

# Dell EMC Unity™ Family

Version 5.x

## Configuring Replication

H14974

REV 07

Copyright © 2016-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published June 2019

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.  
Published in the USA.

Dell EMC  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.DellEMC.com](http://www.DellEMC.com)

# CONTENTS

<b>Chapter 1</b>	<b>Overview</b>	<b>5</b>
	About this document.....	6
	Additional resources.....	6
	About replication.....	7
	Advanced replication topologies for asynchronous file replication.....	16
	Snapshot replication in advanced replication topologies.....	17
	Dell EMC MetroSync for Unity.....	18
	Using replication for disaster recovery.....	19
	Using replication for planned downtime.....	21
	Failback a replication session.....	23
	Asynchronous replication of a synchronous file replication to a third site...	24
	Cascade mode for third site asynchronous replication.....	28
<b>Chapter 2</b>	<b>Replication workflow</b>	<b>31</b>
	Configure replication.....	32
	Configure local replication.....	32
	Configure asynchronous replication.....	32
	Configure synchronous replication.....	33
	Configure asynchronous replication of synchronous replication to a third site.....	34
	Related concepts.....	35
	Replication interfaces.....	35
	About replication connections.....	36
	About replication sessions.....	37
	About replication support for thin storage resources with data reduction.....	39
<b>Chapter 3</b>	<b>Considerations for synchronous replication</b>	<b>41</b>
	Synchronization states.....	42
	Fibre Channel configuration for synchronous replication.....	43
<b>Chapter 4</b>	<b>Configure replication using Unisphere</b>	<b>47</b>
	Configure protection and mobility interfaces.....	48
	Configure replication connections.....	49
	Create a replication session.....	49
	Manage replication sessions.....	51
	Configure an override network address for file replication .....	56
<b>Chapter 5</b>	<b>Configure replication using the CLI</b>	<b>59</b>
	Manage network interfaces.....	60
	Create interfaces.....	61
	View interfaces.....	62
	Change interface settings.....	64
	Delete interfaces.....	65
	Manage remote storage systems.....	65
	Create remote system configurations.....	67
	View settings for remote storage systems.....	68

Verify settings for remote storage systems.....	69
Change settings for remote storage systems.....	69
Delete remote system configurations.....	70
Cabinet level unplanned failover of replication sessions.....	71
Manage replication sessions.....	72
Create replication sessions.....	79
View replication sessions.....	85
Change replication session settings.....	87
Pause replication sessions.....	91
Resume replication sessions.....	91
Manually synchronize replication sessions.....	93
Delete replication sessions.....	94
Fail over replication sessions.....	95
Fail back replication sessions.....	96
Preserve asynchronous replication sessions.....	98
Associated CLI commands.....	98
The show action command.....	98
Manage Ethernet ports.....	100
Manage FC ports.....	104
Manage link aggregations.....	106

<b>Appendix A</b>	<b>Use case for CLI</b>	<b>111</b>
	Replication configuration use case.....	112
	Configure local replication.....	112
	Configure asynchronous replication.....	112
	Configure synchronous replication.....	113
	Create a replication interface.....	114
	Create a replication connection.....	116
	Create a replication session for block storage.....	117
	Create an asynchronous replication session for file storage.....	117
	Create a synchronous replication session for file storage.....	119
	View replication sessions.....	120

# CHAPTER 1

## Overview

This chapter addresses the following topics:

- [About this document](#).....6
- [Additional resources](#).....6
- [About replication](#).....7
- [Advanced replication topologies for asynchronous file replication](#).....16
- [Dell EMC MetroSync for Unity](#).....18
- [Using replication for disaster recovery](#).....19
- [Using replication for planned downtime](#).....21
- [Failback a replication session](#).....23
- [Asynchronous replication of a synchronous file replication to a third site](#).....24

## About this document

This document provides information you can use to configure and manage replication on your Unity storage system. Along with relevant concepts and instructions to configure replication using the Unisphere GUI, this document also include information on the CLI commands associated with configuring replication.

---

### Note

For more information on other Unisphere features or CLI commands, refer to the *Unisphere online help* and *CLI User Guide*.

---

## Additional resources

As part of an improvement effort, revisions of the software and hardware are periodically released. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features. Contact your technical support professional if a product does not function properly or does not function as described in this document.

### Where to get help

Support, product, and licensing information can be obtained as follows:

#### Product information

For product and feature documentation or release notes, go to Unity Technical Documentation at: [www.emc.com/en-us/documentation/unity-family.htm](http://www.emc.com/en-us/documentation/unity-family.htm).

#### Troubleshooting

For information about products, software updates, licensing, and service, go to Online Support (registration required) at: <https://Support.EMC.com>. After logging in, locate the appropriate **Support by Product** page.

#### Technical support

For technical support and service requests, go to Online Support at: <https://Support.EMC.com>. After logging in, locate **Create a service request**. To open a service request, you must have a valid support agreement. Contact your Sales Representative for details about obtaining a valid support agreement or to answer any questions about your account.

### Special notice conventions used in this document



Indicates a hazardous situation which, if not avoided, will result in death or serious injury.

---



Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

---

**⚠ CAUTION**

Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

---

**NOTICE**

Addresses practices not related to personal injury.

---

**Note**

Presents information that is important, but not hazard-related.

---

## About replication

Data replication is one of the many data protection methodologies that enable your data center to avoid disruptions in business operations. It is a process in which storage data is duplicated to a remote or local system. It provides an enhanced level of redundancy in case the main storage backup system fails. It minimizes the downtime-associated costs of a system failure and simplifies the recovery process from a natural disaster or human error.

The system supports asynchronous and synchronous replication of all storage resources, including file systems, NAS servers, LUNs, LUN groups (also known as consistency groups (CGs)), VMware VMFS datastores, VMware NFS datastores, and thin clones. The asynchronous replication feature leverages the Unified Snapshots technology to produce a read-only, point-in-time copy of source storage data and periodically updates the copy to keep it consistent with the source data. It leverages crash consistent replicas to provide remote data protection of storage resources. The synchronous replication feature leverages the MirrorView/Synchronous technology to mirror data in real time between local and remote storage resources.

---

**Note**

In general, Unity OE versions 4.0 or later support replication interoperability. However, the exception to this occurs when the source system is configured with features that are not compatible with an earlier Unity OE version that is running on the destination side of the replication session. For example:

- Inline compression is only supported for block objects with OE versions 4.1 or later, while file objects are supported with OE versions 4.2 or later.
  - Block objects support asynchronous and synchronous replication with all Unity OE versions, while file supports only asynchronous replication with OE versions 4.0 or later, and synchronous replication with OE versions 4.4 or later.
  - Block and file objects support asynchronous replication of user snapshots with OE versions 4.2 or later, while only file supports synchronous replication of snapshots with OE versions 4.4 or later.
  - Advanced replication topologies, that is, fan-out or 1-N replication and cascading replication for asynchronous replication, are only supported with OE versions 5.x and only for file objects.
-

**NOTICE**

Do not disable the automatic failback policy setting, otherwise the following issues may occur:

- The operations to pause or delete a synchronous replication session may not complete and appear not to make any progress.
- If SP reboots, synchronous replication sessions may not be recovered to a synchronization state of `Consistent` or `In Sync`. Instead, they may remain in a synchronization state of `Out of Sync`.

In Unisphere under Settings, ensure the checkbox **Management > Automatic failback policy** is selected or the `Automatic failback` system general attribute in CLI is set to `on`.

---

**Replication modes**

Replication for both block and file storage can operate in the following modes:

- **Asynchronous**—Use this mode when you want the data between the source and destination storage resources synchronized automatically at a specific interval, based on the Recovery Point Objective (RPO).
- **Synchronous**—Use this mode when you want the data between the source and destination storage resources to always remain in sync.
- **Manual**—Use this mode when you want to manually synchronize changes in the source storage resource to the destination storage resource. When you choose this mode, ensure that you periodically synchronize the session to avoid excessive pool space consumption.

**Recovery Point Objective**

Recovery Point Objective (RPO) is an industry accepted term that indicates the acceptable amount of data, measured in units of time, that may be lost in a failure. When you set up an asynchronous replication session, you can configure automatic synchronization based on the RPO. You can specify an RPO from a minimum of 5 minutes up to a maximum of 1440 minutes (24 hours). The default RPO is set at 60 minutes (1 hour) interval. In the case of synchronous replication, RPO is set to 0. You can use the Unisphere CLI or Unisphere Management REST API to specify a more granular RPO.

---

**Note**

Although a smaller time interval provides more protection and lesser space consumption, it also has a higher performance impact, and results in more network traffic. A higher RPO value may result in more space consumption. This may affect the snapshot schedules and space thresholds.

---

**Source and destination storage resources**

In Unisphere, for all replicated storage resources except for thin clones, once replication is configured, the destination storage resource is automatically created. In CLI, you must manually create the destination storage resource and then create the replication session between the source and destination storage resources.



**NOTICE**

For file synchronous replication, the model-to-model source and destination replication pair allowed is as follows:

- Unity 300(F) to a Unity 300(F)
- Unity 350F, 380(F), 400(F), or 500(F) to a Unity 350F, 380(F), 400(F), or 500(F)
- Unity 450F, 480(F), 550F, 600(F), 650F, 680(F), or 880(F) to a Unity 450F, 480(F), 550F, 600(F), 650F, 680(F), or 880(F)

---

You can convert a thin LUN to a non-thin (thick) LUN, or a thick LUN to a thin LUN with a LUN move operation. To enable data reduction on a thin LUN requires an All-Flash pool on the destination system. For thick file systems, the replication process matches the destination storage resource to the source. In this case, thin and data reduction cannot be selected for file systems. For thin file systems, the following rules apply for replication:

- If the source file system is thin, then the destination file system is also thin.
- If the source file system is thick, then the destination file system is also thick.
- If both the source and destination systems support data reduction, then the source or destination can have data reduction enabled or not.
- However, if you change the source file system data reduction attribute, the change is not replicated to the destination system (retains the original setting) regardless of using asynchronous or synchronous replication.

The following file asynchronous replication rules also apply to synchronous replication:

- The type of NAS server is set by the user during server creation.
- Users can change the type of a NAS server when there are no replication sessions on any of the file systems using that NAS server.
- In destination mode, only the local configuration is enabled and all attributes in the local configuration are active.

**Note**

A storage resource in destination mode can be or is used as a replication destination and data access is restricted. For asynchronous replication, the storage resource is read-only. For synchronous replication, the storage resource is not accessible. When replication is active, the storage resource in the destination site is in destination mode. When replication is failed over, the destination storage resource becomes the source while the original source storage resource is set to destination mode.

- Override configuration specifies a set of attributes which are enabled when the NAS server is changed from destination mode to source. The override configuration of a NAS server is not replicated as part of NAS server replication.
- Local configuration (also known as back up) specifies a set of attributes that are related to enabling back up or local test through NFS or NDMP protocols. The local configuration of a NAS server is not replicated as part of NAS server replication.

**Table 1** on page 10 classifies the Global, Override, and Local attributes for NAS server. Yes means the attribute exists in the configuration, and No means the attribute does not exist in the configuration.

**Table 1** Classification of configuration attributes for NAS server

Attribute	Global Configuration	Override Configuration	Local Configuration
Production IP interface	Yes	Yes	No
Backup IP interface	No	No	Yes
DNS	Yes	Yes	Yes
NIS	Yes	Yes	Yes
CIFS server (name, domain name, NetBIOS name, LDAP org string)	Yes	No	No
NFS export	Yes	No	No
NFS export for Local snapshots	No	No	Yes
CAVA	Yes	No	No
NDMP user/password	No	No	Yes
ASA user/password	Yes	No	No

**Table 2** on page 10 lists file system level mount options that are also saved into the NAS server configuration, and all of these are Global configuration:

**Table 2** File system level mount options saved into NAS server configuration

UEMCLI Option	RESTful Option	Mount Option
-cifsNotifyOnWrite	isCIFSNotifyOnWriteEnabled	FS_PROPERTY_NOTIFYONWRITE_SYNC
-cifsNotifyDirDepth	cifsNotifyOnChangeDirDepth	FS_PROPERTY_TRIGGERLEVEL
-cifsNotifyOnAccesses	isCIFSNotifyOnAccessEnabled	FS_PROPERTY_NOTIFYONACCESS_SYNC
-cifsOpLocks	isCIFSOpLocksEnabled	FS_PROPERTY_NOOPLOCK
-cifsSyncWrites	isCIFSsyncWritesEnabled	FS_PROPERTY_NP_CIFSSYNCWRITE
-accessPolicy	accessPolicy	FS_PROPERTY_NP_ACCESSPOLICY
-folderRenamePolicy	folderRenamePolicy	FS_PROPERTY_NP_RENAMEPOLICY
-lockingPolicy	lockingPolicy	FS_PROPERTY_NP_LOCKINGPOLICY
-eventProtocols	fileEventSettings	FS_PROPERTY_NP_CEPPPOLICY

A user can set these properties through Unisphere and the CLI. If a synchronous or asynchronous replication session is created on the file system, these properties will be replicated through the NAS server's replication session to the destination and can be seen on the destination NAS server after the configuration view is refreshed.

**Table 3** on page 11 lists other file system properties not related to mount options.

**Table 3** File system properties not related to mount options

UEMCLI Option	RESTful Option	Description
-id	id	Not replicated, the source and destination can be different.
-name	name	Destination must be the same as the source. Unisphere creates the destination file system with the same value as the source.
-desc	description	Not replicated, the source and destination can be different.
-size	size	Replicated.
-thin	isThinEnabled	Destination must be the same as the source. Unisphere creates the destination file system with the same value as the source.
-compression	isCompressionEnabled	Not replicated, in Unisphere the system shows the source value as default for destination. It can be changed.
-type	supportedProtocols	Destination must be the same as the source. Unisphere creates the destination file system with the same value as the source.
-fastvpPolicy	fastVPPParameters	Not replicated, in Unisphere the system shows the source value as default for destination, but it can be changed.
-poolFullPolicy	poolFullPolicy	Not replicated, the source and destination can be different.

**NOTICE**

If both the source and the destination file systems are legacy file systems that have been upgraded from OE version 4.2.x and a synchronous replication session has been created on them, the file system setting of `minSizeAllocated` will not be replicated between them.

For properties that are not replicated when a failover occurs, ensure that you modify the attributes of the associated destination storage resource to match the attributes of the source storage resource.

When a thin clone is replicated, the destination resource is automatically created with the same attributes as the source thin clone, except that the destination resource is a full copy, rather than a thin clone.

**Snapshots**

Block and file objects support asynchronous replication of user snapshots at OE version 4.2 and later, while only file objects support synchronous replication of snapshots at OE version 4.4 and later. Also, asynchronous replication supports the replication of existing user snapshots during the initial replication session configuration, while synchronous replication does not. For synchronous replication,

source snapshots will be replicated to the destination only after the replication sessions are established.

---

**Note**

To do snapshot replication from a source system running OE version 4.0, 4.1, 4.2, 4.3, or 4.4 to a destination system running OE version 5.x, requires upgrading the source system to OE version 4.5 first. Upgrading to OE version 4.5 is not required but recommended if you want to do LUN or file system replication from OE version 4.0, 4.1, 4.2, 4.3, or 4.4 to OE version 5.x without any snapshot replication.

---

Asynchronous replication supports the replication of read-only user snapshots to either a local or a remote site along with the resource data. Both scheduled snapshots and user created snapshots can be replicated. Snapshots are supported for all resources that support asynchronous replication (that is, file system, LUN, LUN group (also known as a consistency group (CG)), VMware VMFS, and VMware VMNFS).

---

**Note**

User snapshots do not apply to the NAS server resource type.

---

Asynchronous replication of scheduled snapshots can be enabled during session creation, or enabled or disabled at any time in the lifetime of the replication session. User snapshots can be replicated with a remote retention policy that is different than that of the source.

To support asynchronous replication of snapshots, both the source and destination systems must be running Unity OE version 4.2 or later. Snapshot replication can be enabled on an existing Unity OE version 4.0.x- or OE version 4.1.x-based session after both the production and the remote systems have been upgraded to Unity OE version 4.2 or later. Only read-only snapshots are eligible for replication, and they can only be replicated to the disaster recovery site where the replication destination storage resource is located. Any snapshots that are writable, such as attached block snapshots or file snapshots with shares or exports, are not replicated.

Snapshots that exist prior to asynchronous replication session creation can be selected for replication during replication session creation. Snapshots that are older than the last sync (RPO) time can be manually selected for replication and included in the next RPO sync.

A user snapshot can have one of the following asynchronous (async) replication state attributes:

- Not marked for replication (No)-snapshot is not marked for replication
- Pending sync (Pending)-snapshot is marked for replication but is awaiting transfer
- Replicated (Yes)-snapshot has successfully transferred to the disaster recovery resource
- Failed to replicate (Failed)-snapshot failed to replicate

When the operational status of a synchronous replication session is Active, checkpoint snapshots that are created on the source array (either manually through Unisphere, CLI, RESTful command, or by a snapshot schedule) are synchronously created on the destination system. The data is consistent between the source and destination snapshots with the destination snapshot having the same content, name, description and retention policies as the source snapshot. If a checkpoint snapshot create operation fails on either the source or destination system due to any reason (such as out of space) when a synchronous replication session is Active, the snapshot is deleted on the other system.

You can associate a snapshot schedule to any file system; however, if the file system is under synchronous replication, an association with a snapshot schedule is allowed only on the source side. The association will be propagated to the destination side according to the following rules:

- Association of the source file system with cluster schedule (synchronously replicated Schedule00) makes the destination file system associated with the same cluster schedule.
- De-association of the source file system from the cluster schedule (change to local schedule or no-schedule) de-associates the destination file system from any schedule.
- Change from one local schedule or no-schedule to another local schedule on the source file system does not affect the association on the destination file system.

---

#### Note

The destination mode file system is not always on the destination side. For example, when you perform an unplanned failover from the destination site, the destination file system is no longer in destination mode. The snapshot schedule will execute on the destination site. When the source site is up, the source file system will be changed to destination mode automatically. If the network connection between the source and destination is OK, the role will not be switched until the session is failed back or resumed from the destination side.

---

If a synchronous replication session is not Active (a condition in which I/O is not synchronously mirrored to the destination system due to a syncing in progress, a connectivity issue, or other system problem), checkpoint snapshots created on the source are marked as not-replicated. It is possible to delete the checkpoint snapshots on the source system but leave them on the destination system. It is also possible to change the retention policy settings or snapshot description on the source system without replicating them to the destination system.

#### NOTICE

Restoring a snapshot on a file system under synchronous replication session is not allowed.

---

When the operational status of a synchronous replication session is Active, checkpoint snapshots that are deleted from the source system are synchronously deleted from the destination system. Also, changes that are made to the retention policy or snapshot description on the source system are replicated to the destination system. It is possible to create snapshots on the destination system, but such snapshots are not replicated to the source system.

Replicated snapshot deletion or modification of its retention settings or snapshot description generate a warning on the destination system. Operations on the destination system only affect the local objects and are not reflected to the source system. Restoring a snapshot on a file system under a synchronous replication session is not allowed.

A user snapshot can have one of the following synchronous (sync) replication state attributes:

#### Sync Replicated snapshot schedule replication

Sync Replicated snapshot schedule is synchronized to the peer site while the two sites are connected (source site connected to a synchronous replication destination site). If the two sites are disconnected, you cannot create a new Sync Replicated schedule. If

a synchronously replicated schedule is updated, the schedule on the peer is also updated.

A Sync Replicated schedule can be configured from either site using Unisphere, or CLI or REST commands. However, the schedule can be associated to a storage resource (file system or VMware NFS datastore) only from the source for each replication session. A change to a Sync Replicated schedule on either site is synchronized to the other and updates the Sync Replicated schedule with the matching name when the two sites are connected through the management interface. If the two sites are disconnected, you cannot modify the Sync Replicated schedule, however, you can associate the production file system with a local schedule.

Sync Replicated schedule deletion must be synchronized to the peer site. You can delete a Sync Replicated schedule when it is not associated to any resources on any site. If the peer sites are management-fractured, you cannot delete a Sync Replicated schedule. If a system is not participating in a cluster, you can delete the Sync Replicated schedule which has no associations with the file system.

### File-based replication session actions

On Unity systems running OE version 4.2, the following asynchronous replication actions affect both the NAS server and its associated file systems when run at the NAS server level:

- Failover
- Failover-with-sync
- Failback
- Pause
- Resume

On Unity systems running OE version 4.4, the following synchronous replication actions affect both the NAS server and its associated file systems when run at the NAS server level:

- Failover
- Failback
- Pause
- Resume

#### NOTICE

NDMP backup/restore operations on a synchronous replicated source file system are not preserved after failover.

If an unplanned failover occurs, it is necessary to check whether the file systems on the original destination system have the expected sizes. For example, in the following scenario, the resize action, either expand or shrink, triggered on the source system could not be applied to the destination system, therefore, the sizes are different:

1. The source file system is manually resized while communication is disconnected (only the source could be resized but the destination could not be resized in this situation).
2. An unplanned failover is performed on the destination system before communication is fully recovered. The destination file system size could have an unexpected value after the unplanned failover.

Check and resize the destination file system to the expected value after the unplanned failover and before executing a resume or failback of the synchronous replication

session. This action will help to avoid an unexpected size being updated to the source after a failback or resume operation, which may cause a potential issue.

