths through the switch are never used in this configuration, despite the fact that raw connectivity does exist. The only paths that are used are the following:

For the host I/O:

- H1 to/from SA1-A1 and to/from SA1-B1
- H2 to/from SA1-A1 and to/from SA1-B1
- H3 to/from SA2-A1 and to/from SA2-B1
- H4 to/from SA2-A1 and to/from SA2-B1

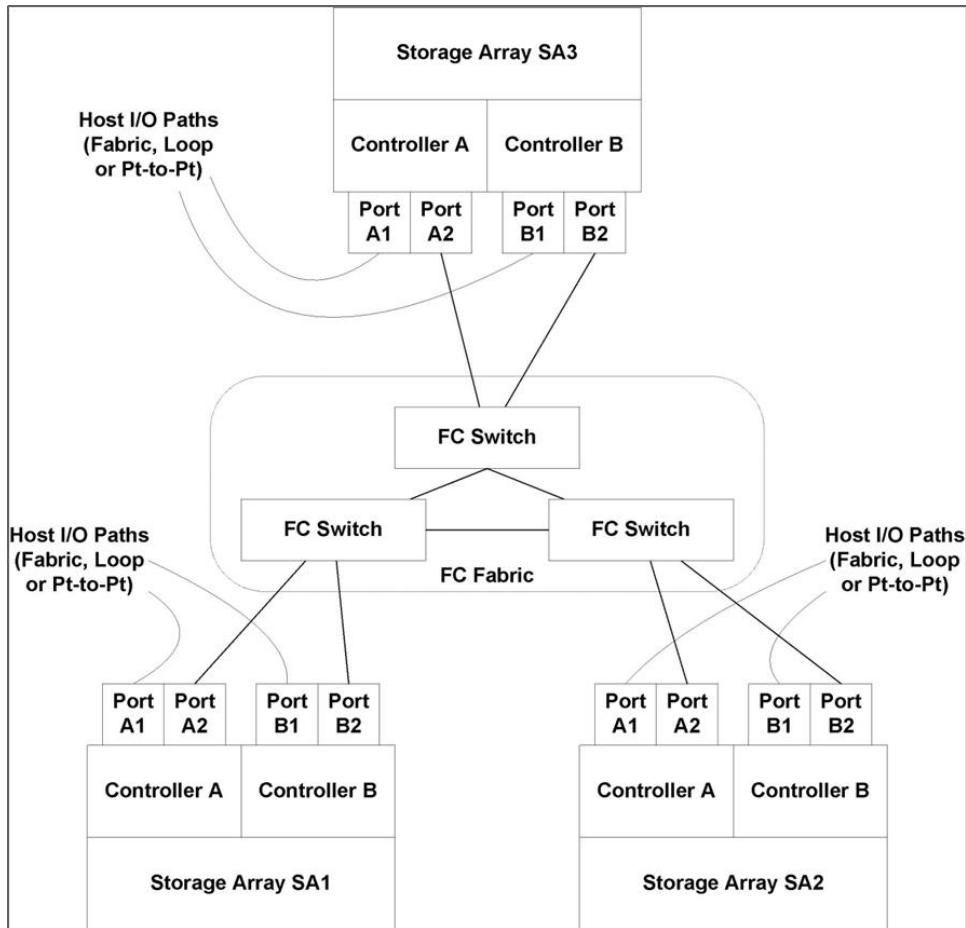
For mirror operations:

- SA1-A2 to/from SA2-A2
- SA1-B2 to/from SA2-B2

In this configuration, H1 and H2 constitute one partition of the compute environment, and H3 and H4 constitute a separate partition. These two partitions do not share any common storage devices, so SA1 is dedicated to the H1/H2 partition, and SA2 is dedicated to H3/H4. However, SA1 and SA2 use a mirroring link that causes one or more volumes to be mirrored between the two arrays. This approach allows either partition to take over operations for the other if an outage occurs, using the mirrored volume or volumes as necessary during the outage.

In a configuration with a fabric dedicated to mirroring, as in Figure 6, the host traffic connects to the storage arrays by using fabric, FC-AL, or point-to-point configurations that are completely independent of the dedicated mirroring fabric. Note that the host I/O connections can be made only to the A1 and B1 ports of each controller, because the A2 and B2 ports are dedicated to asynchronous mirroring.

Figure 6) Fabric dedicated to mirroring.



Note that fabric connectivity paths exist that are never used. The only paths that are used in the illustrated fabric for synchronous mirroring operations are the following:

- SA1-A2 to/from SA2-A2
- SA1-A2 to/from SA3-A2
- SA2-A2 to/from SA3-A2
- SA1-B2 to/from SA2-B2
- SA1-B2 to/from SA3-B2
- SA2-B2 to/from SA3-B2

iSCSI interface (asynchronous only)

The following guidelines apply to mirroring through an iSCSI interface (asynchronous only):

- Unlike FC, iSCSI does not require a dedicated port. The controller maintains a list of remote storage arrays with which the iSCSI initiator attempts to establish a session. The first port that successfully establishes an iSCSI connection is used for all subsequent communication with that remote storage array. If communication fails, a new session is attempted by using all available ports.
- You configure iSCSI ports at the array level on a port-by-port basis. Intercontroller communication for configuration messaging and data transfer uses the global settings, including settings for:
 - Virtual LAN (VLAN; both local and remote systems must have the same VLAN setting to communicate)

- iSCSI listening port
- Jumbo frames
- Ethernet priority

Note: The iSCSI intercontroller communication must use a host connect port and not the management Ethernet port.

Note: VLAN is not a requirement for configuring asynchronous replication.

In addition, asynchronous mirroring uses the storage array's host-side I/O ports to convey mirrored data from the primary side to the secondary side. Because asynchronous mirroring is intended for higher-latency, lower-cost networks, iSCSI (and thus TCP/IP-based) connections are a good fit for it. When you use asynchronous mirroring in iSCSI environments, you do not have to dedicate any of the array's front-end iSCSI ports for use with asynchronous mirroring. Those ports are shared for both asynchronous mirror traffic and host-to-array I/O connections.

Distances and speed

Synchronous mirroring is limited to a maximum of 6.2 miles (10km) in distance between arrays (measured by total cable length). See Table 6 for a summary of cable distance by type and speed. Asynchronous mirroring can go much farther by extension over the WAN, but it is subject to a maximum latency of 120ms round-trip time.

Table 6) Cable distances by cable type and speed.

Speed	Twinax	OM1	OM2	OM3	OM4	OS1
1Gbps FC	-	300m	500m	860m	-	10km
2Gbps FC	-	150m	300m	500m	-	10km
4Gbps FC	-	70m	150m	380m	400m	10km
8Gbps FC	-	21m	50m	150m	190m	10km
16Gbps FC	-	15m	35m	100m	125m	10km
32Gbps FC	-	-	20m	70m	100m	10km
10Gb Ethernet	15m	-	82m	300m	400m	10km
25Gb Ethernet	3m	-	-	70m	100m	10km

WAN bandwidths for selected carrier lines are shown in Table 7.

Table 7) Transmission rates by carrier type.

DS1	DS3	OC-1	OC-3	OC-12	OC-48	OC-192
1.544Mbps	44.736Mbps	51.48Mbps	155.52Mbps	622.08Mbps	2.4Gbps	9.6Gbps

Host connections

A controller can receive host I/O through one protocol and use a different protocol for mirroring with a remote storage array. For example, a host might be attached to the controller through a SAS connection, and the intercontroller mirroring data is sent over FC to the remote system.

Controller-to-controller interactions for mirroring I/O occur only as follows:

- Controller A on one end of a mirrored pair interacts only with controller A on the other end.
- Controller B on one end of a mirrored pair interacts only with controller B on the other end.

No mirroring interactions are attempted between controller A on one end of a mirrored pair and controller B on the other end.

Capacity and volume requirements

Mirroring operations include the capacity and volume requirements that are outlined in this section.

Mirror volume requirements

Mirrored volumes must meet the following requirements and characteristics:

- The reported capacity of a volume that participates in mirroring must be the lesser of the primary or the secondary capacity.
- The secondary volume capacity must be at least as large as the primary volume at creation time; however, after a role reversal, the new primary could be smaller than the new secondary.
- A volume is allowed to participate in only one mirroring relationship, whether asynchronous or synchronous. There is no support for multiple data center mirroring.
- Pool and volume group, RAID level, caching parameters, and segment size can be different on the two mirrored volumes.
- If the primary is configured to enable data assurance, the secondary must also have data assurance enabled.
- Drive types should be the same on primary and secondary, either both NVMe drives or both non-NVMe drives. Mixing of 4K NVMe and SAS based SSD/HDD drive groups between primary and secondary is not supported.
- Regarding drive security, a best practice is for the primary and secondary volumes to be of the same drive type (full disk encryption, FIPS, secure-capable). They should also have the same security setting (secure enabled or not). Each type of mirroring has the following drive security requirements:
 - **Synchronous mirroring.** The controller does not enforce any drive security rules with respect to the synchronous mirroring feature.
 - **Asynchronous mirroring.** If the primary is not secure-capable, the secondary must also not be secure-capable. If the primary is secure-enabled, the secondary must also be secure-enabled. If the primary is secure-capable (but not enabled), the secondary can be either secure-capable or secure-enabled. If the primary ever attains a higher security level than the secondary, the controller raises a Needs Attention condition. The controller allows a mix of FDE and FIPS volumes as mirrored pairs, but as noted earlier, a best practice is for the drive types to match.

