

Dell Unity™ Family Configuring Replication

Version 5.3

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Additional resources.....	6
Chapter 1: Overview.....	7
About this document.....	7
About replication.....	7
Required ports for replication.....	14
Snapshot replication in advanced replication topologies.....	14
Dell MetroSync for Unity.....	15
Using replication for disaster recovery.....	15
Using replication for planned downtime.....	17
Fail back a replication session.....	19
Advanced replication topologies for file replication.....	21
Asynchronous replication of a synchronous file replication to a third site.....	23
Preserve replication sessions in advanced topologies.....	26
Limitations for advanced replication topology	27
Upgrade a system that has an advanced replication topology.....	29
Chapter 2: Replication workflow.....	31
Configure replication.....	31
Configure local replication.....	31
Configure asynchronous replication.....	31
Configure synchronous replication.....	32
Configure asynchronous replication of synchronous replication to a third site.....	33
Related concepts.....	34
Replication interfaces.....	34
About replication connections.....	34
About replication sessions.....	35
About replication support for thin storage resources with data reduction.....	36
Chapter 3: Considerations for synchronous replication.....	37
Synchronization states.....	37
Code page translation of NFSv3 and FTP clients.....	38
Fibre Channel configuration for synchronous replication.....	38
Chapter 4: Configure replication using Unisphere.....	40
Configure protection and mobility interfaces.....	40
Configure replication connections.....	41
Create a replication session.....	41
Delta synchronization after failover	43
Prevent a full synchronization after a failover	43
Re-create a replication session.....	44
Manage replication sessions.....	44
Configure an override network address for file replication	47
Configure failover settings.....	47

Chapter 5: Configure replication using the CLI.....	48
Manage network interfaces.....	48
Create interfaces.....	49
View interfaces.....	50
Change interface settings.....	51
Delete interfaces.....	52
Manage remote storage systems.....	53
Create remote system configurations.....	54
View settings for remote storage systems.....	55
Verify settings for remote storage systems.....	56
Change settings for remote storage systems.....	57
Delete remote system configurations.....	58
Cabinet level unplanned failover of replication sessions.....	59
Manage replication sessions.....	59
Create replication sessions.....	64
View replication sessions.....	68
Change replication session settings.....	69
Pause replication sessions.....	72
Resume replication sessions.....	72
Manually synchronize replication sessions.....	74
Delete replication sessions.....	74
Fail over replication sessions.....	75
Fail back replication sessions.....	76
Preserve asynchronous replication sessions.....	78
Enable asynchronous snapshot replication after a replication session is created.....	78
Disable asynchronous snapshot replication after a replication session is created.....	80
Cascade a replicated snapshot in an asynchronous replication session.....	82
Provision an asynchronous replication session to autoreplicate user snapshots.....	83
Common base snapshots.....	87
Run a precheck for existing common base snapshots.....	87
Re-create a replication session.....	88
Create a replication session with a user snapshot as a common base.....	90
Associated CLI commands.....	90
The show action command.....	90
Manage Ethernet ports.....	92
Manage FC ports.....	95
Manage link aggregations.....	97
 Appendix A: Use case for CLI.....	 102
Replication configuration use case.....	102
Configure local replication.....	102
Configure asynchronous replication.....	102
Configure synchronous replication.....	103
Create a replication interface.....	104
Create a replication connection.....	105
Create a replication session for block storage.....	106
Create an asynchronous replication session for file storage.....	107
Create a synchronous replication session for file storage.....	107

View replication sessions.....	108
--------------------------------	-----

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 described below.

Product information

For product and feature documentation or release notes, go to Unity Technical Documentation at: dell.com/unitydocs.

Troubleshooting

For information about products, software updates, licensing, and service, go to Support (registration required) at: dell.com/support. After logging in, locate the appropriate product page.

Overview

Topics:

- [About this document](#)
- [About replication](#)
- [Snapshot replication in advanced replication topologies](#)
- [Dell MetroSync for Unity](#)
- [Using replication for disaster recovery](#)
- [Using replication for planned downtime](#)
- [Fail back a replication session](#)
- [Advanced replication topologies for file replication](#)

About this document

This document provides information that 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 user interface, this document also includes information about the CLI commands that are associated with configuring replication.

NOTE: For more information about other Unisphere features or CLI commands, see the *Unisphere online help* and *CLI User Guide*.

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 vStorage VMFS datastores, VMware NFS datastores, and thin clones.

NOTE: VMFS datastores that are being replicated are treated as CGs by replication and therefore are subject to those limits which are lower than the regular LUN limits.

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.



NOTE: If data is replicated to an unsupported earlier version of OE, that replication could result in data corruption on the replica.

- For fan-out and cascade topologies with synchronous sessions, the systems on both sides of the synchronous session must be running OE 5.2 or later.

NOTE: It is a best practice to enable 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 may not appear to make any progress.
- If an 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 to synchronize 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 synchronized.
- **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, which is 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. For synchronous replication, RPO is fixed at 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, which 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.



NOTE: For file synchronous replication, the model-to-model source and destination replication pair that is 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 nonthin (thick) LUN, or a thick LUN to a thin LUN with a LUN move operation. For more information on LUN move operations, see *Dell Unity Family Configuring and managing LUNs*.

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 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 that is either enabled or not 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.

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

- The user sets the type of NAS server 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.
- i 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 backup) specifies a set of attributes that are related to enabling backup or local test through NFS or NDMP protocols. The local configuration of a NAS server is not replicated as part of NAS server replication.

[Classification of configuration attributes for NAS server](#) classifies the Global, Override, and Local attributes for NAS server. Yes means that the attribute exists in the configuration, and No means that 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

[File system level mount options saved into NAS server configuration](#) lists file system-level mount options that are also saved into the NAS server configuration, and all are Global configurations:

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
-cifsNotifyOnAccess	isCIFSNotifyOnAccessEnabled	FS_PROPERTY_NOTIFYONACCESS_SYNC
-cifsOpLocks	isCFSOpLocksEnabled	FS_PROPERTY_NOOLOCK
-cifsSyncWrites	isCIFSSyncWritesEnabled	FS_PROPERTY_NP_CIFSSYNCWRITE
-accessPolicy	accessPolicy	FS_PROPERTY_NP_ACCESSPOLICY
-folderRenamePolicy	folderRenamePolicy	FS_PROPERTY_NP_RENAMEPOLICY
-lockingPolicy	lockingPolicy	FS_PROPERTY_NP_LOCKINGPOLICY

Table 2. File system level mount options saved into NAS server configuration (continued)

UEMCLI Option	RESTful Option	Mount Option
<i>-eventProtocols</i>	<i>fileEventSettings</i>	<i>FS_PROPERTY_NP_CEPPPOLICY</i>

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.

[File system properties not related to mount options](#) lists other file system properties that are not related to mount options.

Table 3. File system properties not related to mount options

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

NOTE: 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` is not 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 are 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 the following resources that support asynchronous replication:

- File system
- LUN
- LUN group (also known as a consistency group, or CG)
- VMware vStorage VMFS
- VMware NFS

 **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 either 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.
- In-Progress (Yes)-snapshot is currently syncing.
- 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.

 **NOTE:** 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:

- No-snapshot was created on the destination system and will not be replicated on the source system, or the snapshot was created before synchronous replication session setup
- Yes-successfully replicated
- Failed to replicate (Failed)-snapshot was created on the source system while the replication session was in fracture and will not be replicated on the destination system

Sync Replicated snapshot schedule replication

A 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. You can configure or change the schedule time zone that is used in your storage environment. The selected schedule time zone reflects the Universal Time Coordinated (UTC) adjusted by an offset for the local time zone. The schedule time zone applies only to snapshot scheduling, and asynchronous replication throttling scheduling. It is not a general system time zone setting.

The schedule can only be associated to a storage resource (file system or VMware NFS datastore) 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 with 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

 **NOTE:** 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 performing 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 the source file system has been resized after a NAS server failover, the sessions are fractured until the NAS server is failed back.