---

#### Note

If you need to perform a graceful/planned asynchronous replication failover, use `Failover-with-sync` from the source system. If you need to perform an emergency/unplanned synchronous replication failover when there is no network connectivity to the source system, use `Failover` from the destination system without switches, then `Resume` to restart the sessions on the destination system. This action also reverses the direction from destination to source. Or, instead of doing `Resume`, you could simply `Failback` to go back to the source. When there is network connectivity to the source system, CLI requires you to use the `failover -sync no switch` issued from the destination system before it allows an unplanned failover. If you need to perform a graceful/planned synchronous replication failover, use `Failover` from the source system, where sessions remain running and reversed back from the destination to the original source system. To move back to the original configuration, run another `Failover`, this time from the destination system, to gracefully fail back to the original source system. A `Resume` is not required in this scenario because the sessions remain running.

---

Each of these actions triggers a group operation towards the NAS server replication session and its associated file system replication sessions. A NAS server replication as a group is available for local and remote asynchronous replication.

#### NOTICE

Do not perform a group operation at both sides of a replication session at the same time. This action is not prohibited by the storage system, however, a group operation performed at the same time at both sides of a replication session can cause the group replication session to enter an unhealthy state. Also, `failover-with-sync` for asynchronous replication is not a transparent operation. During the `failover-with-sync` process, hosts' write/read requests may be rejected.

---

A group replication session operation on a NAS server supports up to 500 file system replication sessions in such a way that those sessions look like one replicated unit. If group operations are conducted on a group session whose file system replication session numbers exceeds 500, the group replication session may enter an unhealthy state, along with some file system replication sessions.

---

#### Note

Although a group replication session looks like one operation, each file system is replicated individually. If any of the individual file system replication sessions fail, you can resolve the issue and then select the individual file system to replicate.

---

Those same replication actions towards a file system remain at the file system level. Those actions are still individual operations toward file system replication sessions.

The following replication actions affect only the NAS server when run at the NAS server level or are still individual operations toward file system replication sessions:

- Create
- Sync
- Delete

- Modify

A destination file system changes from read only (RO) mounted to half-mounted when a synchronous replication session is created on it. It changes back to RO mounted when the synchronous replication session is deleted. Any file system functionality that relies on a RO mounted state will not work on the destination file system under a synchronous replication session (for example, disaster recovery access on the destination file system through Proxy NAS Server). In this case for such access, a snapshot should be created on the file system instead.

#### Updating the view of the destination NAS server configuration

While file system data and NAS server configuration are synchronously maintained between the source and destination systems, by default, the view of the NAS server configuration at the destination system from the management interface is only updated automatically every 15 minutes. However, if you need or want to see whether any changes to the NAS server configuration have occurred before the default update runs, you can manually issue an on demand update of the view of the NAS server configuration at the destination system from either Unisphere or the CLI.

---

#### Note

Since the synchronous replicated configuration file system is unreadable during initial synchronization or synchronizing after a fracture, it is not possible to update the view of the NAS server configuration at the destination system at those times.

---

## Advanced replication topologies for asynchronous file replication

Unity supports advanced replication topologies, that is, fan-out (1 to many) replication and cascading (multi-hop) replication for asynchronous file replication only. Fan-out supports a maximum of four asynchronous replication sessions on the same file storage object between two remote systems, including the local system, whether or not the storage object is in destination mode. Cascaded replication replicates to another tier or level from an already replicated resource. Each cascade level can use fan-out replication for up to three additional sites. Each replication session can have an independent Recovery Point Objective (RPO).


The existing replication operations are supported with some restrictions:

- This feature only supports file storage objects and does not support block storage objects.
- All systems joining the multiple sessions, either in fan-out (star) or cascaded mode, must be running OE version 5.x.
- A user snapshot replication can only be supported for one asynchronous session among all the sessions associated with the same storage object.
- Only one loopback asynchronous session is supported per storage object.

The following figure shows an example of a possible configuration with both fan-out and cascaded replication. Each lettered box represents a system running OE version 5.x. A represents a production site. The source object is not in destination mode and all of the replication sessions on this resource act as the source. B, D, and E represent cascaded remote sites. Sessions on these resources act as the source for one session and act as the destination for another session. C, F, G, H, I, and J represent end remote sites. All of the sessions on these resources act as the destination and are in destination mode.




**Figure 1** Fan-out and cascade replication topologies



### Snapshot replication in advanced replication topologies

For snapshot replication support, only one asynchronous replication session among all the sessions associated with the same storage resource can be enabled to transfer the snapshot from the source to the destination. The following figure and table show an example of how basic snapshot replication is supported along with independent RPO settings on each session. In the figure, A represents a production site with a NAS server and associated file system and user snapshot, and B, C, D, E, F, G, H, and I represent remote sites. Solid lines represent replication sessions that have a replicated snapshot while dashed lines represent replication sessions that do not have a replicated snapshot. The table lists the systems for each asynchronous replication session, whether a snapshot is replicated, and the RPO for each session.

**Figure 2** Replication sessions with and without snapshots to remote sites



**Table 4** Snapshot replication and RPO for replication sessions

Asynchronous Replication Session	User Snapshot Replication	RPO (in minutes)
A to B	Yes	60
A to C	No	180
A to D	No	240
A to E	No	480
B to F	No	360
C to G	Yes	720
D to H	Yes	960
E to I	Yes	1440

## Dell EMC MetroSync for Unity

Dell EMC MetroSync for Unity is a collection of features that provide a file resource disaster recovery solution. To use all of the features for this solution, the systems must be running OE versions 4.4 or later. The underlying features of MetroSync for Unity are:

- NAS server and file system synchronous replication
- File synchronous replication support for snapshots
- Sync replicated snapshot schedule replication
- Cabinet-level unplanned failover of file-based replication sessions

- Asynchronous replication of synchronous file replication to a third site

## Using replication for disaster recovery

In a disaster recovery scenario, the primary (source) system is unavailable due to a natural or human-caused disaster. Data access is still available because a replication session was configured between the primary and destination systems, and the destination system contains a full copy, or replica, of the production data. The replica is up-to-date in accordance with the last time the destination synchronized with the source, as specified by the automatic synchronization recovery point objective (RPO) setting. By issuing a session failover on the destination system, you make the destination system the new production system, using the replica of the primary system's data that resides on the destination system. Using replicas for disaster recovery minimizes potential data loss. The amount of potential data loss is affected by the RPO that is configured when setting up the replication session. In synchronous replication configuration, where the RPO is set to 0, the amount of potential data loss will be minimal.

The asynchronous failover operation always restores the destination resource to the replication common base snapshot. If failing over to the common base is not sufficient and replicated user snapshots exist, the destination resource should be manually restored to any of the replicated user snapshots.

Once the session is failed over to the destination system, the destination storage resource becomes read-write. At this point, ensure that the storage resource has the correct access permissions to the host and share. When originally establishing a replication session between the primary and destination systems, create the proper host access on the destination system ahead of time to reduce downtime in an event of a disaster.

To resume the operations on the destination and switch the roles, resume the replication session. To resume the operations on the source, fail back the replication session.

### File-based replication consideration

Switch over the NAS server replication session using the **Failover** option. This action triggers a group operation towards the NAS server replication session and its associated file system replication sessions.

The NAS server replication session should be in one of the following states in order for it to be failed over to the destination system:

- Idle
- Auto Sync Configured
- Lost Communication
- Lost Sync Communication
- Non Recoverable Error

If the NAS server replication session is in one of the following states, it cannot be failed over to the destination system:

- Paused

---

#### Note

The Paused state only affects NAS server replication session failover in systems with OE versions 4.x. It does not affect synchronous or asynchronous NAS server replication session failover in OE versions 5.x.

---

- Error states other than Lost Communication, Lost Sync Communication, or Non Recoverable Error

**NOTICE**

(This notice does not apply to asynchronous file replication in systems running OE versions 5x.) When a source site has a power outage and file replication sessions are failed over to the destination site, after power is restored and when the source site is restarted with the destination site well connected, a duplicate IP issue (the production IP addresses of the source and destination NAS servers are the same and both are in service) can be avoided. The duplicate IP issue may not be avoided for other cases which include but not are limited to:

- Source site is alive when failover is executed on the destination site.
- Remote system IP connection is broken during the source site restarting.
- SP failover and file synchronous replication failover are executed at the same time.

---

To resume the operations on the destination and switch the roles, resume the NAS server replication session. To resume the operations on the source, fail back the NAS server replication session.

**NOTICE**

If a synchronous replication session is created on an import target NAS server with file systems on it, do not execute a failover before the import is committed. In this case, the system will reject the failover from the source (planned failover), but the system will not reject failover from the destination (unplanned failover) if the destination is disconnected from the source (import target). Ensure the source site is permanently down before failing over.

**Cabinet level unplanned failover of file-based replication sessions**

In case the source system is not available for any reason, you can execute a failover of all NAS server synchronous replication sessions from the remote system. This operation also automatically fails over the replication sessions of file systems created on the affected NAS servers. The cabinet level operation must be executed using the `/remote/sys/ failover` CLI command and must be run from the destination system. You must run the `/remote/sys show` CLI command from the destination system to obtain the remote system ID of the source system in order to perform the cabinet level failover. When the destination system detects that the source system is actually still online, you can run the command with the `-force` option. Once the source system has been recovered, there is no option to perform a cabinet level failback. Instead, resume and fail back each NAS server session to the original source using the graceful failover operation.

In case a NAS server session failover operation fails as part of a cabinet level failover, the system still continues to failover the other NAS server sessions. For any NAS server session failover that fails as part of a cabinet level failover operation, you can switch over the individual NAS server after the cabinet level failover completes using either the **Failover** action option in Unisphere or the `/prot/rep/session failover` CLI command.

**Unplanned failover in advanced file asynchronous replication topology**

With the advanced file asynchronous replication topologies introduced in OE version 5.x, multiple unplanned failover operations on the same resource could be executed serially or concurrently. For the end remote sites, the unplanned failover behavior does

not change from earlier OE releases. However, the behavior of a cascade mode site depends on when the processing on the upstream and downstream sessions occur with regard to the planned failover command.

---

#### Note

Unplanned failover may cause duplicate IP addresses for the NAS server. If the NAS server is running with CIFS support, it may cause a duplicate CIFS or SMB server as well. IP addresses need to be well planned and assigned to avoid this issue.

---

- If the processing of the upstream session unplanned failover command on the cascade mode occurs before the processing of the downstream session unplanned failover command, then the cascade site will be in destination mode.
- If the processing of the upstream session unplanned failover command on the cascade mode occurs after the processing of the downstream session unplanned failover command, then the cascade site will not be in destination mode.

With the advanced file asynchronous replication topologies, the destination of a resume operation could be functioning as a read/write (RW) production source resource for another replication session. The resume operation checks for this condition and, if discovered, causes the resume operation to fail. In this RW to RW (source to source) case, the local data change could exist in both sides. The resume operation includes an option to overwrite data on the destination resource to complete the operation successfully:

- In Unisphere, selecting **Resync the remote and overwrite any data written to the remote** discards the data changes in the destination resource. The data in the local source resource is retained and synchronized with the data in the destination resource. The local source resource is changed to destination mode and replication resumes in the original direction.
- In the CLI, `-forceSyncData` forces data transfer from the local source resource to the destination resource, even if the destination resource has data that is not replicated from the local source resource. The data in the local source resource is retained and synchronized with the data in the destination resource. The local source resource is changed to destination mode and replication resumes in the original direction.

In another case, the destination of a resume operation could be functioning as an active destination resource for another replication session. The resume operation checks for this condition and, if discovered, causes the resume operation to fail. A Resume operation from a local resource to keep local data changes is not allowed to a remote resource that is an active destination.

## Using replication for planned downtime

Unlike a disaster, in which the primary (source) system is lost due to an unforeseen event, planned downtime is a situation for which you plan and take the source system offline for maintenance or testing purposes on the destination system. Prior to the planned downtime, both the source and destination are running with an active replication session. When you want to take the source offline in this scenario, the destination system is used as the production system for the duration of the maintenance period. Once maintenance or testing completes, return production to the original source system. Planned downtime does not involve data loss.

To initiate a planned downtime, use the **Failover with sync** option (for synchronous replication, use **Failover** option in Unisphere) on the source system. When you fail over a replication session from the source system, the destination system is fully

synchronized with the source to ensure that there is no data loss. The session remains active with roles switched for synchronous replication, and paused for asynchronous replication, while the source becomes Read-Only and the destination becomes Read-Write. The destination storage resource can be used for providing access to the host.

Performing a failover with sync operation on an asynchronous replication session results in replication copying all the data, including any snapshots that have been created or marked for copy since the last sync occurred, to the destination site. Once the copy is finished, the destination is an exact replica of the source site and the roles are switched similar to the failover operation.

To restore operations on the source, fail back the replication session.

For synchronous replication, both roles and operations switch sides. To resume the operations on the original source, perform a failover again. For asynchronous replication, to resume the operations on the destination and switch the roles, resume the replication session. To resume the operations on the source, fail back the replication session.

#### **File-based replication consideration**

The NAS server replication session should be in one of the following states in order to do a planned fail over to the destination system:

- Idle
- Auto Sync Configured
- Active

If the NAS server replication session is in one of the following states, you cannot do a planned fail over to the destination system:

- Paused
- Error states

For asynchronous replication to minimize disruption during a planned downtime window, ensure that the NAS server and associated file system replication sessions are manually synchronized first and then failed over. Follow these steps:

1. Synchronize the NAS server replication session using the **Sync** option.
2. Synchronize the replication sessions for each of the file systems associated with the NAS server using the **Sync** option. This ensures that the destination file systems have the latest data and minimal data will need to be transferred when the replication sessions switch over.
3. Inform file system users and quiesce I/O operations from hosts and applications using the file systems in the NAS server.
4. Switch over the NAS server replication session using the **Failover with sync** option. This action triggers a group operation towards the NAS server replication session and its associated file system replication sessions.
5. Once all replication sessions have successfully failed over, resume I/O operations with the relevant applications and hosts.

---

#### **Note**

Any I/O attempted when the failover is occurring may result in read/write errors or stale file handle exceptions.

---

#### **Planned failover in advanced file asynchronous replication topology**

With the advanced file asynchronous replication topologies introduced in OE version 5.x, multiple planned failover operations on the same resource could be executed serially or concurrently. For the end remote sites, the planned failover behavior does

not change from earlier OE releases. However, the behavior of a cascade mode site depends on when the remount task on the upstream and downstream sessions occur with regard to the planned failover command.

---

#### Note

Planned failover may cause duplicate IP addresses for the NAS server. If the NAS server is running with CIFS support, it may cause a duplicate CIFS or SMB server as well. IP addresses need to be well planned and assigned to avoid this issue.

---

- If the remount task of the upstream session planned failover command on the cascade mode occurs before the remount task of the downstream session planned failover command, then the cascade site will be in destination mode.
- If the remount task of the upstream session planned failover command on the cascade mode occurs after the remount task of the downstream session planned failover command, then the cascade site will not be in destination mode.

With the advanced file asynchronous replication topologies, the destination of a resume operation could be functioning as a read/write (RW) production source resource for another replication session. The resume operation checks for this condition and, if discovered, causes the resume operation to fail. In this RW to RW (source to source) case, the local data change could exist in both sides. The resume operation includes an option to overwrite data on the destination resource to complete the operation successfully:

- In Unisphere, selecting **Resync the remote and overwrite any data written to the remote** discards the data changes in the destination resource. The data in the local source resource is retained and synchronized with the data in the destination resource. The local source resource is changed to destination mode and replication resumes in the original direction.
- In the CLI, `-forceSyncData` forces data transfer from the local source resource to the destination resource, even if the destination resource has data that is not replicated from the local source resource. The data in the local source resource is retained and synchronized with the data in the destination resource. The local source resource is changed to destination mode and replication resumes in the original direction.

In another case, the destination of a resume operation could be functioning as an active destination resource for another replication session. The resume operation checks for this condition and, if discovered, causes the resume operation to fail. A Resume operation from a local resource to keep local data changes is not allowed to a remote resource that is an active destination.

## Failback a replication session

To resume operations on a source system, the associated replication session needs to be failed back. To fail back a replication session, use the **Failback** option on the original destination system. For asynchronous replication, failback synchronizes the original source with the changes made on the original destination after failover, including any snapshots that have been created since the failover operation occurred. It then restores the source as the production system and restarts the replication session in the original direction.

For synchronous replication, failback is only used after an unplanned failover (use the failover operation again if a planned failover had occurred). Failback synchronizes the original source with the changes made on the original destination after an unplanned failover, not including any snapshot. A full copy is needed (to resume from the

destination) due to a restriction of the MirrorView/Synchronous technology. If there are snapshots on the source storage resource, its size will be increased (about 100%). If those old snapshots are deleted later, the size will shrink back. Failback then restores the source as the production system and restarts the replication session in the original direction.

#### File-based replication consideration

To resume operations on a source system, the associated NAS server replication session needs to be failed back. To fail back a NAS server replication session, use the **Failback** option on the original destination system. This action triggers a group operation towards the NAS server replication session and its associated file system replication sessions.

#### Failback in advanced file asynchronous replication topology

With the advanced file asynchronous replication topologies, fan-out and cascade, introduced in OE version 5.x, the destination of a failback operation could be functioning as a read/write (RW) production source resource for another replication session. The failback operation checks for this condition and, if discovered, causes the failback operation to fail. In this RW to RW (source to source) case, the local data change could exist in both sides. The failback operation includes an option to overwrite data on either the source or destination resource to complete the operation successfully:

- In Unisphere **Keep local data changes by updating the remote resource** or in the CLI `-syncData force` discards the data changes in the destination resource. The data in the local source resource is retained and synchronized with the data in the destination resource. The local source resource is changed to destination mode and replication resumes in the original direction.
- In Unisphere **Keep remote data by discarding all local data changes** or in the CLI `-syncData ignore` discards the data changes in the local source resource. The data in the destination resource is retained and synchronized with the data in the local source resource. The local source resource is changed to destination mode and replication resumes in the original direction.

In another case, the destination of a failback operation could be functioning as an active destination resource for another replication session. The failback operation checks for this condition and, if discovered, causes the failback operation to fail. Local data changes cannot be synchronized back to the remote session when that session is already an active destination. The failback option, in Unisphere **Keep remote data by discarding all local data changes** or in the CLI `-syncData ignore`, can be used to overwrite data on the source resource to complete the operation successfully.

## Asynchronous replication of a synchronous file replication to a third site

MetroSync for Unity supports configuring synchronous replication sessions of a NAS server and its file systems to one Unity system destination site and asynchronous replication sessions of the same NAS server and its file systems to a different Unity system destination site. These systems must have OE version 4.4 or later.

When creating a new NAS server or when a NAS server does not have an associated replication session, you can configure one synchronous replication session and one asynchronous replication session for that NAS server. In the case of an existing NAS server, if one asynchronous replication session is already associated with it, only a synchronous replication session can be created for it. If one synchronous replication



session is already associated with the NAS server, only an asynchronous replication session can be created for it.

---

#### Note

The asynchronous replication destination NAS server is selected as **Used as backup only** in Unisphere by default when synchronous replication is already enabled. When using CLI, it must be specified, otherwise, the asynchronous replication session creation will fail.

---

When creating a new file system or when a file system does not have an associated replication session, you can configure one synchronous replication session and one asynchronous replication session for that file system. In the case of an existing file system, if one asynchronous replication session is already associated with it, only a synchronous replication session can be created for it. If one synchronous replication session is already associated with the file system, only an asynchronous replication session can be created for it.

---

#### Note

By default, the new file system replication session is created with the same attributes as the associated NAS server replication session.

---

With asynchronous replication, internal checkpoint snapshots are routinely taken using the manual or automatic Recovery Point Objective (RPO) policy, and then replicated to the destination system. In addition, the internal checkpoint snapshots are synchronously replicated to the partner MetroSync system. This action ensures that a common-base snapshot will be available on the source and asynchronous and synchronous destination sites in this topology. After a MetroSync failover of synchronously replicated NAS servers and associated file systems, run a preserve operation on the new source synchronous NAS server replication session. This operation restores the asynchronous NAS server and file system replication sessions by using the replicated internal checkpoint snapshots on the new source system, as a common-base snapshot, without requiring a full synchronization.

---

#### Note

Internal asynchronous replication snapshots are refreshed on the source and asynchronous and synchronous destination sites on every asynchronous session sync operation. If the snapshot create or refresh operation failed on the synchronous replication destination site, the internal snapshot create or refresh operation on the synchronous replication destination site will be retried as part of the next asynchronous session sync operation. If snapshots cannot be created or refreshed on the synchronous replication destination site because the synchronous session is not Active or the connection is broken, the internal snapshot create or refresh operation on the synchronous replication destination site will be retried as part of the next asynchronous session sync operation after the synchronous replication session is restored.

---

#### Concurrent operations compatibility

A majority of the synchronous and asynchronous replication operations for coexisting synchronous and asynchronous replication sessions can be run concurrently with the following exceptions:

- Create and delete - You cannot run a create operation for an asynchronous replication and a delete operation for a synchronous replication or the reverse concurrently. The operation will be rejected.

- Failover or failback - Both of these operations are not supported on an asynchronous replication destination NAS server that is selected as **Used as backup only**.

### Preserve asynchronous replication sessions

When the synchronous replication sessions of a NAS server and its file systems are failed over (planned or unplanned) or failed back, the associated asynchronous replication sessions can be switched manually from the production site. This switching operation on the synchronous replication sessions preserves the asynchronous replication sessions with the active production site.