Table 8 lists the allowed and restricted candidates for a mirror volume.

Table 8) Candidates for mirror volumes.

Mirror volume type	Allowed candidates	Disallowed candidates
Asynchronous primary	<ul style="list-style-type: none"> • Standard volume (only if the secondary is also standard) • NetApp Snapshot™ base volume • Volume copy source • Thin volume (only if the secondary is also thin) 	<ul style="list-style-type: none"> • Asynchronous or synchronous mirror primary • Asynchronous or synchronous mirror secondary • Volume copy target • Snapshot volume
Asynchronous secondary	<ul style="list-style-type: none"> • Standard volume (only if the primary is also standard) • Thin volume (only if the primary is also thin) 	<ul style="list-style-type: none"> • Asynchronous or synchronous mirror primary • Asynchronous or synchronous mirror secondary • Volume copy source

Mirror volume type	Allowed candidates	Disallowed candidates
	<p>Note: A secondary can become a volume copy source if a role reversal occurs after a volume copy on the original primary is complete. If the role reversal occurs during a copy in progress, the copy fails and cannot be restarted.</p>	<ul style="list-style-type: none"> Volume copy target Snapshot base volume Snapshot volume
Synchronous primary	<ul style="list-style-type: none"> Standard volume Snapshot base volume Volume copy source Volume copy target 	<ul style="list-style-type: none"> Asynchronous or synchronous mirror primary Asynchronous or synchronous mirror secondary Snapshot volume Thin volume
Synchronous secondary	<ul style="list-style-type: none"> Standard volume <p>Note: A secondary can become a volume copy source or target if a role reversal occurs after a volume copy on the original primary is complete. If the role reversal occurs during a copy in progress, the copy fails and cannot be restarted.</p>	<ul style="list-style-type: none"> Asynchronous or synchronous mirror primary Asynchronous or synchronous mirror secondary Volume copy source Volume copy target Snapshot base volume Snapshot volume Thin volume
Reserved capacity (repository)	<ul style="list-style-type: none"> Standard volume 	<ul style="list-style-type: none"> Asynchronous or synchronous mirror primary Asynchronous or synchronous mirror secondary Volume copy source Volume copy target Snapshot base volume Snapshot volume Thin volume

Note: A secondary volume can become a Snapshot base volume after the mirroring relationship is established.

Free capacity

For mirroring, free capacity must meet the following requirements:

- Enough free capacity must be available on the remote storage array to create a secondary volume at least as large as the primary volume.
- Enough free capacity must be available for the reserved capacity (repository) volumes, as follows:
 - For asynchronous mirroring, this capacity is one volume for each primary volume on the local system and one volume for each secondary volume on the remote array. The default for the reserved capacity volumes is 20% of the respective mirror volume, but the default can be changed.
 - For synchronous mirroring, this capacity is one reserved capacity volume for each controller on the local and remote systems regardless of how many mirrored pairs are created. The size of each reserved capacity volume is 128MiB if it is created on a volume group and 4GiB if it is created on a pool.

Reserved capacity

Reserved capacity is required for logging write information to recover from controller resets and other temporary interruptions. You must ensure that there is enough free capacity available on both storage arrays that participate in the mirror relationship.

For asynchronous mirroring, reserved capacity volumes are created when a mirrored pair is created, one for the primary volume and one for the secondary volume. The characteristics of the reserved capacity volumes are as follows:

- The two volumes are not required to be of the same capacity.
- A reserved capacity volume can be on a different pool or different volume group with a different RAID level from the associated primary or secondary volume.
- Reserved capacity volumes must have the same data assurance attribute as their associated mirror volumes.
- Mixing of 4K NVMe and SAS-based SSD/HDD drive groups between base and repository is not supported.
- If a mirror volume has security enabled, its associated reserved capacity volume must also have security enabled. The type of security should be the same, also (FDE or FIPS), but this requirement is not enforced by the controller.
- Mixing of SSD and HDD drive groups between a primary or secondary volume and its repository volume is supported as long as both drive groups share the same DA and security attributes.

For synchronous mirroring, one reserved capacity volume is created for each controller. The volumes can be on a pool or a volume group of any RAID level except 0.

Connectivity and volume ownership

The controller that owns the primary volume determines the current owner of the secondary volume. The primary volume and the secondary volume in a mirrored pair use the following ownership rules:

- If the primary volume is owned by controller A on the primary side, the secondary volume is owned by controller A on the secondary side.
- If the primary volume is owned by controller B on the primary side, the secondary volume is owned by controller B on the secondary side.
- If primary controller A cannot communicate with secondary controller A, controller ownership changes do not take place. A primary controller attempts to communicate only with its matching controller in the secondary synchronous mirror relationship.

The next remote write that is processed automatically triggers a matching ownership change on the secondary side if one of these conditions exists:

- An I/O path error causes a volume ownership change on the primary side.
- The storage administrator changes the current owner of the primary volume.

For example, a primary volume is owned by controller A and the controller owner is changed to controller B. In this case, the next remote write changes the controller owner of the secondary volume from controller A to controller B. Because controller ownership changes on the secondary side are controlled by the primary side, they do not require any special intervention.

Asynchronous mirroring configuration and management

This section describes some additional considerations for asynchronous mirroring, configuration instructions, and management tasks.

Determine synchronization times and performance impacts

Before you configure asynchronous mirroring, review the following sections for additional recommendations on calculating synchronization times and balancing synchronization with other factors.


Balance recovery point objectives with synchronization

The recovery point objective (RPO) for asynchronous mirroring is the maximum tolerable time during which changes that are made on the primary side are subject to loss if a primary-side disaster occurs. For a single interval, the potential loss period is from the beginning of the synchronization interval to when the secondary contains a point-in-time image that matches the contents of the primary at the beginning of the next interval.

Note the blue line in Figure 7. At time T1, the primary takes a point-in-time image and begins resynchronizing the secondary. From then on, new writes to the primary site are not reflected at the secondary site. Thus, they are not protected from a disaster at the primary site until T2 plus the time to resynchronize that interval. The RPO is the maximum allowable for all of these out-of-synchronization time periods.

Note: In Figure 7, PiT refers to “point-in-time.”

Figure 7) Time from one protection point to the next.



Several factors affect the ability to meet the RPO. These factors include the synchronization interval, the size of the dataset and the rate of change, and the bandwidth of the communications between the primary and secondary sites.

To determine how to set the synchronization interval, and eventually the required bandwidth, it is useful to begin by analyzing the relationship between resynchronization time and RPO by using the following values:

- **SI:** synchronization interval
- **MRT:** maximum resynchronization time
- **MRTC:** maximum resynchronization time with contingency

The relationship between synchronization interval and maximum resynchronization time can be expressed as follows:

$$SI + MRT \leq RPO$$

Also:

$$SI \geq MRT \text{ and } SI \leq 10 \text{ minutes}$$

Selecting values is a balancing act. If $SI = MRT$, then the system spends significant time resynchronizing, which affects the performance of the primary. The system also runs the risk that resynchronization will not be complete by the next interval, thus affecting RPO achievement. If for very short synchronization intervals, $SI \gg MRT$, then the bandwidth of the communications link must be greater, leading to potential unnecessary expense. This situation leads to two general rules:

- Use a contingency on the expected MRT in case something unexpected happens, such as an increase in network latencies. A good starting point is a 50% contingency, or $MRTC = 1.5MRT$.
- $RPO = 2 \times MRTC$.

These rules provide contingency, while balancing the cost of inter-array communications components. You can select other values for the contingency if you prefer. For example, select smaller values if the amount of change during peak periods is well known, or if a bit more risk in RPO achievement is acceptable.

Taking the contingency into account, the synchronization interval is:

$$SI = RPO - MRTC = MRTC$$

For example, if the RPO is 60 minutes, then the maximum resynchronization time (with contingency) would be 30 minutes. The maximum expected actual resynchronization time is 20 minutes, and the synchronization interval would be set to 30 minutes.

Size the network to meet RPO

Before you configure mirroring, make sure that the communications link can handle the bandwidth that is required to resynchronize the maximum expected changes within the synchronization interval. Bandwidth is limited to the slowest segment in the network. Also, latencies can reduce the maximum achievable bandwidth on IP networks.

After you determine the maximum actual resynchronization time, you can calculate the maximum required bandwidth if you can estimate how much data will change during any given synchronization interval. One way to arrive at this value is to estimate the portion of the total dataset that is changed on a peak day or in peak hour. For example, if the dataset is 20TB, and 10% of that data changes on a peak day, then 2TB must be transferred to the secondary site in a day. If there is a peak hour during the peak day that gets much of the activity, then you must take that extra traffic into account. For example, if 15% of the activity on the peak day occurs during an hour, then 300GB would be written in that hour.

In the previous example, the network must be sized so that all changes over a 30-minute period can be transferred in 20 minutes to the secondary site. If 300GB must be written in the peak hour, and it is approximately uniform during that hour, then 150GB would be the maximum that must be transferred to the secondary site in 20 minutes to meet the RPO. In this example, the network must have a bandwidth of 125MBps, or about 1.25Gbps.

