

FUJITSU Storage  
ETERNUS AB series All-Flash Arrays,  
ETERNUS HB series Hybrid Arrays

Technical Report

SANtricity Synchronous and Asynchronous Mirroring



Feature Descriptions and Deployment Guide

Fujitsu Limited  
July 2023

**Abstract**

Fujitsu ETERNUS AB/HB series systems provide rock-solid local and remote data mirroring capabilities using the SANtricity synchronous and asynchronous mirroring features. This document provides descriptions of the features and specific setup and operations guidance, including the use of SANtricity Unified Manager when the customer only has new generation ETERNUS AB/HB series systems.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b>	<b>5</b>
1.1	Intended Use	5
1.2	Overview of the Mirroring Features	5
1.3	Terminology	6
<b>2</b>	<b>SANtricity Unified Manager – Usage and Configuration</b>	<b>8</b>
2.1	Installing Unified Manager	8
2.2	Logging in to Unified Manager	9
2.3	Discovering and Adding Storage Arrays	10
2.4	Remote Mirroring with SANtricity Unified Manager	11
2.5	SANtricity Unified Manager Security	11
2.6	Web Services Certificates and Recommended Browsers	11
<b>3</b>	<b>Comparing Asynchronous and Synchronous Mirroring</b>	<b>12</b>
3.1	Typical Use Cases	13
<b>4</b>	<b>Requirements and Limitations</b>	<b>13</b>
4.1	Feature Enablement and Activation	13
4.2	Prerequisites	14
4.3	Supported Connections	14
4.4	Distances and Speed	17
4.5	Mirror Volume Requirements	18
4.6	Reserved Capacity Volume Requirements	19
4.7	Connectivity and Volume Ownership	19
4.8	Storage System Limits	20
<b>5</b>	<b>Asynchronous Mirroring Operational Model</b>	<b>20</b>
5.1	Initial Synchronization	20
5.2	Resynchronization Methods for Asynchronous Mirroring	22
5.3	Periodic Resynchronization	22
5.4	Using Thin-Provisioned Volumes in Asynchronous Mirroring	24
5.5	Suspend and Resume	25
5.6	Orderly Role Reversal	25
5.7	Forced Promotion of Secondary	26
5.8	Forced Demotion of Primary	27
5.9	Example Failures and How to Handle	27
<b>6</b>	<b>Considerations for Setting Up Asynchronous Mirroring</b>	<b>29</b>
6.1	RPO and Synchronization Interval	29
6.2	Sizing the Network to Meet RPO	30
6.3	Latency Considerations with iSCSI	31

6.4	Time to Complete Initial or Full Synchronization.....	31
6.5	Effect of Asynchronous Mirroring Operations on Performance.....	32
<b>7</b>	<b>Synchronous Mirroring Operational Model .....</b>	<b>33</b>
7.1	Initial Synchronization.....	33
7.2	Write Process .....	33
7.3	Resynchronization Methods .....	34
7.4	Suspend and Resume .....	35
7.5	Role Changes on a Synchronous Mirror Pair .....	35
<b>8</b>	<b>Considerations for Setting Up Synchronous Mirroring .....</b>	<b>36</b>
8.1	Latency .....	36
8.2	Distance and Number of IOPS .....	36
8.3	Bandwidth.....	37
<b>9</b>	<b>Configuring Through Management GUI .....</b>	<b>37</b>
9.1	Use Cases .....	37
9.2	Management Station Port Numbers .....	38
9.3	Web Services Certificates.....	38
9.4	Mirroring and Role-Based Access Control (RBAC).....	39
9.5	Asynchronous with System Manager as Primary .....	40
9.6	Synchronous with System Manager as Primary .....	47
9.7	Testing the Communication Link .....	52
	<b>Summary .....</b>	<b>54</b>
	<b>Related Information and Documents .....</b>	<b>55</b>
	<b>Version History .....</b>	<b>55</b>

## LIST OF TABLES

Table 1) Comparison between mirroring features. ....	12
Table 2) Cable distances by cable type and speed. ....	17
Table 3) Transmission rates by carrier type. ....	17
Table 4) Allowed candidates for mirror volumes. ....	18
Table 5) Storage system limits for mirroring. ....	20
Table 6) Approximate hours needed for initial or full synchronization by pipe size. ....	31
Table 7) System configuration for performance example. ....	32
Table 8) Example performance results. ....	32
Table 9) Management use cases for mirroring. ....	37
Table 10) Port numbers used by SANtricity OS releases. ....	38

## LIST OF FIGURES

Figure 1) WSP installation wizard with the link to launch Unified Manager. ....	8
Figure 2) SANtricity Unified Manager login page. ....	9
Figure 3) SANtricity Unified Manager landing page. ....	10
Figure 4) SANtricity Unified Manager showing discovered arrays. ....	10
Figure 5) Single-fabric configuration. ....	15
Figure 6) Fabric dedicated to mirroring. ....	16
Figure 7) Periodic resynchronization in asynchronous mirroring. ....	23
Figure 8) Time from one protection point to the next. ....	29
Figure 9) Synchronous mirroring write process. ....	34
Figure 10) Web services certificate management in SANtricity Unified Manager. ....	38
Figure 11) Completing CSR for web services. ....	39
Figure 12) SANtricity System Manager: invoke test communication on mirror consistency group. ....	53
Figure 13) SANtricity System Manager: mirror consistency group test communication results. ....	53
Figure 14) SANtricity System Manager: invoke test communication for synchronous mirrored pair. ....	54
Figure 15) SANtricity System Manager: synchronous mirrored pair test communication results. ....	54

# 1 Introduction

This document describes the two mirroring features available with Fujitsu ETERNUS AB/HB series storage systems and provides guidance on how to deploy them. The mirroring features enable administrators to replicate data online over near or far distances to protect against disasters or to centrally back up data at a remote location.

To accomplish this replication, the administrator creates mirrored volume pairs on the local and remote storage systems. For each pair, the volume on the local site is known as the primary, and the remote volume is known as the secondary. Any given storage system might have both primary and secondary volumes, enabling companies to have multiple active sites that protect data for each other.

If a disaster or other failure causes a set of primary volumes on one site to become unavailable, the administrator can promote the corresponding secondary volumes to the role of primary to take over responsibility for maintaining IT operations.

## 1.1 Intended Use

This information is for Fujitsu customers and partners. For the information and procedures described in this document to be useful, it is assumed that the reader has:

- A minimal knowledge of ETERNUS AB/HB series platforms and products, particularly in the area of data protection
- General knowledge of disaster recovery (DR) solutions

## 1.2 Overview of the Mirroring Features

ETERNUS AB/HB series storage systems offer two types of mirroring: asynchronous and synchronous:

- **Asynchronous mirroring.** Replicates changed data from the primary to the secondary at discrete points in time, providing a recovery point at the secondary that is behind the current primary. Asynchronous mirroring can function over long distances and includes the following attributes:
  - 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 ETERNUS AB/HB series storage systems.
  - A standard Internet Protocol (IP) or Fibre Channel (FC) network can be used for replication.
- **Synchronous mirroring.** Continuously replicates data so that the secondary is always up to date. Synchronous mirroring supports shorter distances and includes the following attributes:
  - 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 ETERNUS AB/HB series storage systems.
  - An FC network is used for replication.

Both types of mirroring have some effect on performance, but it is important to be aware that, with synchronous mirroring, each write I/O is not completed until the write to the remote volume has completed. This is an important consideration for high performance environments given the additional write latency associated with remote write confirmation messages.

## 1.3 Terminology

There are a number of terms used in this document and in the industry that bear description to help reduce confusion.

### SANtricity System Manager

SANtricity System Manager is the web-based management GUI that is bundled with all SANtricity OS versions running on new generation ETERNUS AB/HB series arrays (ETERNUS HB1100/HB1200/AB2100/AB5100/HB2000/HB5000). To access the SANtricity System Manager GUI, point your browser to the fully qualified domain name (FQDN) or IP address of a management port on one of the two storage array controllers.

**Note:** You can also launch SANtricity System Manager from SANtricity Unified Manager when the new generation array has been discovered by the central management interface.

### SANtricity Web Services Proxy

The SANtricity Web Services Proxy (WSP) is the set of API endpoints that are used to configure, monitor, and manage ETERNUS AB/HB series storage arrays. Many of the endpoints support older generation and new generation arrays.

**Note:** Some API endpoints do not apply to older generation arrays when those arrays do not support the associated features. For example, endpoints associated with managing web server certificates do not apply to older generation ETERNUS AB/HB series arrays.

### SANtricity Unified Manager

SANtricity Unified Manager is a web-based GUI that is built on top of the SANtricity WSP. It provides storage array administrators a GUI-based API orchestration tool to monitor and manage ETERNUS HB1100/HB1200/AB2100/AB5100/HB2000/HB5000 arrays from a single interface.

### Initial (or Full) Synchronization

When a mirror relationship is established, all data from the primary is copied to the secondary. During this process, the primary volume is fully accessible for I/O operations. Full synchronization can also be required under certain failure conditions of the communications link or the secondary volume. If a volume being replicated is thin, a full synchronization only copies blocks to the secondary that have been written on the primary. Any blocks on a thin primary volume that have been unmapped are also not replicated to the secondary.

### Mirror Consistency Group (Asynchronous Mirror Group)

A mirror consistency group is a container for one or more volumes participating in asynchronous mirroring. The asynchronous mirroring feature resynchronizes all mirror pairs in a group simultaneously, thus preserving a consistent recovery point at the secondary across all volumes in the group.

**Note:** Synchronous mirroring does not use mirror consistency groups.

### Mirror Pair

Two volumes linked together through a mirroring relationship are called a mirror pair. A mirror pair consists of a primary volume and a secondary volume.

## Mirror Volume

A mirror volume is any volume participating in mirroring. It can have either the primary or secondary role.

## Primary Site

Primary site generally refers to the production site for any given mirror pair or group. Mirroring occurs from the primary site to the secondary site. Note that with SANtricity mirroring features, sites can mirror for each other; that is, one site can be primary for one set of volumes and secondary for another.

## Primary Volume

A mirror volume from which reads and writes are initiated from the host or hosts is said to be in a primary role. The mirroring features copy data from the primary volume to the secondary volume.

## Reserved Capacity Volume

These volumes that are not user accessible are required so that the controller can persistently save information needed to maintain mirroring in an operational state. They contain information such as delta logs and copy-on-write data.

## Resynchronization

For synchronous mirroring, the normal operating state is for the primary and secondary to be synchronized. However, a link failure or other type of failure or a suspend from the administrator causes the primary and secondary to be out of synchronization. When this desynchronization happens, restoring to normal operations requires a resynchronization in which changed data is sent from the primary to the secondary.

For asynchronous mirroring, after initial synchronization, the mirror pair is out of synchronization most of the time. The administrator can select either periodic or manual resynchronization. Every resynchronization creates a recovery point.

## Roles

A volume in a mirror pair is always in either a primary or secondary role. A volume does not always have the same role: roles can be reversed by the administrator, or an individual volume can have its role changed by the administrator in the event both members of a mirror pair are not accessible. Role changes are typically used for recovery when a failure has occurred at the primary site or when communication between members of a pair is lost.

## Secondary Site

This site is where the destination, or secondary, volumes reside for any given mirror pair/group. It is sometimes called the recovery or disaster site. Note that with SANtricity mirroring features, sites can mirror for each other; that is, one site can be primary for one set of volumes and secondary for another.

## Secondary Volume

A mirror volume that holds data copied from the primary volume is known as the secondary volume. It is either an exact copy with synchronous mirroring or a copy that lags in time from the primary with asynchronous mirroring.

## 2 SANtricity Unified Manager – Usage and Configuration

SANtricity Unified Manager is a web-based, central management interface for managing new generation ETERNUS HB1100/HB1200/AB2100/AB5100/HB2000/HB5000 systems. It is bundled with SANtricity WSP 3.0 and later.

SANtricity WSP with the Unified Manager GUI is installed on a management server that has IP access to all arrays that you want to manage from the central management interface and can manage up to 500 arrays. SANtricity Unified Manager provides a higher level of management interface security and adds the following time-saving features:

- Supports a new, central storage array upgrade wizard for arrays managed using the web based SANtricity System Manager GUI.
- Supports Lightweight Directory Access Protocol (LDAP) and role-based access control (RBAC) just like SANtricity System Manager. SANtricity Unified Manager includes a simplified certificate management workflow to manage the Unified Manager and Web Services Proxy server certificates (truststore and keystore certificates).
- Supports organizing arrays by groups you can create, name, and arrange.
- Supports importing common settings from one array to another, saving time from duplicating set-up steps for each array.
- Supports synchronous and asynchronous mirroring for ETERNUS HB1100/HB1200/AB2100/AB5100/HB2000/HB5000 arrays.

**Note:** The legacy Symbol API management interface must be turned on to set up mirroring, but after mirroring is created, the Symbol API interface can be disabled again.

### 2.1 Installing Unified Manager

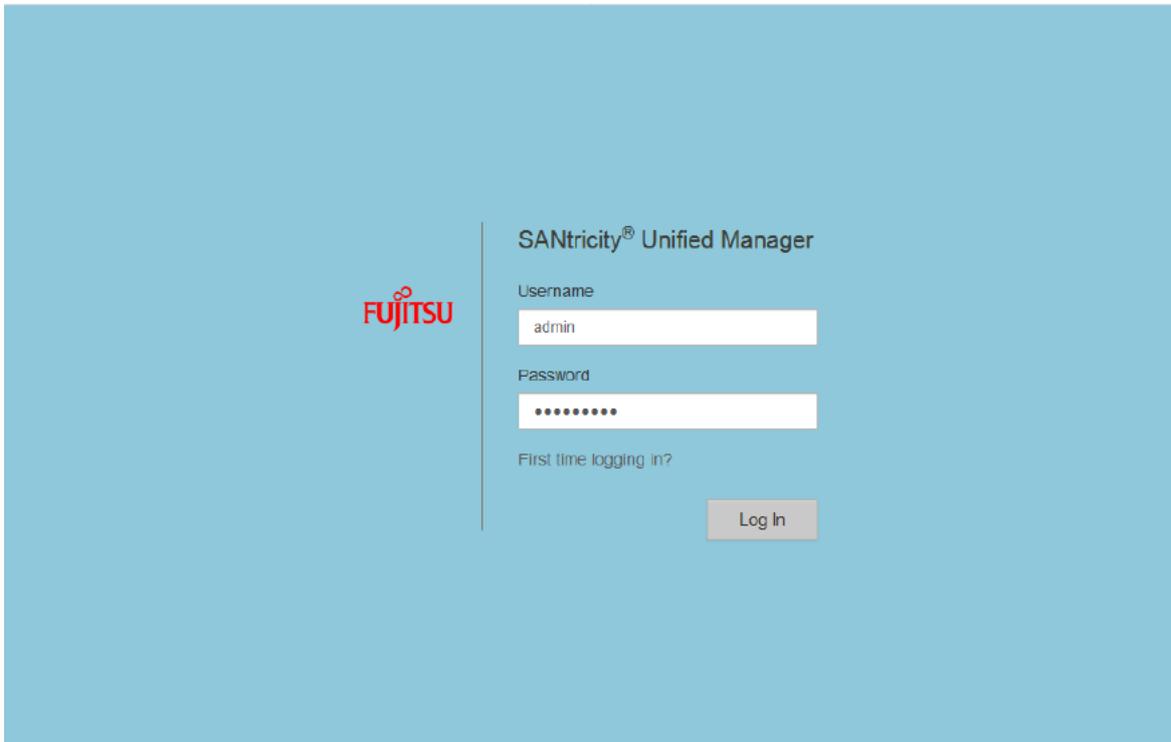
To install SANtricity Unified Manager, install either ETERNUS AB/HB series SANtricity Unified Manager or ETERNUS AB/HB series SANtricity Web Services Proxy from the DVD. Either link downloads the combined package. When the WSP installation wizard completes, you can open Unified Manager as shown in Figure 1.

