



Sync concepts

SANtricity 11.8

NetApp
May 08, 2024

Table of Contents

- Sync concepts 1
 - How synchronous mirroring works 1
 - Synchronous mirroring terminology 3
 - Workflow for mirroring a volume synchronously 4
 - Requirements for using synchronous mirroring 5
 - Synchronous mirroring status 7
 - Volume ownership 7
 - Role change between volumes in a mirrored pair 8

Sync concepts

How synchronous mirroring works

Synchronous mirroring replicates data volumes in real time to ensure continuous availability.



Synchronous mirroring is not available on the EF600 or EF300 storage array.

Synchronous mirroring achieves a recovery point objective (RPO) of zero lost data by having a copy of important data available if a disaster happens on one of the two storage arrays. The copy is identical to production data at every moment because each time a write is done to the primary volume, a write is done to the secondary volume. 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.


This type of mirroring is ideal for business continuity purposes such as disaster recovery.

Synchronous mirroring relationship

A synchronous mirroring relationship consists of a primary volume and a secondary volume on separate storage arrays. The storage array containing the primary volume is usually located at the primary site and serves the active hosts. The storage array containing the secondary volume is usually located at a secondary site and holds a replica of the data. The secondary volume is used if the primary volume's storage array is unavailable because of, for example, a complete power outage, a fire, or a hardware failure at the primary site.

Synchronous mirroring session

The synchronous mirroring configuration process involves configuring volumes into pairs. After you create a mirrored pair, which consists of a primary volume on one storage array and a secondary volume on another storage array, you can start synchronous mirroring. The steps in synchronous mirroring are depicted below.



1. A write comes in from the host.
2. The write is committed to the primary volume, propagated to the remote system, and then committed to the secondary volume.
3. The primary volume's storage array sends an I/O completion message to the host system *after* both write operations have been successfully completed.

Reserved capacity is used to log information about the incoming write request from a host.

When the current controller owner of the primary volume receives a write request from a host, the controller first logs information about the write to the primary volume's reserved capacity. It then writes the data to the primary volume. Next, the controller initiates a remote write operation to copy the affected data blocks to the secondary volume at the remote storage array.

Because the host application must wait for the write to occur on the local storage array and across the network on the remote storage array, a very fast connection between the local storage array and remote storage array is required to maintain the mirror relationship without overly reducing local I/O performance.

Disaster recovery

Synchronous mirroring maintains a copy of data that is physically distant from the site where the data resides. If a disaster occurs at the primary site, such as a power outage or a flood, the data can be quickly accessed from the secondary site.

The secondary volume is unavailable to host applications while the synchronous mirroring operation is in progress, so, in the event of a disaster at the local storage array, you can fail over to the remote storage array. To fail over, promote the secondary volume to the primary role. Then the recovery host is able to access the newly promoted volume, and business operations can continue.

Synchronization settings

When you create a mirrored pair, you also define the synchronization priority and resynchronization policy that the mirrored pair uses to complete the resynchronization operation after a communication interruption.

If the communication link between the two storage arrays stops working, hosts continue to receive acknowledgements from the local storage array, preventing an access loss. When the communication link is working again, any unreplicated data can be automatically or manually resynced to the remote storage array.

Whether data is resynchronized automatically depends on the mirrored pair's resynchronization policy. An automatic resynchronization policy allows the mirrored pair to resynchronize automatically when the link is working again. A manual resynchronization policy requires you to manually resume synchronization after a communication problem. Manual resynchronization is the recommended policy.

You can edit the synchronization settings for a mirrored pair only on the storage array that contains the primary volume.

Unsynchronized data

The primary and secondary volumes become unsynchronized when the primary volume's storage array is unable to write data to the secondary volume. This can be caused by the following issues:

- Network problems between the local and remote storage arrays
- A failed secondary volume
- Synchronization being manually suspended on the mirrored pair

Orphaned mirrored pair

An orphaned mirrored pair volume exists when a member volume has been removed on one side (either the primary side or secondary side) but not on the other side.

Orphaned mirrored pair volumes are detected when inter-array communication is restored and the two sides of the mirror configuration reconcile mirror parameters.

You can remove a mirrored pair to correct an orphaned mirrored pair state.

Configuration and management

To enable and configure mirroring between two arrays, you must use the Unified Manager interface. Once mirroring is enabled, you can manage mirrored pairs and synchronization settings in System Manager.

Synchronous mirroring terminology

Learn how the synchronous mirroring terms apply to your storage array.