To determine the required network configuration for asynchronous mirroring, consider the following:

- Maximum round-trip time latency for asynchronous mirroring is 120ms.
- Datasets grow over time, and growth likely requires more data to be written to the secondary site during each synchronization interval.
- With iSCSI or FC extensions over the WAN, increases in overhead can reduce the effective bandwidth by up to 50%.
- NetApp recommends a minimum iSCSI speed of 10Gbps when you configure asynchronous mirroring.

- The network that you use for mirroring should not be shared with other traffic, which can cause variable bandwidth and prevent adherence to the RPO.

Determine latency issues

Latency refers to any delay that occurs as data is read or written. When you use iSCSI for the communications link, latency is an important factor in the system's ability to support the required RPO. Although the rated bandwidth of the link might be high, the achievable bandwidth could be limited by latency and the TCP window size. The TCP window size is the maximum number of bytes that can be sent without receiving an acknowledgment. A large round-trip time slows down synchronizations and resynchronizations. When you use iSCSI for asynchronous mirroring, you must take this potential limitation into account.

Note: The TCP window size was increased in SANtricity 11.50 from 256KiB to 512KiB and from two simultaneous connections between the local and remote array to eight simultaneous connections.

Current NetApp E-Series systems have a TCP window size of 512KiB. Maximum achievable bandwidth per connection can be calculated by (TCP window size in bits) / (latency in seconds) = throughput in bits per second.

As an example of maximum bandwidth calculation, assume that you have a primary site in Chicago and a secondary site in Salt Lake City. These cities are about 1,400 miles (2253km) apart; at the speed of light, communication can travel the distance in 7.5ms. This situation has a round-trip time of 15ms, not counting any other delays through the actual media and components. If these extra latencies add up to 5ms round trip, then the communications link has a round-trip latency of 20ms. By using these values, you can calculate the maximum bandwidth that can be achieved by asynchronous mirroring over iSCSI as follows:

- TCP window size in bits = 512 x 1024 x 8 = 4,194,304 bits
- Latency in seconds = 20 / 1,000 = 0.02 seconds
- Maximum bandwidth per connection = 4,194,304 bits / .02 seconds = ~210Mbps
- Maximum bandwidth for all connections = 210Mbps x 8 = ~ 1677Mbps per controller

Note: The owning controller is the only one to synchronize data. Therefore, if your bandwidth is constrained, make sure the volumes are distributed across both controllers. In this example, the maximum bandwidth for the array would then be 3355Mbps.

Determine the time to complete initial or full synchronization

When you configure mirroring, you must know how long it takes for normal operations to begin with a valid recovery point at the secondary after post-disaster creation of a mirrored pair or re-creation the primary. Depending on the pipe size and how large the dataset is, this approach could take a very long time. Table 9 shows some examples using the theoretical rate of various WAN links.

Note: Various overheads will reduce these rates. For iSCSI, as described in the section above, the TCP Window size can prevent the link from being fully utilized depending on the link latency.

Table 9) Approximate hours needed for initial or full synchronization by pipe size.

Data to Copy	DS1 (1.544Mbps)	DS3 (44.736Mbps)	OC-1 (51.48Mbps)	OC-3 (155.52Mbps)	OC-12 (622.08Mbps)	OC-48 (2.4Gbps)	OC-192 (9.6Gbps)	10Gb Ethernet
100GB	143	4.97	4.32	1.43	0.357	0.0926	0.0231	0.0222
1TB	1,430	49.7	43.2	14.3	3.57	0.926	0.231	0.222
10TB	14,300	497	432	143	35.7	9.26	2.31	2.22

Determine the effects of asynchronous mirroring on host I/O performance

Performance varies based on workload, the number and type of drives, and system configuration. The following tables provide helpful information about the performance effects when you deploy mirroring operations. Table 10 summarizes a configuration that is used in performance testing to simulate a typical database.

Table 10) System configuration for performance example.

Workload Profile	Number of Drives	Type of Volume Container
8K; 75% read; 85% random; 8 outstanding I/Os	24 10K RPM SAS	SANtricity Dynamic Disk Pools

Table 11 shows the performance results from the preceding example test.

Note: Actual results might vary greatly from this example test and cannot be guaranteed.

Table 11) Example performance results.

Copy Services in Use	I/O Performance (IOPS)	Bandwidth Performance (MBps)	Performance Effect	Notes
None: baseline	1,071	8.37	n/a	–
Snapshot copy only	1,052	8.24	–1.7%	Initial performance drops ~50% while the Snapshot reserved capacity volume is being initialized. The performance shown is after initialization has been completed. During initialization, CPU utilization for this test was ~55%, which dropped to ~2% after completion. With Snapshot copies enabled, CPU utilization during the lometer run was ~21%. Without Snapshot copies, CPU utilization was ~18%.
Mirroring only	1,052	8.24	–1.7%	Initial performance drops ~50% while the reserved capacity volume is being initialized. Performance shown is after initialization has been completed. During initialization, CPU utilization for this test on the primary site was ~53%, which dropped to ~2% after completion. CPU utilization on the remote site was ~12%, which dropped to ~2% after completion. With mirroring enabled, CPU utilization during the lometer run was ~21%. Without Snapshot copies, CPU utilization was ~18%. If lometer is running during initial synchronization, CPU utilization is ~63%, and performance drops ~60%.
Snapshot copy and mirroring	1,052	8.24	–1.7%	As with the other testing, after initialization is complete, performance is minimally affected.

When a point-in-time image is taken, the system must initialize the reserved capacity volume. This situation takes some cycles away from the disks, but CPU overhead that is associated with this operation is minimal.

Configure asynchronous mirroring

Configuration involves the following steps:

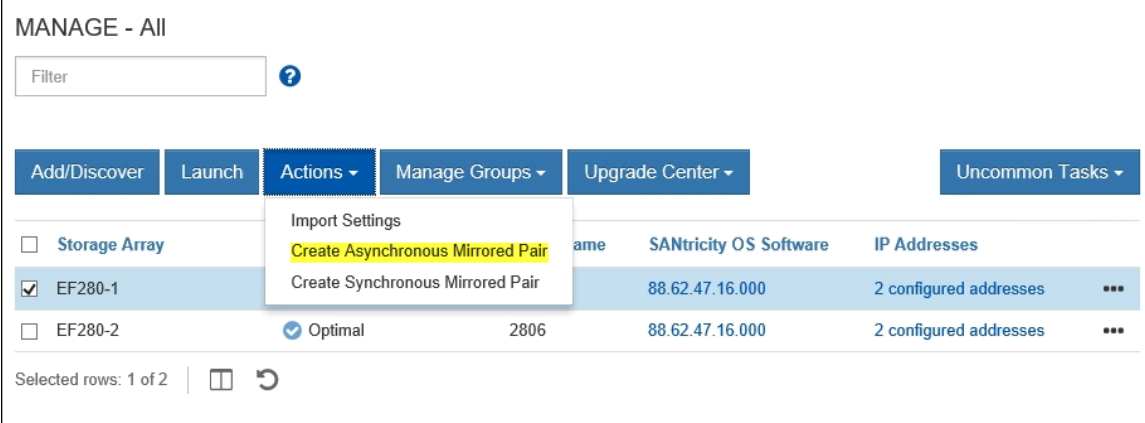
- In Unified Manager, create mirrored pairs by selecting the local and remote storage arrays, creating a consistency group, and then selecting the volumes.
- When initial synchronization begins, check the progress of the synchronization by opening System Manager for the local storage array.

Create mirrored pairs

To configure asynchronous mirroring, you first create a consistency group to hold the mirrored pairs. Next you select a primary volume from the local storage array and a secondary volume from the remote storage array for the mirrored pair.

To configure mirrored pairs, do the following:

1. Open Unified Manager.
2. From the Manage page, select the local storage array that you want to use for the primary site.
Note: Only discovered arrays appear in this window. To discover arrays, select Add/Discover from the Manage panel to open a discovery wizard.
3. Select Actions > Create Asynchronous Mirrored Pair.



The screenshot shows the 'MANAGE - All' interface. At the top, there is a 'Filter' input field with a help icon. Below this are several action buttons: 'Add/Discover', 'Launch', 'Actions', 'Manage Groups', 'Upgrade Center', and 'Uncommon Tasks'. The 'Actions' menu is open, showing three options: 'Import Settings', 'Create Asynchronous Mirrored Pair' (which is highlighted in yellow), and 'Create Synchronous Mirrored Pair'. Below the menu is a table with columns for 'Storage Array', 'Name', 'SANtricity OS Software', and 'IP Addresses'. The table contains two rows: 'EF280-1' and 'EF280-2'. The 'EF280-1' row is selected, and its 'Optimal' status is indicated by a checkmark. The 'Selected rows: 1 of 2' indicator is visible at the bottom left of the table area.

The Create Asynchronous Mirrored Pair wizard opens.

4. In the Select/Create Group panel, select either an existing group or the option to create a new one.
To create a new group, do the following:
 - a. Select A New Mirror Consistency Group, then click Next.

Note: If you previously created a group and want to use that group, select An Existing Mirror Consistency Group from this panel and click the group name from the list.

Create Asynchronous Mirrored Pair
✕

1 Select/Create Group
2 Select Primary Volume
3 Select Secondary Volume

What do I need to know before creating a mirrored pair?

Create a mirrored pair using...