Figure 1) WSP installation wizard with the link to launch Unified Manager.



To launch the Unified Manager GUI after SANtricity WSP installation, open a browser and navigate to the server IP address and secure port number that was reserved during the WSP software installation. The installation wizard prompted for the port # during installation. For example, the URL is in the form of `https://proxy-FQDN or IP address:port #/`. Then select the link for Unified Manager. You can go straight to the Unified Manager login screen by adding `/um` to the URL; for example, `https://proxy-FQDN or IP address:port #/um`, as shown in Figure 2.

Figure 2) SANtricity Unified Manager login page.



## 2.2 Logging in to Unified Manager

SANtricity Unified Manager has a similar look and feel to SANtricity System Manager, but there is one significant difference at the initial login if you are using SANtricity Unified Manager 3.x. SANtricity System Manager requires administrators to set the array admin password as part of the initial login.

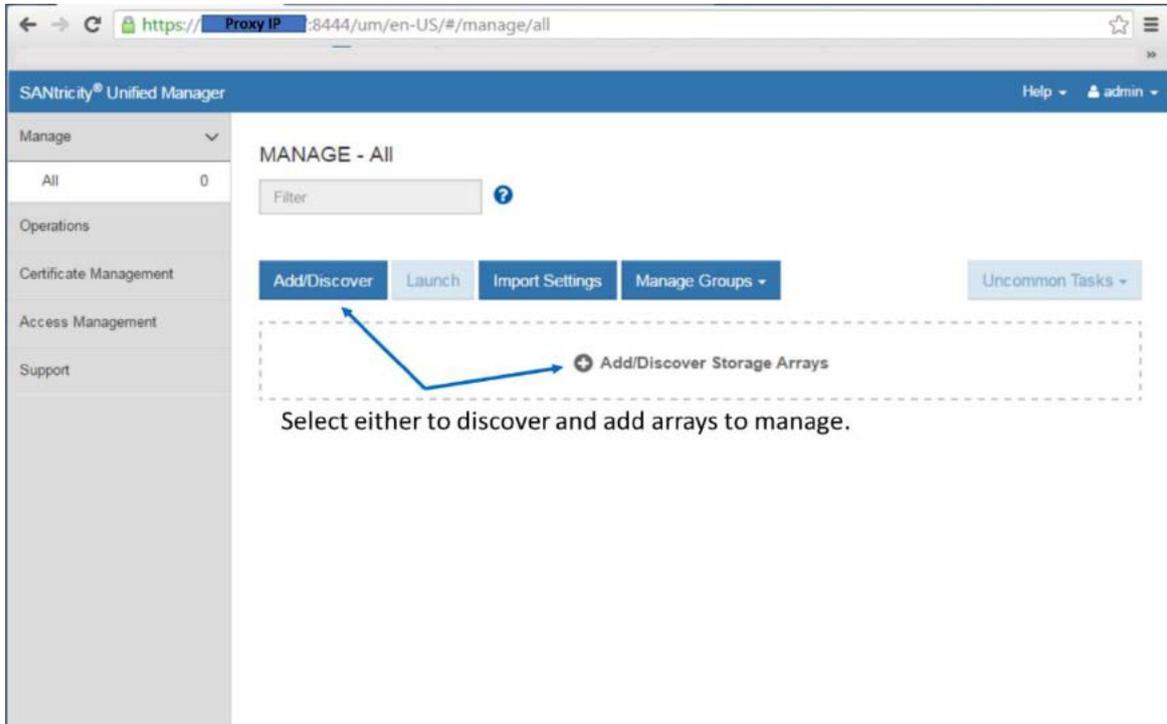
If you are using SANtricity Unified Manager 3.x, it has a factory default account and password: `user=admin` and `password=admin`. The administrator can choose to continue the setup using the default admin settings or the admin password can be changed before additional setup work is completed.

If you are using SANtricity Unified Manager 4.0 or later, you must set the admin password as part of the initial login, the same way as with SANtricity System Manager.

## 2.3 Discovering and Adding Storage Arrays

SANtricity Unified Manager must discover arrays to manage. You can discover a single array or scan a range of IP addresses to discover multiple arrays at once using the Add/Discover arrays wizard shown in Figure 3. After discovering one or more arrays, you can then add the array(s) to be managed by Unified Manager.

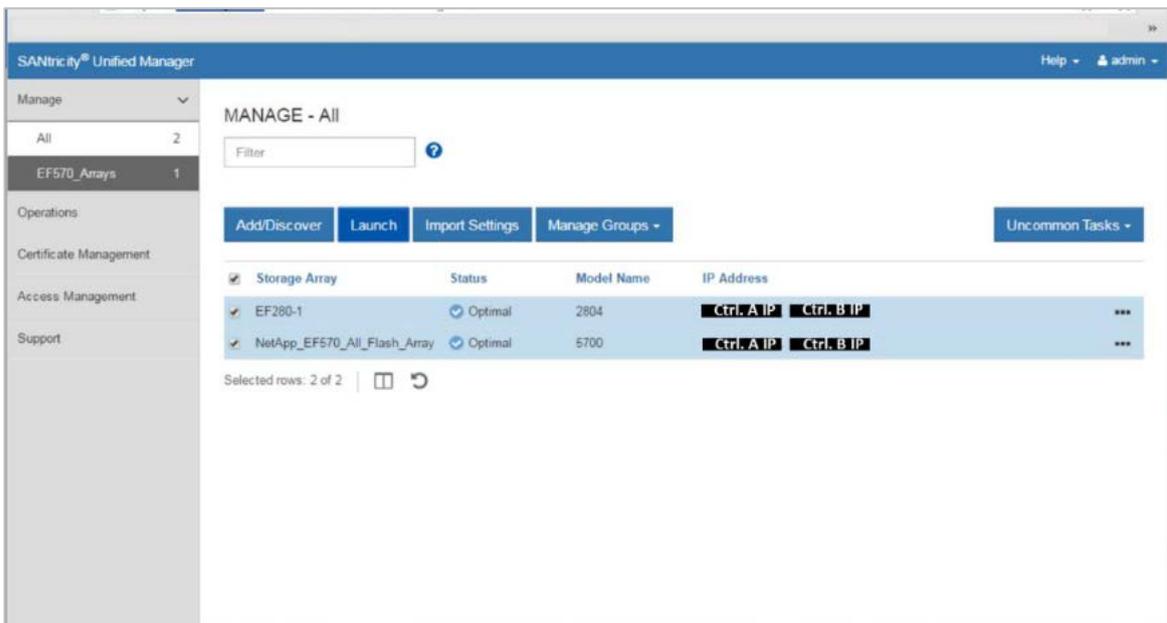
Figure 3) SANtricity Unified Manager landing page.



**Note:** The array discovery process includes either accepting the array self-signed certificates or importing an array's signed web server certificates and entering the array's admin user password.

After arrays are discovered and added, they are displayed on the landing page of Unified Manager as seen in Figure 4.

Figure 4) SANtricity Unified Manager showing discovered arrays.



## 2.4 Remote Mirroring with SANtricity Unified Manager

Both the initiator and target arrays must be discovered and managed by Unified Manager. The System Manager session for the initiator and target arrays must be launched from Unified Manager. You must always launch array management sessions from Unified Manager.

**Note:** To set up mirroring relationships using Unified Manager 3.0 (WSP 3.0), the legacy management SYMBol interface on the arrays must be turned on temporarily until the mirroring configuration is complete. You can turn on the legacy management SYMBol interface using SANtricity System Manager at Settings > System > Change Management Interface. For Unified Manager 3.1 (WSP 3.1) and later, you do not need to enable the legacy management interface on the arrays. The setup is fully supported through the HTTPs interfaces.

Mirroring is configured by launching SANtricity Unified Manager. Any mirroring relationship requires that both the initiator and target arrays are discovered by and listed in SANtricity Unified Manager. You must have the browser-based SANtricity Unified Manager installed, and you must have discovered the two storage arrays you want to mirror data between. Then, from Unified Manager, you select the initiator array and click Launch to open the browser-based SANtricity System Manager and set up mirroring.

## 2.5 SANtricity Unified Manager Security

SANtricity Unified Manager supports the same secure management features as SANtricity System Manager including LDAP, RBAC, and SSL certificates.

## 2.6 Web Services Certificates and Recommended Browsers

Certificates identify website owners for secure connections between clients and servers.

### Trusted Certificates

The SANtricity Unified Manager interface is installed with the WSP on a host system. When you open a browser and try connecting to Unified Manager, the browser attempts to verify that the host is a trusted source by checking for a digital certificate. Unified Manager might prompt you to accept the array's self-signed certificate, or it might show that the array's existing security certificate is untrusted. If the certificate is untrusted, you must import the Certificate Authority (CA) root and, in some cases, the intermediate certificates to the Unified Manager truststore.

To acquire signed, digital certificates from a CA from within SANtricity Unified Manager, complete the following steps:

1. Generate a certificate signing request (CSR) for the host system on which SANtricity Unified Manager is installed.
2. Send the .CSR file(s) to a CA, and then wait for them to send you the certificate files.
3. When the CA sends back the signed certificates, import the signed certificates from the Unified Manager interface.

For mirroring involving systems managed by SANtricity System Manager, Fujitsu recommends importing the trusted certificates for the web services in SANtricity Unified Manager that allow the storage systems to authenticate with the web server.

### Self-Signed Certificates

Self-signed certificates can be used to authenticate client-server communications. If the administrator attempts to configure mirroring without importing signed certificates, SANtricity System Manager displays an error dialog box that allows the administrator to accept the self-signed certificate. In this case, Fujitsu recommends using the latest version of Chrome or Firefox as the browser.

You can accept a self-signed certificate or install your own security certificate using Unified Manager and navigating to Certificate > Certificate Management.

### 3 Comparing Asynchronous and Synchronous Mirroring

Asynchronous mirroring differs from synchronous mirroring in one essential way. It captures the state of the source volume at a particular point in time and copies just the data that has changed since the last image capture.

- With asynchronous mirroring, the remote storage array is not always fully synchronized with the local storage array. Therefore, if the application needs to transition to the remote storage array due to a loss of the local storage array, some transactions could be lost.
- With synchronous mirroring, the state of the source volume is not captured at some point in time, but rather reflects all changes that were made on the source volume to the target volume. The copy is identical to production data at every moment because, with this type of mirror, each time the host writes to the primary volume, the storage system writes the same data to the secondary volume at the remote site. The host does not receive an acknowledgment that the write was successful until the secondary volume is successfully updated with the changes that were made on the primary volume.

Table 1 presents a brief comparison between asynchronous and synchronous mirroring.

Table 1) Comparison between mirroring features.

Attribute	Asynchronous Mirroring	Synchronous Mirroring
Replication method	Point in time: Snapshot copies are taken on demand or at user-defined intervals. Changed data between the current and previous Snapshot 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.  <b>Note:</b> The remote write acknowledgement adds latency to all mirrored-volume write I/O.
Consistency groups	Each mirrored volume is a member of a mirror consistency group (asynchronous mirror 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 done on a volume basis only.
Reserved capacity	Reserved capacity volumes required for each volume in a mirrored pair.	One reserved capacity volume created for each controller with synchronous mirroring activated.
Allowed volume types	Standard or thin	Standard only
Communication between arrays	FC or iSCSI	FC only
Distance	Virtually unlimited: limited only by the capabilities of the network, network components, and required synchronization interval.	10km (6.2 miles) to meet latency and application performance requirements.

### 3.1 Typical Use Cases

Asynchronous mirroring is ideal for satisfying the demand for nonstop operations and, in general, is far more network efficient for periodic processes such as backup and archive.

Although synchronous mirroring is ideal 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 archive.

Asynchronous mirroring is good for the following use cases:

- Remote backup consolidation
- Wide area content distribution
- Long-distance disaster recovery
- Backup of 24x7 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 metro area
- Business continuity during scheduled maintenance of the primary storage system with the secondary acting as the primary

## 4 Requirements and Limitations

This section provides a brief overview of what to know before deploying either or both of the SANtricity mirroring features.

### 4.1 Feature Enablement and Activation

Before using mirroring, the desired feature must be both enabled and activated on each storage system that participates in mirroring operations. Fujitsu ships all new ETERNUS AB/HB series systems with both asynchronous and synchronous mirroring enabled.

#### Activating the Mirroring Features in the GUI

For systems managed by SANtricity System Manager (HB2000, HB5000, and AB5100), no separate activation step is needed; activation occurs behind the scenes while mirror groups/pairs are being set up.

## 4.2 Prerequisites

The following are prerequisites that must be met to configure mirroring:

- Two dual-controller (duplex) ETERNUS AB/HB series storage systems that support the mirroring feature. They do not have to be the same OS version.
- The administrator must know the password for both storage systems.
- Enough free capacity is available on the remote storage system to create a secondary volume at least as large as the primary volume.
- Enough free capacity is available for the reserved capacity volumes:
  - 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 disk pool.

## 4.3 Supported Connections

Asynchronous mirroring can use either FC connections, iSCSI connections, or both for communication between local and remote storage systems. At the time of creating a mirror consistency group (also known as an asynchronous mirror group), the administrator can select either FC or iSCSI for that group if both are connected to the remote storage system. There is no failover from one channel type to the other.

Synchronous mirroring supports only FC for communication between storage systems.

It is possible for a controller to receive host I/O through one protocol and use a different protocol for mirroring with a remote storage system. For example, a host might be attached to the controller through a SAS connection, while 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 mirror pair only interacts with controller A on the other end.
- Controller B on one end of a mirror pair only interacts with controller B on the other end.