If you must perform a graceful/planned asynchronous replication failover, you must use `Failover-with-sync` from the source system. If you must 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 from the destination system before it allows an unplanned failover. If you must 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.

NOTE: If you must perform a graceful/planned asynchronous replication failover, use `Failover-with-sync` from the source system. If you must 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 from the destination system before it allows an unplanned failover. If you must 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.

NOTE: Do not perform a group operation at both sides of a replication session simultaneously. This action is not prohibited by the storage system, however, a group operation that is performed simultaneously 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, the write and read requests of the hosts 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 exceed 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
- Pause/Resume

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 does 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.

Required ports for replication

Communications with the Unisphere and CLI interfaces are conducted through HTTPS on port 443. Attempts to access Unisphere on port 80 (through HTTP) are automatically redirected to port 443.

Replication services are associated with the following ports:

Service	Protocol	Port	Description
Replication services	TCP	5080	Associated with DHSM services.
Replication services	TCP	5085	Associated with replication services.
Replication services	TCP	8888	Used by the replicator (on the secondary side). It is left open by the replicator when some data must be replicated. After it is started, the service cannot be stopped.

See more detailed information about storage system network ports in the *Dell Unity Family Security Configuration Guide*.

Snapshot replication in advanced replication topologies

For snapshot replication support, you can enable all file replication sessions that are associated with the same storage resource to transfer the snapshots 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. The solid lines represent replication sessions that 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 1. 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	Yes	180
A to D	Yes	240
A to E	Yes	480
B to F	Yes	360
C to G	Yes	720
D to H	Yes	960
E to I	Yes	1440

Dell MetroSync for Unity

Dell 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 system. That

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 that is 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 data that resides on the destination system. Using replicas for disaster recovery minimizes potential data loss. The RPO, which you configure when you set up the replication session, affects the amount of potential data loss. In synchronous replication configuration, where the RPO is set to 0, the amount of potential data loss is 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.

NOTE: A snapshot of a consistency group or VMFS datastore that was created on a Unity OE release before 5.1 and replicated by asynchronous replication cannot be used as a common base.

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

NOTE: (This note 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 performed on the destination site.
- Remote system IP connection is broken during the source site restarting.
- SP failover and file synchronous replication failover are executed simultaneously.

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.

NOTE: 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 rejects the failover from the source (planned failover), but the system does 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 perform 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 performed 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 fallback. 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 performed 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 regarding 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 must 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.

Ensuring host access after failover

You can perform checks to ensure that the host can reach the NAS server or LUN after a planned failover.

To ensure host access to the NAS server, verify the following information:

- All the related network ports for the NAS server production network interfaces are connected.
- The source and destination NAS server production networks have the same settings (for example, subnets, gateways). If the settings are the same, the source passes the same settings to the destination.
 - If the settings are the same, the source passes the same settings to the destination.
 - If source and destination NAS server production network settings are different, you must override the destination NAS server network interface or interfaces.

You can ensure host access to the LUN by verifying the following information:

- Create a host on destination array if a host does not exist.
- Add the host to the LUN.

i **NOTE:** If you perform an unplanned failover in an isolated network as part of a test scenario, the failover might cause data loss because both the disaster recovery array and the source array are in read-write mode. The isolation causes both the source and destination arrays to be writable. The failback operation from the disaster recovery array triggers a full synchronization, updating the source array with data written to the disaster recovery array after failover. This full synchronization overwrites data changes that are made on source array from the time when the unplanned failover occurred to the time the failback operation begins.

File-based replication considerations

The NAS server replication session should be in one of the following states in order to do a planned failover 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 failover 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 that are associated with the NAS server using the **Sync** option. This ensures that the destination file systems have the latest data and minimal data needs 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.

i **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 performed 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 remote side is remounted as read-only.
- 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.

NOTE: It is a best practice to avoid writing more data to the source during a planned asynchronous failover operation. If the data that is written to the source during the failover is larger than 100 MB, the failover might continue to run. The still-running failover operation can cause a long wait time.

Fail back a replication session

If you want to resume operations on a source system, the associated replication session must be failed back. Use the **Failback** option on the original destination system to fail back a replication session. For asynchronous replication, failback synchronizes the original source with the changes that are made on the original destination after failover, including any snapshots that have been created since the failover operation occurred. The failback 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 that are made on the original destination after an unplanned failover, not including any snapshots.

Beginning with OE version 5.1, you can avoid a full copy when performing a failback operation to the original source after an unplanned failover. Instead, only the data that are modified after the failover is replicated back to the original source from the destination system. To avoid a full copy when a failback operation is required, you must earlier establish common base snapshots on the source and destination systems either manually or by creating a snapshot schedule.

NOTE: Re-creating NAS server replication sessions using common base snapshots is not supported. For these types of sessions, a full copy must be used. A full copy is needed (to resume from the destination) due to a restriction of the MirrorView/Synchronous technology.

Avoiding a full copy saves space on the source storage system after failback and often reduces the time that is required to synchronize the source and destination sites. The common base snapshot also enables you to re-create a file system replication session that is interrupted by an unplanned failover.

If there are snapshots on the source storage resource, its size is increased (about 100%). If those old snapshots are deleted later, the size shrinks back. Failback restores the source as the production system and restarts the replication session in the original direction.

NOTE: A snapshot of a consistency group or VMFS datastore that was created on a Unity OE release before 5.1 and replicated by asynchronous replication cannot be used as a common base.

File-based synchronous replication

In file-synchronous replication, the source SPs and the destination SPs have connections (SPA to SPA, SPB to SPB) that are independent of each other. Therefore, the source NAS server and the destination NAS server must be in the same SP to stay synchronized. If a disruption occurs on the owner of the source NAS server, the source NAS server is failed over to the peer SP. The destination NAS server remains in the original SP.

Before OE 5.1, the data could not be synchronized when the source NAS server and the destination NAS server were on different SPs. With OE 5.1, the data stays synchronized. The system generates an alert that tells you that the operational status has degraded on each file system of the NAS server. Latency might increase as the operation proceeds.

Synchronous replication is protected in the following cases; in these examples, SPA is the primary SP:

OE version on source SPA	OE version on destination SPA	Synchronous replication protected?
5.1	5.1	Yes
Pre-5.1	5.1	Yes
5.1	Pre-5.1	No

The following restrictions apply when the source and destination NAS servers are on different SPs:

- Snapshots created, refreshed, or deleted on the source are not replicated to the destination system.
- If the size of the file system on the source changes, that change is not synched to the destination. The session enters the "syncing" state, and the data cannot be synchronized. The synchronization is recovered when the source NAS server is failed back to its owner SP.

File-based replication consideration

To resume operations on a source system, the associated NAS server replication session must 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.

Advanced replication topologies for file replication

Unity supports advanced replication topologies, that is, fan-out (1 to many) replication and cascading (multihop) file replication. Fan-out supports a maximum of four replication sessions on the same file storage object between two remote systems, including the local system. The storage object is not required to be 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).

Beginning with OE version 5.2, cascaded topologies allow one of the replication sessions to be synchronous. However, if the replication from the source to the destination is synchronous, the replication sessions from that destination to subsequent sites must be asynchronous. In the following figure, the dotted line from Site A represents a synchronous replication session to site B. The solid lines from Site B and C represent asynchronous replication sessions to destination sites:




Figure 2. Cascaded topology with synchronous replication from the source

If the replication session type from the source is asynchronous, one of the three sessions cascaded from the destination site can be synchronous:




Figure 3. Cascaded topology with asynchronous replication from the source, synchronous replication from the destination

User snapshot replication is supported on all replication sessions that are associated with the same storage object (an individual file system of NFS datastore). However, some restrictions apply to existing replication operations:

- 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.
- Only one local asynchronous session is supported per storage object.
- Four replication sessions are supported for one storage resource. Only one session can be synchronous; the remaining three must be asynchronous.