An existing mirror consistency group

Name	Remote storage array
Reyes_Mirror_Group	EF280-2

A new mirror consistency group

- b. Under Mirror Consistency Group Name, enter a unique name that best describes the data on the volumes that will be mirrored between the two storage arrays. A name can consist only of letters, numbers, and the special characters underscore (_), dash (-), and the hash sign (#). A name cannot exceed 30 characters and cannot contain spaces.

Create Asynchronous Mirrored Pair
✕

1 Select/Create Group
2 Select Primary Volume
3 Select Secondary Volume

Step 1b - Create Mirror Consistency Group

What do I need to know before creating a mirror consistency group?

Mirror consistency group name ?

Gallagher_Mirror_Group ✕

Remote storage array

Select

Synchronization settings

Synchronize the mirrored pairs...

Manually

Automatically, every - 10 + Minutes ▼

Alert settings

Alert me...

When my reserved capacity threshold reaches - 80 % + ?

- c. Under Remote Storage Array, select the secondary storage array on which you want to establish a mirror relationship with the local storage array.

Storage arrays are listed by their storage array name. If you have not named a storage array, it will be listed as Unnamed.

Note: If your remote storage array is password protected, the system prompts for a password.

Create Asynchronous Mirrored Pair

1 Select/Create Group 2 Select Primary Volume 3 Select Secondary Volume

Step 1b - Create Mirror Consistency Group

What do I need to know before creating a mirror consistency group?

Mirror consistency group name ?

Gallagher_Mirror_Group

Remote storage array

Select

EF280-2 (Fibre Channel)

Synchronize the mirrored pairs...

Manually

Automatically, every

Alert settings

Alert me...

When my reserved capacity threshold reaches %

- d. Under Synchronize the Mirrored Pairs..., select whether you want to synchronize the mirrored pairs manually or automatically:

Manually—Select this option to manually start synchronization for all mirrored pairs in this group. When you want to perform a resynchronization later, launch System Manager for the primary storage array, go to Storage > Asynchronous Mirroring, select the group from the Mirror Consistency Groups tab, and then select More > Manually resynchronize.

Automatically—Select the desired interval in minutes, hours, or days, from the beginning of the previous update to the beginning of the next update. For example, if the synchronization interval is set at 30 minutes and the synchronization process starts at 4 p.m., the next process starts at 4:30 p.m., and so on.
 - e. Under Alert Settings, select the alert settings:

For manual synchronizations, specify the threshold (defined by the percentage of the capacity remaining) for when you receive alerts.

For automatic synchronizations, you can set three methods of alerting:

 - When the synchronization has not been completed in a specific length of time
 - When the recovery point data on the remote array is older than a specific time limit
 - When the reserved capacity is nearing a specific threshold (defined by the percentage of the capacity remaining)
5. Click Next to create the group (if you defined a new one), and then move to the Select Primary Volume panel.

If you defined a new mirror consistency group, Unified Manager creates the mirror consistency group on the local storage array first and then creates the mirror consistency group on the remote storage array.

Note: If Unified Manager successfully creates the mirror consistency group on the local storage array but fails to create it on the remote storage array, it automatically deletes the mirror consistency group from the local storage array. If an error occurs while Unified Manager is attempting to delete the mirror consistency group, you must manually delete it.

- In the Select Primary Volume panel, select a volume that you want to use as the primary volume. The system displays a list of all eligible volumes for this mirrored pair. Any volumes that are not eligible do not display in the list.

Create Asynchronous Mirrored Pair
✕

1 Select/Create Group
2 Select Primary Volume
3 Select Secondary Volume

Why don't I see all my volumes?

Select your primary volume from EF280-1...

Name	Capacity (GiB)
VMware1	1024.00
VMware2	1024.00

Total rows: 2

- Click Next to allocate the reserved capacity.
- From the list of eligible candidates, select the reserved capacity for the primary volume.

Create Asynchronous Mirrored Pair
✕

1 Select/Create Group
2 Select Primary Volume
3 Select Secondary Volume

Step 2b - Reserve Capacity

I want to reserve capacity for the mirrored pair that is... ?

-

20

%

+

of the primary volume capacity (1024.00 GiB) = 204.80 GiB.

Why would I change this percentage?

Refresh Candidates
?

Capacity (GiB)	Pool/ Volume Group	RAID Level
208.00	Disk_Pool_1	Pool

Total rows: 1

When you are selecting reserved capacity, keep the following guidelines in mind:

- The default setting for reserved capacity is 20% of the capacity of the base volume, and usually this capacity is sufficient. If you change the percentage, click Refresh Candidates.
- The capacity that you need varies, depending on the frequency and size of I/O writes to the primary volume and how long you need to keep the capacity.
- In general, select a larger capacity for reserved capacity if one or both of these conditions exist:
 - You intend to keep the mirrored pair for a long time.

- A large percentage of data blocks will change on the primary volume because of heavy I/O activity. To help you determine typical I/O activity to the primary volume, use historical performance data or other operating system utilities.
- Click Next to move to the Select Secondary Volume panel.
 - Select a volume that you want to use as the secondary volume in the mirrored pair.
The system displays a list of all eligible volumes for this mirrored pair. Any volumes that are not eligible do not display in the list.

Create Asynchronous Mirrored Pair
✕

1 Select/Create Group
2 Select Primary Volume
3 Select Secondary Volume

Primary volume attributes

Reported capacity: 1024.00 GiB
 Data Assurance (DA) enabled: Yes
 Secure-capable: Yes - FDE (🔒 not enabled)
 Volume type: Thick

[Why don't I see all the volumes on the remote storage array?](#)

Filter

?

Select your secondary volume from EF280-2...

Name	Capacity (GiB)
Test1	1024.00

Total rows: 1

- Click Next to select reserved capacity.
- From the list of eligible candidates, select the reserved capacity for the secondary volume.

Create Asynchronous Mirrored Pair
✕

1 Select/Create Group
2 Select Primary Volume
3 Select Secondary Volume

Step 3b - Reserve Capacity

I want to reserve capacity for the mirrored pair that is... ?

-

20

%

+

of the secondary volume capacity (1024.00 GiB) = 204.80 GiB.

[Why would I change this percentage?](#)

Refresh Candidates

?

Capacity (GiB)	Pool/ Volume Group	RAID Level
208.00	POOL_1	Pool

Total rows: 1

When you are selecting reserved capacity, keep the following guidelines in mind:

- The default setting for reserved capacity is 20% of the capacity of the base volume, and usually this capacity is sufficient. If you change the percentage, click Refresh Candidates.

- The capacity that you needed varies, depending on the frequency and size of I/O writes to the primary volume and how long you need to keep the capacity.
 - In general, select a larger capacity for reserved capacity if one or both of these conditions exist:
 - You intend to keep the mirrored pair for a long time.
 - A large percentage of data blocks will change on the primary volume because of heavy I/O activity. To help you determine typical I/O activity to the primary volume, use historical performance data or other operating system utilities.
13. Select Finish to complete the asynchronous mirroring sequence.

Asynchronous mirroring is now set. The system performs the following actions:

- It creates the mirrored pair and the reserved capacity. When that step is finished, Unified Manager displays Mirrored Pair Was Successfully Created.
- It begins initial synchronization between the local storage array and the remote storage array. For more information, see “Initial synchronization process for asynchronous mirroring.”

Note: If the volume that is being mirrored is a thin volume, only the provisioned blocks (allocated capacity rather than reported capacity) are transferred to the secondary volume during the initial synchronization. This approach reduces the amount of data that must be transferred to complete the initial synchronization.

Check the progress of the initial synchronization

To check the progress of the initial synchronization, do the following:

- In Unified Manager, from the Manage page, select the local storage array and click Launch to open System Manager.

The screenshot shows the 'MANAGE - All' interface. At the top, there is a 'Filter' input field and a help icon. Below this are several action buttons: 'Add/Discover', 'Launch' (highlighted with a yellow box and a hand cursor), 'Actions', 'Manage Groups', and 'Upgrade Center'. A table below these buttons lists storage arrays. The table has columns for 'Storage Array', 'Status', 'Model Name', and 'SANtricity OS Software'. Two rows are visible: 'EF280-1' and 'EF280-2', both with a status of 'Optimal'. At the bottom of the table, it says 'Selected rows: 1 of 2' and includes icons for table manipulation.

Storage Array	Status	Model Name	SANtricity OS Software
<input checked="" type="checkbox"/> EF280-1	Optimal	2804	88.62.47.16.000
<input type="checkbox"/> EF280-2	Optimal	2806	88.62.47.16.000

- In System Manager, view the status of the mirroring operation. Go to Storage > Asynchronous Mirroring > Mirror Consistency Groups.

ASYNCHRONOUS MIRRORING

[Learn More >](#)

Mirror Consistency Groups | Mirrored Pairs

Filter ?

Create Mirrored pair | Test Communication | More ▾

Name	Status	Number of Mirrored Pairs	Local Role	Remote Storage Array	Connection Type
Reyes_Mirror_Group	Optimal	2	Primary	EF280-2	Fibre Channel
Gallagher_Mirror_Group	Synchronizing	1	Primary	EF280-2	Fibre Channel

Total rows: 2 |

- Confirm that mirroring is complete. When mirroring is complete, the status of the mirrored pair is Optimal.