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

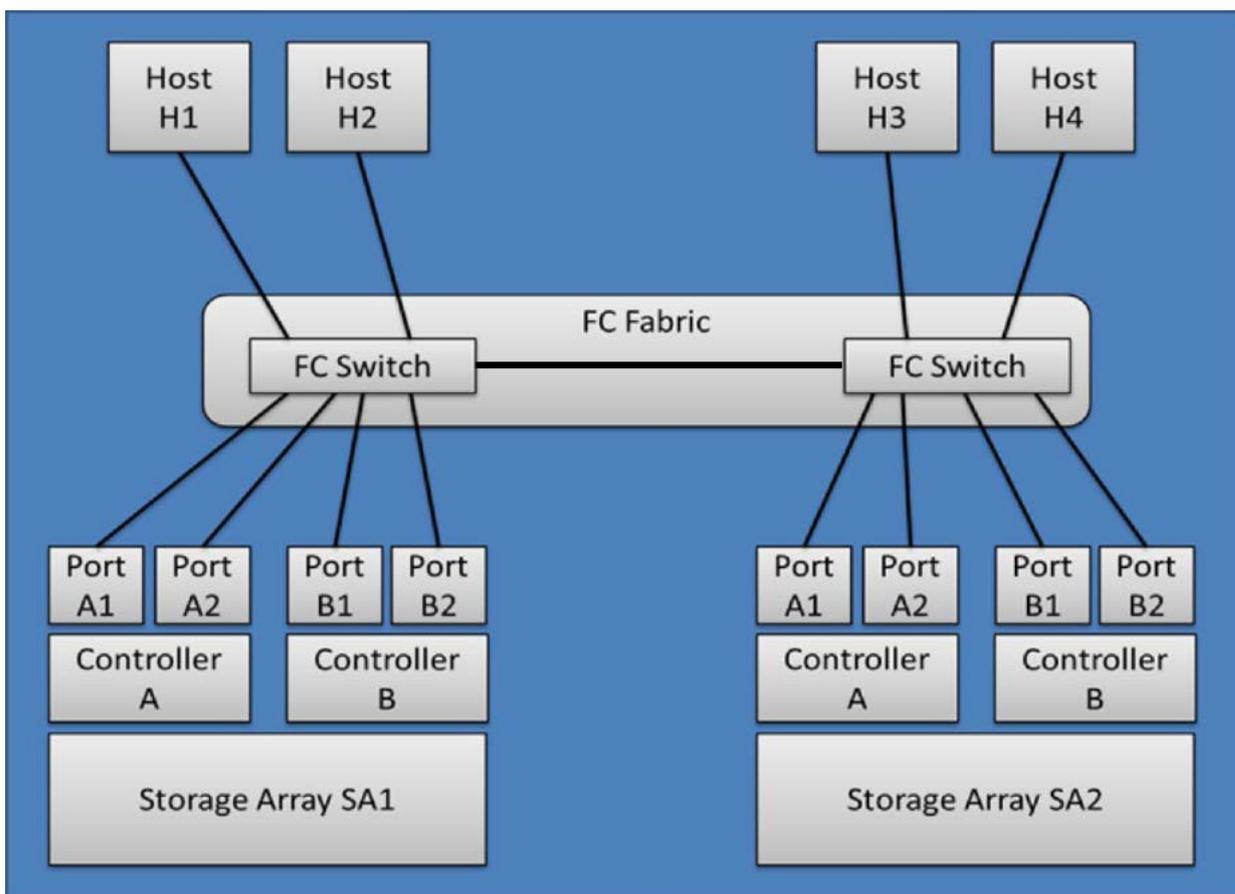
## FC Connection Requirements

When either mirroring feature is activated, each controller of the storage system dedicates its highest numbered FC host port to mirroring operations. Note that if the controller has both base FC ports and host interface card (HIC) FC ports, the highest numbered port is on an HIC. Any host logged onto 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. The following examples illustrate FC configurations supported for mirroring.

In the configuration shown in Figure 5, a single fabric is used to provide total connectivity among all participating devices, including host systems and storage systems.

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

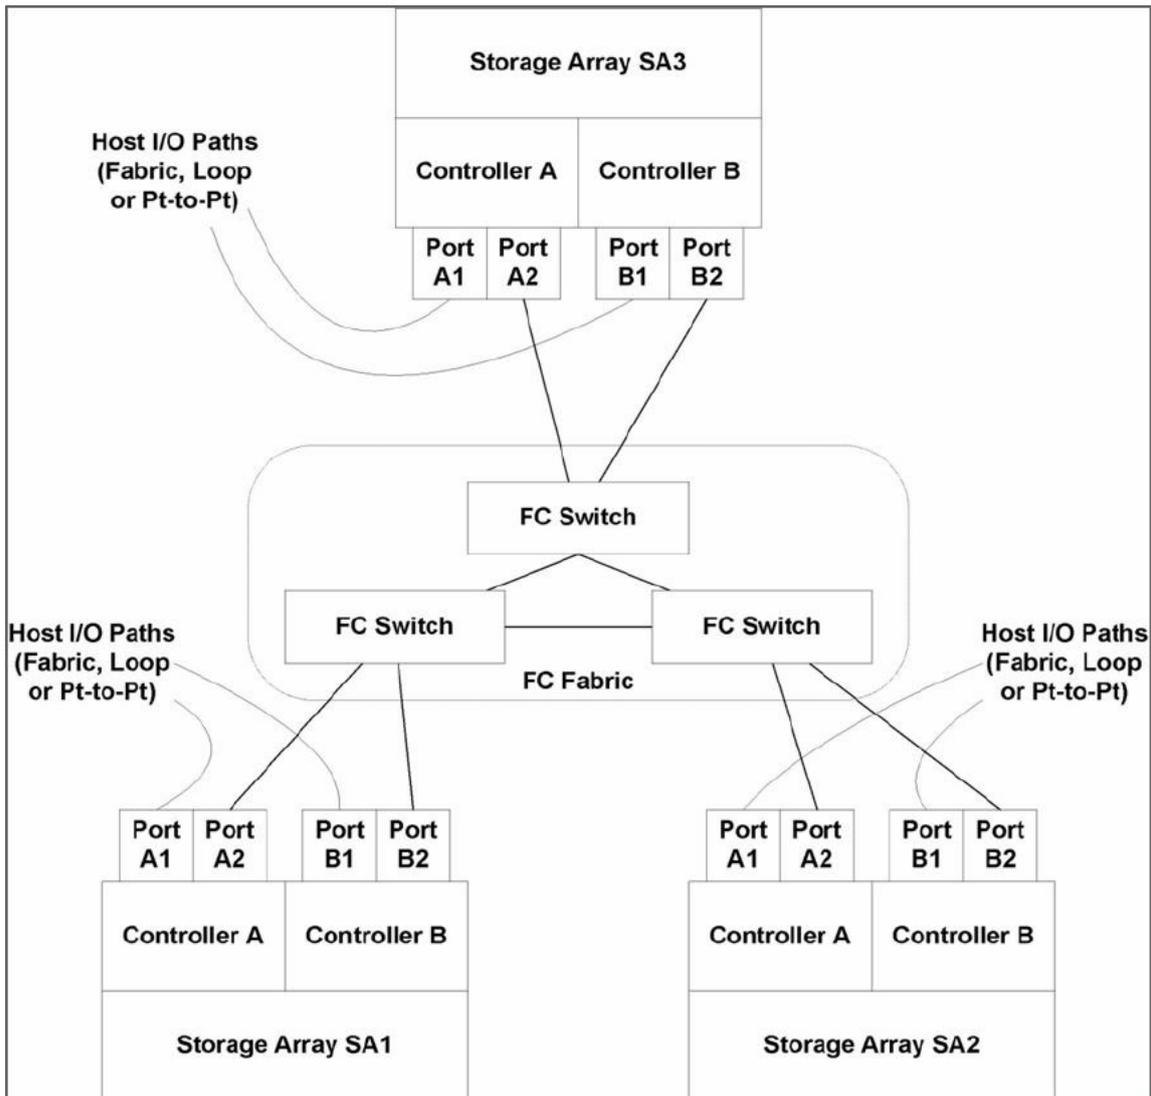
For mirror operations:

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

The assumption here is that H1 and H2 constitute one partition of the computing environment, while 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 in the event of an outage, using the mirrored volume or volumes as necessary during the outage.

In a configuration having a fabric dedicated to mirroring, as in Figure 6, the host traffic is assumed to connect to the storage systems 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 activities.

Figure 6) Fabric dedicated to mirroring.



Again, note that fabric connectivity paths exist that are never used. The only paths that are used in the illustrated fabric for synchronous mirror 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 Connection Requirements (Asynchronous Mirroring Only)

Unlike FC, iSCSI does not require a dedicated port. The controller maintains a list of remote storage systems 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 system. If communication fails, a new session is attempted using all available ports.

iSCSI ports are configured at the system level on a port-by-port basis. Intercontroller communication for configuration messaging and data transfer uses the global settings, including settings for the following:

- VLAN: Both local and remote systems must have the same VLAN setting to communicate.
- iSCSI listening port
- Jumbo frames
- Ethernet priority

Asynchronous mirroring uses the storage system'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 asynchronous mirroring is used in iSCSI environments, it is not necessary 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.

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

**Note:** Activation is not required if the storage system is mirroring only through an iSCSI interface.

### 4.4 Distances and Speed

Synchronous mirroring is limited to a maximum of 10km in distance between arrays (measured by total cable length). 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 2) 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 3.

Table 3) 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

## 4.5 Mirror Volume Requirements

Following are requirements and characteristics of mirrored volumes:

- The reported capacity of a volume participating in mirroring is the lesser of the primary and secondary capacities. 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 three-data center mirroring or n-way 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 (DA), the secondary must also have DA enabled.
- Regarding drive security, a best practice is for the primary and secondary volumes to be of the same drive type (FDE, FIPS, or not security capable). They should also have the same security setting (secure enabled or not):
  - 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 be not secure capable. If the primary is secure enabled, the secondary must 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 primary and secondary volumes should be the same drive type (FDE, FIPS, or not security capable).
- There are a number of allowed and restricted candidates to become a mirror volume (see Table 4).

Table 4) Allowed candidates for mirror volumes.

Mirror Volume Type	Allowed Candidates	Disallowed Candidates
Asynchronous primary	<ul style="list-style-type: none"> <li>• Standard volume (only if the secondary is also standard)</li> <li>• Snapshot base volume</li> <li>• Volume copy source</li> <li>• Thin volume (only if the secondary is also thin)</li> </ul>	<ul style="list-style-type: none"> <li>• Async or sync mirror primary</li> <li>• Async or sync mirror secondary</li> <li>• Volume copy target</li> <li>• Snapshot volume</li> </ul>
Asynchronous secondary	<ul style="list-style-type: none"> <li>• Standard volume (only if the primary is also standard)</li> <li>• Thin volume (only if the primary is also thin)</li> </ul> <p><b>Note:</b> A secondary can become a volume copy source if a role reversal occurs after a volume copy on the original primary completes. If the role reversal occurs during a copy in progress, the copy fails and cannot be restarted.</p>	<ul style="list-style-type: none"> <li>• Async or sync mirror primary</li> <li>• Async or sync mirror secondary</li> <li>• Volume copy source</li> <li>• Volume copy target</li> <li>• Snapshot base volume</li> <li>• Snapshot volume</li> </ul>
Synchronous primary	<ul style="list-style-type: none"> <li>• Standard volume</li> <li>• Snapshot base volume</li> <li>• Volume copy source</li> <li>• Volume copy target</li> </ul>	<ul style="list-style-type: none"> <li>• Async or sync mirror primary</li> <li>• Async or sync mirror secondary</li> <li>• Snapshot volume</li> <li>• Thin volume</li> </ul>

Mirror Volume Type	Allowed Candidates	Disallowed Candidates
Synchronous secondary	<ul style="list-style-type: none"> <li>Standard volume</li> </ul> <p><b>Note:</b> A secondary can become a volume copy source or target if a role reversal occurs after a volume copy on the original primary completes. If the role reversal occurs during a copy in progress, the copy fails and cannot be restarted.</p>	<ul style="list-style-type: none"> <li>Async or sync mirror primary</li> <li>Async or sync mirror secondary</li> <li>Volume copy source</li> <li>Volume copy target</li> <li>Snapshot base volume</li> <li>Snapshot volume</li> <li>Thin volume</li> </ul>
Reserved capacity	<ul style="list-style-type: none"> <li>Standard volume</li> </ul>	<ul style="list-style-type: none"> <li>Async or sync mirror primary</li> <li>Async or sync mirror secondary</li> <li>Volume copy source</li> <li>Volume copy target</li> <li>Snapshot base volume</li> <li>Snapshot volume</li> <li>Thin volume</li> </ul>

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

## 4.6 Reserved Capacity Volume Requirements

For asynchronous mirroring, reserved capacity volumes are created at the time 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 (DA) attribute as their associated mirror volumes.
- If a mirror volume has security enabled, its associated reserved capacity volume must also have security enabled. Note: The type of security should be the same also (FDE vs. FIPS), but this requirement is not enforced by the controller.

For synchronous mirroring, one reserved capacity volume is created for each controller when the feature is activated. The characteristics are as follows:

- The volumes can be on a pool or a volume group of any RAID level except 0.
- From Unified Manager, the pool/volume groups that contain the free space in the system are automatically selected and created the reserve space.

## 4.7 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 processed automatically triggers a matching ownership change on the secondary side if one of these conditions exists:

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

For example, suppose that 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 by the storage administrator.

## 4.8 Storage System Limits

ETERNUS AB/HB series systems have different limits on the number of volumes and mirrors supported, as is shown in Table 5.

Table 5) Storage system limits for mirroring.

Maximum	HB1100, HB1200	HB2000, AB2100	HB5000, AB5100	AB6100
Volumes per system (including reserved capacity volumes)	256	512	2,048	2,048
Mirrors per system (both async and sync)	—	32	128	—
Async mirror consistency groups per system	4	4	4	—
Volumes per async mirror consistency group	32	32	64	—
Async mirrors per system	32	32	128	—
Sync mirrors per system	—	16	128	—

## 5 Asynchronous Mirroring Operational Model

Asynchronous mirroring operates at the volume level and performs incremental, block-based replication as frequently as every 10 minutes.

The first step is initialization, the creation of a one-time transfer of the entire primary volume to the secondary. This step is required before incremental updates can be performed. At the end of the initialization, the source system creates a Snapshot copy and begins tracking changed blocks.

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:

1. The source storage system creates a new Snapshot copy and begins tracking changes after this new Snapshot copy.
2. The changed blocks between Snapshot copies are sent to the destination system.
3. After the update is complete, both systems have the new copy, which becomes the baseline copy for the next update.

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

The following sections provide more depth on how the asynchronous mirroring works.

### 5.1 Initial Synchronization

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 would normally have access to it. Until the initial synchronization completes, the secondary volume has minimal value as a recovery point.

The initial synchronization is completed in multiple phases. The first phase copies all data from the primary to the secondary volume. For a thin volume, this approach includes only blocks that have been written on the primary and not unmapped. However, the first phase does not create a Snapshot copy as the synchronization point. The first phase synchronization 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 Snapshot copy on the primary to use as a synchronization point. Not using a Snapshot copy for the first phase of the synchronization process minimizes the effect on performance during the initialization process.