For example, source production site A has a synchronous replication of a NAS server and its file systems to destination site B and an asynchronous replication of the same NAS server and its file systems on the source production site A to a third site C. When the synchronous replication is failed over from site A to site B, you can manually preserve the asynchronous replication sessions by running a Preserve asynchronous replication operation on the synchronous NAS server replication session on the new source production site B. This operation switches the asynchronous replication sessions from the old source production site A to the new source production site B.

---

#### Note

Only users with Administrator or Storage Administrator roles are allowed to perform this preserve operation.

---

While performing the initial synchronization operation between the new source site and the asynchronous replication destination site, the new source storage system searches for internal snapshots on itself and the asynchronous replication destination site, determines whether there is a common base snapshot, and if so, replicates only a differential copy of the production data and not a full copy. If a common base snapshot is not found, a full data copy will be performed.

When a NAS server synchronous replication session's production site is switched between the source site and the destination synchronous replication site (for example, through a failover or a failback operation), the source NAS server of the asynchronous replication sessions changes to replication destination mode, the data transfer of the source NAS server asynchronous replication sessions is stopped. The operational status for the related sessions in this case is `hibernated`. After a preserve operation, the asynchronous replication sessions between the destination synchronous replication site (new source) and the asynchronous replication site are established. Later, the original source site may become the production site of the synchronous replication sessions again. In this case, the NAS server asynchronous replication sessions between the old destination synchronous replication site and the destination asynchronous replication site will become `hibernated`.

### Restrictions and limitations

The following restrictions and limitations relate to file asynchronous replication to another site for backup:

- The asynchronous replication sessions can be preserved only to a resource on the synchronous replication production site.
- The remote system connection towards the backup (asynchronous replication destination) site needs to be created in advance on the synchronous replication destination site for successful preserve operations.
- The asynchronous replication session can be used for backup only, failover is not allowed on the asynchronous replication session unless the backup only property is removed manually and failover is executed from the asynchronous replication destination site. Such an operation will break the synchronous replication,

therefore, you must ensure that either the synchronous replication session is deleted or both sites of the synchronous replication session are down and will not be recovered before removing the backup only property and failing over.

- If the limit of an asynchronous replication session on the preservation site is reached, the restore will fail with an error message and no more sessions can be restored.
- The preserve asynchronous replication sessions operation establishes the asynchronous replication sessions from the new synchronous production site to the backup site. After the preserve operation successfully completes, the new synchronous production site becomes the asynchronous replication source of the backup site. If the NAS server, file systems, and snapshots of the new synchronous production site are not the same as the old synchronous production site (for example, a snapshot is created or marked for asynchronous replication when a synchronous replication session is not in the Active state or is asynchronously replicated when a synchronous replication session is fractured) the new synchronous production site to the asynchronous replication backup site will continue based on the snapshots and information on the new synchronous production site. Some snapshots may not be replicated to the backup site. In this case, a warning will be generated when marking the snapshot for asynchronous replication when the synchronous replication session is not Active. Some snapshots may result in duplicate storage space.
- Supported on Unity systems with OE versions 4.4 or later.
- The preserve file system asynchronous replication sessions process can avoid the full copy process only when a system snapshot is available for a common baseline. If there are any snapshots on the destination file system when a full copy is initiated, the storage space will be increased (by 100% of the production file system size). If those old snapshots are deleted, the size will shrink back.
- Asynchronous replication session state is always preserved as either Auto Sync Configured or Idle.
- If a preserve asynchronous replication session operation is ongoing, it must be cancelled before a planned failover or fallback synchronous session is executed.

#### **Remove an asynchronous replication session**

You can remove the asynchronous replication session between the source site and the asynchronous replication destination site. While performing this operation, the internal snapshots of the source and asynchronous destination sites that are used for the asynchronous replication session are removed. The internal snapshots of the synchronous destination site will be removed if the deletion happens when the synchronous replication session is Active.

---

#### **Note**

The internal snapshots used for asynchronous replication between the source site and the asynchronous replication site remain existing after asynchronous replication session removal in case a preserve operation needs to be performed.

---

#### **Remove a synchronous replication session**

You can remove the synchronous replication session between the source site and the synchronous replication destination site. While performing this operation, the internal snapshots of the synchronous destination site that are used for the asynchronous replication session are removed. The internal snapshots of the source and asynchronous destination sites that are used for the asynchronous replication session remain existing.

### Remove an internal file system snapshot manually

You can remove a file system internal snapshot manually when no asynchronous replication session exists which uses that snapshot. If an asynchronous replication session does exist which uses that snapshot, the remove operation will fail.

#### Note


Internal snapshots on the source site and its asynchronous replication destination site are removed as part of the asynchronous replication session delete operation. If the asynchronous replication session still exists, the internal snapshots cannot be removed. Internal snapshots on the synchronous replication destination site can be removed either manually using the `-force` delete operation attribute or during a synchronous replication session delete operation.

## Cascade mode for third site asynchronous replication

When one synchronous replication session is created for MetroSync for Unity, a session should not be created from the destination side and only one remote asynchronous session should be created from the source. When a storage object is created as the destination mode and one asynchronous session is created on it as the destination, that storage object can be used to create another asynchronous session when it acts as the source.

The following figure and table show an example of cascade mode in a MetroSync configuration. Each lettered box represents a system running OE version 5.x. The thick solid lines between boxes represent asynchronous replication sessions and the thin solid line between boxes represents a synchronous replication session. A represents a production site. All of the replication sessions on this resource act as the source and are not in destination mode. B represents a synchronous replication destination remote site. C represents an asynchronous replication destination remote site in cascade mode. The resource on C acts as the source for a session on D and acts as the destination for another session from A. D represents an end remote site. The session on this resource acts as the destination and is in destination mode.

**Figure 3** Cascade mode for third site asynchronous replication



**Table 5** RPO for replication sessions in third site asynchronous replication topology

<b>Replication Session</b>	<b>User Snapshot Replication</b>	<b>RPO (in minutes)</b>
A to B (synchronous)	Yes	0
A to C (asynchronous)	Yes	60
C to D (asynchronous)	No	1440



# CHAPTER 2

## Replication workflow

This chapter addresses the following topics:

- [Configure replication](#)..... 32
- [Related concepts](#)..... 35

## Configure replication

Replication can occur locally or between a local storage system and a remote storage system. Use the native Replication feature in Unisphere to configure replication for block or file storage resources. The system supports asynchronous and synchronous replication of all storage resources, including file systems, NAS servers, LUNs, LUN groups, VMware VMFS datastores, VMware NFS datastores, and thin clones.

---

### Note

You can also configure synchronous and asynchronous replication for LUNs and VMware VMFS datastores using the RecoverPoint appliance-based solution.

---

## Configure local replication

---

### Note

Local replication is applicable to asynchronous replication only. Local replication is not applicable to synchronous replication.

---

Replication interfaces and connections do not apply to local replication. When using the CLI or the REST API, once you create the identical source and destination storage resources or NAS servers on the storage system, you can proceed to configure a replication session. When using Unisphere, you only need to create the source storage resources or NAS servers on the storage system. Unisphere does not allow you to create a session with an existing destination. A DR\_ is concatenated onto the resource name for local destinations to ensure that the source and destination names on the same system are unique (that is, LUN names need to be unique).

Note the following:

- For a disaster recovery scenario, it is recommended that the destination storage resource and NAS server are configured on a storage pool other than the pool used for the source storage resource and NAS server.
- For a migration scenario, which means migrating the source storage resource and NAS server to a destination storage resource and NAS server on the same pool, use the CLI to configure local replication. The Unisphere GUI does not allow local replication between storage resources and NAS servers on the same pool.

## Configure asynchronous replication

### Before you begin

If you are configuring asynchronous replication for a tenant, create a pool for the tenant on the destination system that matches the corresponding pool on the source system (if one exists). Then add the tenant to the destination system, using the same UUID and VLANs as the tenant on the source.

If you are configuring asynchronous replication in a coexisting synchronous and asynchronous replication topology, create the asynchronous replication destination NAS server with both the `-replDest` and the `-backupOnly` attributes set to `yes`. These attributes must be set to `yes` on the asynchronous replication destination NAS server when the source NAS server is synchronous replicated; otherwise, the asynchronous replication session cannot be created.



### Procedure

1. Configure the replication interfaces on each SP of the source and destination systems.
2. Configure a replication connection using the **Asynchronous** connection mode.
3. For file storage, create a replication session for the NAS server associated with the file storage.

---

#### Note

Storage resources included in a NAS server automatically get replicated when a replication session is first configured for the NAS server. The replication session for the storage resources will inherit the same attributes as the associated replication session of the associated NAS server. For the storage resources you do not want participating in replication, you can choose to remove the associated replication sessions manually.

---

4. Create a remote replication session for the storage resource.

---

#### Note

You only need to configure replication interfaces and connections for the first replication session between two systems. The same connection can be used again for subsequent replication sessions between the same systems.

---

## Configure synchronous replication

---

#### Note

Only remote replication is supported for synchronous replication.

---

### Procedure

1. Identify the Synchronous Replication Fibre Channel (FC) ports on each system.  
To determine the FC port used for synchronous replication, in the CLI console, run the command `/remote/sys show -detail`. Port information, similar to the following example, will appear in the output:

```
Synchronous FC ports = spb_fc4, spa_fc4
```

For more information, see the *Unisphere CLI User Guide*.

2. Zone the Synchronous Replication FC ports between the systems.  
If the source and destination systems are co-located, instead of zoning, you can choose to use direct-connected FC cables between the SPs.
3. Configure the replication interfaces on each SP of the source and destination systems based on the connection mode you want:
  - For synchronous replication support (**Synchronous** connection mode), use the Synchronous Replication Management Ports on each SP of both the systems.
  - For asynchronous and synchronous replication support (**Both** connection mode), in addition to the replication interfaces for Synchronous Replication

Management Ports, configure additional interfaces using the Ethernet Ports on each SP of the source and destination systems.

4. Configure the replication connection between source and destination systems from the source system only.
  - For synchronous replication support, specify the **Synchronous** connection mode.
  - For asynchronous and synchronous replication support, specify the **Both** connection mode.
5. Create the synchronous replication session.

---

**Note**

You only need to configure replication interfaces and connections for the first replication session between two systems. The same connection can be used again for subsequent replication sessions between the same systems.

---

## Configure asynchronous replication of synchronous replication to a third site

### Before you begin

To configure coexisting synchronous and asynchronous replication, the following preconditions must exist:

1. The source Unity system and two remote Unity systems with OE version 4.4 or later should be installed and powered on.
2. A zoned Fibre Channel connection exists between the source system and the remote synchronous replication destination system.
3. IP connections exist between the source and two remote systems.
4. Each system should have sufficient storage available.

---

**Note**

Only remote replication is supported for synchronous replication.

---

### Procedure

1. Create a synchronous replication management interface between the source system and the remote synchronous replication destination system.
2. Create remote systems to establish a synchronous replication connection between the source system and the synchronous replication destination system.

---

**Note**

The connection mode must be set to **Both**.

---

3. Create asynchronous replication IP interfaces on the source system and the remote destination systems.
4. Create remote systems to establish asynchronous replication connections between the source system and the asynchronous replication destination system and the synchronous replication destination system and the asynchronous replication destination system.

---

**Note**

The connection mode must be set to **Asynchronous**.

---

5. Create a NAS server on the source system.
  6. Create file systems on the NAS server on the source system.
  7. Create a NAS server synchronous replication session on the source system.
- 

**Note**

You only need to configure replication interfaces and connections for the first replication session between two systems. The same connection can be used again for subsequent replication sessions between the same systems.

---

8. Create synchronous replication sessions for the file systems on the NAS server on the source system.
- 

**Note**

Storage resources included in a NAS server automatically get replicated when a replication session is first configured for the NAS server. The replication session for the storage resources will inherit the same attributes as the associated replication session of the associated NAS server.

---

9. Create a NAS server asynchronous replication session on the source system.  
The asynchronous replication destination NAS server is selected as **Used as backup only** in Unisphere by default when synchronous replication is already enabled. When using CLI, it must be specified, otherwise, the asynchronous replication session creation will fail.
  10. Create asynchronous replication sessions for the file systems on the NAS server on the source system.
- 

**Note**

Storage resources included in a NAS server automatically get replicated when a replication session is first configured for the NAS server. The replication session for the storage resources will inherit the same attributes as the associated replication session of the associated NAS server.

---

## Related concepts

### Replication interfaces

Each SP must include at least one replication interface configured on it. Each replication interface on the SP of the source system should be able to communicate to the replication interface on the SP of the destination system. To establish this, you must set up a switched network for remote replication. You can configure a replication interface with an IPv4 or IPv6 address.

Replication is supported between like type IP addresses. For example, you can set up a replication connection between a replication interface configured with IPv4 address on

the source system and a replication interface configured with IPv4 address on the remote system.

### Configuration for asynchronous replication

When you create replication interfaces for asynchronous replication, use any of the available Ethernet ports to configure an IP address on each SP. These IP addresses are then used for asynchronous replication between the remote systems.

### Configuration for synchronous replication

When you create replication interfaces for synchronous replication, you must use the Ethernet Port titled 'Sync Replication Management Port' to configure an IP address on each SP. This port is specifically used for management of the synchronous replication session, leveraging the same physical interface as the storage system's Unisphere management port, whereas the Fibre Channel (FC) ports are used for replicating data between systems. Before creating replication interfaces and connections for synchronous replication, ensure that you complete FC zoning and configuration.

---

#### Note

You can choose to configure Asynchronous, Synchronous, or both connection modes for block storage resources. If you choose to configure both connection modes, ensure that you configure an additional Ethernet Port IP interface on each SP of the source and destination systems.

---

## About replication connections

Replication requires a configured connection between the source system and destination system. This connection is called a replication connection. Replication connections use replication interfaces on each SP from the local system and pairs them up with the replication interfaces on each SP in the remote system.

For synchronous replication, a Fibre Channel (FC) connection is required for the data communication, and an IP interface for the management communication. Only one FC port can be configured and used for synchronous replication. Once an FC port is assigned, it will continue to act as the synchronous replication port and can never be changed.

To determine the FC port used for synchronous replication, in the CLI console, run the command `/remote/sys show -detail`. Port information, similar to the following example, will appear in the output:

```
Synchronous FC ports = spb_fc4, spa_fc4
```

For more information, see the *Unisphere CLI User Guide*.

---

#### Note

Before creating a replication interface or connection for synchronous replication, it is recommended that you configure the FC zoning.

---

### Connection modes

The following connection modes are supported:

- **Asynchronous**—Supports asynchronous replication of file and block storage resources. It uses the Ethernet Port IP interfaces configured on Storage

Processor A (SP A) and Storage Processor B (SP B) on the source and destination systems. Manual replication uses asynchronous connection as well.

- Synchronous—Supports synchronous replication of file and block storage resources. It uses the Sync Replication Management Port with IP addresses configured on SP A and SP B on the source and destination systems.
- Both—Supports synchronous replication of file and block storage resources using the Sync Replication Management Port and asynchronous replication of file and block storage resources. Ensure that you set up the additional Ethernet Port IP interfaces on SP A and SP B on the source and destination systems.

## About replication sessions

A replication session leverages a configured replication connection and associated interfaces to establish an end-to-end path for replicating data between the source and destination storage resources. For asynchronous replication, once a replication session is established, the system creates two replication system snapshots on the source storage resource and two replication system snapshots on the destination storage resources. The snapshots are used to keep point-in-time copies of data and maintain a common base between the source and destination storage resources.

The snapshots are refreshed based on the replication mode:

- For automatic synchronization, the snapshots are refreshed based on the Recovery Point Objective (RPO) you specified. You can adjust the RPO interval based on the amount of protection you want.

---

### Note

For asynchronous replication, you can specify an RPO from a minimum of 5 minutes up to a maximum of 1440 minutes (24 hours). Once set, the RPO value cannot be reset from non-zero (for asynchronous replication) to zero (the automatic static setting of RPO for synchronous replication) or from zero to non-zero.

- For manual synchronization, it is recommended that you synchronize the replication periodically.

---

### Note

The snapshots consume protection space based on the snapshots feature. Although the internal snapshots may appear listed in the Snapshots tab of a storage resource (source or destination), you cannot manage or modify them. The snapshots also do not participate in automatic delete operations.

The following steps provide a high level overview of how the asynchronous replication session works:

1. When the session and the internal snapshots are created, an initial synchronization occurs where all the data from the first snapshot of the source storage resource is replicated to the destination storage resource.

---

**Note**

For automatic synchronization, initial synchronization is enabled by default. For manual synchronization, there is no initial synchronization. An option is available for asynchronous replication to replicate existing snapshots on the storage resource as part of the initial synchronization. This option is only available when creating a replication session. In this case, replication transfers the snapshots in the order in which they were created on the source to the destination site. If you opt not to synchronize existing snapshots as part of the initial synchronization, the replication copies only the source data from the internal system snapshot.

---

The system supports up to 16 concurrent initial synchronization operations per SP.

2. Once data is replicated to the destination storage resource, the first snapshot on the destination storage resource is refreshed.  
The first snapshot on both systems then becomes your common base at that point in time. In event of a disaster, the common base is used to restore data from the last time it was synchronized.
  3. When the session is synchronized next, manually or automatically (based on the RPO), the second internal snapshot on the source system gets refreshed.
- 

**Note**

If a replication session is created with auto synchronization, any snapshots created or marked for replication since the previous RPO are replicated upon RPO expiry to the destination site. In case of a session configured with manual synchronization, any snapshots created or marked for replication since the last synchronization occurred are replicated whenever the session is manually synchronized.

---

4. Once the snapshot is refreshed, data is replicated to the destination storage resource.
- 

**Note**

Only the differential data between the first and second snapshots gets copied over.

---

5. The second snapshot on the destination storage resource is then refreshed.
6. Once this occurs, at that point-in-time, the second snapshots on both the systems now become the common base.

These steps occur each time a synchronization occurs. When a disaster occurs, replication uses the common base to go back to a good point in time copy of data. If failing over to the common base is not sufficient and replicated user snapshots exist, manually restore the destination resource to any of the replicated user snapshots.

Synchronous file replication works differently than asynchronous file replication. Once the operational status of a synchronous file replication session is Active, checkpoint snapshots that are created on the source system (either manually through Unisphere, CLI, RESTful command, or by a snapshot schedule) are synchronously created on the destination system. Snapshots that were created before the synchronous replication session was setup are not replicated.

For synchronous replication, the RPO is set to 0 automatically and the NAS server configuration is synchronously replicated to the destination system so that after failover, the destination NAS server can reflect the latest configuration. However, the view of the NAS server configuration at the destination system from the management interface is not always up to date, it is only updated automatically every 15 minutes.

An on demand update of the view of the NAS server configuration at the destination system can be issued from either Unisphere (**Refresh Configuration**) or CLI (`net/nas/server {-id <value> | -name <value> update -confView}`). Since the synchronous replicated configuration file system is unreadable during initial synchronization or synchronizing after a fracture, it is not possible to update the view of the NAS server configuration at the destination side at those times.

## About replication support for thin storage resources with data reduction

A thin destination LUN or LUN Group can be created with data reduction during replication session creation when the following criteria is met:

- The source LUN or LUN group is thin.
- The OE version is 4.2 or later.
- The pool is All-Flash.

If the source LUN or LUN group meet this criteria, thin is selected by default for the destination LUN or LUN group. You can choose to unselect it to create a thick destination LUN or LUN group. If thin is not selected for the destination storage resource, data reduction cannot be selected. Data reduction can only be selected when thin can be and is selected for the destination storage resource.

---

### Note

For file systems, replication matches the destination storage resource to the source. The destination is converted to match with the source after the synchronization operation (initial synchronization or manual synchronization). Therefore in this case, thin and data reduction cannot be selected.

---





# CHAPTER 3

## Considerations for synchronous replication

This chapter addresses the following topics:

- [Synchronization states](#).....42
- [Fibre Channel configuration for synchronous replication](#).....43

## Synchronization states

Replication sessions configured with synchronous replication (RPO set to zero) can have the following synchronization states. Each of these describe the state of the user data on the destination storage resource. They are useful for troubleshooting. However, for replication session management, it is recommended to check the operational status. For accurate information on the synchronization state of a replication, refer to the synchronization state on the source system.

---

### Note

The synchronization state on the destination system is not updated after a synchronous replication session is fractured.

---

- **Consistent** – The data on the destination system is consistent from the host perspective. The destination LUN is the latest or point in time copy of the source LUN.
- **In Sync** – The data on the destination system is consistent from the host perspective. The destination LUN is an exact copy of the source LUN.
- **Syncing** – The data on the destination system is getting updated from the source system. The updates represent the pending changes on the source system. These changes are not yet replicated to the destination system because of the following reasons:
  - A new replication session requires a full synchronization.
  - Replication session requires a full synchronization after an unplanned failover operation, when a failback or resume operation is performed.
  - Replication session was previously in a Consistent or In Sync state. A connection issue has created a backlog of changes on the source system. The backlog will be updated to the destination system once the connection issue is resolved.
  - Replication session was paused resulting in a backlog of changes on the source system. The backlog will be updated to the destination system once the replication session is resumed.
- **Out of Sync** – The state of the data on the destination LUN cannot be determined by the replication session. The update to the destination system may be incomplete. This can occur for any of the following reasons:
  - A connection issue disrupts a replication that is in a Syncing state.
  - The replication session is paused on the source system.
  - The replication session was deleted locally on the destination system.