Home / Storage / Asynchronous Mirroring

ASYNCHRONOUS MIRRORING

[Learn More >](#)

Mirror Consistency Groups | Mirrored Pairs

Filter ?

Create Mirrored pair | Test Communication | More ▾

Name	Status	Number of Mirrored Pairs	Local Role	Remote Storage Array	Connection Type
Reyes_Mirror_Group	Optimal	2	Primary	EF280-2	Fibre Channel
Gallagher_Mirror_Group	Optimal	1	Primary	EF280-2	Fibre Channel

Total rows: 2 |

Test the asynchronous communications link

If you need to diagnose communication problems between the local storage array and the remote storage array that is associated with a mirror consistency group, you can test the communications link.

Asynchronous mirroring has four tests:

- **Connectivity** verifies that the two controllers have a communication path. The connectivity test sends an intercontroller message between the storage arrays and then validates that the corresponding mirror group on the remote storage array exists. It also validates that the member volumes of the group on the remote array match the member volumes of the group on the local array.
- **Latency** sends a SCSI test unit command to each mirrored volume on the remote storage array that is associated with the mirror group to test the minimum, average, and maximum latency.

- **Bandwidth** sends two intercontroller messages to the remote storage array to test the minimum, average, and maximum bandwidth, as well as the negotiated link speed of the port on the controller that is performing the test.
- **Port connections** shows the port that is being used for mirroring on the local storage array and the port that is receiving the mirrored data on the remote storage array.

To test the communications link, do the following:

1. In System Manager, select Storage > Asynchronous Mirroring.
2. Select the Mirror Consistency Groups tab and then select the mirror consistency group that you want to test.
3. Select Test Communication.

Home / Storage / Asynchronous Mirroring

ASYNCHRONOUS MIRRORING

[Learn More >](#)

Mirror Consistency Groups | Mirrored Pairs

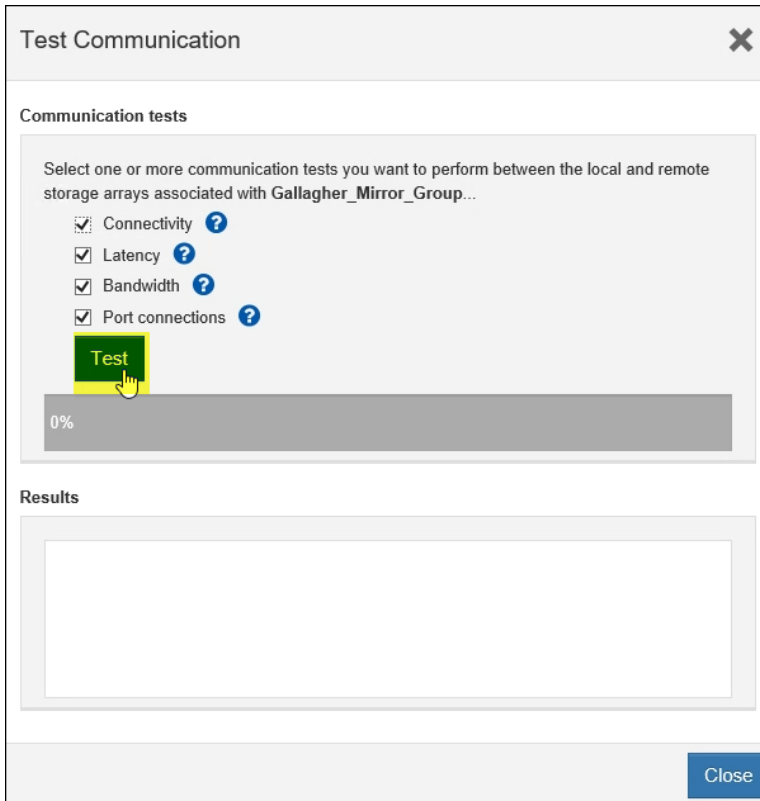
Filter ?

Create Mirrored pair | **Test Communication** | More >

Name	Status	Number of Mirrored Pairs	Local Role	Remote Storage Array	Connection Type
Reyes_Mirror_Group	Optimal	2	Primary	EF280-2	Fibre Channel
Gallagher_Mirror_Group	Optimal	1	Primary	EF280-2	Fibre Channel

Total rows: 2 | ↻

4. In the Test Communication dialog box, select one or more communication tests to perform between the local and the remote storage arrays that are associated with the selected mirror consistency group. Then click Test.



5. Review the information displayed in the Results window.
 - Normal. The mirror group is communicating correctly.
 - Passed. Check possible network or connection problems and retry the test.
 - Failed. The reason is indicated. To correct the problem, refer to the Recovery Guru.
 - Port connection error. This error might be because the local array is not connected or because the remote array cannot be contacted. To correct the problem, refer to the Recovery Guru.

Manage asynchronous mirroring

You have several options for managing the mirroring process in System Manager, including:

- Suspending or resuming the synchronization for a mirror consistency group.
- Changing the role assignments (primary and secondary) of the mirrored pairs in a mirror consistency group.

Suspend or resume synchronization

You can suspend synchronization on a mirror consistency group basis. With this suspension, the controller immediately halts any synchronization or resynchronization activity on that group. No attempt is made to contact the secondary volumes for that group. A recovery point on the secondary volumes remains valid, and no alerts are issued for the age of the recovery point. While the group is suspended, writes to the primary volumes continue as normal and are logged.

When you resume synchronization for the mirror consistency group, a resynchronization of all mirrored pairs in the group begins immediately. After the resynchronization is complete, the group resumes the normal periodic resynchronization schedule if the group is configured for periodic resynchronization.

To suspend or resume mirroring, do the following:

1. Open System Manager.
2. Go to the Storage > Asynchronous Mirroring > Mirror Consistency Groups tab. The Mirror Consistency Group table displays all the mirror consistency groups that are associated with the storage array.
3. Select the mirror consistency group, select More, and then select either Suspend or Resume. Note that if synchronization is currently set, Resume is unavailable from the More menu. If synchronization is currently suspended, Suspend is unavailable.

ASYNCHRONOUS MIRRORING

[Learn More >](#)

Mirror Consistency Groups | Mirrored Pairs

Filter ?

Create Mirrored pair | Test Communication | More ▾

Name	Status	Remote Storage Array	Connection Type
Reyes_Mirror_Group	Optimal	N/A	Fibre Channel
Gallagher_Mirror_Group	Optimal	N/A	Fibre Channel

Total rows: 2 | ↻

More ▾

- Suspend
- Resume
- Manually resynchronize
- Change role
- Edit settings
- Update remote IP address
- View unsynchronized data amount

Perform an orderly role reversal

At certain times, you might want to reverse the roles of the primary and secondary sites. Examples are if you anticipate a natural disaster such as a hurricane that is headed for the primary site and if you want to perform scheduled maintenance on the primary site. In these cases, you can reverse roles to preserve continuity of service. Later, you can reverse the roles again to restore the original site to normal operations.

The role-reversal change affects all asynchronous mirrored pairs in the selected mirror consistency group. For example, when a primary group is demoted to a secondary role, all the primary volumes of the asynchronous mirrored pairs in that group are also demoted to secondary volumes.

When demoting a primary group to a secondary role, if the current secondary group can be contacted, it is automatically promoted to a primary role in the mirror relationship. Likewise, when promoting a secondary group to a primary role, if the current primary group can be contacted, it is automatically demoted to a secondary role in the mirror relationship.

Note: If the environment has a suspended asynchronous mirror group operation, it resumes during the role-change operation.

Before you reverse roles, keep these guidelines in mind:

- When the primary group becomes a secondary, hosts that have been mapped to the mirrored volumes in the group no longer have write access to them.
- When the secondary group becomes a primary, any hosts that are accessing the secondary volumes in the group can then write to the asynchronous mirrored pairs.

- If a communication problem between the local and remote sites prevents the promotion of the secondary group, an error message appears. However, you can still force the secondary group to a primary role. This forced promotion leads to a dual primary asynchronous mirroring condition.

To change role assignments, do the following:

1. Open System Manager.
2. Go to the Storage > Asynchronous Mirroring > Mirror Consistency Groups tab. The Mirror Consistency Group table displays all the mirror consistency groups that are associated with the storage array.
3. Select the mirror consistency group for which you want to change the role and then select More > Change Role.

ASYNCHRONOUS MIRRORING

[Learn More >](#)

Mirror Consistency Groups | Mirrored Pairs

Filter ?

Create Mirrored pair | Test Communication | More ▾

Name	Status	Remote Storage Array	Connection Type
Reyes_Mirror_Group	Optimal	N/A	Fibre Channel
Gallagher_Mirror_Group	Optimal	N/A	Fibre Channel

Total rows: 2 | ↻

More ▾ menu items: Suspend, Resume, Manually resynchronize, **Change role**, Edit settings, Update remote IP address, View unsynchronized data amount

4. Confirm the operation in the Change Role dialog box.

Change Role to Secondary ✕

The Selected Mirror Consistency Group Will Be Demoted

Note the following before confirming this operation:

- If the mirror group is currently in a suspended state, it will be resumed
- Hosts assigned to volumes in the newly promoted group will now have Write access to those volumes
- Hosts assigned to volumes in the newly demoted group will no longer have Write access to those volumes

Important: It is highly recommended that you elect to resynchronize your mirror consistency group as a part of this operation. However, resynchronization may take a long time. If you choose not to resynchronize your mirror consistency group, you will lose all data written to the primary volumes since the last synchronization.