When an asynchronous mirrored pair is first established, the following steps are taken on the primary array:

1. The mirror state is set to initializing.
2. One of the delta logs (log A) in the primary side's reserved capacity is initialized to a clean state and then activated to track updates made 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 out of sync.
4. A background synchronization process transfers the contents of the primary volume to the secondary. The process iterates through delta log B, copying any data 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 may be copied to the secondary, resulting in several bits in the delta log being marked complete in the synchronization process.
5. A Snapshot copy 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 data regions have been copied to the secondary, a message is sent to the remote array so that a Snapshot copy can be created to capture the recovery point on the secondary volume.
7. The Snapshot copy that was created to capture the primary synchronization image is deleted.
8. Steps 5 through 7 are repeated as long as the synchronization process does not complete within the configured synchronization interval.
9. When the background synchronization process completes 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 Snapshot copy 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 side. Because this synchronization is the initial one, the updates are written directly to the secondary's base volume, without first creating a Snapshot copy.
3. The secondary side receives a message from the primary side that the background synchronization process is complete. A Snapshot copy is created to capture the synchronized image.

## 5.2 Resynchronization Methods for Asynchronous Mirroring

After the initial synchronization is complete, and as the primary changes, the system needs to resynchronize every so often to create a new recovery point that meets the desired SLA. Two resynchronization methods, manual and automatic, are available for asynchronous mirroring.

### Manual

- Manual resynchronization causes immediate resynchronization of data on all of the asynchronous mirrored pairs in the asynchronous mirror group. Select this option to manually start resynchronization for all mirrored pairs in the group.
- If the administrator chooses manual as the synchronization method, the administrator must use the manual resynchronization option to send updates of modified data from the local storage array to the remote storage array.

### Automatic (Periodic)

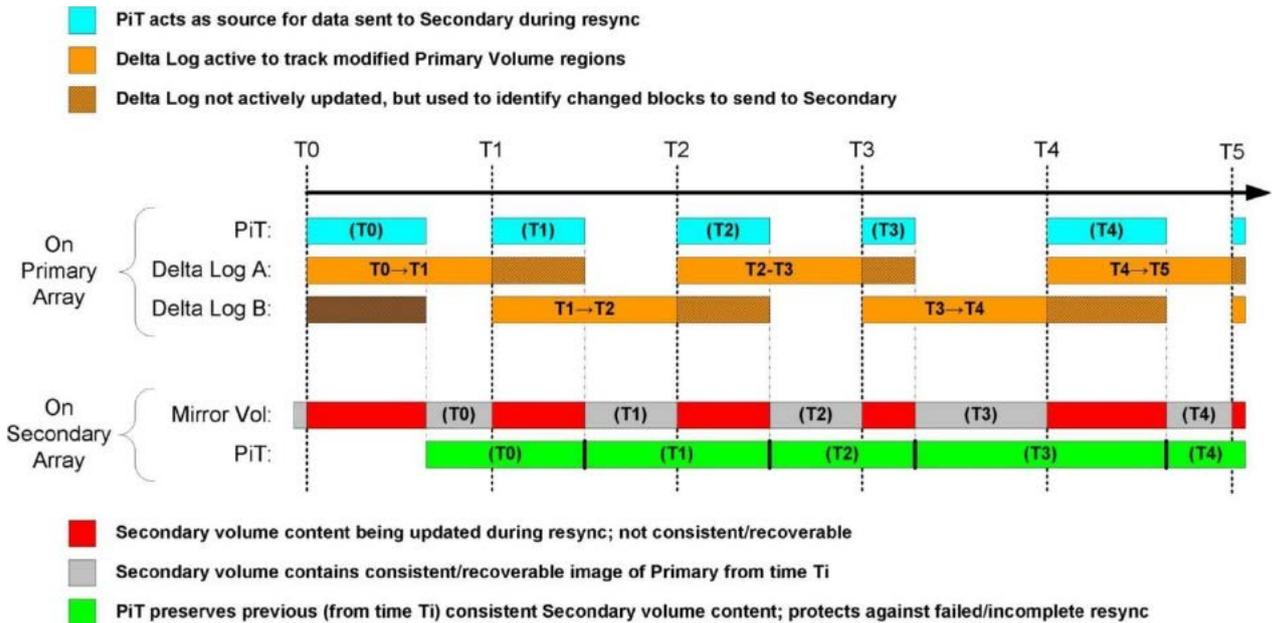
- Specifies the time (in minutes) from the beginning of the previous update to the beginning of the next update. For example, if the synchronization interval is set at 30 minutes, and the synchronization process starts at 4:00 p.m., the next process is started at 4:30 p.m.
- To change the automatic synchronization interval from the default of every 10 minutes, the administrator can edit the interval value, which is defined in minutes.
- If a communication failure occurs between the local storage array and the remote storage array and a synchronization interval was missed during that time period, the controller owner of the primary asynchronous mirror group starts the resynchronization process immediately after communication has been restored.
- If the administrator chooses automatic synchronization, the administrator can also execute a manual resynchronization at any time by using the manual resynchronization option.

## 5.3 Periodic Resynchronization

Each mirror consistency group (asynchronous mirror group) has a configurable attribute that specifies the resynchronization interval for all mirror pairs in the group. This interval is the amount of time between automatically sending updates of modified primary-side data to the secondary side. The value represents the time between the starting points of sending updates from primary to secondary.

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. The actual amount of time to complete a resynchronization varies due to differing quantities of modified data on the primary volume and due to performance variations in the intercontroller communication link. For example, the first resynchronization cycle may 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 7) Periodic resynchronization in asynchronous mirroring.



**Note:** In Figure 7, PiT refers to a point-in-time image, which is the same as a Snapshot copy.

When a resynchronization interval arrives (for example at time  $T_1$ , as shown in Figure 7), the controller on the primary side of the mirror pair manages the synchronization process with the following steps:

1. The mirror synchronization activity state changes to active.
2. A Snapshot copy is created to capture the  $T_1$  image of the primary volume.
3. Delta log B (DL-B) is cleared and activated to track primary volume updates that occur after  $T_1$ .
4. Delta log A (DL-A) is deactivated to no longer track writes. A background synchronization process transfers the regions of the primary  $T_1$  Snapshot 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 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 resulting from write requests to the primary volume are minimized by performing the copy-on-write only for regions 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 affects a region that is not flagged, there is no need to save off the original data because the synchronization process never accesses that region.
6. When the background synchronization process completes (all flagged regions from delta log A have been copied to the secondary), the mirror synchronization activity state is set to idle.
7. The Snapshot copy that was created to capture the initial synchronization image is deleted, leaving no active Snapshot copies on the primary mirror repository.

The following steps occur on the secondary side:

1. The secondary side 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 Snapshot copy was created at the end of the previous synchronization cycle, the consistent secondary image is preserved during the synchronization process.
3. The secondary side receives a mirror synchronization activity state change to idle from the primary side. A Snapshot copy is created to protect the newly established consistent image; any old Snapshot copy 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 that the use of DL-A and DL-B is reversed. On the secondary side, the processing is the same as just described for the first resynchronization interval. This pattern repeats for all subsequent resynchronization intervals.

## 5.4 Using Thin-Provisioned Volumes in Asynchronous Mirroring

Thin-provisioned volumes are allowed to participate in a mirror consistency group (asynchronous mirror group) only when paired with another thin-provisioned volume. Thin-provisioned volumes that are in a group follow the same rules as normal volumes (for example, security settings, data assurance, capacity, and so on). 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. (Note: This approach is the default in the management GUIs when creating a thin volume.) Thin volumes with automatic expansion disabled are not shown as candidates for mirroring in the GUI.
- The secondary thin volume parameters controlling its growth are set to match the primary thin volume parameters. When selecting an existing thin volume for the secondary, the user is warned of these impending changes and allowed to cancel the operation.
- The alert thresholds may only be changed on the primary side of the mirror pair. Any changes to these parameters on the primary 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 controlling its growth.

## 5.5 Suspend and Resume

The administrator can suspend synchronization on a mirror consistency group basis. When this suspension happens, 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 the administrator resumes the mirror consistency group, a resynchronization of all mirror 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.

## 5.6 Orderly Role Reversal

At certain times, an administrator might want to reverse the roles of primary and secondary. This reversal could be when it looks like a disaster at the primary site is possible and the company wants to migrate operations to the recovery site. Another example might be when conducting scheduled maintenance at the primary site, the administrator can reverse roles to preserve continuity of service. Later, the roles can be reversed again to restore the original site to normal operations.

The role reversal change affects all asynchronous mirrored pairs in the selected mirror consistency group (asynchronous mirror 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 change role operation.

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

## 5.7 Forced Promotion of Secondary

In the event of a catastrophic failure at the primary site, the administrator 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 primary.

Forced promotion first transitions the mirror consistency group to a role-change-in-progress state on the local storage system. Because the remote storage system containing the primary group is inaccessible, this state change is only made locally. Next, the mirror consistency group role is set to primary, which allows 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 mirror pairs prior to 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 is due to a communication link failure and not the primary being destroyed, and if the original primary is 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, the administrator is notified by a needs-attention alert. The administrator needs to follow the Recovery Guru procedures to resolve this conflict.

**Note:** The preceding description of the role change applies when the controller at the secondary site cannot communicate with the primary. If the controller is able to 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 reversing roles, resulting in any data that had been written to the primary since the last resynchronization being lost.

Forced promotion is not allowed if any of the following conditions exist:

- Not all of the mirror consistency group member volumes have been previously synchronized and have a consistent recovery point.
- Any member volumes do not have a Snapshot 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.

## 5.8 Forced Demotion of Primary

Forced demotion of a mirror consistency group in the primary role to function in the secondary role allows the storage administrator to protect the volumes from receiving new write requests while disconnected from the remote storage system. This approach can be useful in the event that the corresponding group on the remote system had been promoted with the force option to primary.

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 system. Because the remote storage system is inaccessible, this state change is only made locally. If the original primary was actively synchronizing the mirror pairs prior to the forced role change, the synchronization operation is canceled, and the Snapshot copies 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, the administrator is notified by a needs-attention alert. The administrator needs to follow the Recovery Guru procedures to resolve this conflict.

**Note:** The preceding description of the role change applies when the controller at the primary site cannot communicate with the secondary. If the controller is able to 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 reversing roles, resulting in any data that had been written to the primary since the last resynchronization being lost.

Forced demotion of the primary mirror consistency group is not allowed 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.

## 5.9 Example Failures and How to Handle

Following are a few failure modes with guidelines on how an administrator might handle them.

**Note:** These are not the only steps required to handle site failures. The steps listed here are related only to the two storage systems involved in a mirroring relationship. They do not discuss moving applications over from a primary site to a secondary, reestablishing communications, and so on.

## Communication Lost: Secondary Not Promoted

This situation could occur under a number of 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.

## Communication Lost: Secondary Promoted

If communication between the storage systems is lost, and host systems also cannot access the primary mirror consistency group, the administrator might want to promote the secondary so that operations can continue. Following are high-level steps of operation and recovery:

1. If possible, administrator demotes the primary to prevent any further writes.
2. Administrator promotes the secondary, and 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. The newly promoted primary is now read/write accessible, and all writes are logged.
4. After communication is restored, periodic resynchronization from the newly promoted primary is automatically resumed. Note that it may take about 30 minutes to resume.
5. When ready, the administrator can initiate an orderly role reversal to restore normal operation.

## Failure at Primary Site with Storage System Recoverable

A failure such as a prolonged power outage at the primary site that renders the primary storage system inoperable but can be recovered later can be handled as follows:

1. Administrator promotes the secondary, and 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. The newly promoted primary is now read/write accessible, and all writes are logged.
3. When the primary site becomes operable again, before restoring communication with the remote site, the administrator should force-demote the original primary mirror consistency groups to secondary.
4. After communication is restored, periodic resynchronization from the newly promoted primary is automatically resumed. Note that it may take about 30 minutes to resume.
5. When ready, the administrator can initiate an orderly role reversal to restore normal operation.

## Failure at Primary Site, Storage System Destroyed

This situation is the classic disaster scenario in which the primary site is sufficiently damaged that it needs new storage equipment as part of the rebuilding process.

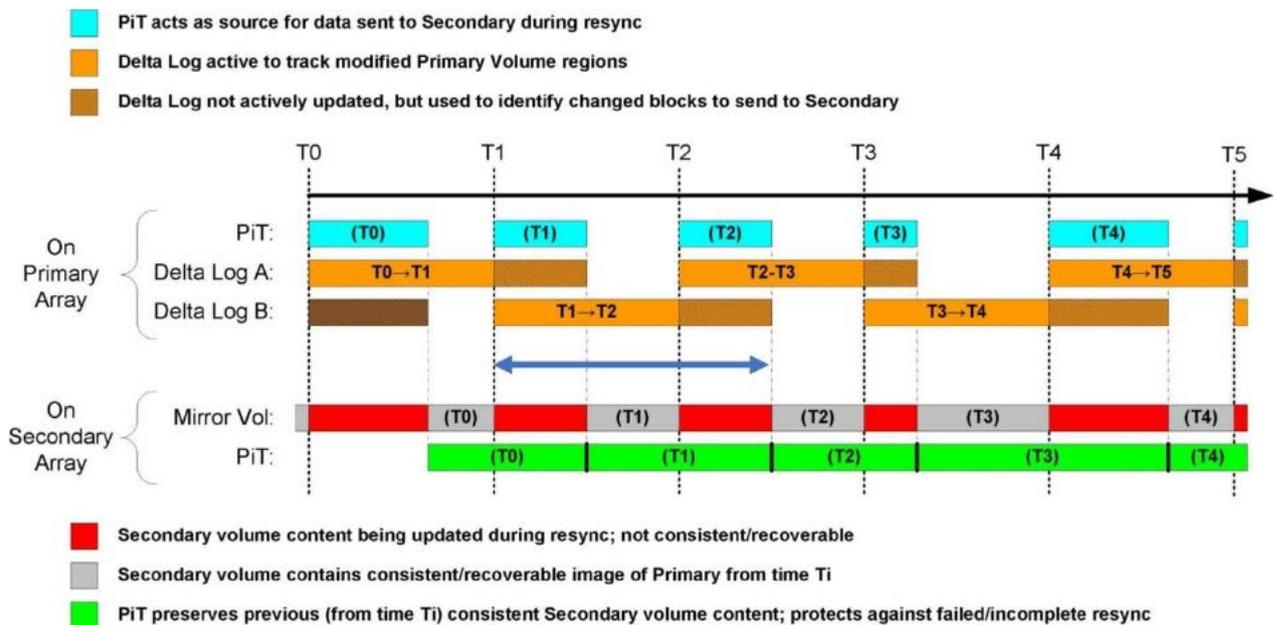
1. Administrator promotes the secondary, and operations proceed from the last recovery point. Note that any writes that have occurred on the original primary since the last recovery point are not reflected at the newly promoted primary.
2. Administrator removes member volumes from the mirror consistency groups, then force-deletes the groups.
3. After a storage system is available at the primary site (or a third site), the administrator creates new mirror consistency groups and adds the volumes back in.
4. At this point, an initial synchronization takes place, after which normal periodic resynchronizations resume.
5. When ready, the administrator can initiate an orderly role reversal to restore normal operation.