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 5.2. A represents a production site. The source object is not in destination mode and all the replication sessions on this resource act as the source.

The replication session from Site A to a destination can be synchronous. The replication sessions to the other three cascaded remote sites must be asynchronous.

B, C, 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. F, G, H, I, J, K, and L represent end remote sites. All the sessions on these resources act as the destination and are in destination mode.

In the figure, the dotted arrows represent synchronous sessions, and the solid arrows represent asynchronous sessions:

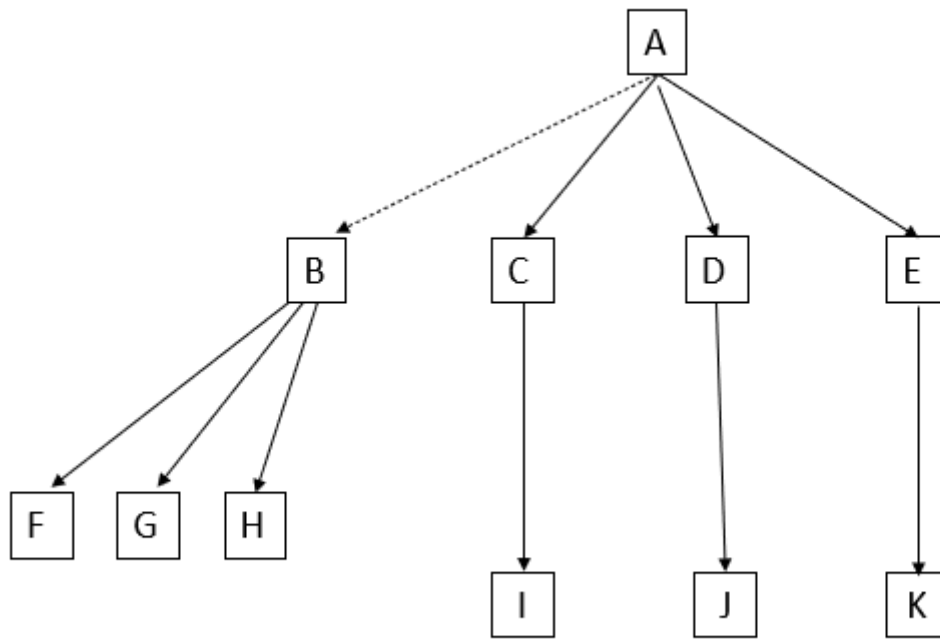


Figure 4. Example of Replication with Fan-Out and Cascade

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 NAS server or when a NAS server does not have an associated replication session, you can configure one synchronous replication session and three asynchronous replication sessions for that NAS server for a total of four sessions.

For 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. If you are using the CLI and you want the asynchronous replication destination NAS server to be used as backup only, you must set the **-backupOnly** attribute to **yes**.

When creating a file system or when a file system does not have an associated replication session, you can configure one synchronous replication session and three asynchronous replication sessions for that file system.

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 is 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.

NOTE: Beginning with OE 5.2, you must create a snapshot when you create the replication sessions if you want to run a preserve operation later.

The preserve 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 destination sites on every asynchronous session sync operation. If the snapshot create operation or refresh operation fails on the synchronous replication destination site, the internal snapshot create operation or refresh operation on the synchronous replication destination site is 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 operation or refresh operation on the synchronous replication destination site is retried as part of the next asynchronous session sync operation after the synchronous replication session is restored.

Concurrent operations compatibility

Most of the synchronous and asynchronous replication operations for co-existing 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 is 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 to use as a common base 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 is performed.

Beginning with OE 5.2, the system snapshot from the site A to site C asynchronous session is not synced to site B and cannot be used as a common base snapshot. The system snapshot is automatically transferred from Site A to Site B during the synchronous replication session. However, you must manually create a snapshot on source production Site A and sync that snapshot to both Site B and Site C. This snapshot is used as the common base snapshot for the Site B to Site C preserve operation.

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.

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.
- Asynchronous replication to another site for backup is supported on Unity systems with OE versions 4.4 or later.
- The remote system connection towards the backup (asynchronous replication destination) site must 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 performed from the asynchronous replication

destination site. Such an operation breaks the synchronous replication. Ensure that either the synchronous replication session is deleted, or that both sites of the synchronous replication session are down and are not 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 fails 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 continues based on the snapshots and information about the new synchronous production site. Some snapshots may not be replicated to the backup site. In this case, a warning is generated when marking the snapshot for asynchronous replication when the synchronous replication session is not Active. Some snapshots may result in duplicate storage space.
- 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 is increased (by 100% of the production file system size). If those old snapshots are deleted, the size shrinks.
- 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 canceled before a planned failover or failback synchronous session is performed.

Remove an asynchronous replication session

You can remove the asynchronous replication session between the source site and the asynchronous replication destination site. When you perform this operation, the internal snapshots of the source and asynchronous destination sites that are used for the asynchronous replication session are removed. If the deletion happens when the synchronous replication session is Active, the internal snapshots of the synchronous destination site are removed.

NOTE: The internal snapshots that are 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 must be performed.

Remove a synchronous replication session

You can remove the synchronous replication session between the source site and the synchronous replication destination site. When you perform 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 that uses that snapshot. If an asynchronous replication session exists and uses that snapshot, the remove operation fails.

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 replication

When a storage object is created for MetroSync for Unity as the destination, that storage object can be used to create additional replication sessions when that object acts as the source.

The following figure and table show a simple 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 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 5. Cascade mode for third site asynchronous replication

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

Replication Session	User Snapshot Replication	RPO (in minutes)
A to B (synchronous)	Yes	0
A to C (asynchronous)	Yes	60
C to D (asynchronous)	Yes	1440

Preserve replication sessions in advanced topologies

In OE 5.2 and later releases, after a synchronous replication session is failed over in a fan-out or cascaded topology, the asynchronous session is not automatically hibernated.

You can still preserve asynchronous replication sessions. Create a user snapshot for any new synchronous replication sessions to use the preserve function after a failover or failback operation.

To use the preserve function after a failover or failback operation, you must create a system snapshot for any new synchronous replication sessions.

If you want to preserve all asynchronous sessions, you must perform preserve operations for each session to the destination site.

NOTE: If a user snapshot pair existed before the system was upgraded to OE 5.2, the user snapshot pair can be used as a common base. The common base snapshot prevents the need for a full copy after failover or failback.

For example, in OE 5.2, the system snapshot from the Site A to Site C asynchronous replication session is not automatically synced to Site B. The A to C snapshot cannot be used as a common base for a B to C preserve operation. Instead, you must create a snapshot on A that is synchronized to B by a synchronous replication session and C by an asynchronous replication session. This new A to B snapshot pair is used as the common base snapshot for a B to C preserve operation as shown in the following figure. The thin lines represent normal synchronous and asynchronous sessions. The thick line shows the B to C preserve operation made possible by the creation of the new snapshots:




Figure 6. New snapshots for a preserve operation

NOTE: It is a best practice to use the snapshot pair that is created in OE 5.2 after a nondisruptive upgrade (NDU).

Preserve replication options

In Unisphere, there are two options for preserving an asynchronous replication session.

The first option, **Preserve first async replication** option, switches the first asynchronous replication session to the production site after the synchronous replication session has been failed over or failed back. Use this option when only one asynchronous replication session exists on the storage resource.

The second option, **Preserve specified async replication**, enables you to choose which asynchronous replication session is switched to the active production site after the synchronous replication session has been failed over or failed back. Use this option when more than one asynchronous replication session exists on the storage resource.

Limitations for advanced replication topology

There are limitations with fan-out and cascade replication topologies that support both synchronous and asynchronous replication sessions.

NOTE: A topology with inbound asynchronous replication and outbound synchronous replication is never a recommended topology. Such a configuration, as shown in the following example, should only occur after certain disaster recovery cases:




Figure 7. Abnormal replication topology

After all systems are recovered, Dell strongly recommends that you immediately return your system to a normal topology. Such a normal topology is, for example, one synchronous session to a second system plus one asynchronous session to another system.