Resynchronize mirror consistency group (recommended)

Type CHANGE to confirm that you want to perform this operation.

Type change

Change Role | Cancel

5. Check the Local Role status to make sure that the role changed.

ASYNCHRONOUS MIRRORING

[Learn More >](#)

Mirror Consistency Groups Mirrored Pairs

Filter ?

Create Mirrored pair Test Communication More ▾

Name	Status	Number of Mirrored Pairs	Local Role	Remote Storage Array	Connection Type
Reyes_Mirror_Group	Optimal	2	Primary	N/A	Fibre Channel
Gallagher_Mirror_Group	Optimal	1	Secondary	N/A	Fibre Channel

Total rows: 2 | ↻

Force promotion of a secondary mirror consistency group

If a catastrophic failure occurs at the primary site, you might decide to promote the secondary mirror consistency group to primary so that business operations can resume from the current recovery point. In this case, the controller at the secondary site that receives the command to promote attempts to communicate with the primary to coordinate a role change. If communication is not possible, as would likely be the case in a disaster scenario, the controller begins the process of promoting the secondary mirror consistency group to the primary.

Forced promotion first transitions the mirror consistency group to a role-change-in-progress state on the local storage array. Because the remote storage array containing the primary group is inaccessible, this state change is made only locally. Next, the mirror consistency group role is set to primary. This step enables all member volumes in the group to function as mirror primary volumes, allowing write access and tracking writes to resynchronize with the original primary when it is available again. If the original primary mirror consistency group was actively synchronizing the mirrored pairs before the forced role change, the data on the original secondary (the new primary) volumes is in an inconsistent state. In this case, a rollback is performed on all members of the group as an online background operation to restore the data back to the last known consistent state as preserved by the recovery point.

Note: If the inability to communicate was because of a communications link failure and not because the primary had been destroyed, and if the original primary was not force-demoted to secondary, when communication is restored, it is likely that the mirror consistency group will be in a dual primary role conflict. If this conflict occurs, a Needs Attention alert is issued. To resolve this conflict, you must follow the Recovery Guru procedures.

Note: The preceding description of the role change applies when the controller at the secondary site cannot communicate with the primary. If the controller can communicate with its counterpart at the primary site when it receives the forced promotion command, the two controllers proceed with an immediate role reversal. In this case, the primary volumes are not synchronized with the secondary before they reverse roles, so any data that had been written to the primary since the last resynchronization is lost.

Forced promotion is disallowed if any of the following conditions exist:

- Not all the mirror consistency group member volumes have been previously synchronized and have a consistent recovery point.

- Any member volumes do not have a point-in-time image at the recovery point. This situation could occur if the reserved capacity volume is full, for example.
- The mirror consistency group has no member volumes.
- The mirror consistency group is in the failed, role-change-pending, or role-change-in-progress state.
- There is a dual role conflict.

To force promotion of the secondary mirror consistency group, follow the same instructions in the previous section, “Perform an orderly role reversal.”

Force demotion of a primary mirror consistency group

Forced demotion of a mirror consistency group in the primary role to function in the secondary role helps you protect the volumes from receiving new write requests while they are disconnected from the remote storage array. This approach can be useful if the corresponding group on the remote system has been promoted to primary by using the force option.

If the original primary is force-demoted before communication is restored, then after communication is restored, the mirror consistency groups on both sides complete the role reversal and begin periodic resynchronization as normal (with roles reversed). If any writes occurred to the original primary after the communication loss but before the demotion, they are discarded. As in the disaster recovery scenario, operations continue from the recovery point with the original secondary now as primary.

Forced demotion first transitions the mirror group to a role-change-in-progress state on the local storage array. Because the remote storage array is inaccessible, this state change is made only locally. If the original primary was actively synchronizing the mirrored pairs before the forced role change, the synchronization operation is canceled, and the point-in-time images used for the resynchronization are deleted. The group role is set to secondary, which prevents all member volumes in the group from receiving new write requests.

Note: If the original secondary is not force-promoted to primary while unable to communicate, then when communication is restored, it is likely that the mirror consistency group is in a dual secondary role conflict. If this situation occurs, a Needs Attention alert is issued. To resolve this conflict, you must follow the Recovery Guru procedures.

Note: The preceding description of the role change applies when the controller at the primary site cannot communicate with the secondary. If the controller can communicate with its counterpart at the secondary site when it receives the forced demotion command, the two controllers proceed with an immediate role reversal. In this case, the primary volumes are not synchronized with the secondary before they reverse roles, so any data that had been written to the primary since the last resynchronization is lost.

Forced demotion of the primary mirror consistency group is disallowed if any of the following conditions exist:

- Any of the member volumes are failed.
- Any of the associated reserved capacity volumes are failed.
- The mirror consistency group is in the failed, role-change-pending, or role-change-in-progress state.
- There is a dual role conflict.

To force demotion of the primary mirror consistency group, follow the same instructions that are in the previous section, “Perform an orderly role reversal.”

Troubleshooting of common failures

Table 12 lists some common failure modes, along with guidelines on how to troubleshoot the issue.

Note: These actions are not the only steps that are required to handle site failures. The steps listed here relate only to the two storage arrays that are involved in a mirroring relationship. They do not

discuss moving applications over from a primary site to a secondary, reestablishing communications, and so on.

Table 12) Troubleshooting of common failures.

Problem	Description and Suggested Steps
Communication Lost: Secondary Not Promoted	<p>This situation could occur in several scenarios, such as the secondary system temporarily going offline or a short-term interruption in the communication between sites. In this case, the primary mirror consistency group is still available to the host systems for writing and reading. The administrator does not promote the secondary. The primary continues operations as normal but continues to log changes until communication is restored and normal resynchronization with the secondary resumes.</p>
Communication Lost: Secondary Promoted	<p>If communication between the storage arrays is lost, and host systems also cannot access the primary mirror consistency group, promote the secondary so that operations can continue, as follows:</p> <ol style="list-style-type: none"> 1. If possible, demote the primary to prevent any further writes. 2. Promote the secondary so that operations proceed from the last recovery point. Note that any writes that have occurred on the original primary since the last recovery point are not reflected at the newly promoted primary. 3. Wait while the newly promoted primary is read/write accessible and all writes are logged. After communication is restored, the mirror consistency groups on both systems complete the role reversal, and periodic resynchronizations resume. 4. When ready, initiate an orderly role reversal to restore normal operation.
Failure at Primary Site with Storage System Recoverable	<p>To handle a failure such as a prolonged power outage at the primary site that renders the primary storage array inoperable but recoverable later, take the following steps:</p> <ol style="list-style-type: none"> 1. Promote the secondary so that operations proceed from the last recovery point. Note that any writes that have occurred on the original primary since the last recovery point are not reflected at the newly promoted primary. 2. Wait while the newly promoted primary is read/write accessible and all writes are logged. 3. When the primary site becomes operable again, before you restore communication with the remote site, you must force-demote the original primary mirror consistency groups to secondary. After communication is restored, the mirror consistency groups on both systems complete the role reversal, and periodic resynchronizations resume. 4. When ready, initiate an orderly role reversal to restore normal operation.
Failure at Primary Site, Storage System Destroyed	<p>This situation is a common disaster scenario in which the primary site is sufficiently damaged and needs new storage equipment as part of the rebuilding process. Take the following steps:</p> <ol style="list-style-type: none"> 1. Promote the secondary so that operations proceed from the last recovery point. Note that any writes that have occurred on the original primary since the last recovery point are not reflected at the newly promoted primary. 2. Remove member volumes from the mirror consistency groups, then force-delete the groups. 3. After a storage array is available at the primary site (or a third site), create new mirror consistency groups and add the volumes back in. At this point, an initial synchronization takes place, after which normal periodic resynchronizations resume. 4. When ready, initiate an orderly role reversal to restore normal operation.

Synchronous mirroring configuration and management

This section describes some additional considerations for synchronous mirroring, configuration instructions, and management tasks.

Balance link speeds and data traffic

Before you configure synchronous mirroring, review the following sections for additional considerations about latency, IOPS, and bandwidth.

Consider latency issues

Latency refers to any delay that occurs as data is read or written. The longer the distance between the two sites, the longer the round-trip time. Distance is the limiting factor for the number of I/Os that a mirror link can handle per second. Therefore, latency determines the I/O rate that can be achieved in a mirrored solution.

In addition to distance, latency is influenced by the FC network; for example, by the number of hops and by the equipment that is used for the link between the two sites. These latencies are the result of delays that are introduced by the equipment (switches, protocol converters for IP, firewalls) and the transport medium itself. The best method to determine the latency is to use the link test tool that is included in System Manager.

Calculate the distance and number of IOPS

In a synchronous mirroring solution, all writes to the primary site are mirrored to the secondary site before the acknowledgment is sent to the host. The link must be sized so that the required IOPS performance on the primary side can be achieved.

To calculate the correct link between the two sites, you must know what factors influence the effective IOPS rate compared with the nominal link speed. As described earlier, latency is the most important factor in designing the solution.