## 6 Considerations for Setting Up Asynchronous Mirroring

### 6.1 RPO and Synchronization Interval

The recovery point objective (RPO) for asynchronous mirroring is the maximum tolerable time during which changes made on the primary side are subject to loss in the event of a primary side disaster. For a single interval, the time subject to loss is from the beginning of the synchronization interval to the time the secondary contains a Snapshot image consistent with the contents of the primary at the beginning of the next interval. Note the blue line in Figure 8. At time T1, the primary takes a Snapshot copy and begins resynchronizing the secondary. From that point on, new writes to the primary are not reflected at the secondary, and thus not protected from a disaster at the primary, until T2 plus the time to resynchronize that interval. The RPO is the maximum allowable for all of these out-of-synchronization time periods.

Figure 8) Time from one protection point to the next.



Several factors affect the ability to meet the RPO. These include the synchronization interval, the size of the dataset and rate of change, and the bandwidth of the communication between primary and secondary. 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. Some shorthand notation:

- **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:

$$SI + MRT \leq RPO$$

Also:

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

Choosing values for these 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 of resynchronization not completing by the next interval, thus affecting RPO achievement. If, for very short synchronization intervals,  $SI \gg MRT$ , then the bandwidth of the communication link must be greater, leading to potential unnecessary expense. This situation leads to two general rules:

- It is a good general rule to 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 interarray communication components. The administrator can choose other values for the contingency if desired; for example, the administrator might choose smaller values if the amount of change during peak periods is well known, or a bit more risk in RPO achievement is acceptable.

Taking the contingency into account, the synchronization interval would be:

$$SI = RPO - MRTC = MRTC$$

For example, suppose 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.

## 6.2 Sizing the Network to Meet RPO

It is very important to make sure that the communication link can handle the bandwidth required to resynchronize the maximum expected changes within the synchronization interval chosen. Bandwidth is limited to the slowest segment in the network. Also, latencies can reduce the maximum achievable bandwidth on IP networks as well.

After the maximum actual resynchronization time is known, the maximum required bandwidth can be calculated if the administrator knows or can estimate how much data will change during any given synchronization interval. One way to arrive at this situation is to estimate the portion of the total dataset that is changed on a peak day or peak hour. For example, if the dataset is 20TB, and 10% changes on a peak day, then 2TB would have to be transferred to the secondary in a day. If there is a peak hour during the peak day that gets much of the activity, then that must be taken into account. Suppose that 15% of the activity on the peak day occurs during one hour. Then 300GB would be written in that hour.

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

To determine the required network, the administrator needs to consider the following:

- Maximum round-trip time latency for asynchronous mirroring is 120ms.
- Plan for how much the dataset grows over time. Growth likely requires more data to be written to the secondary during each synchronization interval.
- If using iSCSI or extending FC over the WAN, expect increases in overhead that can reduce the effective bandwidth by up to 50%.
- Fujitsu recommends minimum iSCSI speed of 10Gbps when configuring asynchronous mirroring.
- It is recommended not to share the network with other traffic, which can cause bandwidth to be variable. This situation could prevent adherence to the RPO.

### 6.3 Latency Considerations with iSCSI

Another important factor in the system's ability to support the required RPO is the latency over the communications link, especially when using iSCSI. Although the rated bandwidth of the link may 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 the sender can send without receiving an acknowledgement. A large round-trip time slows down synchronizations and resynchronizations.

Current ETERNUS AB/HB series systems have a TCP window size of 256KiB.

Maximum achievable bandwidth can be calculated by (TCP window size in bits) / (latency in seconds) = throughput in bits per second.

For example, suppose the primary site is in Chicago, and the secondary site is in Salt Lake City, Utah. These cities are about 1,400 miles apart; at the speed of light, communication can travel the distance in 7.5ms. This situation would have a round-trip time of 15ms, not counting any other delays through the actual media and components. If these additional latencies add up to 5ms round trip, then the communication link has a round-trip latency of 20ms. Then the maximum bandwidth that can be achieved by asynchronous mirroring over iSCSI between two current ETERNUS AB/HB series systems is calculated as follows:

TCP window size in bits = 262,144 x 8 = 2,097,152 Latency in seconds = 20 / 1,000 = 0.02

Maximum bandwidth = 2,097,152 / .02 = ~105Mbps

When using iSCSI for asynchronous mirroring, the administrator must take this potential limitation into account.

It is possible to increase the TCP window size in some instances. Contact Fujitsu Support for assistance with this situation.

### 6.4 Time to Complete Initial or Full Synchronization

When configuring mirroring, the administrator needs to know how long it takes after creating a mirrored pair or recreating the primary after a disaster before normal operations can begin with a valid recovery point at the secondary. Depending on the size of pipe and how large the dataset is, this approach could take a very long time. Some examples are shown in Table 6.

Table 6) Approximate hours needed for initial or full synchronization by pipe size.

Data to Copy	DS1 (1.544 Mbps)	DS3 (44.736 Mbps)	OC-1 (51.48 Mbps)	OC-3 (155.52 Mbps)	OC-12 (622.08 Mbps)	OC-48 (2.4 Gbps)	OC-192 (9.6 Gbps)	10Gb Ethernet
100GB	143	4.97	4.32	1.43	.357	.0926	.0231	.0222
1TB	1430	49.7	43.2	14.3	3.57	.926	.231	.222
10TB	14,300	497	432	143	35.7	9.26	2.31	2.22

## 6.5 Effect of Asynchronous Mirroring Operations on Performance

Performance varies based on workload, number and type of drives, and how the system is set up. This section offers some guidance about the performance effects of deploying the mirroring features.

This example was performed on 24 10k RPM SAS drives in a single SANtricity Dynamic Disk Pool, using Iometer. This profile was chosen to simulate a typical database. Table 7 summarizes the configuration used in performance testing.

Table 7) 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 Pool

Performance results from the example test are shown in Table 8.

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

Table 8) 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 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 completed. During initialization, CPU utilization for this test was ~55%, which dropped to ~2% after completion. With Snapshot copies enabled, CPU utilization during the Iometer 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 completed. During initialization, CPU utilization for this test on the primary site was ~53%, which dropped to ~2% after completion. CPU on the remote site was ~12%, which dropped to ~2% after completion. With mirroring enabled, CPU utilization during the Iometer run was ~21%. Without Snapshot copies, CPU utilization was ~18%. If Iometer is running during initial synchronization, CPU is ~63%, and performance drops ~60%.
Snapshot and mirroring	1,052	8.24	-1.7%	As with the other testing, after initialization is complete, performance is minimally affected.

When a Snapshot copy is taken, the system must initialize the reserved capacity volume. This situation takes some cycles away from the disks, but not much. The CPU overhead associated with this operation is minimal.

## 7 Synchronous Mirroring Operational Model

Synchronous mirroring operates at the volume level and performs incremental, block-based replication. It is called synchronous because the secondary volume is updated before the controller at the primary acknowledges the write completion to the host.

The first step is the creation of a one-time, baseline transfer of the entire dataset. This step is required before incremental updates can be performed and is called initial synchronization.

After the initialization 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:

1. Changed data blocks are written to the primary volume.
2. The primary storage system sends the changed blocks to the secondary system.
3. The secondary system sends an acknowledgement after the changed blocks were written to cache (or disk in case of writethrough caching).
4. The primary storage system sends an acknowledgement to the host.
5. After the operation is complete, both systems are in 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.

### 7.1 Initial Synchronization

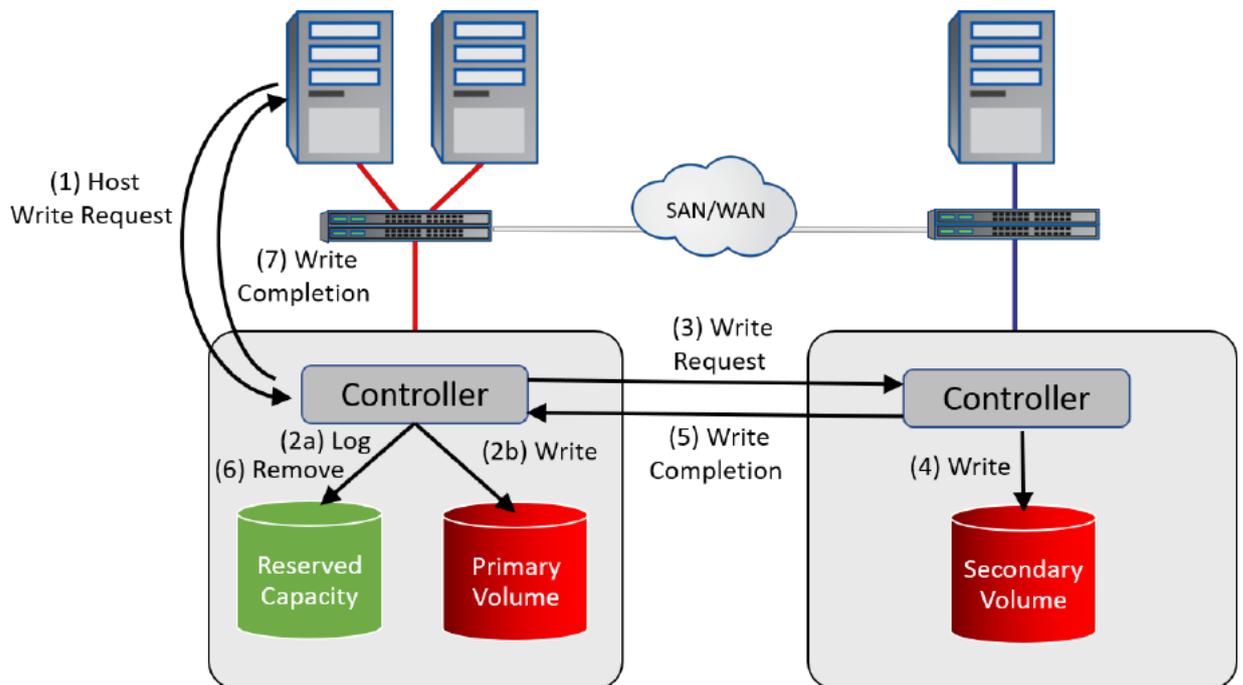
When the mirror relationship is first established, the data from the primary volume is copied in its entirety to the secondary volume. The owning controller on the primary side directs this process. During the copy or mirror initialization, the primary volume is fully accessible for I/O operations by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

During the initial synchronization, all write requests on the primary volume are synchronized with the secondary volume. This situation has to be taken into account for planning the mirror link between the two sites. The link has to be able to support the initial copy load and the write requests at the same time.

### 7.2 Write Process

This process is illustrated in Figure 9. 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 mirror reserved capacity volume (the information is placed in a queue), (2a). In parallel, it writes the data to the primary volume (2b). The controller then initiates a remote write operation to copy the affected data blocks to the secondary volume at the remote site (3). When the remote write operation is complete (4) (5), the primary controller removes the log record from the mirror reserved capacity volume (deletes it from the queue) (6). Finally, the controller sends an I/O completion indication back to the host system (7).

Figure 9) Synchronous mirroring write process.



### 7.3 Resynchronization Methods

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 and they are completed, but remote writes to the secondary volume do not take place. In this case, the mirrored volumes' pair or pairs are running out of synchronization. Some examples of events that cause a mirrored pair to become unsynchronized included 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 takes place. 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 were written to a delta bitmap table, which is part of the reserved capacity volume.

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

Two resynchronization methods are available with synchronous mirroring.

#### Manual

The manual resynchronization option is the recommended method because it allows the administrator to manage the resynchronization process in a way that provides best opportunity for recovering data. When enabled, the administrator can manually start resynchronization of the data on the primary volume and the secondary volume after communication has been restored to the unsynchronized mirrored pair.

When this option is selected and a communication failure occurs between the primary volume and the secondary volume, the remote mirrored changes to unsynchronized status. Any write requests to the primary volume are logged, and a needs attention status appears for the storage array.

After the controller owner of the primary volume detects that communication has been restored, the remote mirror stays in unsynchronized status until the administrator issues a resume command. The resume starts the resynchronization process.

## Automatic

When enabled, automatic resynchronization starts immediately after the controller detects that the communication has been restored for an unsynchronized mirrored pair.

When the automatic resynchronization option is selected and a communication failure occurs between the primary storage array and the secondary storage array, the controller owner of the primary volume starts resynchronizing the primary volume and the secondary volume. This action occurs immediately after the controller owner detects that communication has been restored.

**Note:** Any communication disruptions between the primary storage array and the secondary storage array while resynchronization is under way 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, Fujitsu recommends using manual resynchronization and creating a Snapshot copy of the secondary volume before beginning the resynchronization. Then, if a failure occurs during the resynchronization, there is a crash-consistent image at the secondary at the point where synchronization was lost. Asynchronous mirroring is not subject to this resynchronization issue.

## 7.4 Suspend and Resume

A storage administrator can stop mirror synchronization with a remote array by issuing a configuration request to suspend the mirror. This mechanism allows end users to implement backup and other strategies that are based on the split-mirror model used throughout the industry. When a mirror is in a suspended state, no attempt is made to contact the secondary volume for any reason. While suspended, writes to the mirror primary volume are persistently logged so that upon resumption, only the regions of the primary known to have changed are written to the secondary. The state of the mirror remains suspended until the administrator issues a configuration request to resume synchronization activity.

Suspend and resume commands can only be issued on the primary array.

## 7.5 Role Changes on a Synchronous Mirror Pair

The administrator can perform a role reversal on a synchronous mirrored volume pair. This reversal can be done either by promoting the selected volume to a primary role or by demoting the selected volume to a secondary role.

- The role reversal change affects only the selected synchronous mirrored pair.

**Note:** Make sure that the mirror is synchronized before using the role reversal command to have a consistent dataset.

- 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 mirror pair operation, it resumes during the change role operation.

Keep these guidelines in mind:

- Any hosts that are accessing the primary volume in a synchronous mirror 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 are now able to 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, the administrator 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. Use the Recovery Guru to resolve this condition.
- Likewise, the administrator can force the primary volume to a secondary role if communication is not possible between the two storage systems. When communication is restored, the mirror pair is in a dual secondary condition, which can also be resolved by using the Recovery Guru.

## 8 Considerations for Setting Up Synchronous Mirroring

For a successful deployment of a synchronous mirroring solution, the right networking design between the two sites is critical.

### 8.1 Latency

Latency is the time needed for the mirrored I/O and the acknowledgement frame to get over the network link and back. The longer the distance between the two sites, the longer the round-trip time. The 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, the 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, which is included in the management GUIs.

### 8.2 Distance and Number of IOPS

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

To calculate the right link between the two sites, we need to know what influences the effective IOPS rate comparing to 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 has to travel the distance between the two sites twice (to the remote site and for the acknowledgement 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:  $1\text{sec} / (2 \times \text{distance} / 200,000\text{km/sec}) = \text{max number of IOPS}$

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 (QoS).