Term	Description
Local storage array	The local storage array is the storage array that you are acting upon. When you see Primary in the Local Role column, it indicates that the storage array contains the volume that holds the primary role in the mirror relationship. When you see Secondary in the Local Role column, it indicates that the storage array contains the volume that holds the secondary role in the mirror relationship.
Mirrored pair	A mirrored pair is comprised of two volumes, a primary volume and a secondary volume.
Primary volume	The primary volume of a mirrored pair is the source volume to be mirrored.
Recovery point objective (RPO)	Recovery Point Objective (RPO) represents an objective that indicates the difference considered acceptable between the primary volume and secondary volume in a mirrored pair. An RPO of zero indicates that no difference between the primary volume and secondary volume can be tolerated. An RPO greater than zero indicates that the secondary volume is less current or lags behind the primary volume.
Remote storage array	The remote storage array is usually designated as the secondary site, which usually holds a replica of the data in a mirroring configuration.
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.
Role change	Role change is assigning the primary role to the secondary volume and vice versa.
Secondary volume	The secondary volume of a mirrored pair is usually located at a secondary site and holds a replica of the data.
Synchronization	Synchronization occurs at initial synchronization between the local storage array and the remote storage array. Synchronization also occurs when the primary and secondary volumes become unsynchronized after a communication interruption. When the communication link is working again, any unreplicated data is synchronized to the secondary volume's storage array.

Workflow for mirroring a volume synchronously

You configure synchronous mirroring using the following workflow.



This feature is not available on the EF600 or EF300 storage system.

1. Perform the initial configuration in Unified Manager:
 - a. Select a local storage array as the source for the data transfer.
 - b. Select a primary volume from the local storage array.

- c. Select a remote storage array as the destination for the data transfer, and then select a secondary volume.
 - d. Select synchronization and resynchronization priorities.
 - e. Begin the initial data transfer from the primary volume to the secondary volume. Depending on the volume size, this initial transfer could take several hours.
2. Check the progress of the initial synchronization:
 - a. In Unified Manager, launch System Manager for the local array.
 - b. In System Manager, view the status of the mirroring operation. When mirroring is complete, the status of the mirrored pair is "Optimal." The two arrays attempt to stay synchronized through normal operations. Only new and changed blocks are transferred from the primary volume to the secondary volume.
 3. **Optional:** You can change synchronization settings in System Manager.



Because synchronous replication is continuous, the replication link between the two sites must provide sufficient bandwidth capabilities.

Requirements for using synchronous mirroring

If you plan to use synchronous mirroring, keep the following requirements in mind.

Unified Manager

To enable and configure mirroring between two arrays, you must use the Unified Manager interface. Unified Manager is installed on a host system along with the Web Services Proxy.

- The Web Services Proxy service must be running.
- Unified Manager must be running on your local host through an HTTPS connection.
- Unified Manager must be showing valid SSL certificates for the storage array. You can accept a self-signed certificate or install your own security certificate using Unified Manager and navigating to **Certificate > Certificate Management**.

Storage arrays



Synchronous mirroring is not available on the EF300 or EF600 storage array.

- You must have two storage arrays.
- Each storage array must have two controllers.
- The two storage arrays must be discovered in Unified Manager.
- Each controller in both the primary array and secondary array must have an Ethernet management port configured and must be connected to your network.
- The storage arrays have a minimum firmware version of 7.84. (They can each run different OS versions.)
- You must know the password for the local and remote storage arrays.
- You must have enough free capacity on the remote storage array to create a secondary volume equal to or greater than the primary volume that you want to mirror.

- Your local and remote storage arrays are connected through a Fibre Channel fabric.

Supported connections

Communication for synchronous mirroring is supported only on controllers with Fibre Channel (FC) host ports.

Synchronous mirroring uses the highest numbered host port on each controller on both the local storage array and the remote storage array. Controller host bus adapter (HBA) host port 4 is typically reserved for mirror data transmission.

Mirrored volume candidates

- RAID level, caching parameters, and segment size can be different on the primary and secondary volumes of a synchronous mirrored pair.
- The primary and secondary volumes in a synchronous mirrored pair must be standard volumes. They cannot be thin volumes or snapshot volumes.
- The secondary volume must be at least as large as the primary volume.
- Only the primary volume may have snapshots associated with it and/or be the source or target volume in a volume copy operation.
- A volume can participate in only one mirror relationship.
- There are limits to the number of volumes that are supported on a given storage array. Make sure that the number of configured volumes on your storage array is less than the supported limit. When synchronous mirroring is active, the two reserved capacity volumes that are created count against the volume limit.

Reserved capacity

- Reserved capacity is required for a primary volume and for a secondary volume for logging write information to recover from controller resets and other temporary interruptions.
- The reserved capacity volumes are created automatically when synchronous mirroring is activated. Because both the primary volume and the secondary volume in a mirrored pair require reserved capacity, you must ensure that you have enough free capacity available on both storage arrays that are participating in the synchronous mirror relationship.

Drive Security feature

- If you are using secure-capable drives, the primary volume and the secondary volume must have compatible security settings. This restriction is not enforced; therefore, you must verify it yourself.
- If you are using secure-capable drives, the primary volume and the secondary volume should use the same drive type. This restriction is not enforced; therefore, you must verify it yourself.
 - If the primary volume uses Full Disk Encryption (FDE) drives, the secondary volume should use FDE drives.
 - If the primary volume uses Federal Information Processing Standards 140-2 (FIPS) validated drives, the secondary volume should use FIPS 140-2 validated drives.
- If you are using Data Assurance (DA), the primary volume and the secondary volume must have the same DA settings.