The latency in the medium is determined by the speed of light in the fiber cable. The effective speed of light in glass is 200,000km/sec. This speed becomes more relevant as the distance between two sites increases. Keep in mind that in a synchronous mirroring scenario, the signal must travel the distance between the two sites twice (to the remote site and for the acknowledgment back again) to complete an I/O successfully and to acknowledge the I/O to the host. Round-trip time determines the maximum possible I/O rate for the solution.

For a given distance, the calculation for the theoretical maximum number of IOPS is:

$$1\text{sec} / (2x \text{ distance} / 200,000\text{km/sec}) = \text{maximum number of IOPS}$$

At the maximum distance of 10km for synchronous mirroring, the theoretical maximum IOPS is 10,000. Keep in mind that components in the path, such as routers, add to the latency and reduce maximum IOPS. The quality of the link can also add latency. Make sure that the network delivers the required service level (quality of service, or QoS).

Determine bandwidth

Bandwidth becomes an important factor during the initial synchronization or resynchronization of a mirrored pair. Additional bandwidth is required to synchronize the copy pairs in a reasonable amount of time after a suspended mirror is resumed. Especially in shared network environments, sizing for synchronization is important because mirroring operations might need more or even the entire bandwidth for the resynchronization task. Table 9 shows some example times for synchronization pipe size. It is important to size the network with plenty of headroom to complete synchronization tasks within the required time.

Configure synchronous mirroring

Configuration involves the following steps:

- In Unified Manager, create the synchronous mirrored pair by selecting the local and remote storage arrays and then selecting the volumes.
- When initial synchronization begins, check the initial synchronization by opening System Manager for the local storage array.

Create a synchronous mirrored pair

To configure synchronous mirroring, you select a primary volume on the local array and a secondary volume on the remote array for the mirrored pair. Perform the following steps:

1. Open Unified Manager.
2. From the Manage page, select the local storage array that you want to use for the source.
3. Select Actions > Create Synchronous Mirrored Pair.

The screenshot shows the 'MANAGE - All' interface with a table of storage arrays. The 'Actions' menu is open, and 'Create Synchronous Mirrored Pair' is highlighted. The table has columns for Name, SANtricity OS Software, and IP Addresses.

Name	SANtricity OS Software	IP Addresses
EF280-1	88.62.47.16.000	2 configured addresses
EF280-2	88.62.47.16.000	2 configured addresses

The Create Synchronous Mirrored Pair wizard opens.

4. From the Select Primary Volume panel, select a volume that you want to use as the primary volume in the mirror.

The system displays a list of all eligible volumes for that mirrored pair. Any volumes that are not eligible do not display in the list.

The screenshot shows the 'Create Synchronous Mirrored Pair' wizard. Step 1 is 'Select Primary Volume'. It includes a filter box and a table of available volumes.

Name	Capacity (GiB)
Test1	1024.00

5. Click Next to move to the Select Secondary Volume panel.

- Select the remote storage array on which you want to establish a mirror relationship with the local storage array. If your remote storage array is password protected, the system prompts for a password. Storage arrays are listed by their storage array name. If you have not named a storage array, it will be listed as Unnamed.

If the storage array that you want to use is not in the list, make sure that it has been discovered in Unified Manager. To discover arrays, select Add/Discover from the Manage panel to open a discovery wizard.

Create Synchronous Mirrored Pair
✕

1 Select Primary Volume
2 Select Secondary Volume
3 Edit Settings

Primary volume attributes

Reported capacity: 1024.00 GiB
 Data Assurance (DA) enabled: Yes
 Secure-capable: Yes - FDE (🔒 not enabled)
 Volume type: Thick

Select a storage array to view secondary volume candidates...

Select

- From the list of eligible volumes, select a volume that you want to use as the secondary volume in the mirror.

Note: If you select a secondary volume with a capacity that is larger than the primary volume, the usable capacity is restricted to the size of the primary volume.

Create Synchronous Mirrored Pair
✕

1 Select Primary Volume
2 Select Secondary Volume
3 Edit Settings

Primary volume attributes

Reported capacity: 1024.00 GiB
 Data Assurance (DA) enabled: Yes
 Secure-capable: Yes - FDE (🔒 not enabled)
 Volume type: Thick

Select a storage array to view secondary volume candidates...

EF280-1

Why don't I see all the volumes on the remote storage array?

?

Select your secondary volume from EF280-1...

Name	Capacity (GiB)
VMware1	1024.00
VMware2	1024.00

Total rows: 2

8. Click Next to move to the Edit Settings panel.
9. Use the slider bar to set the synchronization priority.

The synchronization priority determines how much of the system resources to use to complete initial synchronization and the resynchronization operation after a communication interruption (how much for these actions as compared with service I/O requests).

The Synchronization Priority setting in this dialog box applies to both the primary volume and the secondary volume. You can modify the rate on the primary volume later by going to System Manager and by selecting Storage > Synchronous Mirroring > More > Edit Settings.

There are five synchronization priority rates:

- Lowest
- Low
- Medium
- High
- Highest

If the synchronization priority is set to the lowest rate, I/O activity is prioritized, and the resynchronization operation takes longer. If the synchronization priority is set to the highest rate, the resynchronization operation is prioritized, but I/O activity for the storage array might be affected.

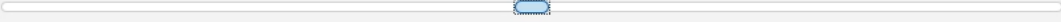
Create Synchronous Mirrored Pair
✕

1 Select Primary Volume
 2 Select Secondary Volume
 3 Edit Settings

Synchronization priority

Synchronization priority ?

Medium (default)



Why is it recommended to use a manual resynchronization policy?

Resynchronization policy

After a communication interruption...

Require synchronization to be manually resumed (recommended)

Automatically resume synchronization

What impact does synchronization priority have on synchronization rates?

10. Select whether you want to resynchronize the mirrored pairs on the remote storage array manually or automatically after a communication interruption:
 - Manually (the option that NetApp recommends)—Select this option to require synchronization to be manually resumed after communication is restored to a mirrored pair. This option provides the best opportunity to recover data.
 - Automatically—Select this option to start resynchronization automatically after communication is restored to a mirrored pair.

To manually resume synchronization, go to System Manager and select Storage > Synchronous Mirroring, highlight the mirrored pair in the table, and select Resume under More.

11. Click Finish to complete the synchronous mirroring sequence.

Note: This operation can be lengthy and could affect system performance.

Synchronous mirroring is now set. The system performs the following actions:

- a. It creates the mirrored pair. When that step is finished, Unified Manager displays Mirrored Pair Was Successfully Created.
- b. It begins initial synchronization between the local storage array and the remote storage array. For more information, see “Initial synchronization process for synchronous mirroring.”
- c. It sets the synchronization priority and resynchronization policy.
- d. It reserves the highest-numbered port of the controller's HIC for mirror data transmission.
- e. I/O requests that are received on this port are accepted only from the remote preferred controller owner of the secondary volume in the mirrored pair. (Reservations on the primary volume are allowed.)
- f. It creates two reserved capacity volumes, one for each controller, which are used to log write information to recover from controller resets and other temporary interruptions.

The capacity of each reserved capacity volume is 128MiB. However, if the volumes are placed in a pool, 4GiB is reserved for each volume.

Check the progress of the initial synchronization

To check the progress of the initial synchronization, do the following:

1. In Unified Manager, from the Manage page, select the local storage array and click Launch to open System Manager.

<input type="checkbox"/>	Storage Array	Status	Model Name	SANtricity OS Software
<input checked="" type="checkbox"/>	EF280-1	Optimal	2804	88.62.47.16.000
<input type="checkbox"/>	EF280-2	Optimal	2806	88.62.47.16.000

2. In System Manager, view the status of the mirroring operation. To view the progress, Select Home > View Operations in Progress.

EF280-1 Preferences | Help | admin | Log Out

✓ Your storage array is optimal. [View Operations in Progress >](#)

Auto-Configuration. [Ignore](#)

[View Performance Details >](#)

5 m 1 h 8 h 1 d 7 d 30 d [Compare IOPS with... ▾](#)

3. In the Operations in Progress panel, check the Status column to see the percentage completed.

Home / Operations in Progress

OPERATIONS IN PROGRESS

In-Progress: 1 | Pending: 0 | Suspended: 0

Operation	Object	Status
Synchronizing	VMware1, Test1	1%

Test the synchronous communications link

If you need to diagnose communication problems between the local storage array and the remote storage array that is associated with a mirrored pair, you can test the communications link as follows:

1. In System Manager, select Storage > Synchronous Mirroring.
2. Select the mirrored pair that you want to test and then select Test Communication.

SYNCHRONOUS MIRRORING

[Learn More >](#)

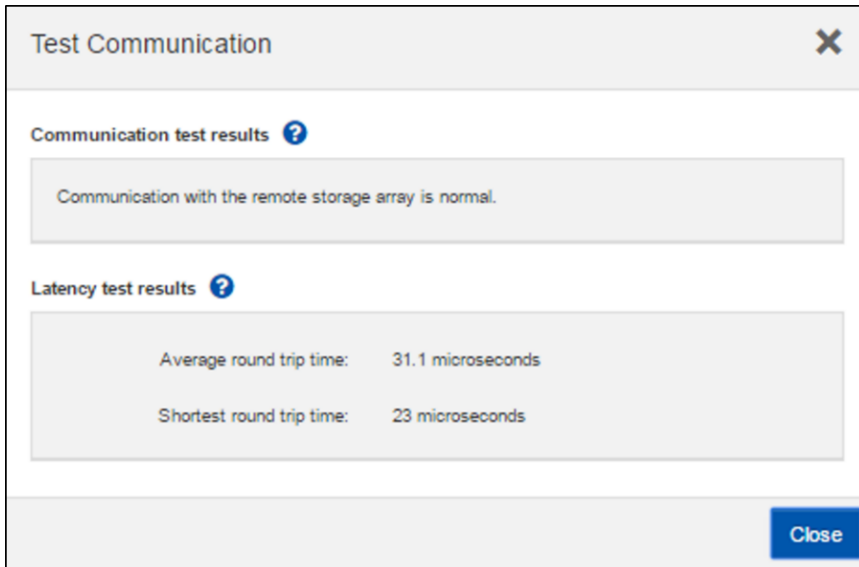
Filter ?