### NOTICE

If a synchronous replication session is fractured or paused and there are still data writes on the source storage resource, recover or resume the session as soon as possible, since the re-synchronizing is by changed extent. If data writes are sparsely distributed on the source storage resource, re-synchronizing may transport much more data than the original data writes and, if there are snapshots on the destination storage resource, they will occupy more storage space size. If those old snapshots are deleted, the size will shrink back.

---

- **Inconsistent** – The state is reported only when replication is failed over. It appears when the synchronization state of the session was not In Sync or Consistent, prior to failover. In this case, it is recommended that you check for the consistency of the destination storage resource.

**Synchronization state transitions**

The following table lists the synchronization state transitions for some of the replication actions:

Action	Synchronization state transition
Initial synchronization	<b>Unknown &gt; Out of Sync &gt; Syncing &gt; Consistent &gt; In Sync</b>
Pause	<b>In Sync &gt; Consistent or Syncing &gt; Out of Sync</b>
Resume	<b>Consistent or Out of Sync &gt; Syncing &gt; Consistent &gt; In Sync</b>
Fail over	<p><b>Syncing &gt; Inconsistent or Out of Sync &gt; Inconsistent</b></p> <hr/> <p><b>Note</b></p> <p>If the state was <b>In Sync</b> or <b>Consistent</b> before a failover occurred, that state is retained.</p> <hr/>
Replication connection down	<b>Syncing &gt; Out of Sync or In Sync &gt; Consistent (if there is I/O)</b>

**NOTICE**

For synchronous replication, I/O performance on a storage resource will be negatively impacted after a synchronous replication session is created on it. Also, initial synchronization takes quite a long time to complete depending on the storage object size and thin or thick type. For a thick storage resource, synchronous replication always copies the complete storage corresponding to its size even when it is empty.

## Fibre Channel configuration for synchronous replication

Every replication session is tied to a Storage Processor (SP) via a replication interface. For synchronous replication, a Fibre Channel (FC) connection is used for data communication, and an IP interface used for management commands. Before creating a replication interface or connection, it is recommended that you configure the FC zoning.

**Note**

When using Single Mode SFP (to support a long range connection), SM cables along with an SM SFP connection on the switch are required as well as manually increasing the credits on the switch. In the case of direct connection between storage systems, both Synchronous Replication ports should use SM SFP along with SM cables. Refer to the *Unity Hardware Information Guide* for more information about SM SFP.

Only one FC port can be configured and used for synchronous replication. Once an FC port is assigned, it will continue to act as the synchronous replication port and can never be changed.

To determine the FC port used for synchronous replication, in the CLI console, run the command `/remote/sys show -detail`. Port information, similar to the following example, will appear in the output:

```
Synchronous FC ports = spb_fc4, spa_fc4
```

For more information, see the *Unisphere CLI User Guide*.

---

#### Note

RecoverPoint and the native Replication feature cannot simultaneously operate on the same volume or Fibre Channel port of the storage system. Even if you do not configure or use synchronous replication, do not configure the synchronous replication port for RecoverPoint. There is no such restriction on the Ethernet ports for iSCSI. You can configure RecoverPoint to use any available Ethernet port.

---

#### Zoning configuration

Ensure that you complete the following:

- Configure the port as an F-port to avoid any potential issues, and remove other types of ports from the list of allowed port types.
  - For a synchronous replication connection, connect SPA port of the source system to the SPA port of the destination system. Similarly, connect SPB port of the source system to the SPB port of the destination system.
- 

#### Note

Only a one to one FC connection is allowed for synchronous replication. One to more than one FC connection is not supported.

---

For example,

- OB\_H1450 SPA Port 4 <--> OB\_H1443 SPA Port 4
  - OB\_H1450 SPB Port 4 <--> OB\_H1443 SPB Port 4
- 

#### Note

If the FC connections are incorrectly crossed configured, for example:

- OB\_H1450 SPA Port 4 <--> OB\_H1443 SPB Port 4
- OB\_H1450 SPB Port 4 <--> OB\_H1443 SPA Port 4

connect the FC connections correctly. After correcting the FC connections, reboot both SPA and SPB of the source system and the destination system.

---

- The SPA and SPB ports on the same side (source or destination) must have the identical port configuration. However, the configuration of ports on the source system can be different from the port configuration on the destination system. For example, SPA Port 4 <--> SPA I/O Module 0 Port 0 and SPB Port 4 <--> SPB I/O Module 0 Port 0 are supported.
- It is recommended to have a separate FC port for FC host access. But, the synchronous replication FC port can be shared with host access.
- Although setting the port speeds to Auto works, it is recommended that you set the port speeds for the FC ports on both systems and the FC switch port.

#### High availability consideration

To ensure high availability for block synchronous replication, consider the following:

- If the FC connection for an SP is disrupted, set the SP ownership for the source LUN to the other SP. SP ownership on the destination LUN automatically switches to the corresponding SP on the destination system.
- If an SP on the destination system shuts down or is faulted (not in a normal state), the system will automatically change the SP ownership on the destination LUN. Ensure that you manually set the SP ownership for the source LUN to the other SP.
- For file synchronous replication, if an SP failover, failback, or NAS server local move results in a synchronously replicated source NAS server and file systems' current SP that is different from the destination NAS server and file systems' current SP, the sessions will be fractured until both are back to the same source SP and destination SP. Unlike block synchronous replication, no automatic switch is triggered. You can perform a NAS server local move manually to make the current SPs of the source and destination match.



# CHAPTER 4

## Configure replication using Unisphere

This chapter addresses the following topics:

- [Configure protection and mobility interfaces](#)..... 48
- [Configure replication connections](#)..... 49
- [Create a replication session](#)..... 49
- [Manage replication sessions](#)..... 51
- [Configure an override network address for file replication](#) ..... 56

## Configure protection and mobility interfaces

### Before you begin

Protection and mobility (import) interfaces can be shared between replication and import. For import, only VDM imports require interfaces. Block imports do not require interfaces.

Protection and mobility (import) interfaces are configured to support VDM imports and must be created prior to creating an import connection. A mobility interface IP address is assigned to SPA and SPB on the target Unity system. Once the mobility interface is configured, you can create the import connection between the Unity system and the VNX system. Mobility interfaces are not used for block import sessions.

Ensure the following:

- The interface port is cabled and connected to a network switch.
- Both SPs are up and running.

Obtain the following information for each Storage Processor (SP):

- IP address associated with the interface (replication or import). Although you can specify an IPv4 or IPv6-based address, ensure that you specify the same type of address for both SPs.
- IP address mask or prefix length that identifies the associated subnet.
- Gateway IP address associated with the interface.
- If applicable, the VLAN ID (between 1 and 4095) you want to associate the interface with.

---

### Note

For the network to continue functioning properly, ensure that you set the VLAN ID only when you have configured the network switch port to support VLAN tagging of multiple VLAN IDs.

---

### Procedure

1. Under **Protection & Mobility**, select **Interfaces**.
2. Perform one of the following actions:
  - To create an interface, select the **Add** icon. On the **Create Interface** window, specify the relevant information:
    - For asynchronous replication or import, from the **Ethernet Port** list, select an available Ethernet port.
    - For synchronous replication, from the **Ethernet Port** list, select **Sync Replication Management Port**.

---

### Note

Do not use **Sync Replication Management Port** for asynchronous replication or import interfaces.

---

- To modify an interface, select the interface, and then select the **Edit** icon. On the **Interface Properties** window, specify the relevant information.
- To delete an interface, select the interface, and then select the **Delete** icon.



**NOTICE**

Before you delete an interface, ensure that the interface is not being used by any replication or import session.

---

## Configure replication connections

### Before you begin

Ensure that you have set up relevant replication interface pairs, one on each SP, on the source and destination systems. Obtain the following information:

- For remote replication, the IP address and associated user authentication credentials to connect to the remote system.
- For local replication, the password associated with your user account.
- The connection mode you want to use for the replication: Asynchronous, Synchronous, or Both.

**NOTICE**

If a replication connection already exists and you plan to add a different mode of file replication, do not attempt to create a new connection. Change the existing replication connection mode to Both. Also, ensure that you have the appropriate interface types configured to support both asynchronous replication (eth2, eth3) and synchronous replication (sync replication mgmt port).

---

### Procedure

1. Under **Protection & Mobility**, select **Replication > Connections**.
2. Perform one of the following actions:
  - To create a replication connection, select the **Add** icon. On the **Create Replication Connection** window, specify the relevant information.
  - To modify a replication connection, select the replication connection, and then select the **Edit** icon. On the **Replication Connection Properties** window, specify the relevant information.
  - If new replication interfaces were added, or existing replication interfaces were deleted, the remote system connection may become outdated. Select the relevant replication, and then select **Verify and Update** to update the remote system connection to pick up the latest replication interface changes on the local and remote systems.

## Create a replication session

### Before you begin

Ensure that you have first created relevant replication interfaces and connections, and then determine the following:

- The replication mode you want. You can choose from the following options:
  - Asynchronous
  - Manual
  - Synchronous

- The Recovery Point Objective (RPO) for the replication session.
  - The system you want to assign as the replication destination system. This is based on the replication connections configured on the storage system.
  - The name, pool, storage provisioning, and tiering policy you want to use for the destination storage resource. The system will automatically create a destination storage resource as part of this process.
- 

#### Note

You can convert a thin LUN to a non-thin (thick) LUN, or a thick LUN to a thin LUN with a LUN move operation. To enable data reduction on a thin LUN requires an All-Flash pool on the destination system. For thick file systems, the replication process matches the destination storage resource to the source. In this case, thin and data reduction cannot be selected for file systems. For thin file systems, the following rules apply for replication:

- If the source file system is thin, then the destination file system is also thin.
  - If the source file system is thin and either the source or destination system does not support data reduction, then the destination file system is also thin.
  - If both the source and destination systems support data reduction, then the source or destination can have either data reduction enabled or no data reduction enabled.
  - However, if you change the source file system data reduction attribute, the change is not replicated to the destination system (retains the original setting) regardless of using asynchronous or synchronous replication.
- 
- Replicate scheduled snapshots to the destination and to use the same retention policy or a different policy. Also, when a replication session is created some time after the resource was created, determine whether to select the one-time option to replicate all existing snapshots.
- 

#### Note

Snapshots that have been created and attached as well as read/write (share) snapshots (as opposed to read-only checkpoint snapshots) are not eligible for replication. Only unattached (read-only) snapshots are eligible for replication. For asynchronous replication, you can replicate existing snapshots and snapshots created from snapshot schedules. For synchronous file replication, you cannot replicate existing snapshots or snapshots created from snapshot schedules. You can only replicate those snapshots and snapshots created from snapshot schedules after you have established the synchronous replication session and it is Active. For synchronous block replication, you cannot replicate any snapshot.

---

#### Note

To do snapshot replication from a source system running OE version 4.0, 4.1, 4.2, 4.3, or 4.4 to a destination system running OE version 5.x, requires upgrading the source system to OE version 4.5 first. Upgrading to OE version 4.5 is not required but recommended if you want to do LUN or file system replication from OE version 4.0, 4.1, 4.2, 4.3, or 4.4 to OE version 5.x without any snapshot replication.

---

With the advanced file asynchronous replication topologies, fan-out (star) and cascaded mode, feature introduced in OE version 5.x, up to four asynchronous

replication sessions can be created on the same storage object. The following restrictions apply to this feature:

- This feature only supports file storage objects and does not support block storage objects.
- Only one session can be created between two remote systems.
- All systems joining the multiple sessions, either in fan-out (star) or cascaded mode, must be running OE version 5.x.
- Only one local or loopback asynchronous session is supported and the associated resource does not have to be in destination mode.
- When one synchronous session is created, a session should not be created from the destination side and only one remote asynchronous session should be created from the source.
- When the storage object is created as the destination mode and one session is created on it as the destination, the storage object can be used to create another session when it acts as the source.
- A user snapshot replication can only be supported for one asynchronous session among all the sessions associated with the same storage object.

#### Procedure

1. Under **Storage**, select the storage resource, and then select **Edit**.

---

#### Note

For file storage, you must first configure replication on the NAS server.

---

2. On the **Properties** window, select the **Replication** tab.
3. Select **Configure Replication**.
4. In the **Create a Session** window, specify the relevant information.

Once you specify the necessary information for the destination storage resource or NAS server, the system will automatically create the storage resource or NAS server on the destination system.

---

#### Note

On a NAS server that is protected by replication, you must create a replication session for each file system on it. Otherwise, file system related configurations like shares and exports may be lost after a NAS server replication session failover.

- When you use Unisphere to create replication for a NAS server, the system automatically creates a replication session for each file system on it.
  - When you use Unisphere to create a new file system, Unisphere creates a replication by default.
- 

## Manage replication sessions

A replication session establishes an end-to-end path for a replication operation between a source and destination. The replication source and destination may be

remote or local, so the session establishes the path that the data follows as it moves from source to destination.

If you want to set up a replication session, see [Configure replication](#) on page 32.

With the advanced file asynchronous replication topologies, fan-out (star) and cascaded mode, feature introduced in OE version 5.x, up to four asynchronous replication sessions can be created on the same storage object. The following restrictions apply to the managing replication sessions operations:

- **Delete**

- In fan-out or star mode:

- On the source side of the replication, concurrent delete for four replication sessions is not supported. If four replication session delete operations are triggered at the same time, the operations are queued to be done one by one.
    - If the remote system is in the lost communication state, the replication session delete operation is done only on one side (source or destination). The residual replication session is visible on the remote side.
    - When the replication session on the destination side is deleted, it is removed only on the destination side, and the source side replication session state is changed to `Non recoverable error`.
    - If the previous replication session delete operation failed, that replication session remains visible and the delete operation can be retried. The residual replication session could be recovered and shown after a system reboot.
    - If the replication session delete failed on the source side, the replication session is already deleted on the destination side and the residual replication session is only visible on the source side. If the replication session failed on the destination side, the replication session is already deleted on the source side and the residual replication session is only visible on destination side.
    - If the replication session to be deleted is already in a `Non recoverable error` state caused by a previous unexpected operation, it is recommended to clean-up the replication session residual components by using the `svc_dataprotection` service command. It can be used to clean up a single session with the session ID or clean up the orphan components for the sessions associated to a single storage resource. If the resource is offline, it needs to recover the resource to be online, then continue the deletion.

---

**Note**

For more information about the `svc_dataprotection` service command, see the *Unity Service Commands Technical Notes*.

---

- In cascade mode:

- With A->B and B->C cascade mode replication session configured, session A->B and B->C concurrent delete is supported.
      - With A->B and B->C cascade mode replication session configured, when deleting A->B first, the destination mode of the storage object on B is not changed.

- **Pause**

- When pausing NAS server related sessions, only the associated file system session to the same remote system will be paused together with the NAS server session.

- Other existing replication operations are supported as in earlier OE versions.
- **Resume**
  - Only one active destination replica should exist.
  - When resuming the NAS server related sessions, only the associated file system session to the same remote system will be resumed together with the NAS server session. When there are multiple sessions associated with one destination object, it is already acting as a source for some sessions, or it is in destination mode while at least one session is not failed over or switched over.

### Procedure

1. Access the **Replication Session Details** screen in one of the following ways:
  - If you know the name of the replication session, go to **Protection & Mobility > Replication > Sessions**.
  - If you know the name of the storage resource or NAS server associated with the replication session, go to the **Replication** tab on the **Properties** window of the relevant storage resource or NAS server.
2. Perform one of the following actions:

---

#### Note

On Unity systems running OE version 4.2, the following asynchronous replication actions affect both the NAS server and its associated file systems when run at the NAS server level:

- Failover
- Failover-with-sync
- Failback
- Pause
- Resume

On Unity systems running OE version 4.4, the following synchronous replication actions affect both the NAS server and its associated file systems when run at the NAS server level:

- Failover
- Failback
- Pause
- Resume
- Preserve

Those same replication actions towards a file system, with the exception of Preserve, remain at the file system level. The following asynchronous replication actions affect only the NAS server when run at the NAS server level or are still individual operations toward file system replication sessions:

- Create
  - Sync (applicable to asynchronous replication only)
  - Delete
  - Modify
-

Action	Description
<b>Modify a session</b>	Select the <b>Edit</b> icon to modify the name and synchronization settings of the session.
<b>Delete a session</b>	<p>Select <b>Delete</b>. Deleting a replication session also deletes the internal snapshots used by the session. If the connection between the systems is healthy, the session on the destination system gets automatically deleted when you delete the session on the source system. You can individually delete replication sessions at the source and destination systems, however, this is not recommended unless the connection is down.</p> <hr/> <p><b>Note</b></p> <p>You cannot delete a replication when a synchronization is in progress.</p>
<b>Pause and resume a session</b>	<p>Select <b>Pause</b> or <b>Resume</b>. Pause stops a replication session without deleting the replication session. Resuming a session from the destination system after a failover or failover with sync reverses the direction of the replication. Resume is not available on the destination if the session is not failed over. Changes on the destination system are copied back to the source system at the next RPO interval after a connection is restored.</p> <hr/> <p><b>Note</b></p> <p>Pausing a replication session while a snapshot is being copied results in stopping the transfer. The destination resource is rolled back to the common base snapshot so that the destination is brought to a consistent state. Resuming a paused session restarts the previously stopped transfer. Replication restores the destination resource to the common base snapshot first, then restarts the snapshot copy by copying the differences between the common base snapshot and the snapshot to be transferred.</p>
<b>Manually synchronize a session (for asynchronous replication only)</b>	<p>Select <b>Sync</b>. Changes since the last sync are replicated to the destination system.</p> <hr/> <p><b>Note</b></p> <p>If an associated synchronous replication session exists between the source site and another destination site when performing this operation, the internal snapshots that are refreshed are replicated on the synchronous replication destination site.</p>
<b>Fail over a session in an unplanned emergency</b>	Select <b>Failover</b> . For asynchronous replication, in response to an unplanned emergency scenario in which the source has become unavailable, failover should switch the status of the source and destination storage resources. After the

Action	Description
	<p>failover operation, the destination object's status becomes read/write, however, the source and destination do not switch roles until a resume operation is run.</p> <hr/> <p><b>Note</b></p> <p>In this type of failover, any active data transfers are terminated, causing a potential loss of data. If the source site is still available when you perform a failover, the system attempts to change the source object from read/write to read-only.</p> <hr/> <p>For synchronous replication, in response to an emergency/unplanned synchronous replication failover when there is no network connectivity to the source system, use <code>Failover</code> from the destination system without switches, then <code>Resume</code> to restart the sessions on the destination system. This action also reverses the direction from destination to source. Or, instead of doing <code>Resume</code>, you could simply <code>Failback</code> to go back to the source. Also, if an associated asynchronous replication session exists between the source site and another destination site, perform the preserve async operation on the synchronous replication destination site.</p>
<p><b>Fail over a session in preparation for a planned downtime (for asynchronous replication only)</b></p>	<p>For asynchronous replication, select <b>Failover with Sync</b>. The destination synchronizes with the source storage resource, the replication session is stopped, and the status of the storage resources are switched. After the failover operation, the destination object's status becomes read/write and the source object's status is shown as Restricted Replication Access = Yes.</p> <hr/> <p><b>Note</b></p> <p>To perform this type of failover, both the source and destination objects of the session must be available.</p> <hr/> <p>For synchronous replication, use <code>Failover</code> from the source system, where sessions remain running and reversed back from the destination to the original source system. After the failover operation, the destination object's status becomes read/write and the source object's status is shown as Restricted Replication Access = Yes. To move back to the original configuration, run another <code>Failover</code>, this time from the destination system, to gracefully fail back to the original source system. A <code>Resume</code> is not required in this scenario because the sessions remain running.</p>
<p><b>Fail back a session</b></p>	<p>Select <b>Failback</b>. A failback operation switches the status of the storage resources following a failover operation. Once the destination synchronizes with the source again, the source storage resource's status returns to read/write</p>

Action	Description
	<p>and the destination storage resource's status is shown as Restricted Replication Access = Yes.</p> <hr/> <p><b>Note</b></p> <p>You can perform this action from the destination system only. Both the source and destination storage resources must be available to perform a successful failback operation.</p> <hr/>
<p><b>Preserve an asynchronous replication session</b></p>	<p>Select <b>Preserve</b>. The preserve operation switches the asynchronous replication session related to the synchronous replication session to the active production site after the synchronous replication session (planned or unplanned) has been failed over or failed back.</p>

**Note**

Any modifications to the attributes of the source storage resource are not automatically synchronized over to the destination storage resource. When a failover occurs, ensure that you modify the attributes of the associated destination storage resource to match the attributes of the source storage resource.

## Configure an override network address for file replication

If you have file replication configured with a destination site on a different physical network, to ensure minimal downtime during a failover, ensure that you modify the destination NAS server properties to include an override address for the network interface.

**Note**

You may still need to manually fail over individual file systems associated with the NAS server.

If you have configured the source NAS server with LACP or FSN, or both, you must do one of the following:

- Configure the destination system with the same LACP or FSN configuration. When the LACP/FSN environment is mirrored on the destination system exactly, no further action is required and the destination NAS server will automatically select the LACP/FSN device as part of its configuration.
- Include an override address for the network interface after first creating the replication session.