### 8.3 Bandwidth

Bandwidth becomes an important factor during the initial synchronization or resynchronization of a mirror pair. Additional bandwidth is required to get the copy pairs in sync in a reasonable period 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 all bandwidth for the resynchronization task. Table 6 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.

## 9 Configuring Through Management GUI

Mirroring can be configured on ETERNUS AB/HB series storage systems using management GUI, CLI, or SANtricity Web Services REST API. This section describes the graphical management interface SANtricity System Manager. The workflows differ slightly depending on which type of systems are involved; for proper operation, it is critical that the administrator pay careful attention to which case applies to the administrator's situation.

### 9.1 Use Cases

SANtricity System Manager. Element manager for HB2000, HB5000, and AB5100. It is embedded in the storage system itself, is browser-based, and manages only the system on which it resides. System Manager can be launched by pointing a browser at a controller.

Table 9) Management use cases for mirroring.

Primarily Managed By	Secondarily Managed By	Asynchronous Mirroring High-Level Workflow	Synchronous Mirroring High-Level Workflow
System Manager (SANtricity OS 11.62 and later)	System Manager (SANtricity OS 11.62 and later)	<ol style="list-style-type: none"><li>1. From Unified Manager, select the storage system that contains the primary volume(s).</li><li>2. Create a mirror consistency group.</li><li>3. Select the primary volume.</li><li>4. Allocate reserved capacity on the primary.</li><li>5. Select the secondary volume.</li><li>6. Allocate reserved capacity on the secondary.</li></ol>	<ol style="list-style-type: none"><li>1. From Unified Manager, select the storage system that contains the primary volume(s).</li><li>2. Select the primary volume.</li><li>3. Select the secondary volume.</li><li>4. Select the synchronization settings.</li></ol>

Later sections show these steps in greater detail.

## 9.2 Management Station Port Numbers

Fujitsu included a web server to enable Unified Manager to discover and handle management operations (including mirroring) for storage systems managed by SANtricity System Manager. For the web server to start and thus be able to configure mirroring, the administrator must make sure that the required ports are open on the management station. These ports are hard-coded and are listed in Table 10.

Table 10) Port numbers used by SANtricity OS releases.

SANtricity Release (Beginning with 11.62.12)	Ports Used by Web Server
SANtricity OS: 11.62.12 and later	49480, 49494, 49443, and 61616

Unified Manager must discover, manage, and facilitate the configuration of mirroring involving systems managed by System Manager.

## 9.3 Web Services Certificates

For mirroring involving systems managed by SANtricity System Manager, Fujitsu recommends importing trusted certificates for the web services in the Unified Manager that allow the storage systems to authenticate with the Unified Manager web server. The steps within Unified Manager are as follows:

1. Generate a certificate signing request (CSR) for the machine where Unified Manager is installed.
2. Send the CSR to a certificate authority (CA).
3. When the CA sends back the signed certificates, import in Unified Manager.

Figure 10 and Figure 11 show the beginning of the process to generate a CSR in Unified Manager.

Figure 10) Web services certificate management in SANtricity Unified Manager.

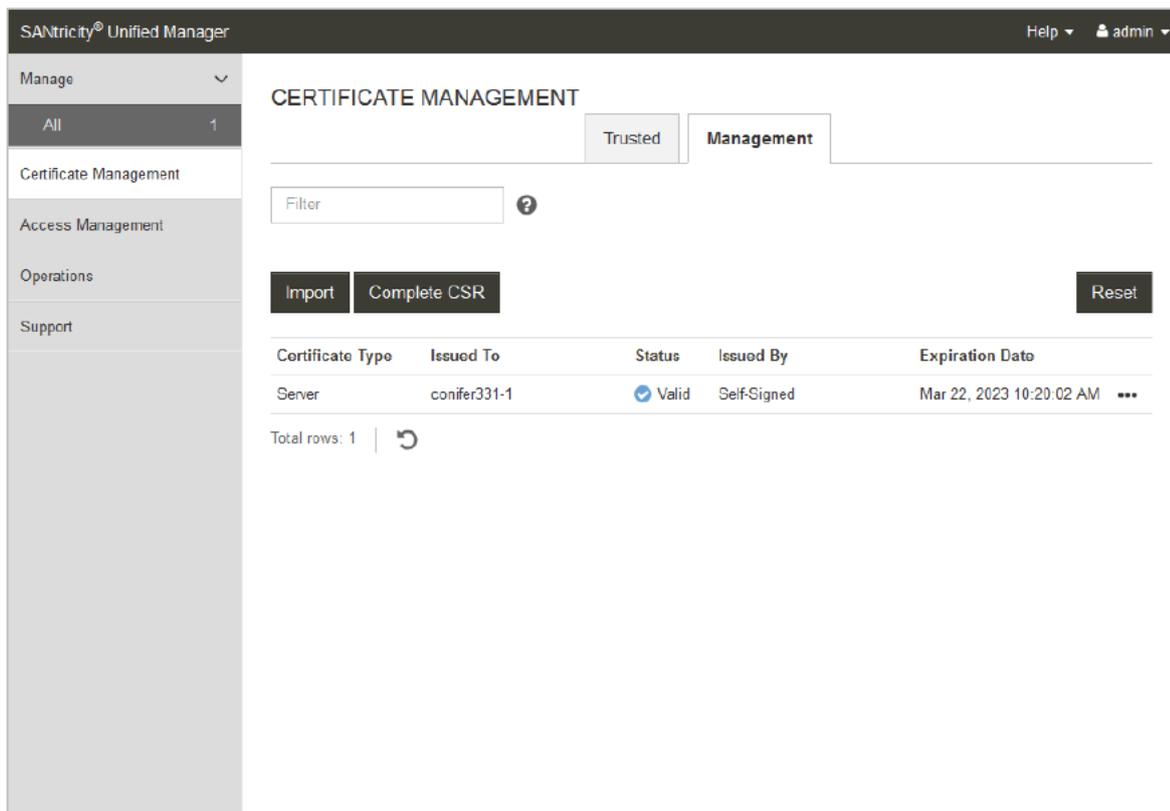


Figure 11) Completing CSR for web services.

**Complete & Download a Certificate Signing Request**

**1 Complete General Information**    **2 Complete System Information**

This information will be saved to a .CSR file. After you obtain the appropriate certificates, you can import them by going to **Settings Certificate Management** and selecting **Import** in the **Management** tab. Because a CSR is associated with a particular management server certificate, do not create another CSR before you import the certificate or that certificate will not be valid.

Organization ?

Organizational unit (optional) ?

City/Locality

State/Region (optional) ?

Country ISO code ?

Cancel    Next >

Self-signed certificates can also be used. If the administrator attempts to configure mirroring without importing signed certificates, System Manager displays an error dialog box that allows the administrator to accept the self-signed certificate. In this case, Fujitsu recommends using the latest version of Chrome or Firefox as the browser.

## 9.4 Mirroring and Role-Based Access Control (RBAC)

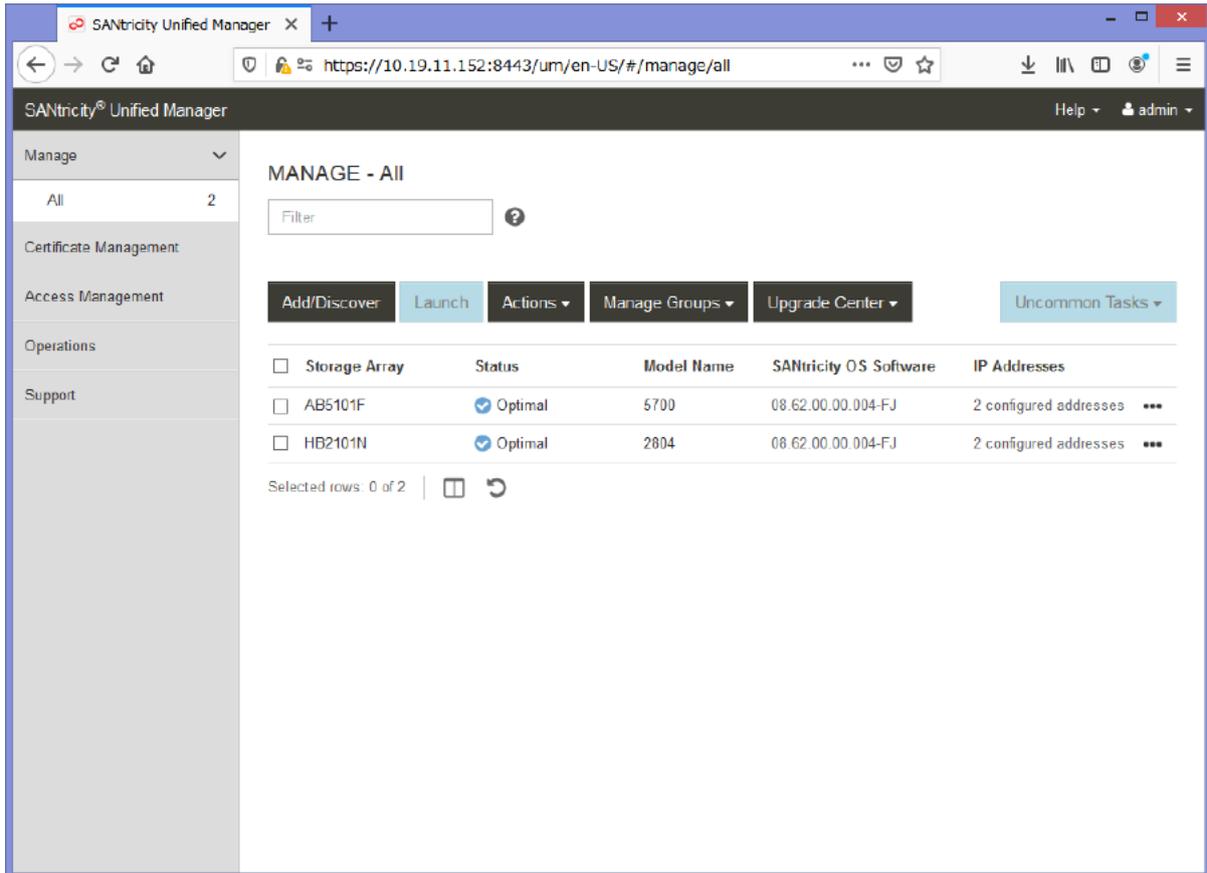
To configure mirroring, note the following requirements when RBAC is being used:

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

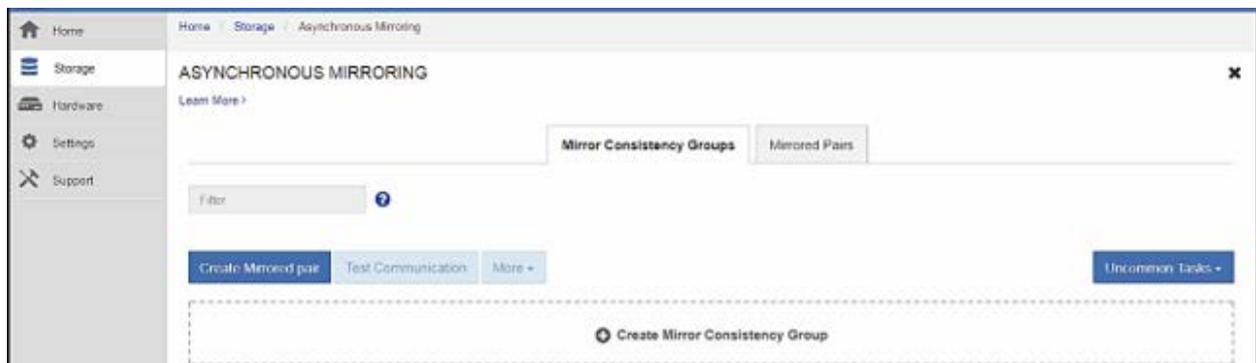
## 9.5 Asynchronous with System Manager as Primary

In this case, the storage system containing the primary volumes is managed by System Manager (HB2000, HB5000 or AB5100), and the secondary is any ETERNUS AB/HB series system that supports asynchronous mirroring. The configuration begins with the SANtricity Unified Manager. Following are the steps:

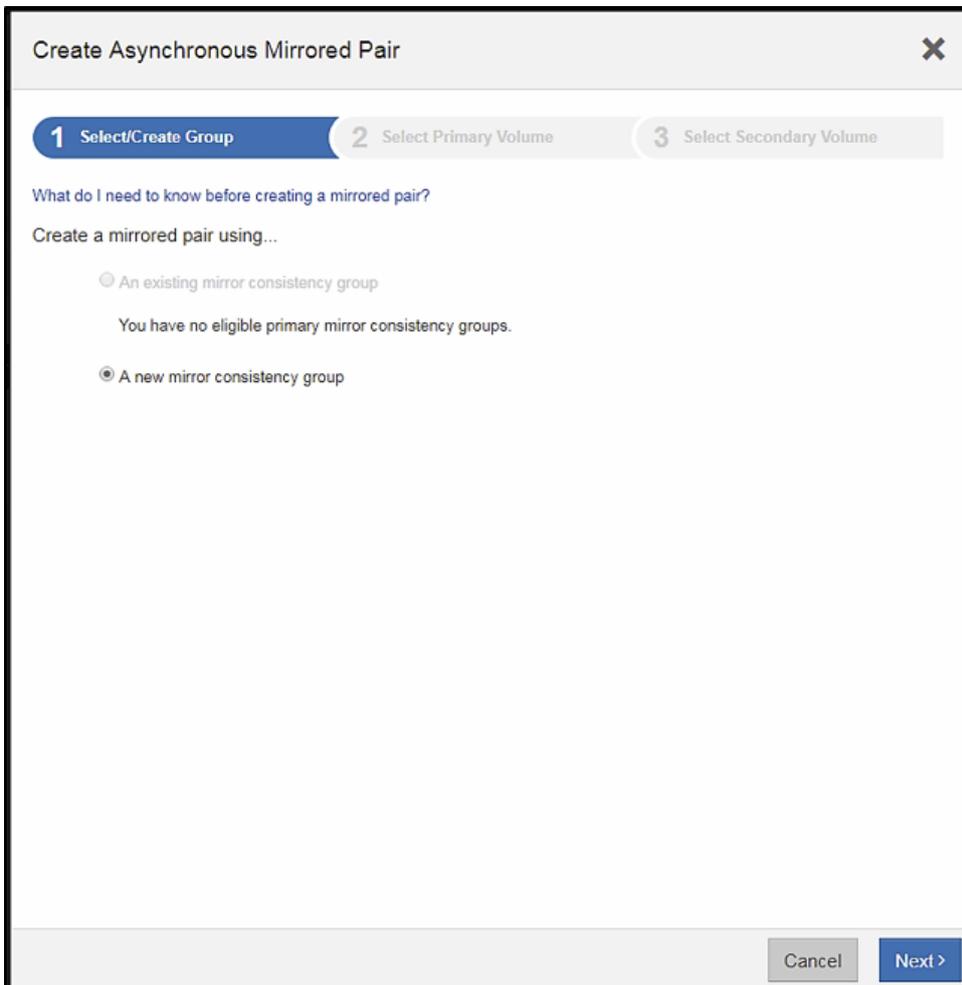
1. Make sure the versions are compatible as described earlier so that the ports are correct.
2. Make sure that both arrays are discovered and listed in the Unified Manager.
3. Recommended: Import trusted certificates in Unified Manager as described earlier.
4. Check the box, and then click the Launch button of the array that contains the primary.