Other limitations include the following:

- Multiwrite topologies are not supported.
- If synchronous replication is used, both the primary site and the destination site must be running OE 5.2.

- Only one synchronous session is supported for each storage resource.
A total of four replication sessions is supported for each storage resource. The remaining three replication sessions must be asynchronous.
- A planned failover of fan-out and cascaded topologies is supported unless the failover results in a multideestination topology. For example, in an A to B to C cascade topology, site A is the primary site. A planned failover from Site A to B is supported. However, a planned failover from Site B, the destination site, to Site C, another destination site, is not supported. Failing over from Site B to Site C would create a multideestination topology.
- Snapshots cannot be taken or refreshed until the initial synchronization between the primary and backup site is finished.
i NOTE: If an initial synchronization for a synchronous session is not complete, creating an asynchronous session on the same resource might cause a nonrecoverable error.
- Snapshots that are sent asynchronously from a primary site might not be sent to downstream to a third site. This problem might occur if there is a synchronous replication session between the initial destination site and the third destination site that is downstream. This problem is shown in the following figure:




Figure 8. Snapshot not synchronously sent to third site

Downstream snapshot replication

Beginning with OE 5.2, a synchronous replication session destination can have a downstream asynchronous session. Snapshot replication continues after a planned failover and is preserved automatically.

For example, in the following figure, snapshots are replicated asynchronously from site A to Site D. After a planned failover and preserve operation, those snapshots are asynchronously replicated from Site B to Site D:

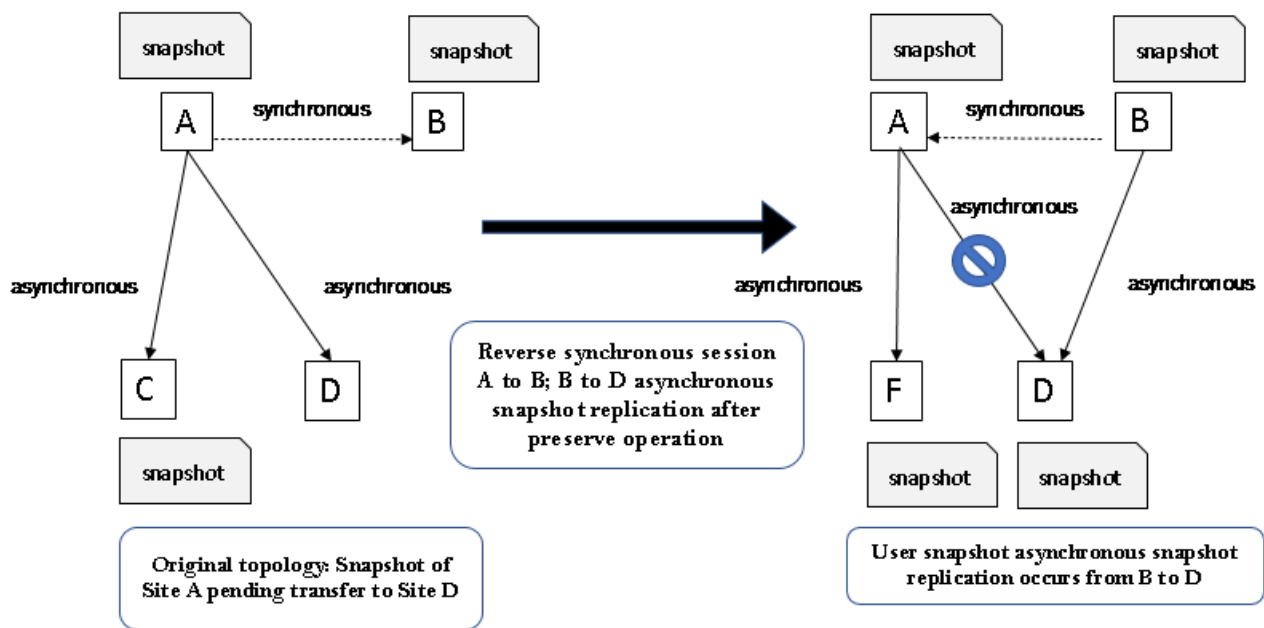


Figure 9. Downstream snapshot replication after failover

Snapshots and deleted replication sessions

If you delete a synchronous replication session in an advanced topology, the user snapshot asynchronous replication information that the synchronous session replicated is also deleted.

If you delete an asynchronous replication session in an advanced topology, the destination side of the synchronous session deletes snapshot information that the synchronous session replicated. The snapshots are maintained.

Upgrade a system that has an advanced replication topology

If you have an advanced replication topology (fan-out or cascade), you must follow a specific sequence when performing a nondisruptive upgrade (NDU). When you upgrade such a system to OE 5.2, you must upgrade the backup sites before the primary production site.

Back up the destination site that has an asynchronous replication session with the primary production site first.

For example, primary production Site A has replication sessions with two backup sites, B and C. Site A has synchronous session with Site B as the destination. Site A also has asynchronous session with Site C. To upgrade all these sites to OE 5.2, you must first upgrade Site C, followed by Site B. Finally, you upgrade the primary production Site A.

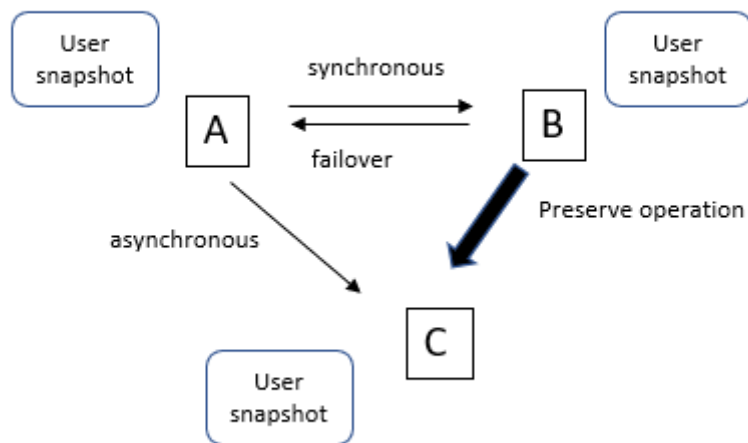


Figure 10. The preserve operation in an advanced topology

If you upgrade the primary site A first, the upgrade might make the asynchronous session from the primary site A to backup site C unrecoverable.

If you upgrade Site B first, Site B no longer receives the system snapshot from Site A. Also, Site C cannot use a user-created snapshot as a common base during a preserve operation.


Replication workflow

Topics:

- [Configure replication](#)
- [Related concepts](#)


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

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.

 **NOTE:** Local replication is applicable to asynchronous replication only. Local replication is not applicable to synchronous replication.

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

Prerequisites

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 a NAS server is being replicated, the destination NAS server must have a matching tenant configuration. For example, you cannot replicate a non-tenanted NAS server to a tenanted NAS server. Tenants must be created on the target system using the same UUID as on the source system.

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.

Steps

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.
5. (Optional) Create one or more bandwidth schedules to control when replication occurs, and how much bandwidth should be used during replication.
Bandwidth schedules are configured to replication connections to the remote system. All asynchronous replication sessions with active data transfer on their replication connections evenly share the bandwidth assigned for the specified days and hours.
NOTE: If no schedule is set, all available bandwidth will be used during replication.

Configure synchronous replication

About this task

Only remote replication is supported for synchronous replication.

If a NAS server is being replicated, the destination NAS server must have a matching tenant configuration. For example, you cannot replicate a non-tenanted NAS server to a tenanted NAS server. Tenants must be created on the target system using the same UUID as on the source system.

Steps

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 that 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

Prerequisites

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.

About this task


 **NOTE:** Only remote replication is supported for synchronous replication.

Steps


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.