**NOTICE**

After the replication session has been created and the source NAS server is using an LACP/FSN configuration, you must do the following:

1. Manually edit the **Network** settings in the **Properties** of the NAS server on the destination system.
  2. Select Override address and an interface to use. If you do not do this, a DU after a failover from the source system to the destination system will occur.
- 

To configure an override address on the destination NAS server:

**Procedure**

1. Under **Storage**, select **File > NAS Servers**.
2. On the **Network** tab, select the network interface, and click **Edit**.
3. On the **Edit NAS Server Network Interface** screen, select the **Override address** checkbox, and specify relevant address information.



# CHAPTER 5

## Configure replication using the CLI

This chapter addresses the following topics:

- [Manage network interfaces](#)..... 60
- [Manage remote storage systems](#)..... 65
- [Manage replication sessions](#)..... 72
- [Associated CLI commands](#)..... 98

## Manage network interfaces

Create interfaces to enable and control access between the system, the network, and configured hosts. Specify the IP address for the interface as well as the IP addresses for the subnet mask and gateway.

You can create the following types of interfaces:

- iSCSI interfaces for controlling access to iSCSI storage. You assign the interface to an iSCSI node.
- Replication interfaces for replication-related data or management traffic.

The system configures each interface on a pair of symmetrical SP ports. The interface can be moved between SPs. You have the option of indicating which SP the interface will use, either a physical port or a link aggregation port. You also have the option of specifying a virtual LAN (VLAN) ID, for communicating with VLAN networks.

Each interface is identified by an ID.

The following table lists the interface attributes with a description of each.

**Table 6** Interface attributes

Attribute	Description
ID	ID of the interface.
Type	Interface type. Value is one of the following: <ul style="list-style-type: none"> <li>• <code>iscsi</code> — Interface for iSCSI storage.</li> <li>• <code>replication</code> — Interface for replication-related data or management traffic.</li> </ul>
Port	ID of the physical port or link aggregation on an SP on which the interface is running. The ID includes the port name and SP name.
VLAN ID	Virtual local area network (VLAN) ID for the interface. The interface uses the ID to accept packets that have VLAN tags. The value range is 1-4095. <hr/> <p><b>Note</b></p> <p>If no VLAN ID is specified, which is the default, packets do not have VLAN tags. The Unisphere online help provides more details about VLANs.</p> <hr/>
IP address	IPv4 or IPv6 address.
Subnet mask	IPv4 subnet mask.
Gateway	IPv4 or IPv6 gateway.
MAC address	MAC address of the interface.
SP	SP that uses the interface.

**Table 6** Interface attributes (continued)

Attribute	Description
Health state	<p>A numerical value indicating the health of the system. Value is one of the following:</p> <ul style="list-style-type: none"> <li>Unknown (0)</li> <li>OK (5)</li> <li>OK BUT (7)</li> <li>Degraded/Warning (10)</li> <li>Minor failure (15)</li> <li>Major failure (20)</li> </ul>
Health details	Additional health information.

## Create interfaces

Create an interface.

### Format

```
/net/if create [ -async ] [-vlanId <value>] -type { iscsi |
replication} -port <value> -addr <value> [-netmask <value>] [-
gateway <value>]
```

### Action qualifier

Qualifier	Description
-async	Run the creation operation in asynchronous mode.
-type	<p>Specify the interface type. Value is one of the following:</p> <ul style="list-style-type: none"> <li><code>iscsi</code> — Interface for iSCSI storage.</li> <li><code>replication</code> — Interface for replication-related data or management traffic.</li> </ul>
-port	<p>Specify the ID of the SP port or link aggregation that will use the interface.</p> <hr/> <p><b>Note</b></p> <p>For systems with two SPs, a file interface is created on a pair of symmetric Ethernet ports rather than on a single specified port. Its current port is defined by NAS server SP and may differ from the specified port. For example, if the user specifies port <code>spa_eth2</code>, but the NAS server is on SP B, the interface is created on port <code>spb_eth2</code>.</p> <hr/>
-vlanId	Specify the virtual LAN (VLAN) ID for the interface. The interface uses the ID to accept packets that have VLAN tags. The value range is 1–4095.

Qualifier	Description
	<p><b>Note</b></p> <p>If no VLAN ID is specified, which is the default, packets do not have VLAN tags. The Unisphere online help provides more details about VLANs.</p>
-addr	Specify the IP address for the interface. The prefix length should be appended to the IPv6 address and, if omitted, will default to 64. For IPv4 addresses, the default length is 24. The IPv4 netmask may be specified in address attribute after slash.
-netmask	Specify the subnet mask for the interface.
	<p><b>Note</b></p> <p>This qualifier is not required if the prefix length is specified in the -addr attribute.</p>
-gateway	Specify the gateway for the interface.
	<p><b>Note</b></p> <p>This qualifier configures the default gateway for the specified port's SP.</p>

**Example**

The following command creates a replication interface. The interface receives the ID IF\_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/if create -type replication -port eth1_spb -addr 10.0.0.1 -netmask 255.255.255.0 -gateway 10.0.0.1
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = IF_1
Operation completed successfully.
```

## View interfaces

View a list of interfaces on the system. You can filter on the interface ID.

**Note**

[The show action command](#) on page 98 explains how to change the output format.

**Format**

```
/net/if [ {-id <value> | -port <value> | -type <value>} ] show
```

**Object qualifier**

Qualifier	Description
-id	Type the ID of an interface.

Qualifier	Description
-port	Type the port the interface is associated with.
-type	Specify the type of the interface. Valid values are: <ul style="list-style-type: none"> <li>• iscsi</li> <li>• replication</li> </ul>

### Example

The following command displays the details of all interfaces on the system.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/if show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                = if_0
      Type                = file
      NAS server          = nas_0
      Port                = eth0_spa
      VLAN ID            = 0
      IP address          = 3ffe:80c0:22c:4e:a:0:2:7f/64
      Subnet mask        =
      Gateway            = fe80::20a8bff:fe5a:967c
      IPv4 mode          =
      IPv4 address       =
      IPv4 subnet mask   =
      IPv4 gateway       =
      IPv6 mode          = static
      IPv6 address       = 3ffe:80c0:22c:4e:a:0:2:7f/64
      IPv6 link-local address =
      IPv6 gateway       = fe80::20a8bff:fe5a:967c
      MAC address        = EA:3E:22:3F:0C:62
      SP                 = spa
      Preferred          = yes

2:      ID                = if_1
      Type                = file
      NAS server          = nas_1
      Port                = eth1_spb
      VLAN ID            = 1
      IP address          = 192.168.1.2
      Subnet mask        = 255.255.255.0
      Gateway            = 192.168.1.254
      IPv4 mode          = static
      IPv4 address       = 192.168.1.2
      IPv4 subnet mask   = 255.255.255.0
      IPv4 gateway       = 192.168.1.254
      IPv6 mode          =
      IPv6 address       =
      IPv6 link-local address =
      IPv6 gateway       =
      MAC address        = EA:3E:22:21:7A:78
      SP                 = spa
      Preferred          = yes

3:      ID                = if_2
      Type                = replication
      NAS server          =
      Port                = eth1_spb
      VLAN ID            =
      IP address          = 10.103.75.56
      Subnet mask        = 255.255.248.0
      Gateway            = 10.103.72.1
```

```

IPv4 mode           = static
IPv4 address        = 10.103.75.56
IPv4 subnet mask    = 255.255.248.0
IPv4 gateway        = 10.103.72.1
IPv6 mode           =
IPv6 address        =
IPv6 gateway        =
MAC address         = EA:3E:22:6D:BA:40
SP                  = spb
Preferred           = no
    
```

## Change interface settings

Change the settings for an interface.

### Format

```

/net/if -id <value> set [-vlanId <value>] [-addr <value>] [-
netmask <value>] [-gateway <value>]
    
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the interface to change.

### Action qualifier

Qualifier	Description
-vlanId	<p>Type the virtual LAN (VLAN) ID for the interface. The interface uses the ID to accept packets that have VLAN tags. The value range is 1–4095.</p> <hr/> <p><b>Note</b></p> <p>If no VLAN ID is specified, which is the default, packets do not have VLAN tags. The Unisphere online help provides more details on VLANs.</p> <hr/>
-addr	<p>Specify the IP address for the interface.</p> <hr/> <p><b>Note</b></p> <p>The prefix length should be appended to the IPv6 address. The IPv4 netmask may be specified in address attribute after the slash.</p> <hr/>
-netmask	Specify the IPv4 subnet mask for the interface.
-gateway	<p>Specify the gateway for the interface.</p> <hr/> <p><b>Note</b></p> <p>The gateway is optional for both IPv4 and IPv6. This qualifier configures the default gateway for the specified port's SP.</p> <hr/>

### Example

The following command changes the gateway address for interface IF\_1:



```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456!/net/if -id IF_1 set
-gateway 2001:db8:0:170:a:0:2:70
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = IF_1
Operation completed successfully.
```

## Delete interfaces

Delete an interface.

### NOTICE

Deleting an interface can break the connection between systems that use it, such as configured hosts.

### Format

```
/net/if -id <value> delete
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the interface to delete.

### Example

The following command deletes interface IF\_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/if -id IF_1
delete
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

## Manage remote storage systems

Configure remote storage systems that connect to the system to which you are logged in. The system uses the configuration to access and communicate with the remote system. For example, to use remote replication, create a configuration that specifies the remote system to use as the destination for the replication session.

Each remote system configuration is identified by an ID.

The following table lists the attributes for remote storage systems:

**Table 7** Remote system attributes

Attribute	Description
ID	ID of the remote system.
Model	Model number of the remote system.

**Table 7** Remote system attributes (continued)

Attribute	Description
Serial number	Serial number of the remote system.
Address	Network name or management IP address of the remote system.
Alternate management address	An alternative management IP address of the remote system.
Health state	<p>Health state of the storage resource. The health state code appears in parentheses. Value is one of the following:</p> <ul style="list-style-type: none"> <li>• <code>OK (5)</code> —Resource is operating normally.</li> <li>• <code>Degraded/Warning (10)</code> —Working, but one or more of the following may have occurred: <ul style="list-style-type: none"> <li>▪ One or more of its storage pools are degraded.</li> <li>▪ Its replication session is degraded.</li> <li>▪ Its replication session has faulted.</li> <li>▪ It has almost reached full capacity. Increase the primary storage size, or create additional resources to store your data, to avoid data loss.</li> </ul> </li> <li>• <code>Minor failure (15)</code> —One or both of the following may have occurred: <ul style="list-style-type: none"> <li>▪ One or more of its storage pools have failed.</li> <li>▪ The associated iSCSI node has failed.</li> </ul> </li> <li>• <code>Major (20)</code> —One or both of the following may have occurred: <ul style="list-style-type: none"> <li>▪ Resource is unavailable.</li> <li>▪ One or more of the associated storage pools have failed.</li> </ul> </li> <li>• <code>Critical failure (25)</code> —One or more of the following may have occurred: <ul style="list-style-type: none"> <li>▪ One or more of its storage pools are unavailable.</li> <li>▪ Resource is unavailable.</li> <li>▪ Resource has reached full capacity. Increase the primary storage size, or create additional resources to store your data, to avoid data loss.</li> </ul> </li> <li>• <code>Non-recoverable error (30)</code> —One or both of the following may have occurred: <ul style="list-style-type: none"> <li>▪ Resource is unavailable.</li> <li>▪ One or more of the associated storage pools are unavailable.</li> </ul> </li> </ul>
Health details	Additional health information.
Source user name	For storage systems that are the source in a replication session, the username that is used to access the system.
Source user password	For storage systems that are the source in a replication session, the user password that is used to access the system.

**Table 7** Remote system attributes (continued)

Attribute	Description
Local interfaces	The list of local interface identifiers used to create the interconnection between the two systems.
Remote interfaces	The list of remote interface identifiers used to create the interconnection between two systems.
Destination user name	For storage systems that are the destination in a replication session, the username that is used to access the system.
Destination user password	For storage systems that are the destination in a replication session, the user password that is used to access the system.
Connection type	The type of connection with the remote system. Valid values are: <ul style="list-style-type: none"> <li>• sync</li> <li>• async</li> <li>• both</li> </ul>
Synchronous FC ports	The fibre channel ports enabled for synchronous replication.  <b>Note</b> For a local system (RS_0), this field will appear empty only when there are no FC ports. For remote systems, this will be empty when the connection type is asynchronous.

## Create remote system configurations

Configures a remote system configuration for the local system to access.

### Note

For a source VNX system with two control stations, the home directory of the sysadmin user, which is used in configuring the import connection, must exist on the primary control station of the VNX.

### Format

```
/remote/sys create -addr <value> [-type VNX] -srcUsername <value> {-srcPassword <value> | -srcPasswordSecure} -dstUsername <value> {-dstPassword <value> | -dstPasswordSecure} [-connectionType {sync | async | both}]
```

### Action qualifiers

Qualifier	Description
-addr	Specify the network name or IP address of the remote system.
-type	Specify the remote system type. Valid values are: <ul style="list-style-type: none"> <li>• VNX</li> </ul>
-srcUsername	For systems that are the source in a replication, type the username that is used to access the system.

Qualifier	Description
-srcPassword	For systems that are the source in a replication, type the user password that is used to access the system.
-srcPasswordSecure	Specify the password in secure mode. Once you run the command with this qualifier, you will be asked to type the password separately.
-dstUsername	For systems that are the destination in a replication session or VNX in an import session, specify the username that is used to access the system.
-dstPassword	For systems that are the destination in a replication session or VNX in an import session, specify the user password that is used to access the system.
-dstPasswordSecure	Specify the password in secure mode. Once you run the command with this qualifier, you will be asked to type the password separately.
-connectionType	Specify this qualifier to indicate the type of replication connection. Valid values are async, sync, or both.

### Example

The following command creates a remote system configuration with these settings:

- Network address is 10.64.75.10.
- Includes access credentials for when the system is the source or destination.

The configure remote system receives the ID RS\_65536:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys create -
addr 10.64.75.10 -type VNX -dstUsername admin1 -dstPassword
Password789!
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = RS_65536
Operation completed successfully.
```

## View settings for remote storage systems

View the configuration for a remote system on the local system. You can filter on the configuration ID of the remote system.

---

### Note

[The show action command](#) on page 98 explains how to change the output format.

---

### Format

```
/remote/sys [-id <value>] show
```

**Object qualifier**

Qualifier	Description
-id	Type the ID of a remote system configuration.

**Example**

The following command lists all configurations for remote storage systems:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                = RS_1
      Address            = 10.2.3.1
      Alternate Management Address =
      Model              = Unity 300
      Serial number      = FCNC987654321
      Connection type    = async
      Local interfaces   = N/A
      Remote interfaces  = N/A
      Operational status = OK (0x2)
      Health state       = OK (5)
      Health details     = "Communication with the
replication host is established. No action is required."
      Synchronous FC ports = spb_fc4, spa_fc4
```

**Verify settings for remote storage systems**

Verify the configuration settings for a remote system to ensure that the source storage resource can connect to the remote storage resource.

**Format**

```
/remote/sys -id <value> verify
```

**Object qualifier**

Qualifier	Description
-id	Type the ID of a remote system configuration to verify the settings.

**Example**

The following command verifies remote system configuration RS\_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys -id RS_1 verify
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

**Change settings for remote storage systems**

Changes the configuration settings for a remote system.

**NOTICE**

If a replication connection already exists and you plan to add a different mode of file replication, do not attempt to create a new connection. Change the existing replication connection mode to Both. Also, ensure that you have the appropriate interface types configured to support both asynchronous replication (eth2, eth3) and synchronous replication (sync replication mgmt port).

**Format**

```
/remote/sys -id <value> set [ -addr <value> ] [ -dstUsername <value> { -dstPassword <value> | -dstPasswordSecure } ] [ -connectionType {sync | async | both}]
```

**Object qualifier**

Qualifier	Description
-id	Type the ID of the remote system configuration to change.

**Action qualifiers**

Qualifier	Description
-addr	Type the network name or management IP address of the remote system.
-dstUsername	Type the username that is used to access the remote system.
-dstPassword	Type the user password that is used to access the remote system.
-dstPasswordSecure	Specify the password in secure mode - the user will be prompted to input the password.
-connectionType	Specify this qualifier to indicate the type of replication connection. Valid values are async, sync, or both.

**Example**

The following command changes the name, IP address, and access credentials for remote system configuration RS\_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys -id RS_1 set -addr "10.64.74.2" -dstUsername Local/joe -dstPassword Password456!
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = RS_1

Operation completed successfully.
```

## Delete remote system configurations

Deletes the configuration for a remote system.

**Note**

Before deleting a remote system configuration, ensure that all I/O operations on the system, such as active replication sessions, have completed to avoid data loss.

**Format**

```
/remote/sys -id <value> delete
```

**Object qualifier**

Qualifier	Description
-id	Type the ID of the remote system configuration to delete.

**Example**

The following command deletes remote system configuration RS\_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys -id RS_1 delete
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

## Cabinet level unplanned failover of replication sessions

Execute a failover of all NAS server synchronous replication sessions from the remote system to the local system (unplanned failover). Replication sessions of file systems created on the affected NAS servers will also fail over automatically.

**Format**

```
/remote/sys -id <value> failover [-force]
```

**Object qualifiers**

Qualifier	Description
-id	Type the ID of the remote system from which to failover its NAS server synchronous replication sessions.

**Action qualifiers**

Qualifier	Description
-force	Specifies whether to skip checking the network connection to the remote system. Required when the network connection is healthy. No values are allowed.

**Example**

The following command executes a cabinet level unplanned failover replication operation issued for a Unity system:

```
uemcli /remote/sys -id RS_1 failover
```

```
Storage system address: 10.0.0.1
Storage system port: 443
```

```
HTTPS connection
ID = RS_1
Operation completed successfully.
```

**Note**

After an unplanned failover, the NAS servers and file systems on the original source system must be updated to reflect the new status. If there is a large number of NAS servers and file systems, this change may take several minutes to complete. During this period, resume and failback operations of the synchronous replication sessions will not work. It is recommended to wait for all of the updates to complete before running a resume or failback operation. There is no impact to data access while this update is occurring.

## Manage replication sessions

Storage replication is a process in which storage data is duplicated either locally or to a remote network device. Replication produces a read-only, point-in-time copy of source storage data and periodically updates the copy, keeping it consistent with the source data. Storage replication provides an enhanced level of redundancy in case the main storage backup system fails. As a result:

- Downtime associated cost of a system failure is minimized.
- Recovery process from a natural or human-caused disaster is facilitated.

Each replication session is identified by an ID. The Unisphere online help provides more details about storage replication.

It is important to note that when replicating from a Unity system running a later OE version (for example, OE 4.1.x) to a Unity system running an older version (for example, OE 4.0.x), you cannot have new OE version features enabled on the source.

**Note**

At any given point in time, only one command is supported on a replication session. Before running a new command, ensure that you wait for the existing action on the replication session to complete.

The following table lists the attributes for replication sessions:

**Table 8** Replication session attributes

Attribute	Description
ID	ID of the session.
Name	Name of the replication session.
Session type	Storage type of the session. Valid values are: <ul style="list-style-type: none"> <li>• lun</li> <li>• block</li> <li>• file</li> <li>• nas server</li> </ul>



**Table 8** Replication session attributes (continued)

Attribute	Description
Synchronization type	<p>Type of synchronization. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>auto</code>—Data is replicated automatically in order to maintain the desired Recovery Point Objective (RPO).</li> <li>• <code>manual</code>—Data will only be replicated when manually initiated.</li> <li>• <code>sync</code>—Data is synchronously replicated with RPO=0.</li> </ul> <hr/> <p><b>Note</b></p> <p>For asynchronous replication, the potential for data loss increases as the RPO increases, as well as the amount of required protection space. Lowering the RPO will reduce the amount of potential data loss, but will also increase network traffic and may negatively impact performance. The default RPO is one hour.</p>
RPO	Recovery Point Objective (RPO) interval for automatic synchronization. For synchronous replication, the RPO is set to 0 automatically.
Resource type	<p>Type of storage resource to which the replication session is applied. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>LUN</code></li> <li>• <code>LUN group</code></li> <li>• <code>File System</code></li> <li>• <code>VMware VMFS</code></li> <li>• <code>VMware NFS</code></li> <li>• <code>NAS Server</code></li> </ul>
Sync State	<p>Additional state of the replication session, specific to the replication mode.</p> <ul style="list-style-type: none"> <li>• For asynchronous replication, valid values are: <ul style="list-style-type: none"> <li>▪ <code>idle</code>— No active syncing.</li> <li>▪ <code>manual</code>— User initiated syncing.</li> <li>▪ <code>auto syncing</code>— System initiated syncing.</li> </ul> </li> <li>• For synchronous replication, valid values are: <ul style="list-style-type: none"> <li>▪ <code>unknown</code>—Unknown sync state.</li> <li>▪ <code>out of sync</code>—Destination is out of sync with the source.</li> <li>▪ <code>in sync</code>—Destination is an exact copy of the source.</li> <li>▪ <code>consistent</code>—Destination is a point in time copy of the source.</li> <li>▪ <code>syncing</code>—System initiated syncing.</li> <li>▪ <code>inconsistent</code>—Destination is not a point in time copy of the source.</li> </ul> </li> </ul>

**Table 8** Replication session attributes (continued)

Attribute	Description
Health state	<p>Health state of the session. Valid values are:</p> <ul style="list-style-type: none"> <li>• Unknown—Session health cannot be determined.</li> <li>• OK—Session is operating normally.</li> <li>• Degraded/Warning—An error has caused one or more of the following: <ul style="list-style-type: none"> <li>▪ Session has been paused.</li> <li>▪ Session has failed over, likely due to the source storage resource becoming unavailable. The destination storage resource is now in a read/write state. Review the state of the source and check your network connections for any problems. Once the source is back online, you can fail back the session to return it to normal operation.</li> <li>▪ Session is syncing.</li> </ul> </li> <li>• Minor failure—Communication with the replication host has been lost. It is likely that the system is either powered down or there is a network connectivity issue between the systems. A change in the network configuration on either side could also interrupt communication.</li> <li>• Critical failure— Session has encountered an error that has halted the session.</li> </ul> <hr/> <p><b>Note</b></p> <p>If the replication session is in an error state, in addition to resolving the issue (for example, destination pool out of space), try pausing, and then resuming the replication session. If the problem persists, delete and then create the replication session again.</p>
Health details	Additional health information.
Operational status	<p>Operational status of the session. The operational status code appears in parentheses.</p> <ul style="list-style-type: none"> <li>• Unknown (0x0)</li> <li>• Non Recoverable Error (0x7)</li> <li>• Lost Communication (0xd)</li> <li>• Failed Over with Sync (0x8400)</li> <li>• Failed Over (0x8401)</li> <li>• Manual Syncing (0x8402)</li> <li>• Paused (0x8403)</li> <li>• Idle (0x8404)</li> <li>• Auto Sync Configured (0x8405)</li> <li>• Destination Extend Failed Not Syncing (0x840B)</li> <li>• Destination Extend In Progress (0x840C)</li> </ul>

**Table 8** Replication session attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> <li>• Active (0x840D)</li> <li>• Lost Sync Communication (0x840E)</li> <li>• Syncing (0x8411)</li> </ul>
Source status	<p>Status of the source storage resource in the session. Valid values are:</p> <ul style="list-style-type: none"> <li>• Unknown—Source status is unknown.</li> <li>• OK—Source is operating normally.</li> <li>• Paused—Replication session for the source is currently paused.</li> <li>• Fatal replication issue—Source has experienced a critical error and the replication session has stopped. Delete the replication session and re-create it.</li> <li>• Lost communication—Communication with the replication host has been lost. It is likely that the system is either powered down or there is a network connectivity issue between the systems. A change in the network configuration on either side could also interrupt communication.</li> <li>• Failed over—The replication session has failed over to the destination site. In a failed over state, the destination object is read/write. When communication is reestablished between the source and destination, the source is shown as Restricted Replication Access = Yes. To resume operations on the source site, the replication session needs to be failed back.</li> <li>• Switched over—The replication session has switched over to the source site. In a switched over state, the source object is read/write. When communication is reestablished between the source and destination, the destination is shown as Restricted Replication Access = Yes. To resume operations on the destination site, the replication session needs to be failed over.</li> </ul>
Destination status	<p>Status of the destination storage resource in the session. Valid values are:</p> <ul style="list-style-type: none"> <li>• Unknown—Status of the destination resource is unknown.</li> <li>• OK—Destination resource is operating normally.</li> <li>• Paused—Replication session for destination resource is currently paused.</li> <li>• Fatal replication issue—Destination has experienced a critical error and the replication session has stopped. Delete the replication session and re-create it.</li> <li>• Lost communication—Communication with the replication host has been lost. It is likely that the system is either powered down or there is a network connectivity issue between the systems. A change in the network configuration on either side could also interrupt communication.</li> <li>• Failed over—The replication session has failed over to the destination site. In a failed over state, the destination object is</li> </ul>

**Table 8** Replication session attributes (continued)

Attribute	Description
	<p>read/write. When communication is reestablished between the source and destination, the source is shown as Restricted Replication Access = Yes. To resume operations on the source site, the replication session needs to be failed back.</p> <ul style="list-style-type: none"> <li>• <code>Switched over</code>—The replication session has switched over to the source site. In a switched over state, the source object is read/write. When communication is reestablished between the source and destination, the destination is shown as Restricted Replication Access = Yes. To resume operations on the destination site, the replication session needs to be failed over.</li> </ul>
Network status	<p>Status of the network connection. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>Unknown</code>—Network status is currently unknown. If you continue to see this value, check the network connections.</li> <li>• <code>OK</code>—Network connection is operating normally.</li> <li>• <code>Lost Communication</code>—Communication with the replication host has been lost. It is likely that the system is either powered down or there is a network connectivity issue (lost IP) between the systems. A change in the network configuration on either side could also interrupt communication.</li> <li>• <code>Lost Sync Communication</code>—Fiber Channel communication with the synchronous replication remote system has been lost. It is likely that the Fiber Channel connection has encountered issues.</li> </ul>
Destination type	<p>Type of destination used in the session. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>local</code>—Maintain a full copy of the storage resource on the local system. This has advantages over snapshots in that a full copy, not just a copy of changes, is retained.</li> <li>• <code>remote</code>—Maintain a full copy of the storage resource on a remote system by transferring the data over the network. Remote replication is often used to ensure that a copy is available at a remote site in case of catastrophic data loss, for example, due to natural disaster at the local site.</li> </ul>
Destination system	<p>For remote sessions, the ID of the remote system on which the data is replicated.</p>
Local role	<p>The local system role. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>Unknown</code>—Status of the local system is unknown.</li> <li>• <code>Source</code>—Resource on the local system is replicated to the remote system.</li> <li>• <code>Destination</code>—Resource on the local system is the replication destination of the resource on the remote system.</li> <li>• <code>Loopback</code>—Resources participating in the replication session are located on the same storage system.</li> </ul>

**Table 8** Replication session attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> <li><code>Local</code>—Resources participating in the replication session are located on the different storage processors of the local system.</li> </ul>
Source resource	ID of the storage resource that is the source of the session. The source can be local or remote.
Source SP A interface	ID of the interface on the SPA of the source system for the replication.
Source SP B interface	ID of the interface on the SPB of the source system for the replication.
Destination resource	ID of the storage resource on which the data is replicated.
Destination SP A interface	ID of the interface on the SPA of the destination system for the replication.
Destination SP B interface	ID of the interface on the SPB of the destination system for the replication.
Time of last sync	Date and time of the last replication synchronization.
Sync status	<p>Percentage of the replication synchronization that has completed and the amount of time remaining.</p> <hr/> <p><b>Note</b></p> <p>For synchronous replication, the percentage is reported when the replication is in the Syncing state.</p> <hr/>
Sync transfer rate	<p>Synchronization transfer rate when the session is in the syncing state. For multi-LUN applications there is a comma-separated list of values.</p> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p> <hr/>
Sync transfer size remaining	<p>Remaining size to be transferred during synchronization. For multi-LUN applications there is a comma-separated list of values.</p> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p> <hr/>
Previous transfer rate	<p>Previous average transfer rate for the replication session.</p> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p> <hr/>
Average transfer rate	<p>Average transfer rate for the replication session.</p> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p> <hr/>

**Table 8** Replication session attributes (continued)

Attribute	Description
Element pairs	For consistency group and VMware VMFS datastore replications, the LUN element pairs within the replication.
Hourly snapshot keep for	<p>Amount of time to keep replicated hourly snapshots on the destination. Output can be:</p> <ul style="list-style-type: none"> <li>• Blank when scheduled snapshots are not replicated.</li> <li>• <code>&lt;value&gt;&lt;qualifier&gt;</code>—When a retention duration is specified, where: <ul style="list-style-type: none"> <li>▪ <code>value</code>—An integer value. If the <code>qualifier</code> is <code>h</code> (hours), the valid range is from 1 to 42840. If the <code>qualifier</code> is <code>d</code> (days), the valid range is from 1 to 1785.</li> <li>▪ <code>qualifier</code>—A value qualifier. The valid values are: <ul style="list-style-type: none"> <li>– <code>h</code> (hours)</li> <li>– <code>d</code> (days)</li> </ul> </li> </ul> </li> <li>• <code>Forever</code>—When <code>-keepFor</code> value is not specified and <code>allow auto-delete</code> is requested</li> <li>• <code>Same as source</code>—Keep the destination retention policy the same as the source retention policy</li> </ul> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p>
Hourly snapshot allow auto-delete	<p>Whether or not the destination pool's auto-delete policy allows replicated hourly snapshots on the destination to be deleted. Output can be:</p> <ul style="list-style-type: none"> <li>• Blank when scheduled snapshots are not replicated.</li> <li>• <code>Same as source</code>—Keep the destination retention policy the same as the source retention policy</li> <li>• <code>yes</code>—When <code>-allowAutoDelete</code> is set</li> <li>• <code>no</code>—When <code>-keepFor</code> is set</li> </ul> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p>
Daily snapshot keep for	<p>Amount of time to keep replicated daily snapshots on the destination. Output can be:</p> <ul style="list-style-type: none"> <li>• Blank when scheduled snapshots are not replicated.</li> <li>• <code>value</code>—An integer value. If the <code>qualifier</code> is <code>h</code> (hours), the valid range is from 1 to 42840. If the <code>qualifier</code> is <code>d</code> (days), the valid range is from 1 to 1785.</li> <li>• <code>qualifier</code>—A value qualifier. The valid values are: <ul style="list-style-type: none"> <li>▪ <code>h</code> (hours)</li> </ul> </li> </ul>

**Table 8** Replication session attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> <li>▪ <code>d (days)</code></li> <li>• <code>Same as source</code>—Keep the destination retention policy the same as the source retention policy</li> </ul> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p>
Daily snapshot allow auto-delete	<p>Whether or not the destination pool's auto-delete policy allows the replicated daily snapshots on the destination to be deleted. Output can be:</p> <ul style="list-style-type: none"> <li>• Blank when scheduled snapshots are not replicated.</li> <li>• <code>Same as source</code>—Keep the destination retention policy the same as the source retention policy</li> <li>• <code>yes</code>—When <code>-allowAutoDelete</code> is set</li> <li>• <code>no</code>—When <code>-keepFor</code> is set</li> </ul> <hr/> <p><b>Note</b></p> <p>This attribute is valid for asynchronous replications only.</p>
Allow Async Snap Replication (file system asynchronous replication sessions only)	<p>Indicates whether or not to allow snap replication in asynchronous replication sessions. Values are:</p> <ul style="list-style-type: none"> <li>• <code>yes</code>--Allow snap replication in asynchronous sessions.</li> <li>• <code>no</code>--Disallow snap replication in asynchronous sessions.</li> </ul>

## Create replication sessions

Create a replication session.

### Note

Snapshots that have been created and attached as well as read/write (share) snapshots (as opposed to read-only checkpoint snapshots) are not eligible for replication. Only unattached (read-only) snapshots are eligible for replication. For asynchronous replication, you can replicate existing snapshots and snapshots created from snapshot schedules. For synchronous file replication, you cannot replicate existing snapshots or snapshots created from snapshot schedules. You can only replicate those snapshots and snapshots created from snapshot schedules after you have established the synchronous replication session and it is Active.

### Note

On a NAS server protected by replication, you must create a replication session for each file system on it. Otherwise, file system related configurations like shares and exports may be lost after a NAS server replication session failover.

**Note**

To do snapshot replication from a source system running OE version 4.0, 4.1, 4.2, 4.3, or 4.4 to a destination system running OE version 5.x, requires upgrading the source system to OE version 4.5 first. Upgrading to OE version 4.5 is not required but recommended if you want to do LUN or file system replication from OE version 4.0, 4.1, 4.2, 4.3, or 4.4 to OE version 5.x without any snapshot replication.

**Prerequisites**

Before creating a replication session, complete the following configuration tasks:

- Create the storage resource that provides the replication source.
- For local replication, create a replication destination on a local system.
- For remote replication, create a replication connection to a remote system, and create a replication destination on that remote system.
- For asynchronous replication in a coexisting asynchronous and synchronous replication with one source resource topology, create the asynchronous replication destination NAS server with both the `-replDest` and the `-backupOnly` attributes set to `yes`. These attributes must be set to `yes` on the asynchronous replication destination NAS server when the source NAS server is synchronous replicated; otherwise, the asynchronous replication session cannot be created.

**Format**

```
/prot/rep/session create [-async] -srcRes <value> [-srcSPAInterface <value>] [-srcSPBInterface <value>] -dstType {local | remote -dstSys <value>} -dstRes <value> [-dstSPAInterface <value>] [-dstSPBInterface <value>] [-name <value>] [-elementPairs <value>] [-syncType {manual [-autoInitiate {yes | no}] | auto [-rpo <value>]}] [-replicateHourlySnaps {yes [{-keepSameAsSource | -keepFor <value> | -allowAutoDelete}] | no}] [-replicateDailySnaps {yes [{-keepSameAsSource | -keepFor <value> | -allowAutoDelete}] | no}] [-replicateExistingSnaps] [-allowAsyncSnapReplication {yes | no}]
```

**Action qualifiers**

Qualifier	Description
<code>-async</code>	Run an action in asynchronous mode.
<code>-srcRes</code>	Type the ID of the storage resource to use as the source.
<code>-srcSPAInterface</code>	Type the ID of the interface on the SPA of the source system for the replication.  <b>Note</b> This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
<code>-srcSPBInterface</code>	Type the ID of the interface on the SPB of the source system for the replication.



Qualifier	Description
	<p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
-dstType	<p>Specify the type of destination. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>local</code>—Maintain a full copy of the storage resource on the local system. This has advantages over snapshots in that a full copy, not just a copy of changes, is retained.</li> <li>• <code>remote</code>—Maintain a full copy of the storage resource on a remote system by transferring the data over the network.</li> </ul> <p>Remote replication is often used to ensure that a copy is available at a remote site in case of catastrophic data loss, for example, due to natural disaster at the local site.</p>
-dstSys	<p>For remote replication, type the ID of the destination system. <a href="#">View settings for remote storage systems</a> on page 68 explains how to view the IDs of the remote system configuration on the local system.</p>
-dstRes	<p>Type the ID of the destination storage resource.</p> <p><b>Note</b></p> <p>To get the proper ID in the case of remote replication, you should use a command that list resources on a local machine with the <code>-remSys</code> qualifier. For example:</p> <ul style="list-style-type: none"> <li>• <code>uemcli /stor/prov/sf/res -remSys &lt;value&gt; show</code></li> <li>• <code>uemcli /stor/prov/iscsi/res -remSys &lt;value&gt; show</code></li> <li>• <code>uemcli /stor/prov/vmware/nfs -remSys &lt;value&gt; show</code></li> </ul>
-dstSPAInterface	<p>Type the ID of the interface on the SPA of the destination system for the replication.</p>

Qualifier	Description
	<p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
<p><code>-dstSPBInterface</code></p>	<p>Type the ID of the interface on the SPB of the destination system for the replication.</p> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
<p><code>-syncType</code></p>	<p>Specify how the source and destination will synchronize. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>auto</code>—Data is replicated automatically in order to maintain the desired Recovery Point Objective (RPO).</li> <li>• <code>manual</code>—Data will only be replicated when manually initiated.</li> </ul> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p> <p>As the RPO increases, the potential for data loss also increases, as well as the amount of required protection space. Lowering the RPO will reduce the amount of potential data loss, but will also increase network traffic and may negatively impact performance. The default RPO is one hour.</p>
<p><code>-autoInitiate</code></p>	<p>Specify whether the system will perform the first replication synchronization automatically. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>yes</code></li> <li>• <code>no</code></li> </ul> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
<p><code>-rpo</code></p>	<p>Type the time interval for when the synchronization will run. Use the following format:</p>

Qualifier	Description
	<p>&lt;HH&gt; [:MM]</p> <p>where:</p> <ul style="list-style-type: none"> <li>• HH—Type the number of hours. Range is 00-24 hours (1 day).</li> <li>• MM—Type the number of minutes, in 5 minute increments. Range is 05 to 55.</li> </ul> <p>For synchronous replication, specify the value 0. Once set, the value cannot be reset from zero to non-zero or from non-zero to zero.</p>
-replicateHourlySnaps	<p>Specify whether or not to mark hourly scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> <li>• yes</li> <li>• no</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
-replicateDailySnaps	<p>Specify whether or not to mark daily scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> <li>• yes</li> <li>• no</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
-keepSameAsSource	<p>Indicates whether or not to use the same retention policy (expiration time and auto-delete) of the source for the destination. This option propagates changes made to the source retention policy to the destination retention policy (from that point forward for newly created scheduled snapshots, old snapshots are left as is). No values are allowed. This option is enabled by default if <code>-keepFor</code> or <code>-allowAutoDelete</code> are not set.</p> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
-keepFor	<p>Specifies the retention time after which the snapshot is deleted on the destination. The interval can be defined in days or hours. Use the following format:</p>

Qualifier	Description
	<p><code>&lt;value&gt;&lt;qualifier&gt;</code></p> <p>where:</p> <ul style="list-style-type: none"> <li>• <i>value</i>—An integer value. If the <i>qualifier</i> is h (hours), the valid range is from 1 to 42840. If the <i>qualifier</i> is d (days), the valid range is from 1 to 1785.</li> <li>• <i>qualifier</i>—A value qualifier. The valid values are: <ul style="list-style-type: none"> <li>▪ h (hours)</li> <li>▪ d (days)</li> </ul> </li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
<p><code>-allowAutoDelete</code></p>	<p>Specify whether auto delete is allowed on the replicated copy of this snapshot or snapshot set. Valid values:</p> <ul style="list-style-type: none"> <li>• yes</li> <li>• no</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
<p><code>-replicateExistingSnaps</code></p>	<p>Indicates whether or not to replicate snapshots already existing on the source resource. This is a one-time option available during session creation that will replicate snapshots existing on the source at that moment in time. All eligible snapshots are replicated and have the source retention policy applied for the destination retention policy. For a snapshot to be eligible for this option, it must meet these 3 criteria:</p> <ul style="list-style-type: none"> <li>• The snapshot is created by either the user or a snapshot schedule.</li> <li>• The snapshot is read-only (file resource snapshot must be a checkpoint snapshot; block resource snapshot must not be attached).</li> <li>• The snapshot is not currently undergoing deletion.</li> </ul>

Qualifier	Description
	<p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
- allowAsyncSnapReplication	<p>Indicates whether or not to allow snapshot replication in asynchronous session. Restriction: For a storage resource, snap replication is allowed only in one of its asynchronous replication sessions. Valid values are:</p> <ul style="list-style-type: none"> <li>• <b>yes</b></li> <li>• <b>no</b></li> </ul> <p><b>Note</b></p> <p>This qualifier is used for file system asynchronous replications only.</p>

### Example

The following command creates a replication session with these settings:

- Source storage resource is file system RS\_1.
- Destination system type is remote.
- Remote destination system is RS\_2.
- Remote storage resource is file system LUN\_2.
- Synchronization type is automatic.
- RPO is 2 hours and 30 minutes.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session
create -name REP1 -srcRes RS_1 -dstType remote -dstSys RS_2 -dstRes
LUN_2 -syncType auto -rpo 02h30m
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID = 81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000
Operation completed successfully.
```

## View replication sessions

View details about replication sessions. You can filter on the session ID.

### Note

The show action command explains how to change the output format.

### Format

```
/prot/rep/session [{-id <value> | -name <value> | -res
<value>}] show
```

### Object qualifiers

Qualifier	Description
-id	Type the ID of the replication session.
-name	Type the name of the replication session.
-res	Type the ID of a local storage resource on the system to view the sessions associated with it.

#### Example 1

The following command displays all replication sessions on the system:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1: ID =
42949672967_FNM00134400082_0000_42949672967_FNM00131800278_0000
Name = REP1
Session type = nas server
Synchronization type = auto
Resource type = NAS Server
Destination type = remote
```

#### Example 2

The following command displays all replication sessions on the system and their details:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1: ID =
42949672967_FNM00134400082_0000_42949672967_FNM00131800278_0000
Name = REP1
Session type = nas server
Synchronization type = auto
RPO =
Resource type = LUN
Sync State = idle
Health state = OK (5)
Health details = "This replication session
is operating normally. No action is required."
Operational status = Idle (0x8404)
Source status = OK
Destination status = OK
Network status = OK
Destination type = local
Destination system = local
Local role = Local
Source resource = sv_1
Source SP-A interface = N/A
Source SP-B interface = N/A
Destination resource = sv_2
Destination SP-A interface = N/A
```

```

Destination SP-B interface      = N/A
Time of last sync              = N/A
Sync status                    =
Sync transfer rate             = N/A
Sync transfer size remaining   = 0
Previous transfer rate         = N/A
Average transfer rate          = N/A
Element pairs                  = N/A
Hourly snapshot keep for       = 3h
Hourly snapshot allow auto-delete = no
Daily snapshot keep for        = same as source
Daily snapshot allow auto-delete = same as source

```

## Change replication session settings

Change the settings for a replication session.

### Format

```

/prot/rep/session {-id <value> | -name <value>} set [-async] [-
newName <value>] [-srcSPAInterface <value>] [-dstSPAInterface
<value>] [-srcSPBInterface <value>] [-dstSPBInterface <value>]
[-syncType {manual | auto -rpo <value>}] [-replicateHourlySnaps
{yes [{-keepSameAsSource | -keepFor <value> | -
allowAutoDelete}] | no}] [-replicateDailySnaps {yes [{-
keepSameAsSource | -keepFor <value> | -allowAutoDelete}] | no}]
[-allowAsyncSnapReplication {yes | no}]

```

### Object qualifiers

Qualifier	Description
-id	Type the ID of the replication session to change.
-name	Type the name of the replication session to change.

### Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.
-newName	Type the new name of the replication session.
-srcSPAInterface	Type the ID of the interface on the SPA of the source system for the replication.  <b>Note</b>  This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-srcSPBInterface	Type the ID of the interface on the SPB of the source system for the replication.