5. Navigate in System Manager to Storage > Asynchronous Mirroring. Click Create Mirrored Pair.



6. The Create Asynchronous Mirrored Pair wizard appears. If there is not already a mirror consistency group created, or if you would like a new one, select the A New Mirror Consistency Group radio button.



7. You can now create the new group. On this screen, complete the following steps:
  - a. Name the mirror consistency group.
  - b. Select the remote storage array on which the secondary volumes reside. System Manager shows all the eligible arrays in the drop-down menu.
  - c. Select the synchronization settings desired to control how often the mirrored pairs in the group resynchronize. Manual is unusual, but sometimes used when a recovery point is only needed, for example, once a day and the admin wants to resynchronize only during periods of lower I/O activity to minimize any effects on performance. For automatic, set the synchronization interval. Ten minutes is the minimum.
  - d. Set the alert for synchronization time. If desired, this time can be set to alert when any periodic resynchronization takes longer than the set value. This approach might be useful to learn when peak time resynchronizations have grown beyond expectations. If the alert is not needed, set to 0.
  - e. Set the alert for recovery point age. This warning can alert the administrator when the RPO is not met for any given interval. It is triggered when the current recovery point on the secondary is older than the specified time. It should be set to the RPO time, but not less than twice the synchronization interval. Set to 0 if using manual resynchronization.

- f. Set the reserved capacity threshold. This threshold alerts when the amount of change during synchronization intervals grows so much that the reserved capacity volume must hold more than the threshold amount. The administrator can then increase the amount of reserved capacity on the primary.

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 mirror consistency group?

Mirror consistency group name ?

Remote storage array

creG1f-SH-A4 (Fibre Channel) ▼

**Synchronization settings**

Synchronize the mirrored pairs...

Manually

Automatically, every - 10 + Minutes ▼ ?

**Alert settings**

Alert me...

If synchronization lasts more than - 10 + Minutes ▼ ?

When my remote recovery point image is older than - 20 + Minutes ▼ ?

When my reserved capacity threshold reaches - 80 % + ?

< Back
Cancel
Next >

8. Select the primary volume. System Manager shows all eligible candidates on the remote array.

### Create Asynchronous Mirrored Pair ✕

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

Why don't I see all my volumes? **Step 2a - Select Volume**

Select your primary volume from creG1f-SH-A1...

Name	Capacity (GiB)
arvm-source-volume	500.00

< Back    Cancel    Next >

9. Select reserved capacity for the primary volume. It can be on any pool or volume group that meets the requirements for protection, capacity, and drive types.

### Create Asynchronous Mirrored Pair ✕

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

I want to reserve capacity that is... ? **Step 2b - Reserve Capacity**

%  of the primary volume capacity (500.00 GiB) = 100.00 GiB.

Why would I change this percentage?

?

Proposed Reserved Capacity (GiB)	Pool/Volume Group	RAID Level
100.00	arvm-source-pool	Pool

10. Select the secondary volume. All valid candidates are shown. This approach is still using System Manager on the primary.

### Create Asynchronous Mirrored Pair ✕

1 Select/Create Group2 Select Primary Volume3 Select Secondary Volume

**Primary volume attributes** Step 3a - Select Volume

Reported capacity: 500.00 GiB  
Data Assurance (DA) Yes  
enabled:  
Secure-capable: Yes - FDE (🔒 not enabled)  
Volume type: Thick

Why don't I see all the volumes on the remote storage array? Show non-eligible volumes

Select your secondary volume from creG1f-SH-A4...

Name	Capacity (GiB)
arvm-remote-volume	500.00

< BackCancelNext >

- Select reserved capacity for the secondary volume. It can be on any pool or volume group that meets the requirements for protection, capacity, and drive types.

### Create Asynchronous Mirrored Pair ✕

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

I want to reserve capacity that is... ? **Step 3b - Reserve Capacity**

of the secondary volume capacity (500.00 GiB) = 100.00 GiB.

Why would I change this percentage? ?

?

Proposed Reserved Capacity (GiB)	Pool/Volume Group	RAID Level
100.00	arvm-remote-pool	Pool

< Back
Cancel
Finish

The mirror pair is now created. The system begins initial synchronization.

Home > Storage > Asynchronous Mirroring
✕

## ASYNCHRONOUS MIRRORING

Learn More >

Mirror Consistency Groups
Mirrored Pairs

Why do I see Not Available values displayed in the table?

Create Mirrored pair
Test Communication
More >
Uncommon Tasks >

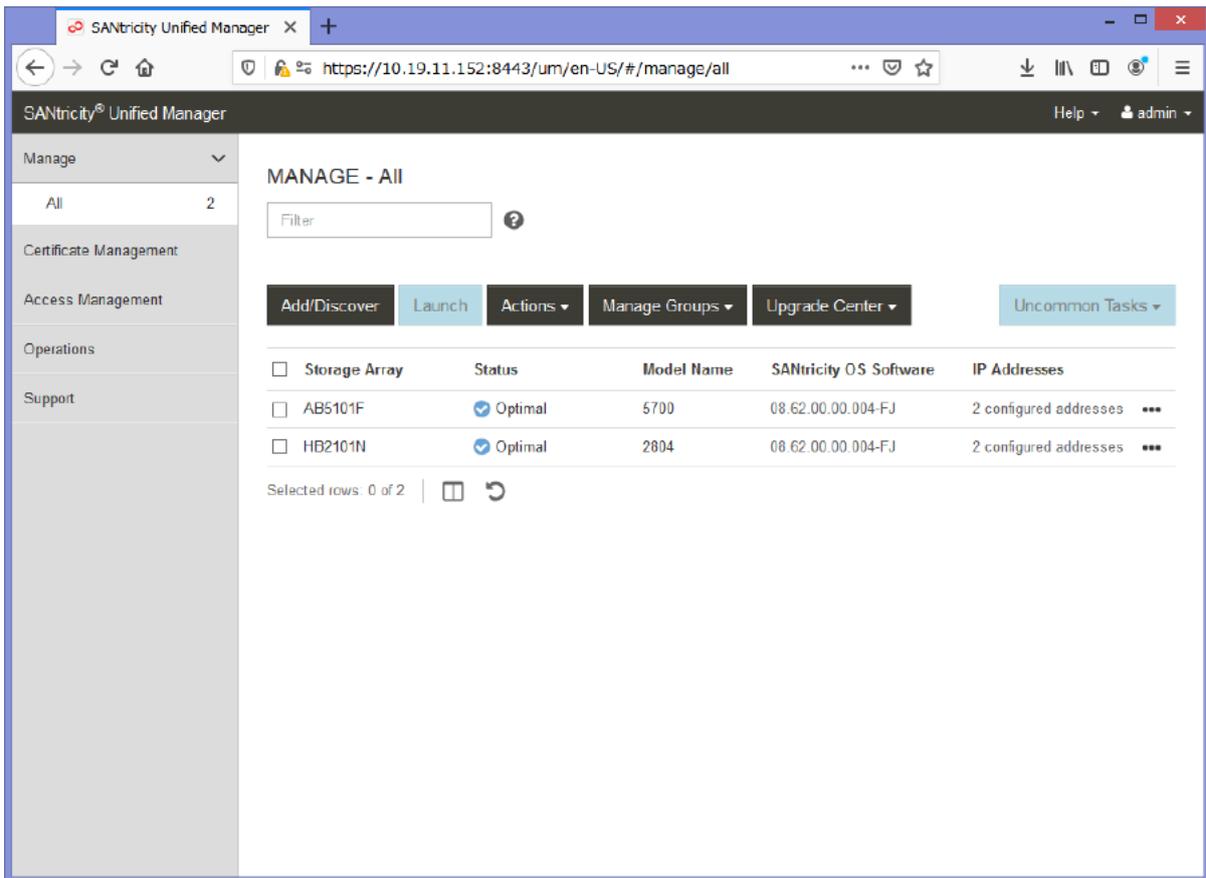
Name	Status	Number of Mirrored Pairs	Local Role	Remote Storage Array	Connection Type	Synchronization Interval	Last Known Recovery Point	Edit
arvm-cg	Synchronizing	1	Primary	crsG11SH-A4	Fibre Channel	10m	N/A	✎

Total rows: 1 ↻

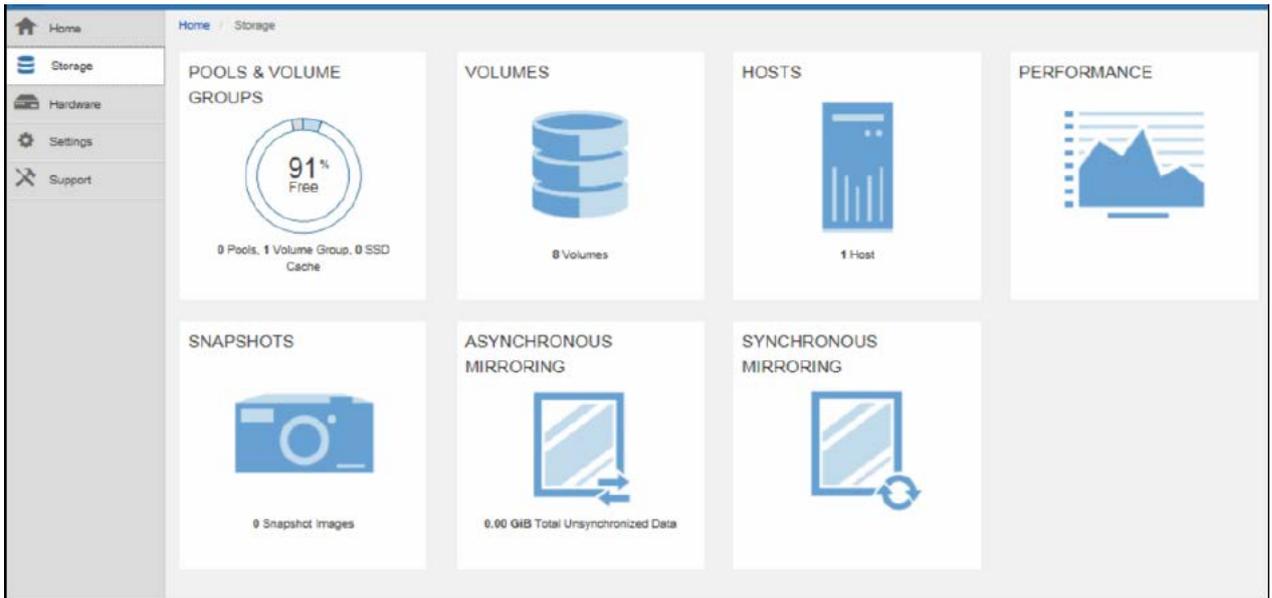
## 9.6 Synchronous with System Manager as Primary

In this case, the storage system containing the primary volumes is managed by System Manager (HB2000 or HB5000), and the secondary is any ETERNUS AB/HB series system that supports asynchronous mirroring. The configuration begins with the SANtricity Unified Manager. To perform configuration, complete the following steps:

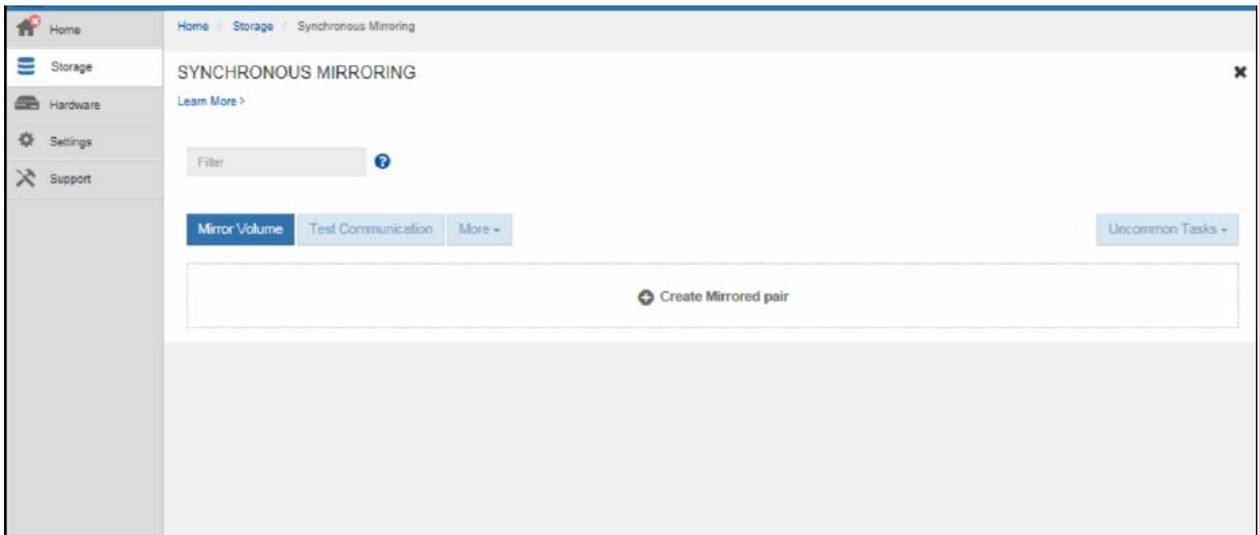
1. Make sure the versions of SANtricity System Manager and SANtricity OS are compatible as described earlier so that the ports are correct to launch System Manager.
2. Make sure that both arrays are discovered and listed in the Unified Manager.
3. Recommended: Import trusted certificates in Unified Manager as described earlier.
4. Check the box, and then click the Launch button of the array that contains the primary.



5. In System Manager for the HB2000, navigate to Storage/Synchronous Mirroring.



6. Select Mirror Volume or Create Mirrored Pair.



7. Select the primary volume.

### Create Synchronous Mirrored Pair ✕

**1** Select Primary Volume    **2** Select Secondary Volume    **3** Edit Settings

What do I need to know before creating a mirrored pair?  
Why don't I see all my volumes?  
Select your primary volume from E2800...

Name	Reported Capacity (GiB)	Pool/Volume Group
Volume1	500.00	VG1
Volume2	500.00	VG1
Volume3	500.00	VG1
Volume4	500.00	VG1
Volume5	100.00	VG1
Volume6	100.00	VG1
Volume7	100.00	VG1
Volume8	100.00	VG1

Cancel Next >

8. Select the secondary array and volume. In this case, it is the HB2000-2, volume1.

**Create Synchronous Mirrored Pair**

1 Select Primary Volume    2 Select Secondary Volume    3 Edit Settings

**Primary volume attributes**

Reported capacity: 500.00 GiB  
Data Assurance (DA) enabled: No  
Secure-capable: No  
Volume type: Thick

Select your secondary volume from remote storage array...