Before OE version 5.1, when creating cascading or fan-out asynchronous replication sessions, all interfaces in a source/destination pair required a valid interface path. In OE version 5.1, replication interface pairing only requires valid interface paths between systems. Interfaces that can be connected between the source and destination systems are used for the replication connection, which allows asynchronous replication sessions to span across different networks or VLANs.

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.

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.


Considerations for synchronous replication

Topics:

- [Synchronization states](#)
- [Code page translation of NFSv3 and FTP clients](#)
- [Fibre Channel configuration for synchronous replication](#)

Synchronization states

Replication sessions configured with synchronous replication (RPO set to zero) have synchronization states that describe the state of the user data on the destination storage resource. These states are useful for troubleshooting. However, for replication session management, it is recommended that you check the operational status. For accurate information about the synchronization state of a replication, see 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.

Synchronous replication sessions can have the following states:

- **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 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.
 - The replication session requires a full synchronization after an unplanned failover operation, when a failback or resume operation is performed.

If there is no common base snapshot between the source system and the destination system, a full synchronization is required after an unplanned failover operation.


- The 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 is updated to the destination system once the connection issue is resolved.
 - The replication session was paused resulting in a backlog of changes on the source system. The backlog is updated to the destination system once the replication session is resumed.
- **Out of Sync** – The replication session cannot determine the state of the data on the destination LUN. The update to the destination system may be incomplete. This issue 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.


 **NOTE:** If a synchronous replication session is fractured or paused and there are still data writes on the source storage resource, you should recover or resume the session as soon as possible. If data writes are sparsely distributed on the source storage resource, resynchronizing may transport more data than the original data writes and, if there are snapshots on the destination storage resource, they occupy more storage space size. If those old snapshots are deleted, the size shrinks 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 before 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	Unknown > Out of Sync > Syncing > Consistent > In Sync
Pause	In Sync > Consistent or Syncing > Out of Sync
Resume	Consistent or Out of Sync > Syncing > Consistent > In Sync
Fail over	Syncing > Inconsistent or Out of Sync > Inconsistent  NOTE: If the state was In Sync or Consistent before a failover occurred, that state is retained.
Replication connection down	Syncing > Out of Sync or In Sync > Consistent (if there is I/O)

 **NOTE:** 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 might take 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.


Code page translation of NFSv3 and FTP clients

All Unity systems (OE version 4.2 and earlier) are configured by default to support NFSv3 and FTP clients configured for UTF-8. The code page on these Unity systems cannot be changed.

With Unity OE version 4.3 and later, the code page of the Unity system can be changed through the `svc_nas {<NAS_server_name> | all} -param -facility vdm -modify codepage -value <value>` service command to UTF-8, 8859-1, or 8859-15. Code page 8859-15, which includes the most used western European characters, is an extension of the 8859-1 code page.


 **NOTE:** The default code page in Unity systems does not need to be changed for either NFSv4 or SMB clients. NFSv4 clients support only UTF-8 and SMB clients support only Unicode.

With Unity systems OE version 5.0.3 and later, the code page of the Unity system can be changed through the same service command to CP949 (EUC-KR Korean). The code page on earlier Unity systems (OE version 5.0.2 and earlier) cannot be changed to CP949. Refer to *Unity Service Commands Technical Notes* for more information.

 **NOTE:** A failover from Unity systems OE version 5.0.3 with the code page set to CP949 to a system with an earlier OE version will revert the codepage to UTF-8. On failback, the code page must be reconfigured on the version 5.0.3 system.

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*.

- i 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.

- i NOTE:** Only a one to one FC connection is allowed for synchronous replication. One to more than one FC connection is not supported. Also, residual zones to an earlier source system used in synchronous replication should be removed.

For example,

- OB_H1450 SPA Port 4 <--> OB_H1443 SPA Port 4
- OB_H1450 SPB Port 4 <--> OB_H1443 SPB Port 4

- i 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.

Configure replication using Unisphere

Topics:

- [Configure protection and mobility interfaces](#)
- [Configure replication connections](#)
- [Create a replication session](#)
- [Manage replication sessions](#)
- [Configure an override network address for file replication](#)
- [Configure failover settings](#)

Configure protection and mobility interfaces

Prerequisites

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.

Steps

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.

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

Configure replication connections

Prerequisites

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**.

NOTE: If a replication connection already exists and you plan to add a different mode of replication, do not attempt to create a new connection. Change the existing replication connection mode to Both.

Steps

1. Under **Protection & Mobility**, select **Replication > Connections**.
2. Perform one of the following actions:
 - To create a replication connection, select the **Add** icon. Complete the **Create Replication Connection** wizard.
 - 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

Prerequisites

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

- The replication mode that you want. You can choose from the following options:
 - Asynchronous
 - Manual
 - Synchronous
- The Recovery Point Objective (RPO) for the replication session.
- The system that you want to assign as the replication destination system. This selection is based on the replication connections that are configured on the storage system.
- The name, pool, storage provisioning, and tiering policy you want to use for the destination storage resource. The system automatically creates a destination storage resource as part of this process.

NOTE: You can convert a thin LUN to a nonthin (thick) LUN, or a thick LUN to a thin LUN with a LUN move operation.

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, 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, the destination file system is also thin.
 - If both the source and destination systems support data reduction, the source or destination can have either data reduction that is 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 that are created from snapshot schedules. For synchronous file replication, you cannot replicate existing snapshots or snapshots that are created from snapshot schedules. You can only replicate those snapshots and snapshots that are 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.

i NOTE: To perform 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, whether those sessions are either in fan-out (star) or cascaded mode, must be running OE version 5.x.
- In Unisphere, NAS servers, file systems, and NFS Datastores that receive replicated data and also replicate data to local or other systems only appear as destinations on the **File > NAS servers** page and the **File > File Systems/VMware > Datastores** page.

Despite being part of a cascade topology, these resources do not appear when the source filter is selected. However, these replication sessions appear correctly on the **Replication > Sessions** page when either the source or destination filter is selected based on the replication session role.

- 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.

Steps

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

i 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 automatically creates the storage resource or NAS server on the destination system.

i 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 file system, Unisphere creates a replication by default.

5. Optional: Select **Support Asynchronous Replication**.

You can also choose the following options:

- **Cascade replicated snapshots**
- **Replicate all existing snapshots**
- **Replicate all scheduled snapshots**

6. Select **Reuse destination resource** if you want to use an existing destination resource.
7. Select **Automatically search user snap as common base** if you want to avoid an initial full synchronization; the pair of common base user snapshots allows you to avoid a full synchronization.

The initial full synchronization is the synchronization that occurs after the replication session is created. If a common base user snapshot pair exists, you can avoid a full synchronization, if a failover or fallback occurs.

Delta synchronization after failover

Beginning with OE version 5.1, you can prevent a full synchronization from occurring when you perform a failback operation from a destination system to the source system.

The avoid full copy feature performs a delta synchronization that prevents redundant data that is created after failover from being replicated back to the source. In addition, delta synchronization typically reduces the time that is required for a failback operation to complete.

To perform a delta synchronization, a pair of common base snapshots must exist on both the source and destination systems. With OE version 5.1, snapshot replication generates a unique signature for both synchronous and asynchronous replication sessions.


During failback, file-asynchronous replication delivers all the snapshot signatures and searches for a pair of snapshots to use as common base snapshots. With file-synchronous replication, only the last 21 replicated snapshot signatures are delivered back to the source system in the common base snapshot search.

You can use the avoid full copy feature for the following types of replication sessions:

- Synchronous file system
- Asynchronous file system


The avoid full copy feature cannot be used for the following types of replication sessions:

- Asynchronous NAS server
- Synchronous NAS server
- Synchronous block

 **NOTE:** A snapshot of a consistency group or VMFS datastore that was created on a Unity OE release before 5.1 and replicated by asynchronous replication cannot be used as a common base.