Mirror Volume **Test Communication** More ▾

Mirrored Pairs (Primary, Secondary)	Status	Local Role
VMware1, Test1	Suspended	Primary

Total rows: 1 | ↻

3. View the Test Communication result window to see whether the test was passed and to review the average and shortest round-trip times.



Manage synchronous mirroring

You have several options for managing the mirroring process in System Manager, including:

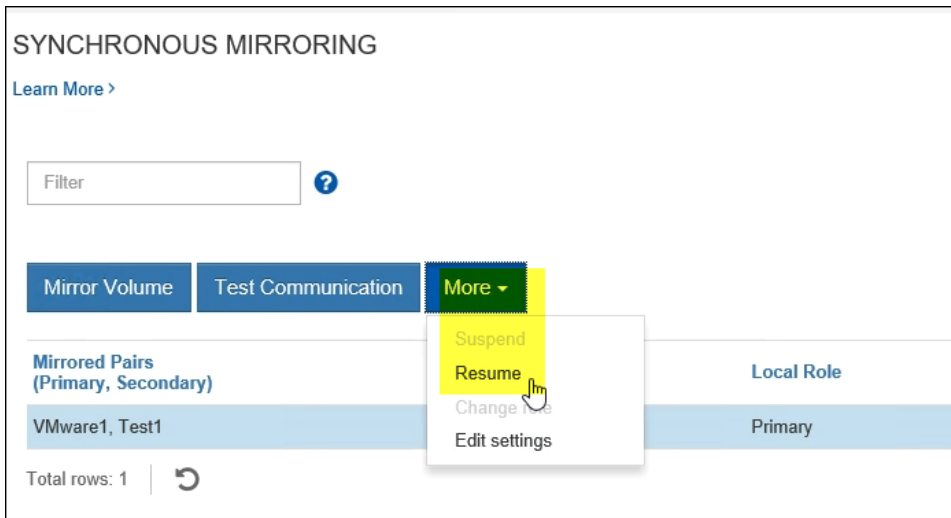
- Suspending or resuming the synchronization for a mirrored pair.
- Changing the role assignments (primary and secondary) of the mirrored pairs.

Suspend or resume synchronization

You can suspend synchronization from the primary storage array. When a mirror is in a suspended state, no attempt is made to contact the secondary volume. Writes to the mirror primary volume are persistently logged so that when synchronization resumes, only the regions of the primary volume that are known to have changed are written to the secondary volume.

To suspend or resume mirroring, do the following:

1. Open System Manager.
2. Go to Storage > Synchronous Mirroring.
3. Select the mirrored pair, select More, and then select either Suspend or Resume. Note that if synchronization is currently set, Resume is unavailable from the More menu. If synchronization is currently suspended, Suspend is unavailable.



Reverse roles on a synchronous mirrored pair

You can perform a role reversal on a synchronous mirrored volume pair. You can either promote the selected volume to a primary role or demote the selected volume to a secondary role. The role reversal affects only the selected synchronous mirrored pair.

When demoting a primary synchronous mirror volume to a secondary role, if the current secondary volume can be contacted, it is automatically promoted to a primary role in the mirror relationship. Likewise, when promoting a secondary volume to a primary role, if the current primary volume can be contacted, it is automatically demoted to a secondary role in the mirror relationship.

Note: If the environment has a suspended synchronous mirrored pair operation, it resumes during the role-change operation.

When you perform a role reversal, keep these guidelines in mind:

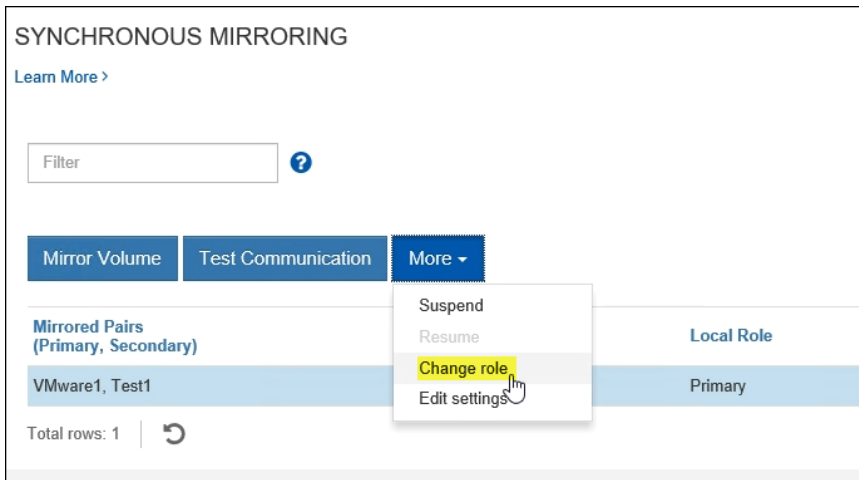
- Any hosts that are accessing the primary volume in a synchronous mirrored pair through a mapping have read/write access to the synchronous mirrored volumes in the mirror relationship. When the primary synchronous mirror volume becomes a secondary, hosts that have been mapped to the mirrored volumes no longer have write access to it.
- When the secondary synchronous mirror volume becomes a primary, any hosts that are accessing that volume can then write to it.
- If a communication problem between the local and remote sites prevents the promotion of the secondary synchronous mirror volume, an error message appears. However, you can still force the secondary volume to a primary role. This forced promotion leads to a dual primary synchronous mirroring condition after communication is restored. To resolve this condition, use the Recovery Guru.
- Likewise, you can force the primary volume to a secondary role if communication is not possible between the two storage arrays. When communication is restored, the mirrored pair is in a dual secondary condition, which can also be resolved by using the Recovery Guru.

To change roles between volumes in a mirrored pair, do the following:

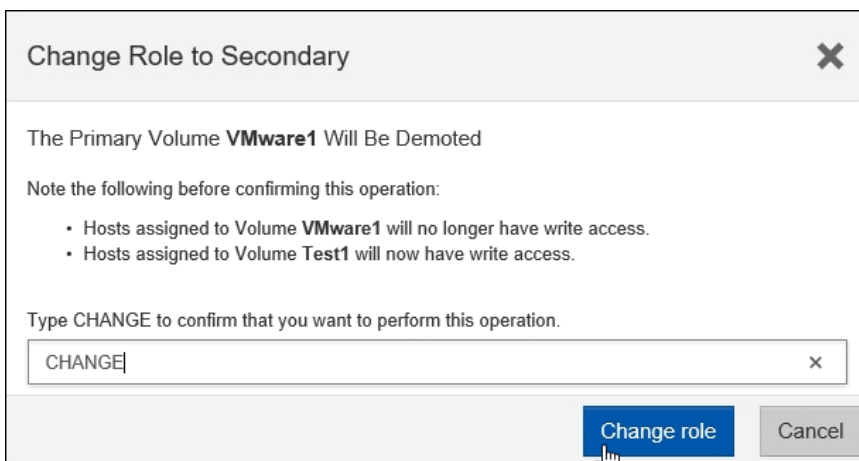
1. Open System Manager.
2. Select Storage > Synchronous Mirroring.

Note: Make sure that the mirrored pair is synchronized before you reverse volume roles.

3. Select the mirrored pair that contains the volumes for which you want to change the role and then select More > Change Role.



4. Confirm the operation in the Change Role dialog box.



5. Check the Local Role status to make sure that the role changed.

Where to find additional information

To learn more about the information that is described in this document, review the following:

- NetApp Product Documentation
[E-Series and SANtricity 11 Documentation Center](#)
- NetApp SANtricity Management Security
[TR-4712: NetApp SANtricity Management Security](#)
- Managing Certificates for NetApp E-Series Storage Systems
[TR 4813: Managing Certificates for NetApp E-Series Storage Systems](#)

Version history

Version	Date	Document Version History
Version 1.0	May 2020	Initial release.
Version 2.0	September 2020	Updated with EF300

Version	Date	Document Version History
Version 2.1	February 2021	Modified the section Latency Considerations with iSCSI to consider changes in SANtricity 11.50 not previously edited. Noted VLAN is not a requirement for iSCSI asynchronous replication.
Version 2.2	June 2022	Minor date and version updates.
Version 3.0	July 2023	Updated for SANtricity 11.80.0 release.
Version 3.1	July 2024	Update to reflect no mirroring between NVMe and SAS drives.
Version 3.2	November 2024	Update for SANtricity OS 11.90.0 release and E4000 release.

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 © 2020–2024 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.

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

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.

TR-4839-0724