E2800-2

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

Name	Reported Capacity (GiB)	Pool/Volume Group
Volume1	500.00	Pool_1
Volume2	500.00	Pool_1
Volume3	500.00	Pool_1
Volume4	500.00	Pool_1

< Back    Cancel    Next >

9. Edit the settings for the mirrored pair:

- Synchronization priority affects the rate that the pair resynchronizes after the connection is lost and then restored. Higher priority provides for faster resynchronization, but it also affects performance of the primary array more. Lower priority takes longer to resynchronize, but it has less impact on the primary performance during resynchronization.
- Automatic synchronization causes the pair to immediately begin resynchronizing after connection is restored. Fujitsu recommends manual synchronization because the administrator has more control. If a problem occurs during resynchronization, the secondary can be rendered useless due to the order of resynchronization. With manual resynchronization, the administrator can take a Snapshot copy of the secondary before starting resynchronization and then there would at least be a valid image at the secondary even if resynchronization is interrupted.

✕
Create Synchronous Mirrored Pair

1 Select Primary Volume
2 Select Secondary Volume
3 Edit Settings

**Synchronization priority**

Synchronization priority ?

Medium (default)

What impact does synchronization priority have on synchronization rates?

**Resynchronization policy**

After a communication interruption...

- Require synchronization to be manually resumed (recommended)
- Automatically resume synchronization

Why is it recommended to use a manual resynchronization policy?

< Back
Cancel
Finish

10. The system prompts for the remote array administrator password.

✕
Enter Remote Storage Array Password

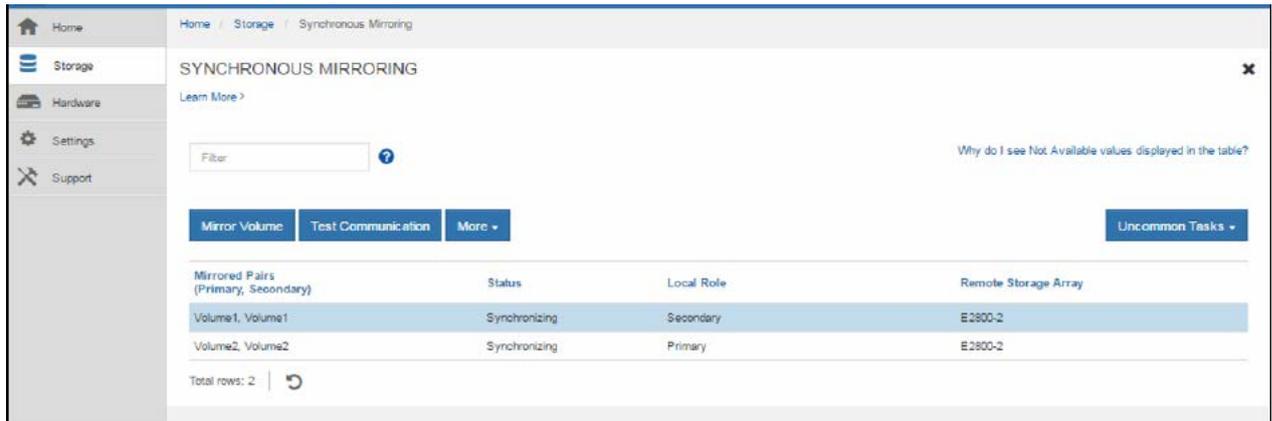
**Authentication is Required**

To perform this mirroring operation, we must communicate with the remote storage array. For security purposes, you must log in using your remote storage array's administrator password.

Storage array password for E2800-2.

Log In
Cancel

The mirrored pair now begins initial synchronization.



## 9.7 Testing the Communication Link

When setting up mirroring, it is good practice to test the communication link between the two sites. This testing can be done for either asynchronous or synchronous mirroring.

### Testing the Asynchronous Mirroring 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 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.** Show 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.

After the test has completed, the window shows the results:

- **Normal.** The mirror group is communicating correctly.
- **Passed.** Check possible network or connection problems and retry.
- **Failed.** The reason is indicated. Refer to the Recovery Guru to correct the problem.
- **Port connection error.** This error might be because the local array is not connected or because the remote array cannot be contacted. Refer to the Recovery Guru to correct the problem.

Figure 12 and Figure 13 show invoking the test and the results in System Manager.

Figure 12) SANtricity System Manager: invoke test communication on mirror consistency group.

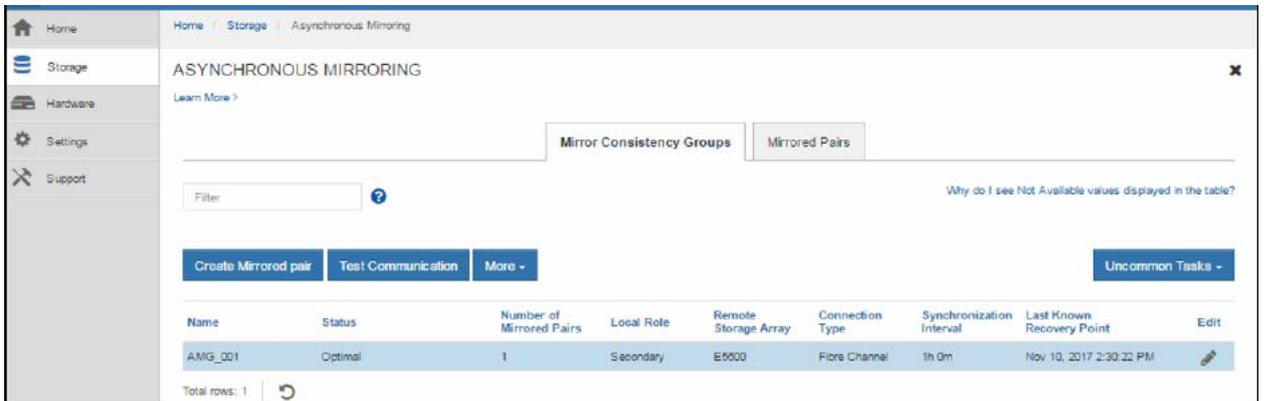
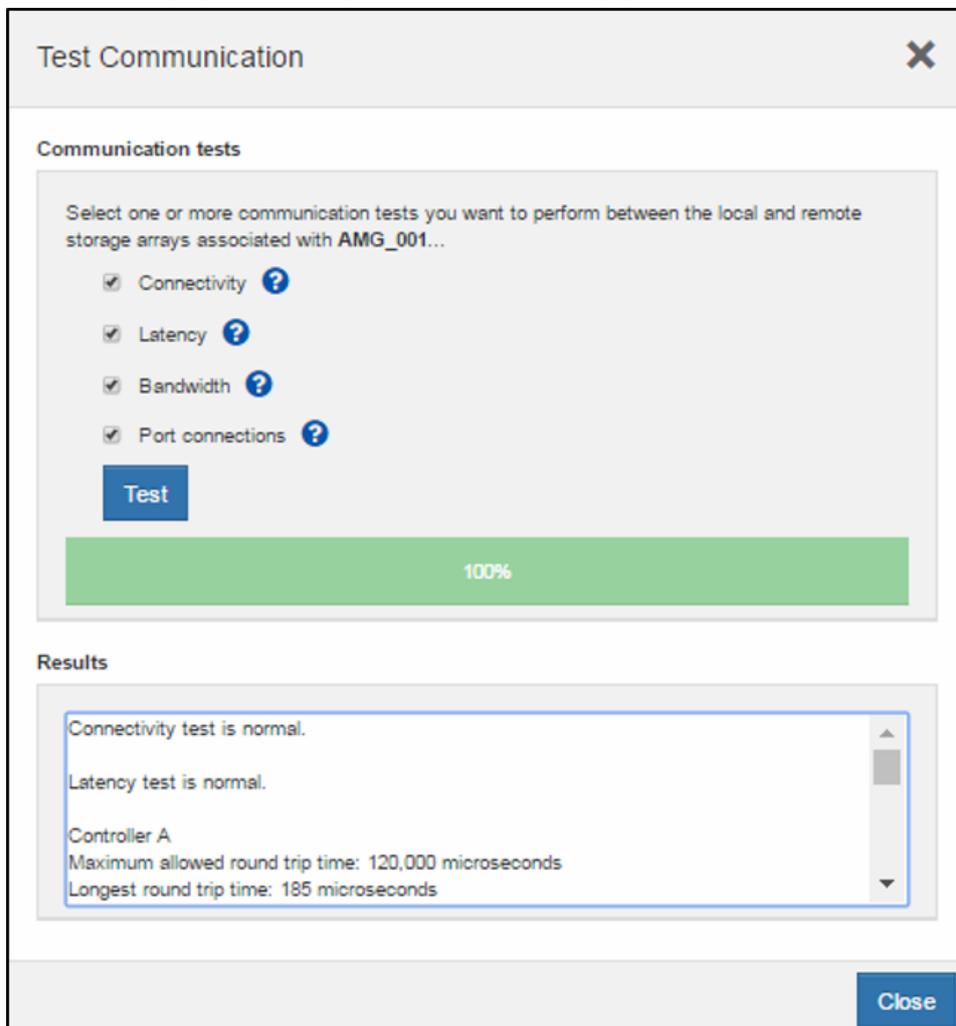


Figure 13) SANtricity System Manager: mirror consistency group test communication results.



## Testing the Synchronous Mirroring Link

To test the communication link, first select the synchronous mirror volume for which the test is desired. Make sure that the volume pair exists on the local and remote storage arrays.

After the communication test has completed, the result window shows if the test was passed, and it also shows the measured round-trip times.

Figure 14 and Figure 15 show examples of this testing in System Manager.

Figure 14) SANtricity System Manager: invoke test communication for synchronous mirrored pair.

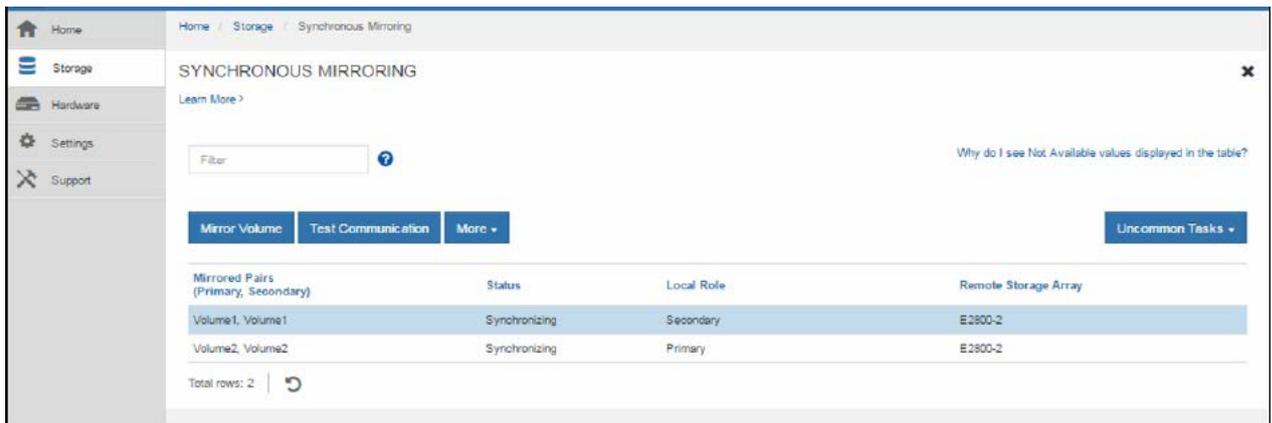
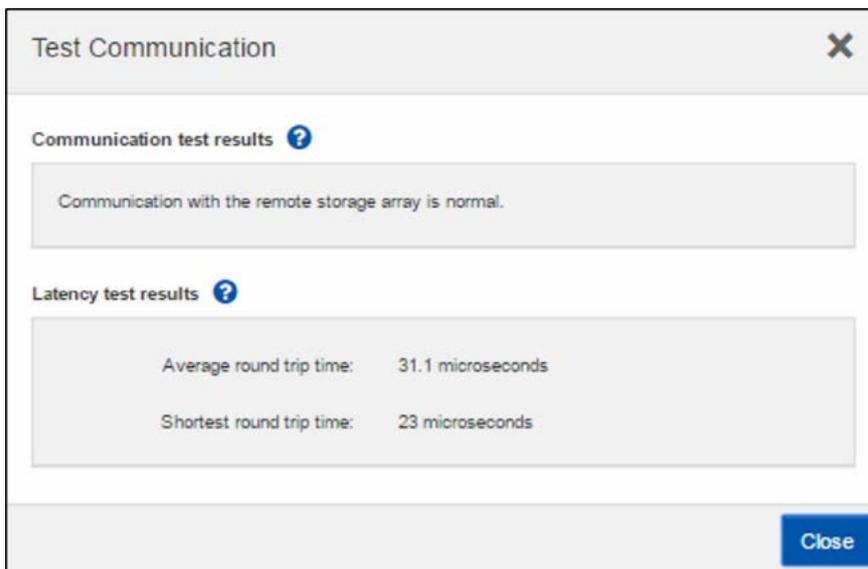


Figure 15) SANtricity System Manager: synchronous mirrored pair test communication results.



## Summary

Designed to deliver worry-free disaster recovery, the Fujitsu SANtricity mirroring features make it possible to quickly and easily protect storage on ETERNUS AB/HB series systems and enhance an enterprise business recovery process. Whether the requirement is an RPO of zero or long-distance replication, Fujitsu ETERNUS AB/HB series systems can play an integral role in an enterprise-wide disaster recovery plan.

Supported across all SANtricity managed arrays, mirroring is the perfect solution for businesses that need to protect their environments while maintaining simplicity, performance, and controlling costs.

## Related Information and Documents

The latest information for the ETERNUS AB/HB series is available at:

<https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html>

## Version History

Version	Date	Document Version History
Version 1.0	July 2020	Newly created
Version 2.0	May 2023	Addition of description of communication restoration
Version 3.0	July 2023	Added description (4.5, 4.6)

## Service Support Matrix

Refer to the Service Support Matrix to validate that the exact product and feature versions described in this document are supported for your specific environment. The Service Support Matrix defines the product components and versions that can be used to construct configurations that are supported by Fujitsu. Specific results depend on each customer's installation in accordance with published specifications.

## Copyright Information

Copyright 2020-2023 Fujitsu Limited. All rights reserved.

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 Fujitsu material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY FUJITSU "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 FUJITSU 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.

Fujitsu reserves the right to change any products described herein at any time, and without notice. Fujitsu assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Fujitsu. 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 Fujitsu.

## Trademarks

Third-party trademark information related to the ETERNUS AB/HB series is available at:

<https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html>