Cascaded and fan-out replication

Using file-asynchronous replication, you can choose to synchronize only the checkpoint snapshots after a synchronous replication session fails over. Checkpoint snapshots are read-only. Data would still be sent through the fan-out asynchronous replication sessions.

 **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.

You can enable the system to automatically send, or cascade, snapshots downstream. In addition, you can fan-out snapshot replication to multiple sites concurrently. Fan-out replication allows snapshot replication on up to four asynchronous replication sessions on the same storage resource.

Prevent a full synchronization after a failover

Prerequisites

A pair of common base snapshots must exist on the source and destination sources.

About this task

You can use the avoid full copy feature for the following types of replication sessions:

Steps

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 that is associated with the replication session, go to the **Replication** tab on the **Properties** window of the relevant storage resource or NAS server.
2. Select **Failback**.

The **Failback Session** window appears and asks if you want to proceed with the failback session the synchronization that is either full or delta.

Re-create a replication session

Beginning with OE 5.1, you can re-create a synchronous or asynchronous replication session that was interrupted because of an unplanned failover.

Prerequisites

A pair of common base snapshots must exist on both the source and destination systems. These snapshots are snapshots that were created according to an established schedule or can be manually created snapshots of both the source and destination systems.

About this task

The avoid full copy feature enables you to re-create the following types of replication sessions:

- Synchronous file system
- Asynchronous file system
- Asynchronous block

The following types of replication sessions cannot be re-created:

- Synchronous NAS server
- Synchronous block
- Asynchronous NAS server

Steps

1. Delete the existing replication session.
2. Under **Storage**, select the storage resource, and then select **Edit**.
3. On the **Properties** window, select the **Replication** tab.
4. Select **Configure Replication**.
5. In the **Create a Session** window, select the checkboxes for **Reuse destination resource** and **Automatically search user snap as common base**.

Manage replication sessions

About this task


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](#).

With the advanced file replication topologies, fan-out (star) and cascaded mode, up to four replication sessions (one synchronous and three asynchronous or four asynchronous) 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 simultaneously, 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 operation 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 that is 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 that are associated with a single storage resource. If the resource is offline, it must recover the resource to be online, and 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 is paused together with the NAS server session.
- Beginning with OE version 5.2, you can pause and resume multiple replication sessions that are associated with remote systems in Unisphere. If any remote destination sites are running OE versions earlier than 5.2, only local replication sessions on systems running OE version 5.2 are paused.
- Other existing replication operations are supported as in earlier OE versions.


- **Resume**

- Only one active destination replica should exist.
- When you resume the NAS server-related sessions, only the associated file system session to the same remote system are resumed together with the NAS server session. When there are multiple sessions that are 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.
- You can resume multiple replication sessions that are paused if the local systems that are the source of the sessions are running OE version 5.2. You cannot resume multiple paused replication sessions on remote systems that are running OE versions earlier than 5.2.

Steps

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 that is 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, except for 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
Modify a session	Select the Edit icon to modify the name and synchronization settings of the session.
Delete a session	<p>Select Delete. 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> <p> NOTE: You cannot delete a replication when a synchronization is in progress.</p>
Pause and resume a session or sessions	<p>Select Pause or Resume. Pause stops a replication session or sessions without deleting the sessions. Resuming sessions 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> <p> NOTE: Pausing a replication session or sessions 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 paused sessions 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>
Manually synchronize a session (for asynchronous replication only)	<p>Select Sync. Changes since the last sync are replicated to the destination system.</p> <p> NOTE: 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>
Fail over a session in an unplanned emergency	<p>Select Failover. 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 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> <p> NOTE: 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> <p>For synchronous replication, in response to 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. 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>
Fail over a session in preparation for a planned downtime (for asynchronous replication only)	<p>For asynchronous replication, select Failover with Sync. 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> <p> NOTE: To perform this type of failover, both the source and destination objects of the session must be available.</p> <p>For synchronous replication, use Failover 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 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.</p>

Action	Description
Fail back a session	<p>Select Failback. 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 and the destination storage resource's status is shown as Restricted Replication Access = Yes.</p> <p> NOTE: 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> <p>Beginning with OE 5.1, you can perform either a full synchronization or a delta synchronization if the source and destination systems have a pair of common base snapshots.</p>
Preserve an asynchronous replication session	<p>Select Preserve. 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


About this task

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.

 **NOTE:** 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:

Steps

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.

Configure failover settings

Beginning with OE 5.1, you can manually set the system to fail over if the NAS server interface becomes inaccessible. Previously, a failover was triggered only when the SRM, array management, and NAS server production interfaces were inaccessible.

Configure replication using the CLI

Topics:

- [Manage network interfaces](#)
- [Manage remote storage systems](#)
- [Manage replication sessions](#)
- [Enable asynchronous snapshot replication after a replication session is created](#)
- [Common base snapshots](#)
- [Associated CLI commands](#)

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"> • <code>iscsi</code> — Interface for iSCSI storage. • <code>replication</code> — Interface for replication-related data or management traffic.
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. <i>i</i> NOTE: 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.
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	A numerical value indicating the health of the system. Value is one of the following: <ul style="list-style-type: none"> • Unknown (0) • OK (5) • OK BUT (7) • Degraded/Warning (10) • Minor failure (15) • Major failure (20)
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	Specify the interface type. Value is one of the following: <ul style="list-style-type: none"> • iscsi — Interface for iSCSI storage. • replication — Interface for replication-related data or management traffic.
-port	Specify the ID of the SP port or link aggregation that will use the interface. <i>i</i> NOTE: 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 spa_eth2, but the NAS server is on SP B, the interface is created on port spb_eth2.
-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. <i>i</i> NOTE: 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.
-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. <i>i</i> NOTE: This qualifier is not required if the prefix length is specified in the -addr attribute.
-gateway	Specify the gateway for the interface. <i>i</i> NOTE: This qualifier configures the default gateway for the specified port's SP.

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.

Format

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

Object qualifier

Qualifier	Description
-id	Type the ID of an interface.
-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">iscsireplication

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          =
        VLAN MTU size    =
        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
VLAN MTU size = 1500
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 =
VLAN MTU size =
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	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.  NOTE: 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.
-addr	Specify the IP address for the interface.  NOTE: The prefix length should be appended to the IPv6 address. The IPv4 netmask may be specified in address attribute after the slash.
-netmask	Specify the IPv4 subnet mask for the interface.
-gateway	Specify the gateway for the interface.  NOTE: The gateway is optional for both IPv4 and IPv6. This qualifier configures the default gateway for the specified port's SP.

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.

 **NOTE:** 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.
Name	Name 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.
Operational status	Operational status of the session. The operational status code appears in parentheses. <ul style="list-style-type: none">• Unknown (0x0)• OK (0x2)• Lost Communication (0xD)• Update Needed (0x8406)• Updating (0x8407)• Failed to Validate (0x840C)• Lost sync Communication (0x840D)
Model	Model number of the remote system.
Serial number	Serial number of the remote system.
Health state	Health state of the storage resource. The health state code appears in parentheses. Value is one of the following: <ul style="list-style-type: none">• OK (5) —Resource is operating normally.• Degraded/Warning (10) —Working, but one or more of the following may have occurred:<ul style="list-style-type: none">◦ One or more of its storage pools are degraded.◦ Its replication session is degraded.◦ Its replication session has faulted.◦ It has almost reached full capacity. Increase the primary storage size, or create additional resources to store your data, to avoid data loss.• Minor failure (15) —One or both of the following may have occurred:<ul style="list-style-type: none">◦ One or more of its storage pools have failed.◦ The associated iSCSI node has failed.• Major (20) —One or both of the following may have occurred:<ul style="list-style-type: none">◦ Resource is unavailable.◦ One or more of the associated storage pools have failed.• Critical failure (25) —One or more of the following may have occurred:<ul style="list-style-type: none">◦ One or more of its storage pools are unavailable.◦ Resource is unavailable.◦ Resource has reached full capacity. Increase the primary storage size, or create additional resources to store your data, to avoid data loss.• Non-recoverable error (30) —One or both of the following may have occurred:

Table 7. Remote system attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> Resource is unavailable. One or more of the associated storage pools are unavailable.
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.
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"> sync async both
Synchronous FC ports	The fibre channel ports enabled for synchronous replication. <i>i</i> NOTE: 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.
Bandwidth schedules	The bandwidth schedules for asynchronous replication sessions with source role having data transfer due to RPO or manual sync on this remote system. <i>i</i> NOTE: This attribute is used for async or both connection types only. If not configured, "not configured" is displayed. For sync connection type, N/A is displayed.
Current bandwidth	Current bandwidth calculated based on the configured bandwidth schedules and the schedule time zone for asynchronous replications on this remote system. <i>i</i> NOTE: This attribute is used for async or both connection types only. If no bandwidth schedules configured, "uses available bandwidth "is displayed. For sync connection type, N/A is displayed.