Synchronous mirroring status

A synchronous mirrored pair's status indicates whether the data on the primary volume and on the secondary volume is synchronized. A mirror status is independent of the component status of the volumes in the mirrored pair.



This feature is not available on the EF600 or EF300 storage system.

Synchronous mirrored pairs can have one of the following statuses:

- **Optimal**

Indicates that the volumes in the mirrored pair are synchronized, which means that the fabric connection between the storage arrays is operational and each volume is in the desired working condition.

- **Synchronizing**

Shows the progress of the data synchronization between the mirrored pairs. This status will also be shown during the initial synchronization.

After a communication link interruption, only the blocks of data that have changed on the primary volume during the link interruption are copied to the secondary volume.

- **Unsynchronized**

Indicates that the primary volume's storage array is unable to write incoming data to the remote array. The local host can continue to write to the primary volume, but remote writes do not take place. Different conditions can prevent the primary volume's storage array from writing incoming data to the secondary volume, such as:

- The secondary volume is not accessible.
- The remote storage array is not accessible.
- The fabric connection between the storage arrays is not accessible.
- The secondary volume cannot be updated with a new World Wide Identifier (WWID).

- **Suspended**

Indicates that the synchronous mirroring operation has been suspended by the user. When a mirrored pair is suspended, no attempt is made to contact the secondary volume. Any writes to the primary volume are persistently logged in the mirror reserved capacity volumes.

- **Failed**

Indicates that the synchronous mirroring operation is unable to operate normally due to a failure with the primary volume, secondary volume, or the mirror reserved capacity.

Volume ownership

You can change the preferred controller owner in a mirrored pair.



This feature is not available for synchronous mirroring on the EF600 or EF300 storage system.

If the primary volume of the mirrored pair is owned by controller A, then the secondary volume will also be owned by controller A of the remote storage array. Changing the primary volume's owner will automatically change the owner of the secondary volume to ensure that both volumes are owned by the same controller. Current ownership changes on the primary side automatically propagate to corresponding current ownership changes on the secondary side.

For example, a primary volume is owned by controller A, and then you change the controller owner to controller B. In this case, the next remote write changes the controller owner of the secondary volume from controller A to 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.

Controller resets

A controller reset causes a volume ownership change on the primary side from the preferred controller owner to the alternate controller in the storage array.

Sometimes a remote write is interrupted by a controller reset or a storage array power cycle before it can be written to the secondary volume. The controller does not need to perform a full synchronization of the mirrored pair in this case.

When a remote write has been interrupted during a controller reset, the new controller owner on the primary side reads information stored in a log file in the reserved capacity volume of the preferred controller owner. The new controller owner then copies the affected data blocks from the primary volume to the secondary volume, eliminating the need for a full synchronization of the mirrored volumes.

Role change between volumes in a mirrored pair

You can change the role between volumes in a mirrored pair. You can do this by demoting the primary volume to the secondary role or promoting the secondary volume to the primary role.



Synchronous mirroring is not available on the EF600 or EF300 storage system.

Review the following information about the role change operation:

- When a primary volume is demoted to the secondary role, the secondary volume in that mirrored pair is promoted to the primary role and vice versa.
- When the primary volume is demoted to the secondary role, hosts that have been assigned to that volume no longer have write access to it.
- When the secondary volume is promoted to the primary role, any hosts that are accessing that volume are now able to write to it.
- If the local storage array is unable to communicate with the remote storage array, you can force the role change on the local storage array.

Force role change

You can force a role change between volumes in a mirrored pair when a communication problem between the local storage array and the remote storage array is preventing the promotion of the secondary volume or the demotion of the primary volume.

You can force the volume on the secondary side to transition to the primary role. Then the recovery host can

access the newly promoted volume, and business operations can continue.



When the remote storage array has recovered and any communication problems have been resolved, a Synchronous Mirroring - Primary Volume Conflict condition occurs. The recovery steps include resynchronizing the volumes. Use the Recovery Guru to recover from this error.

When is a forced promotion allowed and not allowed?

Forced promotion of a volume in a mirrored pair is not allowed under the following conditions:

- Any of the volumes in a mirrored pair are in the process of an initial synchronization.
- The mirrored pair is in the Failed, Role-Change-Pending, or Role-Change-In-Progress states or if any of the associated reserved capacity volumes are failed.

Role change in-progress state

If two storage arrays in a mirroring configuration become disconnected, and the primary volume of a mirrored pair is force demoted to a secondary role, and the secondary volume of a mirrored pair is force promoted to a primary role, then when communication is restored, the volumes on both storage arrays are placed in the Role-Change-In-Progress state.

The system will complete the role change process by transferring the change logs, re-synchronizing, setting the mirrored pair state back to a normal operating state, and continuing with synchronizations.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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