Qualifier	Description
	<p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
<p><code>-dstSPAInterface</code></p>	<p>Type the ID of the interface on the SPA of the destination system for the replication.</p> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
<p><code>-dstSPBInterface</code></p>	<p>Type the ID of the interface on the SPB of the destination system for the replication.</p> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
<p><code>-syncType</code></p>	<p>Specify how the source and destination will synchronize. Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>auto</code>—Data is replicated automatically in order to maintain the desired Recovery Point Objective (RPO).</li> <li>• <code>manual</code>—Data will only be replicated when manually initiated.</li> </ul> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p> <p>As the RPO increases, the potential for data loss also increases, as well as the amount of required protection space. Lowering the RPO will reduce the amount of potential data loss, but will also increase network traffic and may negatively impact performance. The default RPO is one hour.</p>
<p><code>-rpo</code></p>	<p>For automatic synchronization, type the time interval for when the synchronization will run. Use the following format:  <code>&lt;HH&gt; [ :MM]</code></p>



Qualifier	Description
	<p>where:</p> <ul style="list-style-type: none"> <li>• <b>HH</b>—Type the number of hours. Range is 00-24 hours (1 day).</li> <li>• <b>MM</b>—Type the number of minutes, in 5 minute increments. Range is 05 to 55.</li> </ul> <hr/> <p><b>Note</b></p> <p>For synchronous replication, specify the value 0. The value cannot be reset from zero to non-zero or from non-zero to zero.</p>
-replicateHourlySnaps	<p>Specify whether or not to mark hourly scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> <li>• yes</li> <li>• no</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
-replicateDailySnaps	<p>Specify whether or not to mark daily scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> <li>• yes</li> <li>• no</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
-keepSameAsSource	<p>Specify whether or not to use the same retention policy (expiration time and auto-delete) of the source for the destination. This option propagates changes made to the source retention policy to the destination retention policy (from that point forward for newly created scheduled snapshots, old snapshots are left as is). No values are allowed.</p> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
-keepFor	<p>Specify the retention time after which the snapshot is deleted on the destination. The interval can be defined in days or hours. Use the following format:</p>

Qualifier	Description
	<p><code>&lt;value&gt;&lt;qualifier&gt;</code></p> <p>where:</p> <ul style="list-style-type: none"> <li><i>value</i>—An integer value. If the <i>qualifier</i> is h (hours), the valid range is from 1 to 42840. If the <i>qualifier</i> is d (days), the valid range is from 1 to 1785.</li> <li><i>qualifier</i>—A value qualifier. The valid values are: <ul style="list-style-type: none"> <li>h (hours)</li> <li>d (days)</li> </ul> </li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
<p><code>-allowAutoDelete</code></p>	<hr/> <p><b>Note</b></p> <p>Only valid when <code>-replicateHourlySnaps yes</code> or <code>-replicateDailySnaps yes</code>.</p> <hr/> <p>Specify whether auto delete is allowed on the replicated copy of this snapshot or snapshot set. Valid values are:</p> <ul style="list-style-type: none"> <li>yes</li> <li>no</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>
<p><code>-allowAsyncSnapReplication</code></p>	<p>Indicates whether or not to allow snapshot replication in asynchronous session. Restriction: For a storage resource, snapshot replication is allowed only in one of its asynchronous replication sessions. Valid values are:</p> <ul style="list-style-type: none"> <li>yes</li> <li>no</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for file system asynchronous replications only and can only be modified when the session is paused.</p>

**Example**

The following command changes the source interface and destination interface for replication session

81604378625\_FCNC097274B3A\_0000\_81604378627\_FCNC097274B37\_0000:

```
uemcli /prot/rep/session -id
81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000 set -
srcSPAInterface if_1 -srcSPBInterface if_2 -dstSPAInterface if_3 -
dstSPBInterface if_4
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID = 81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000
Operation completed successfully.
```

**Pause replication sessions**

Pause a replication session.

**Format**

```
/prot/rep/session {-id <value> | -name <value>} pause [-async]
```

**Object qualifiers**

Qualifier	Description
-id	Type the ID of the replication session to be paused.
-name	Type the name of the replication session to be paused.

**Action qualifiers**

Qualifier	Description
-async	Run an action in asynchronous mode.

**Example**

The following command pauses replication session

81604378625\_FCNC097274B3A\_0000\_81604378627\_FCNC097274B37\_0000:

```
uemcli /prot/rep/session -id
81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000 pause
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
Operation completed successfully.
```

**Resume replication sessions**

Resumes an existing replication session.

**Format**

```
/prot/rep/session {-id <value> | -name <value>} resume [-async]
[-forceFullCopy] [-srcSPAInterface <value>] [-dstSPAInterface
```

```
<value>] [-srcSPBInterface value>] [-dstSPBInterface <value>]
[-forceSyncData]
```

### Object qualifiers

Qualifier	Description
-id	Type the ID of the replication session to be resumed.
-name	Type the name of the replication session to be resumed.

### Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.
-forceFullCopy	<p>Specifies to force a full synchronization during the resume operation.</p> <hr/> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>If there is no longer a common base for the source storage resource, a full synchronization is required. For such cases, ensure that you specify this qualifier. However, if the resume operation does find the common base, it uses the common base to do the delta synchronization rather than the full synchronization, regardless of this qualifier.</li> <li>This qualifier is used for asynchronous replications only.</li> </ul>
-srcSPAInterface	<p>Type the ID of the interface on the SPA of the source system for the replication.</p> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
-dstSPAInterface	<p>Type the ID of the interface on the SPA of the destination system for the replication.</p> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
-srcSPBInterface	Type the ID of the interface on the SPB of the source system for the replication.

Qualifier	Description
	<p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
<code>-dstSPBInterface</code>	<p>Type the ID of the interface on the SPB of the destination system for the replication.</p> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
<code>-forceSyncData</code>	<p>Specifies to force data transfer from the source system to the remote system even if the remote system has data not replicated from the source system.</p> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>

### Example

The following command resumes replication session

81604378625\_FCNC097274B3A\_0000\_81604378627\_FCNC097274B37\_0000:

```
uemcli /prot/rep/session -id
```

```
81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000 resume
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID = 81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000
Operation completed successfully.
```

## Manually synchronize replication sessions

Manually synchronize a replication session.

### Format

```
/prot/rep/session{-id <value> | -name <value>} sync [-async]
```

### Object qualifiers

Qualifier	Description
<code>-id</code>	Type the ID of the replication session to synchronize.
<code>-name</code>	Type the name of the replication session to synchronize.

### Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.

### Example

The following command initiates a manual resynchronization of replication session REPS\_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session -id REPS_1 sync
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

## Delete replication sessions

Delete a replication session. The deletion process automatically synchronizes the source storage resource with the destination storage resource, makes both read/write, and then deletes the session. You can then connect a host to either storage resource. Deleting the session from the source system automatically removes the destination and source replication sessions. This ensures that you do not have to manually delete the associated storage resources or NAS servers from the destination system.

### Note

Once you delete a replication session, data from the source storage resource will no longer be replicated on the destination, leaving the data unprotected. When deleting a file system synchronous replication session, though the session is deleted, if the initial synchronization does not complete, the destination file system will run into an unrecoverable error. In this case, delete the destination file system.

### Format

```
/prot/rep/session {-id <value> | -name <value>} delete [-async]
```

### Object qualifiers

Qualifier	Description
-id	Type the ID of the replication session to delete.
-name	Type the name of the replication session to delete.

### Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.

### Example

The following command deletes replication session 81604378625\_FCNC097274B3A\_0000\_81604378627\_FCNC097274B37\_0000:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session -
id 81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000
delete
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

## Fail over replication sessions

Run this command on the destination system to perform a failover of a replication session, with possible data loss, in response to an emergency scenario in which the source becomes unavailable.

After the failover, the destination system is read/write. To reestablish communication between the source and destination, fail back the session that has failed over. [Fail back replication sessions](#) on page 96 explains how to fail back a replication session that has failed over.

---

### Note

Failover operations terminate the transfer of data if there is a transfer in progress, causing a potential loss of data. If the source site is still available when you perform a failover, the system attempts to change the source storage resource from read/write to read-only.

---

### Initiate a planned downtime

To initiate a planned downtime, run this command on the source system by specifying the `-sync` option with the value `yes`. When you fail over a replication session from the source system, the destination system is fully synchronized with the source to ensure that there is no data loss. The destination storage resource can be used for providing access to the host.

### Format

```
/prot/rep/session {-id <value> | -name <value>} failover [-
async] [-sync {yes | no}] [-force]
```

### Object qualifiers

Qualifier	Description
<code>-id</code>	Type the ID of the replication session to fail over.
<code>-name</code>	Type the name of the replication session to fail over.

### Action qualifiers

Qualifier	Description
<code>-async</code>	Run an action in asynchronous mode.
<code>-sync</code>	For an asynchronous replication session, specifies whether a synchronization needs to be performed before failing over the replication session. For a synchronous replication session, specifies whether to keep synchronization on the reversed direction after failing over the session. Valid values are:

Qualifier	Description
	<ul style="list-style-type: none"> <li><code>yes</code>—For a planned failover. Can only be issued from the source system. Where <code>-sync</code> is not specified, this is the default value for a local replication session or session where <code>role=source</code>.</li> <li><code>no</code>—For an unplanned failover. Can only be issued from the destination system. Where <code>-sync</code> is not specified, this is the default value for a remote replication session or session where <code>role=destination</code>.</li> </ul> <hr/> <p><b>Note</b></p> <p>If the Network status=OK, the source system is probably OK. The command issued from the destination system without this option will fail. It is recommended to rerun the command using the <code>yes</code> option from the source system. However, in that case, the command issued from the destination system using the <code>no</code> option is still allowed.</p>
<code>-force</code>	Specifies whether to skip a pre-check operation on file systems of a NAS server when a replication failover operation is issued from the source NAS server. No values are allowed.

**Example**

The following command performs a fail over of replication session

81604378625\_FCNC097274B3A\_0000\_81604378627\_FCNC097274B37\_0000:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session -id 81604378625_FCNC097274B3A_0000_81604378627_FCNC097274B37_0000 failover
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

**Fail back replication sessions**

Fail back a replication session that has failed over. A failback results in the following:

- Synchronizes the destination and source storage resources.
- Makes the destination storage resource read-only.
- Makes the source storage resource read/write.

When the failback operation is complete, the replication session will resume and you may connect your hosts to the source storage resource.

**Note**

Ensure that hosts do not write to the destination storage resource, which will become read-only.

**Format**

```
/prot/rep/session {-id <value> | -name <value>} failback [-async] [-forceFullCopy] [-force] [-syncData {force | ignore}]
```



**Object qualifiers**

Qualifier	Description
-id	Type the ID of the replication session to fail back.
-name	Type the name of the replication session to fail back.

**Action qualifiers**

Qualifier	Description
-async	Run an action in asynchronous mode.
-forceFullCopy	<p>Specifies to force a full synchronization before the failback operation occurs.</p> <hr/> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>You may lose the common base on the source storage resource as a result of the event that caused the original failover. If there is no longer a common base for the source storage resource, a full synchronization is required. For such cases, ensure that you specify this qualifier. However, if the failback operation does find the common base, it uses the common base to do the delta synchronization rather than the full synchronization, regardless of this qualifier.</li> <li>This qualifier is used for asynchronous replications only.</li> </ul>
-force	Before failing back a NAS server synchronous replication session, it is checked whether its associated asynchronous file system replication sessions are all preserved when coexisting. When this qualifier is specified, that check is skipped.
-syncData	<p>Specifies the behavior of transferring data from the source system to the remote system. If the qualifier is not specified, the command will fail if the remote system has data not replicated from the source system. Valid values are:</p> <ul style="list-style-type: none"> <li><b>force</b> – Force data transfer from the source system to the remote system, even if the remote system has out-of-sync data.</li> <li><b>ignore</b> – Skip data transfer from the source system to the remote system. Just trigger the replication session reverse from the remote system to the source system.</li> </ul> <hr/> <p><b>Note</b></p> <p>This qualifier is used for asynchronous replications only.</p>

**Example**

The following command performs a fail back of replication session 81604378625\_FCNC097274B3A\_0000\_81604378627\_FCNC097274B37\_0000:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session -
id 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000
failback
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

## Preserve asynchronous replication sessions

Initiates a preserve asynchronous replication sessions operation on a NAS server synchronous replication session. After a NAS server synchronous replication session fails over or fails back with its file system synchronous replication sessions, the asynchronous replication sessions will be switched to the new production site by the preserve asynchronous replication sessions operation.

### Format

```
/prot/rep/session {-id <value> | -name <value>} preserveAsync
```

### Object qualifiers

Qualifier	Description
-id	Identifies the NAS server synchronous replication session.
-name	Identifies the NAS server synchronous replication session by name.

### Example

The following command preserves asynchronous replication sessions for 81604378625\_FCNCH097274B3A\_0000\_81604378627\_FCNCH097274B37\_0000:

```
uemcli /prot/rep/session -id
81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000
preserveAsync
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

## Associated CLI commands

### The show action command

The `show` action command displays a list of objects that exist on the system and the attributes of those objects. You can specify an object qualifier to view the attributes for a single object. The `show` action command provides qualifiers for changing the display of the output, including the format and the attributes to include. The available output formats are name-value pair (NVP), table, and comma-separated values (CSV).

**Format**

```
uemcli [<switches>] <object> [<object qualifier>] show [{-
detail | -brief | -filter <value>}] [-output {nvp | table [-
wrap] | csv}]
```

**Action qualifier**

Qualifier	Description
-output -o	Specify the output format. Value is one of the following: <ul style="list-style-type: none"> <li>nvp — The name-value pair (NVP) format displays output as name=value. <a href="#">Name-value pair format</a> on page 99 provides an example of the NVP format.</li> <li>table — The table format displays output as a table, with column headers and rows. By default, values that are too long to fit in a column are cut off. Add -wrap after the table qualifier, separated by a space, so that the values wrap. <a href="#">Table format</a> on page 99 provides an example of the table format.</li> <li>csv — The comma-separated values (CSV) format is similar to the table format, but the names and values are separated by commas. <a href="#">Comma-separated values format</a> on page 99 provides an example of the CSV format.</li> </ul>
-detail	Display all attributes.
-brief	Display only the basic attributes (default).
-filter	Comma-separated list of attributes which are included into the command output.

**Name-value pair format**

```
1:      ID           = la0_SPA
        SP           = SPA
        Ports        = eth0_SPA,eth1_SPA
        Health state = OK (5)

2:      ID           = la0_SPB
        SP           = SPB
        Ports        = eth0_SPB,eth1_SPB
        Health state = OK (5)
```

**Table format**

```
ID      | SP | Ports                | Health state
-----+---+-----+-----
la0_SPA | SPA | eth0_SPA,eth1_SPA | OK (5)
la0_SPB | SPB | eth0_SPB,eth1_SPB | OK (5)
```

**Comma-separated values format**

```
ID,SP,Ports,Health state
la0_SPA,SPA,"eth0_SPA,eth1_SPA",OK (5)
la0_SPB,SPB,"eth0_SPB,eth1_SPB",OK (5)
```

### Example

The following command modifies the set of attributes in the show action output. For example, if you add `-filter "ID,ID,ID,ID"` to the command, in the output you will see four lines with the "ID" attribute for each listed instance:

```
1:  ID = la_0
    ID = la_0
    ID = la_0
    ID = la_0
```

```
uemcli /net/nas/server show -filter "ID, SP, Health state, ID, Name"
```

### Filter format

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:  ID           = nas_1
    SP           = SPA_
    Health state = OK (5)
    ID           = nas_1
    Name        = Mynas1

2:  ID           = nas_2
    SP           = SPA_
    Health state = OK (5)
    ID           = nas_2
    Name        = Mynas2
```

## Manage Ethernet ports

View and change the settings for the network ports on each SP.

The following table describes the port attributes.

**Table 9** Network port attributes

Attribute	Description
ID	ID of the port.
Name	Name of the port.
SP	Name of the SP on which the port resides. Value is SPA or SPB.
Protocols	Types of protocols the port supports. Value is one of the following: <ul style="list-style-type: none"> <li>• <code>mgmt</code> — Management interface.</li> <li>• <code>file</code> — Network interface for Windows (SMB) and Linux/UNIX (NFS) storage.</li> <li>• <code>iscsi</code> — iSCSI interface for iSCSI storage.</li> </ul> <a href="#">Manage network interfaces</a> on page 60 explains how to configure network interfaces on the system.
MTU size	Maximum transmission unit (MTU) packet size (in bytes) that the port can transmit. Default is 1500 bytes per packet.
Requested MTU size	MTU size set by the user.

**Table 9** Network port attributes (continued)

Attribute	Description
Available MTU size	<p>List of available MTU sizes.</p> <hr/> <p><b>Note</b></p> <p>This can display as either a comma-separated list of exact values (if there is an iSCSI interface on the port), or an interval defined by the minimum or maximum values, such as 1280-9216.</p> <hr/>
Speed	Current link speed of the port.
Requested speed	Link speed set by the user.
Available speeds	List of available speed values.
Health state	<p>Health state of the port. The health state code appears in parentheses. Value is one of the following:</p> <ul style="list-style-type: none"> <li>• Unknown (0) — Status is unknown.</li> <li>• OK (5) — Port is operating normally.</li> <li>• OK BUT (7) — Lost communication, but the port is not in use.</li> <li>• Minor failure (15) — Lost communication. Check the network connection and connected cables.</li> <li>• Major failure (20) — Port has failed. Replace the SP that contains the port.</li> </ul>
Health details	Additional health information. See Appendix A, Reference, for health information details.
Aggregated port ID	If the port is in a link aggregation, the ID of the link aggregation appears. <a href="#">Manage link aggregations</a> on page 106 explains how to configure link aggregations on the SP ports.
Connector type	<p>Physical connector type. Valid values are:</p> <ul style="list-style-type: none"> <li>• unknown</li> <li>• RJ45</li> <li>• LC</li> <li>• MiniSAS_HD</li> <li>• CopperPigtail</li> <li>• NoSeparableConnector</li> </ul>
MAC address	Unique identifier assigned to a network device for communications on a network segment.
SFP supported speeds	List of supported speed values of the inserted Small Form-factor Pluggable.
SFP supported protocols	<p>List of supported protocols of the inserted Small Form-factor Pluggable. Valid values are:</p> <ul style="list-style-type: none"> <li>• unknown</li> <li>• FibreChannel</li> </ul>

**Table 9** Network port attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> <li>Ethernet</li> <li>SAS</li> </ul>

## View Ethernet port settings

View details about the network ports. You can filter on the port ID.

### Note

[The show action command](#) on page 98 explains how to change the output format.

### Format

```
/net/port/eth [-id <value>] show
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the port.

### Example

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/port/eth show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:  ID          = spa_eth2
    Name       = SP A Ethernet Port 2
    SP        = spa
    Protocols  = file, net, iscsi
    MTU size   = 4500
    Requested MTU size = 4500
    Available MTU sizes = 1280-9216
    Linux device name = eth2
    Speed     = 1 Gbps
    Requested speed = auto
    Available speeds = 1 Gbps, 10 Gbps, 100 Mbps, auto
    Health state = OK (5)
    Health details = "The port is operating normally."
    Aggregated port ID = None
    FSN port ID = None
    Connector type = RJ45
    MAC address = 00:60:16:7A:7F:CF
    SFP supported speeds =
    SFP supported protocols =

2:  ID          = spa_eth3
    Name       = SP A Ethernet Port 3
    SP        = spa
    Protocols  = file, net, iscsi
    MTU size   = 1500
    Requested MTU size = 1500
    Available MTU sizes = 1500, 9000
    Linux device name = eth3
    Speed     = 1 Gbps
```

```

Requested speed           = auto
Available speeds         = 1 Gbps, 10 Gbps, 100 Mbps, auto
Health state             = OK (5)
Health details           = "The port is operating normally."
Aggregated port ID      = None
FSN port ID              = None
Connector type           = RJ45
MAC address              = 00:60:16:7A:7F:CE
SFP supported speeds     =
SFP supported protocols =

```

## Change Ethernet port settings

### Note

The new settings are applied to a pair of symmetrical ports on dual SP systems.

Change the maximum transmission unit size and port speed for an Ethernet port.

### Format

```
/net/port/eth -id <value> set [-mtuSize <value>] [-speed <value>]
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the network port.

### Action qualifier

Qualifier	Description
-mtuSize	<p>Type the maximum transmission unit packet size (in bytes) for the port:</p> <ul style="list-style-type: none"> <li>If an Ethernet port carries File interfaces only, the MTU size can be set to a custom value between 1280 and 9216.</li> <li>If an Ethernet port carries iSCSI interfaces, the allowed MTU sizes are 1500 and 9000.</li> </ul> <p>Specific I/O modules may also restrict allowed range for MTU size value. The MTU size values of 1500 bytes (default) and 9000 bytes (jumbo frame) are supported by all interfaces and I/O modules.</p>
-speed	Type the port speed.

### Example

The following command sets the MTU size for Ethernet port 0 (eth0) on SP A to 9000 bytes:

```
uemcli /net/port/eth -id spa_eth0 set -mtuSize 9000
```

```

Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = spa_eth0

```

```
ID = spb_eth0
Operation completed successfully.
```

## Manage FC ports

View and change the settings for the FC ports on each SP.

The following table describes the port attributes.

**Table 10** FC port attributes

Attribute	Description
ID	ID of the port.
Name	Name of the port.
SP	Name of the SP on which the port resides.
WWN	World Wide Name (WWN) of the port.
Speed	Current link speed of the port.
Requested speed	Link speed set by the user.
Available speed	List of available speed values.
Health state	Health state of the port. The health state code appears in parentheses. Value is one of the following: <ul style="list-style-type: none"> <li>Unknown (0) — Status is unknown.</li> <li>OK (5) — Port is operating normally.</li> <li>OK BUT (7) — Lost communication, but the port is not in use.</li> <li>Minor failure (15) — Lost communication. Check the network connection and connected cables.</li> <li>Major failure (20) — Port has failed. Replace the SP that contains the port.</li> </ul>
Health details	Additional health information. See Appendix A, Reference, for health information details.
Connector type	Physical connector type. Valid values are: <ul style="list-style-type: none"> <li>unknown</li> <li>RJ45</li> <li>LC</li> <li>MiniSAS_HD</li> <li>CopperPigtail</li> <li>NoSeparableConnector</li> </ul>
SFP supported speeds	List of supported speed values of the inserted Small Form-factor Pluggable.
SFP supported protocols	List of supported protocols of the inserted Small Form-factor Pluggable. Valid values are:



**Table 10** FC port attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> <li>unknown</li> <li>FibreChannel</li> <li>Ethernet</li> <li>SAS</li> </ul>
Replication capability	<b>Type of replication capability. Valid values are:</b> <ul style="list-style-type: none"> <li>Sync replication</li> <li>RecoverPoint</li> </ul>

## View FC port settings

View details about the FC ports. You can filter on the port ID.

### Format

```
/net/port/fc [-id <value>] show
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the port.

### Example

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/port/fc show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                = spa_fc4
      Name              = SP_A FC Port 4
      SP                = spa
      WWN               = 50:06:BD:01:60:05:8E:
50:06:01:64:3D:E0:05:8E
      Speed              = 1 Gbps
      Requested speed   = auto
      Available speeds  = 4 Gbps, 8 Gbps, 16 Gbps, auto
      Health state      = OK (5)
      Health details    = "The port is operating normally."
      SFP supported speeds = 4 Gbps, 8 Gbps, 16 Gbps
      SFP supported protocols = FibreChannel
      Replication capability = Sync replication
      SFP supported mode = Multimode
```

## Change port settings

Change the speed for an FC port.

### Format

```
/net/port/fc -id <value> set -speed <value>
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the FC port.

### Action qualifier

Qualifier	Description
-speed	Type the port speed.

### Example

The following command sets the speed for FC port fc1 on SP A to 1 Gbps:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/port/fc -id
spa_fc1 set -speed 1Gbps
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = spa_fc1
Operation completed successfully.
```

## Manage link aggregations

Link aggregation lets you link physical ports (for example, port 0 and port 1) on a SP to a single logical port and therefore lets you use up to four Ethernet ports on the SP. If your system has two SPs, and you link two physical ports, the same ports on both SPs are linked for redundancy. For example, if you link port 0 and port 1, the system creates a link aggregation for these ports on SP A and a link aggregation on SP B.

Each link aggregation is identified by an ID.

---

#### Note

The cabling on SP A must be identical to the cabling on SP B, or you cannot configure link aggregation.

Link aggregation has the following advantages:

- Increases overall throughput since two physical ports are linked into one logical port.
- Provides basic load balancing across linked ports since the network traffic is distributed across multiple physical ports.
- Provides redundant ports so that if one port in a linked pair fails, the system does not lose connectivity.

---

#### Note

With link aggregation, both linked ports must be connected to the same switch and the switch must be configured to use link aggregation that uses the Link Aggregation Control Protocol (LACP). The documentation that came with your switch should provide more information on using LACP.

The following table describes the attributes for link aggregation.

**Table 11** Link aggregation attributes

Attribute	Description
ID	ID of the link aggregation. The ID is a combination of the link ID and the SP that contains the linked ports.
Ports	IDs of the linked physical ports. The port names include the name of the SP that contains the ports.
SP	Name of the SP on which the ports are linked. Valid values are: <ul style="list-style-type: none"> <li>• SPA</li> <li>• SPB</li> </ul>
MTU size	Maximum transmission unit (MTU) packet size (in bytes) for the linked ports. Default is 1500 bytes per packet.
Linux device name	Linux network device name.
FSN port ID	ID of the FSN port to which the link aggregation belongs, if it is part of an FSN.
Available MTU size	List of available MTU sizes.  <b>Note</b> This displays as an interval defined by the minimum and maximum values, for example: 1280-9216.
Health state	Health state of the link aggregation. The health state code appears in parentheses. Value is one of the following: <ul style="list-style-type: none"> <li>• Unknown (0) — Status is unknown.</li> <li>• OK (5) — Working correctly.</li> <li>• OK BUT (7) — Lost connection, but the link aggregation is not in use.</li> <li>• Degraded/Warning (10) — Working and performing all functions, but the performance may not be optimum.</li> <li>• Minor failure (15) — Working and performing all functions, but overall performance is degraded. This condition has a minor impact on the system and should be remedied at some point, but does not need to be fixed immediately.</li> <li>• Major failure (20) — Failing and some or all functions may be degraded or not working. This condition has a significant impact on the system and should be remedied immediately.</li> </ul>

**Table 11** Link aggregation attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> <li>Critical failure (25) — Failed and recovery may not be possible. This condition has resulted in data loss and should be remedied immediately.</li> <li>Non-recoverable error (30) — Completely failed and cannot be recovered.</li> </ul>
Health details	Additional health information.

## View link aggregations

View details about link aggregations. You can filter on the link aggregation ID.

### Format

```
/net/la [-id <value>] show
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the link aggregation.

### Example

The following command shows the link aggregations on the system, in this case, for both SPA and SPB:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/la show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:  ID              = spa_la_0_2
   SP              = spa
   Ports           = spa_iom_0_eth2, spa_iom_0_eth3
   FSN port ID    = None
   MTU size        = 3456
   Available MTU sizes = 1280-9216
   Linux device name = bond12
   Health state    = OK (5)
   Health details  = "The component is operating normally.
No action is required."
   Operational status =

2:  ID              = spb_la_0_2
   SP              = spb
   Ports           = spb_iom_0_eth2, spb_iom_0_eth3
   FSN port ID    = None
   MTU size        = 3456
   Available MTU sizes = 1280-9216
   Linux device name = bond12
   Health state    = OK (5)
   Health details  = "The component is operating normally.
```

```
No action is required."
Operational status =
```

## Create link aggregations

Create a link aggregation by linking two physical ports on an SP to create a logical port.

### Format

```
/net/la create -ports <value> [-mtuSize <value>]
```

### Action qualifier

Qualifier	Description
-ports	Type the IDs of the physical ports to link on the SP. Separate the IDs with a comma. For example, to link ports 0 and 1 on SPA, type: eth0_SPA,eth1_SPA.
-mtuSize	Type the MTU size (in bytes) for the linked ports. The MTU size can be set to a custom value between 1280 and 9216. Specific I/O modules may restrict allowed range for MTU size value. The MTU size values of 1500 bytes (default) and 9000 bytes (jumbo frame) are supported by all interfaces and I/O modules.

### Example

The following command links port 0 and port 1 on SPA with the default MTU size. The system has two SPs, so port 0 and port 1 on SPB are also linked, which results in two link aggregation IDs:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/la create -
ports "eth0_SPA,eth1_SPA"
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID = la0_SPA
ID = la0_SPB
Operation completed successfully.
```

## Change link aggregations

Change the settings of a link aggregation.

### Format

```
/net/la -id <value> set [-ports <value>] [-mtuSize <value>]
```

### Object qualifier

Qualifier	Description
-id	Type the ID of the link aggregation to change.

**Action qualifier**

Qualifier	Description
-ports	Type the IDs of the physical ports to link on the SP. Separate the IDs with a comma. For example, to link ports 0 and 1 on SPA, type: eth0_SPA,eth1_SPA
-mtuSize	Type the MTU size (in bytes) for the linked ports. The MTU size can be set to a custom value between 1280 and 9216. Specific I/O modules may restrict allowed range for MTU size value. The MTU size values of 1500 bytes (default) and 9000 bytes (jumbo frame) are supported by all interfaces and I/O modules.

**Example**

The following command changes the MTU size for link aggregation la0\_SPA to 9000 bytes. The system has two SPs, so MTU size is updated for both link aggregation IDs:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/la -id la0_SPA
set -mtuSize 9000
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = la0_SPA
ID = la0_SPB
Operation completed successfully.
```

**Delete link aggregations**

Delete a link aggregation.

**Format**

```
/net/la [-id <value>] delete
```

**Object qualifier**

Qualifier	Description
-id	Type the ID of the link aggregation to delete.

**Example**

The following command deletes link aggregation la0\_SPA. The system has two SPs, so link aggregation la0\_SPB is also deleted:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/la -id la0_SPA
delete
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = la0_SPA
ID = la0_SPB
Operation completed successfully.
```

# APPENDIX A

## Use case for CLI

This appendix addresses the following topics:

- [Replication configuration use case](#)..... 112

## Replication configuration use case

This section describes the use cases for configuring replication for block or file storage resources.

### Before you begin

Before you proceed with configuring replication, ensure that you complete the following:

- Create identical storage resources on the source and destination systems.
- Configure replication interfaces for each SP on the source and destination systems.
- On the destination system, the relevant storage resources and NAS servers are individually created with the `-replDest` attribute set to `yes`.
- For file replication, ensure the following:
  - Start with creating identical NAS servers on both the systems, and then create identical file systems.
  - Configure the NAS server on the destination system with a name other than the NAS server name on the source system.
  - Configure file systems on the destination system with the same name as the file systems on the source system.

## Configure local replication

---

### Note

Local replication is applicable to asynchronous replication only. Local replication is not applicable to synchronous replication.

Replication interfaces and connections do not apply to local replication. When using the CLI or the REST API, once you create the identical source and destination storage resources or NAS servers on the storage system, you can proceed to configure a replication session. When using Unisphere, you only need to create the source storage resources or NAS servers on the storage system. Unisphere does not allow you to create a session with an existing destination. A `DR_` is concatenated onto the resource name for local destinations to ensure that the source and destination names on the same system are unique (that is, LUN names need to be unique).

Note the following:

- For a disaster recovery scenario, it is recommended that the destination storage resource and NAS server are configured on a storage pool other than the pool used for the source storage resource and NAS server.
- For a migration scenario, which means migrating the source storage resource and NAS server to a destination storage resource and NAS server on the same pool, use the CLI to configure local replication. The Unisphere GUI does not allow local replication between storage resources and NAS servers on the same pool.

## Configure asynchronous replication

### Before you begin

If you are configuring asynchronous replication for a tenant, create a pool for the tenant on the destination system that matches the corresponding pool on the source



system (if one exists). Then add the tenant to the destination system, using the same UUID and VLANs as the tenant on the source.

If you are configuring asynchronous replication in a coexisting synchronous and asynchronous replication topology, create the asynchronous replication destination NAS server with both the `-replDest` and the `-backupOnly` attributes set to `yes`. These attributes must be set to `yes` on the asynchronous replication destination NAS server when the source NAS server is synchronous replicated; otherwise, the asynchronous replication session cannot be created.

### Procedure

1. Configure the replication interfaces on each SP of the source and destination systems.
2. Configure a replication connection using the **Asynchronous** connection mode.
3. For file storage, create a replication session for the NAS server associated with the file storage.

---

#### Note

Storage resources included in a NAS server automatically get replicated when a replication session is first configured for the NAS server. The replication session for the storage resources will inherit the same attributes as the associated replication session of the associated NAS server. For the storage resources you do not want participating in replication, you can choose to remove the associated replication sessions manually.

---

4. Create a remote replication session for the storage resource.

---

#### Note

You only need to configure replication interfaces and connections for the first replication session between two systems. The same connection can be used again for subsequent replication sessions between the same systems.

---

## Configure synchronous replication

---

### Note

Only remote replication is supported for synchronous replication.

---

### Procedure

1. Identify the Synchronous Replication Fibre Channel (FC) ports on each system. To determine the FC port used for synchronous replication, in the CLI console, run the command `/remote/sys show -detail`. Port information, similar to the following example, will appear in the output:

```
Synchronous FC ports = spb_fc4, spa_fc4
```

For more information, see the *Unisphere CLI User Guide*.

2. Zone the Synchronous Replication FC ports between the systems.

If the source and destination systems are co-located, instead of zoning, you can choose to use direct-connected FC cables between the SPs.

3. Configure the replication interfaces on each SP of the source and destination systems based on the connection mode you want:
  - For synchronous replication support (**Synchronous** connection mode), use the Synchronous Replication Management Ports on each SP of both the systems.
  - For asynchronous and synchronous replication support (**Both** connection mode), in addition to the replication interfaces for Synchronous Replication Management Ports, configure additional interfaces using the Ethernet Ports on each SP of the source and destination systems.
4. Configure the replication connection between source and destination systems from the source system only.
  - For synchronous replication support, specify the **Synchronous** connection mode.
  - For asynchronous and synchronous replication support, specify the **Both** connection mode.
5. Create the synchronous replication session.

---

#### Note

You only need to configure replication interfaces and connections for the first replication session between two systems. The same connection can be used again for subsequent replication sessions between the same systems.

---

## Create a replication interface

### Before you begin

Protection and mobility (import) interfaces can be shared between replication and import. For import, only VDM imports require interfaces. Block imports do not require interfaces.

Protection and mobility (import) interfaces are configured to support VDM imports and must be created prior to creating an import connection. A mobility interface IP address is assigned to SPA and SPB on the target Unity system. Once the mobility interface is configured, you can create the import connection between the Unity system and the VNX system. Mobility interfaces are not used for block import sessions.

Ensure the following:

- The interface port is cabled and connected to a network switch.
- Both SPs are up and running.

Obtain the following information for each Storage Processor (SP):

- IP address associated with the interface (replication or import). Although you can specify an IPv4 or IPv6-based address, ensure that you specify the same type of address for both SPs.
- IP address mask or prefix length that identifies the associated subnet.
- Gateway IP address associated with the interface.
- If applicable, the VLAN ID (between 1 and 4095) you want to associate the interface with.

**Note**

For the network to continue functioning properly, ensure that you set the VLAN ID only when you have configured the network switch port to support VLAN tagging of multiple VLAN IDs.

Ensure that you create replication interfaces on each SP.

**Procedure**

1. Run the following command to create the interface on SP A:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/if create
-type replication -port eth1_spa -addr 10.0.1.1 -netmask
255.255.255.0 -gateway 10.0.1.0
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = IF_1
Operation completed successfully.
```

2. Run the following command to create the interface on SP B:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/if create
-type replication -port eth1_spb -addr 10.0.1.2 -netmask
255.255.255.0 -gateway 10.0.1.0
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = IF_2
Operation completed successfully.
```

**View interfaces**

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /net/if show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                = IF_0
      Type              = iscsi
      Port              = eth0_spa
      VLAN ID           = 0
      IP address        = 3ffe:80c0:22c:4e:a:0:2:7f/64
      Subnet mask       =
      Gateway           = fe80::20a8bff:fe5a:967c
      SP                = spa

2:      ID                = IF_1
      Type              = replication
      Port              = eth1_spa
      VLAN ID           = 1
      IP address        = 10.0.1.1
      Subnet mask       = 255.255.255.0
      Gateway           = 10.0.1.0
```

```

3:      SP                               = spa
      ID                                 = IF_2
      Type                               = replication
      Port                               = eth1_spb
      VLAN ID                            =
      IP address                         = 10.0.1.2
      Subnet mask                        = 255.255.248.0
      Gateway                            = 10.0.1.0
      SP                                 = spb

```

## Create a replication connection

### Before you begin

Ensure that you have set up relevant replication interface pairs, one on each SP, on the source and destination systems. Obtain the following information:

- For remote replication, the IP address and associated user authentication credentials to connect to the remote system.
- For local replication, the password associated with your user account.
- The connection mode you want to use for the replication: Asynchronous, Synchronous, or Both.

#### NOTICE

If a replication connection already exists and you plan to add a different mode of file replication, do not attempt to create a new connection. Change the existing replication connection mode to Both. Also, ensure that you have the appropriate interface types configured to support both asynchronous replication (eth2, eth3) and synchronous replication (sync replication mgmt port).

Consider that you want to create an asynchronous replication connection to the remote system with the IP address 10.1.1.1.

### Procedure

- Run the following command:

```

uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys
create -addr 10.1.1.1 -srcUsername admin1 -srcPassword Password456!
-dstUsername admin2 -dstPassword Password986! -connectionType async

```

```

Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = RS_1
Operation completed successfully.

```

## View settings for remote storage systems

```

uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys show -
detail

```

```

Storage system address: 10.0.0.1
Storage system port: 443

```

```
HTTPS connection
```

```
1:      ID                = RS_1
      Name                = MyTargetSystem
      Address              = 10.1.1.1
      Model                = Unity 300
      Serial number        = FCNCH01234567A90
      Connection type      = async
      Source interfaces    = N/A
      Local interfaces     = N/A
      Remote interfaces    = N/A
      Operational status   = OK (0x2)
      Health state         = OK (5)
      Health details       = "Communication with the replication
                           host is established. No action is
                           required."
      Synchronous FC ports = spb_fc4, spa_fc4
```

## Create a replication session for block storage

### Before you begin

Complete the following:

- For remote replication:
  - Identify the remote system that will act as the replication destination.
  - Create relevant replication interfaces, replication connection, and a storage resource on the remote system that will act as the destination.
- For local replication, create a storage resource that will act as the destination.
- Determine the replication synchronization mode you want. You can specify asynchronous (async), synchronous (sync), or manual synchronization (manual).
- For asynchronous replication, determine the Recovery Point Objective (RPO) for the replication session.

### Procedure

- Run the following command to create a synchronous replication session between the LUN "LUN\_1" on the source system and the LUN "LUN\_2" located on the remote system "RS\_2":

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session
create -name REP1 -srcRes LUN_1 -dstType remote -dstSys RS_2 -
dstRes LUN_2 -syncType auto -rpo 0
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID =
81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000
Operation completed successfully.
```

## Create an asynchronous replication session for file storage

### Before you begin

Complete the following:

- For remote replication:
  - Identify the remote system that will act as the replication destination.
  - Create relevant replication interfaces, replication connection, and a storage resource on the remote system that will act as the destination. The storage resource on the destination system must have the same size
- For local replication, create a storage resource that will act as the destination.
- For file replication, create a replication session on the NAS server associated with the file storage.
- For asynchronous replication, determine the Recovery Point Objective (RPO) for the replication session.

#### NOTICE

If you are configuring asynchronous replication in a coexisting synchronous and asynchronous replication with one source resource topology, create the asynchronous replication destination NAS server with both the `-replDest` and the `-backupOnly` attributes set to `yes`. These attributes must be set to `yes` on the asynchronous replication destination NAS server when the source NAS server is synchronous replicated; otherwise, the asynchronous replication session cannot be created.

Configure an asynchronous replication session between the NAS servers associated with the file storage, with an RPO set to 2 hours 30 minutes and automatic synchronization. On the source system, the file system "res\_7" is associated with NAS server "nas\_1". And, the file system "res\_8" is associated with NAS server "nas\_2" on the remote system.

#### Procedure

1. Run the following command to create an asynchronous replication session between the NAS servers:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/
session create -async -srcRes nas_1 -dstType remote -dstSys RS_2
-dstRes nas_2 auto -rpo 02h30m
```

```
Job ID = N-86
Operation completed successfully.
```

2. Run the following command to create an asynchronous replication session between the file system "res\_7" on the source system and the file system "res\_8" located on the remote system "RS\_2", with an RPO set to 2 hours 30 minutes and automatic synchronization:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/
session create -async -srcRes res_7 -dstType remote -dstSys RS_2
-dstRes res_8 -syncType auto -rpo 02h30m
```

```
Job ID = N-89
Operation completed successfully.
```

## Create a synchronous replication session for file storage

### Before you begin

Complete the following:

- For remote replication:
  - Identify the remote system that will act as the replication destination.
  - Create relevant replication interfaces, replication connection, and a storage resource on the remote system that will act as the destination. The storage resource on the destination system must have the same size.
- For local replication, create a storage resource that will act as the destination.
- For file replication, create a replication session on the NAS server associated with the file storage.

Configure a synchronous replication session between the NAS servers associated with the file storage.

### Procedure

1. Run the following command to create a synchronous replication session between the NAS servers:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/
session create -name MyNSRepl -srcRes nas_1 -dstType remote -
dstSys RS_1 -dstRes nas_1
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID =
103079215106_FCNC097274999_0000_103079215106_FCNC0972749A9_0
000
Operation completed successfully.
```

2. Run the following command to create a synchronous replication session between file systems on the source system and the remote system:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/
session create -name MyFSRepl -srcRes res_1 -dstType remote -
dstSys RS_1 -dstRes res_1
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID =
171798691844_FCNC097274999_0000_171798691844_FCNC0972749A9_0
000
Operation completed successfully.
```

## View replication sessions

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session  
show
```

```
Storage system address: 10.0.0.1  
Storage system port: 443  
HTTPS connection  
  
1:      ID              =  
81604378635_FCNCH097274B3A_0000_81604378637_FCNCH097274B37_0000  
      Name              = REP2  
      Session type      = nas server  
      Synchronization type = auto  
      Resource type      = NAS Server  
      Destination type   = remote
```