Create remote system configurations

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

i **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}] [ -bandwidthSchedules <value> ]
```

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">• VNX
-srcUsername	For systems that are the source in a replication, type the username that is used to access the system.
-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: <ul style="list-style-type: none">• async• sync• both
-bandwidthSchedules	(async or both connection types only) Specify a schedule during which to cap bandwidth. Type the following to specify the schedule [{Sun Mon Tue Wed Thu Fri Sat}] [HH-HH] [/<bandwidth>] , where: <ul style="list-style-type: none">• Sun Mon Tue Wed Thu Fri Sat - days of the week to run the schedule• HH-HH - start and end time for the schedule• /<bandwidth> - value at which to cap bandwidth, in KBps

Example

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

- Network address is 10.60.10.10.
- Includes access credentials for when the system is the source or destination.
- Limits bandwidth to 2,000 KB/s from 7 A.M. to 6 P.M. Monday through Friday. Otherwise, bandwidth is limited to 8,000 KB/s.

The configure remote system receives the ID RS_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /remote/sys create -addr 10.60.10.10
-srcUsername LocalAdmin -srcPassword LocalAdminPassword -dstUsername RemAdmin -dstPassword
RemAdminPassword -connectionType async -bandwidthSchedules MonTueWedThuFri7-18/2000,/8000
```

```
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

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 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_0
      Name              = FNM00151702099
      Address           = 10.245.252.229
      Model             = Unity 300 (Local System)
      Serial number     = FNM00151702099
      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 = spa_fc4, spb_fc4
      Bandwidth schedules = MonTueWedThuFri7-18/2000,/8000
      Current bandwidth   = 8000

2:      ID              = RS_1
      Name              = FNM00151702100
      Address           = 10.244.205.127
      Model             = Unity 300
      Serial number     = FNM00151702100
      Connection type   = sync
      Source interfaces = 10.245.252.231, 10.245.252.232
      Local interfaces  = 10.245.252.231, 10.245.252.232
      Remote interfaces = 10.244.205.131, 10.244.205.132
      Operational status = OK (0x2)
      Health state      = OK (5)
      Health details    = "Communication with the replication host is established. No
action is required."
      Synchronous FC ports = spa_iom_0_fc0, spb_iom_0_fc0
      Bandwidth schedules = N/A
      Current bandwidth   = N/A
```

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.

i NOTE: 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}] [-bandwidthSchedules <value>]
```

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 is prompted to specify the password.
-connectionType	Specify this qualifier to indicate the type of replication connection. Valid values are: <ul style="list-style-type: none">• async• sync

Qualifier	Description
	<ul style="list-style-type: none"> • both
-bandwidthSchedules	<p>(async or both connection types only) Specify a schedule during which to cap bandwidth. Type the following to specify the schedule [{Sun Mon Tue Wed Thu Fri Sat}] [HH-HH] [/<bandwidth>], where:</p> <ul style="list-style-type: none"> • Sun Mon Tue Wed Thu Fri Sat - days of the week to run the schedule • HH-HH - start and end time for the schedule • /<bandwidth> - value at which to cap bandwidth, in KBps <p>Leave blank to delete an existing bandwidth schedule.</p>

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 earlier 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"> • lun • block • file • nas_server
Synchronization type	Type of synchronization. Valid values are: <ul style="list-style-type: none"> • auto—Data is replicated automatically in order to maintain the desired Recovery Point Objective (RPO). • manual—Data will only be replicated when manually initiated. • sync—Data is synchronously replicated with RPO=0. <p>NOTE: 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	Type of storage resource to which the replication session is applied. Valid values are: <ul style="list-style-type: none"> • LUN • LUN group • File System • VMware VMFS • VMware NFS • NAS Server
Sync State	Additional state of the replication session, specific to the replication mode. <ul style="list-style-type: none"> • For asynchronous replication, valid values are: <ul style="list-style-type: none"> ◦ idle—No active syncing. Beginning with OE 5.x, you can change replication session attributes while the session is in an Idle state. ◦ manual—User initiated syncing. ◦ auto_syncing—System initiated syncing. • For synchronous replication, valid values are: <ul style="list-style-type: none"> ◦ unknown—Unknown sync state. ◦ out_of_sync—Destination is out of sync with the source. ◦ in_sync—Destination is an exact copy of the source. ◦ consistent—Destination is a point in time copy of the source. ◦ syncing—System initiated syncing. ◦ inconsistent—Destination is not a point in time copy of the source.
Health state	Health state of the session. Valid values are: <ul style="list-style-type: none"> • Unknown—Session health cannot be determined. • OK—Session is operating normally. • Degraded/Warning—An error has caused one or more of the following: <ul style="list-style-type: none"> ◦ Session has been paused. ◦ 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.

Table 8. Replication session attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> ○ Session is syncing. ● 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. ● Critical failure—Session has encountered an error that has halted the session. <p>i NOTE: 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"> ● Unknown (0x0) ● Non Recoverable Error (0x7) ● Lost Communication (0xd) ● Failed Over with Sync (0x8400) ● Failed Over (0x8401) ● Manual Syncing (0x8402) ● Paused (0x8403) ● Idle (0x8404) ● Auto Sync Configured (0x8405) ● Destination Extend Failed Not Syncing (0x840B) ● Destination Extend In Progress (0x840C) ● Active (0x840D) ● Lost Sync Communication (0x840E) ● Syncing (0x8411)
Source status	<p>Status of the source storage resource in the session. Valid values are:</p> <ul style="list-style-type: none"> ● Unknown—Source status is unknown. ● OK—Source is operating normally. ● Paused—Replication session for the source is currently paused. ● Fatal replication issue—Source has experienced a critical error and the replication session has stopped. Delete the replication session and re-create it. ● 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. ● 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. ● 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.
Destination status	<p>Status of the destination storage resource in the session. Valid values are:</p> <ul style="list-style-type: none"> ● Unknown—Status of the destination resource is unknown. ● OK—Destination resource is operating normally. ● Paused—Replication session for destination resource is currently paused. ● Fatal replication issue—Destination has experienced a critical error and the replication session has stopped. Delete the replication session and re-create it.

Table 8. Replication session attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> • 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. • 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. • 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.
Network status	<p>Status of the network connection. Valid values are:</p> <ul style="list-style-type: none"> • Unknown—Network status is currently unknown. If you continue to see this value, check the network connections. • OK—Network connection is operating normally. • 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 (lost IP) between the systems. A change in the network configuration on either side could also interrupt communication. • Lost Sync Communication—Fiber Channel communication with the synchronous replication remote system has been lost. It is likely that the Fiber Channel connection has encountered issues.
Destination type	<p>Type of destination used in the session. Valid values are:</p> <ul style="list-style-type: none"> • local—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. • remote—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.
Destination system	For remote sessions, the ID of the remote system on which the data is replicated.
Local role	<p>The local system role. Valid values are:</p> <ul style="list-style-type: none"> • Unknown—Status of the local system is unknown. • Source—Resource on the local system is replicated to the remote system. • Destination—Resource on the local system is the replication destination of the resource on the remote system. • Loopback—Resources participating in the replication session are located on the same storage system. • Local—Resources participating in the replication session are located on the different storage processors of the local system.
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.

Table 8. Replication session attributes (continued)

Attribute	Description
Sync status	Percentage of the replication synchronization that has completed and the amount of time remaining. i NOTE: For synchronous replication, the percentage is reported when the replication is in the Syncing state.
Sync transfer rate	Synchronization transfer rate when the session is in the syncing state. For multi-LUN applications there is a comma-separated list of values. i NOTE: This attribute is valid for asynchronous replications only.
Sync transfer size remaining	Remaining size to be transferred during synchronization. For multi-LUN applications there is a comma-separated list of values. i NOTE: This attribute is valid for asynchronous replications only.
Previous transfer rate	Previous average transfer rate for the replication session. i NOTE: This attribute is valid for asynchronous replications only.
Average transfer rate	Average transfer rate for the replication session. i NOTE: This attribute is valid for asynchronous replications only.
Element pairs	For consistency group and VMware VMFS datastore replications, the LUN element pairs within the replication.
Hourly snapshot keep for	Amount of time to keep replicated hourly snapshots on the destination. Output can be: <ul style="list-style-type: none"> Blank when scheduled snapshots are not replicated. <value><qualifier>—When a retention duration is specified, where: <ul style="list-style-type: none"> value—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. qualifier—A value qualifier. The valid values are: <ul style="list-style-type: none"> h (hours) d (days) Forever—When -keepFor value is not specified and allow auto-delete is requested Same as source—Keep the destination retention policy the same as the source retention policy i NOTE: This attribute is valid for asynchronous replications only.
Hourly snapshot allow auto-delete	Whether or not the destination pool's auto-delete policy allows replicated hourly snapshots on the destination to be deleted. Output can be: <ul style="list-style-type: none"> Blank when scheduled snapshots are not replicated. Same as source—Keep the destination retention policy the same as the source retention policy yes—When -allowAutoDelete is set no—When -keepFor is set i NOTE: This attribute is valid for asynchronous replications only.
Daily snapshot keep for	Amount of time to keep replicated daily snapshots on the destination. Output can be: <ul style="list-style-type: none"> Blank when scheduled snapshots are not replicated. value—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. qualifier—A value qualifier. The valid values are: <ul style="list-style-type: none"> h (hours) d (days) Same as source—Keep the destination retention policy the same as the source retention policy i NOTE: This attribute is valid for asynchronous replications only.

Table 8. Replication session attributes (continued)

Attribute	Description
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"> Blank when scheduled snapshots are not replicated. Same as source—Keep the destination retention policy the same as the source retention policy yes—When <code>-allowAutoDelete</code> is set no—When <code>-keepFor</code> is set <p>NOTE: 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"> yes—Allow snap replication in asynchronous sessions. no—Disallow snap replication in asynchronous sessions.
Cascade Replicated Snap	<p>Whether or not to replicate snap in the cascade replication session.</p> <ul style="list-style-type: none"> yes—When snap is replicated from the inbound session, it will be automatically replicated in this replication session. no—When snap is replicated from the inbound session, it will not be automatically replicated in this replication session. <p>NOTE: This qualifier is used for file system asynchronous replications only.</p>

Create replication sessions

Ensure you have completed the necessary configuration tasks before creating a replication session.

For asynchronous replication, you can replicate existing snapshots and snapshots that are created from snapshot schedules.

For synchronous file replication, you cannot replicate existing snapshots or snapshots that are created from snapshot schedules.

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.

The following conditions apply to snapshot replication:

- Asynchronous replication: You can replicate existing snapshots and snapshots that are created from snapshot schedules.
- Synchronous file replication: You cannot replicate existing snapshots or snapshots that are created from snapshot schedules. You can only replicate those snapshots and snapshots that are 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.

Performing a 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 perform LUN or file system replication from OE version 4.0, 4.1, 4.2, 4.3, or 4.4 to OE version 5.0 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 co-existing asynchronous and synchronous replication with one source resource topology, create the asynchronous replication destination NAS server with both the `-rep1Dest` 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}] [enableUserSnapCommonBase {-autosearch | -srcCommonBaseSnap <value> -dstCommonBaseSnap <value> [-skipSnapSignatureCheck]}]
```

Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.
-srcRes	Type the ID of the storage resource to use as the source.
-srcSPAInterface	Type the ID of the interface on the SPA of the source system for the replication.  NOTE: 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.  NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstType	Specify the type of destination. Valid values are: <ul style="list-style-type: none"> local—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. remote—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 if there is catastrophic data loss, for example, due to natural disaster at the local site.
-dstSys	For remote replication, type the ID of the destination system. View settings for remote storage systems explains how to view the IDs of the remote system configuration on the local system.
-dstRes	Type the ID of the destination storage resource.
-dstSPAInterface	Type the ID of the interface on the SPA of the destination system for the replication.  NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPBInterface	Type the ID of the interface on the SPB of the destination system for the replication.  NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-syncType	This parameter indicates whether to perform the first replication sync (a full sync) automatically. The following example shows the manual option used for the first replication sync: <code>-syncType {manual [-autoInitiate {yes no}] auto [-rpo <value>]}</code> Valid values are: <ul style="list-style-type: none"> manual—Data is replicated when manually initiated. -autoInitiate—The values are for -autoInitiate are as follows:

Qualifier	Description
	<ul style="list-style-type: none"> o yes enables automatic synchronization. o no disables automatic synchronization. • auto [-rpo <value>] —Data is replicated automatically in order to maintain the wanted Recovery Point Objective (RPO). <p>i NOTE: This qualifier is used for asynchronous replications only.</p> <p>As the RPO increases, the potential for data loss also increases, and the amount of required protection space. Lowering the RPO reduces the amount of potential data loss, but also increases network traffic and may negatively impact performance. The default RPO is one hour.</p>
-autoInitiate	<p>Specify whether the system performs the first replication synchronization automatically. Valid values are:</p> <ul style="list-style-type: none"> • yes • no <p>i NOTE: This qualifier is used for asynchronous replications only.</p>
-rpo	<p>Type the time interval for when the synchronization runs. Use the following format: <HH>[:MM]</p> <p>Where:</p> <ul style="list-style-type: none"> • HH—Type the number of hours. Range is 00-24 hours (1 day). • MM—Type the number of minutes, in 5 minute increments. Range is 05 to 55. <p>For synchronous replication, specify the value 0. Once it is set, the value cannot be reset from zero to nonzero or from nonzero to zero.</p>
-replicateHourlySnaps	<p>Specify whether to mark hourly scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> • yes • no <p>i NOTE: This qualifier is used for asynchronous replications only.</p>
-replicateDailySnaps	<p>Specify whether to mark daily scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> • yes • no <p>i NOTE: This qualifier is used for asynchronous replications only.</p>
-keepSameAsSource	<p>Indicate whether to use the same retention policy (expiration time and auto-delete) of the source for the destination. This option propagates changes that are 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> <p>This option is enabled by default if -keepFor or -allowAutoDelete are not set.</p> <p>i NOTE: 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> <p><value><qualifier></p> <p>Where:</p> <ul style="list-style-type: none"> • value—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. • qualifier—A value qualifier. The valid values are: <ul style="list-style-type: none"> o h (hours) o d (days) <p>i NOTE: This qualifier is used for asynchronous replications only.</p>

Qualifier	Description
<code>-allowAutoDelete</code>	Specify whether auto delete is allowed on the replicated copy of this snapshot or snapshot set.  NOTE: This qualifier is used for asynchronous replications only.
<code>-replicateExistingSnaps</code>	Indicates whether to replicate snapshots that exist on the source resource. This qualifier is a one-time option available during session creation that replicates snapshots existing on the source at that moment in time.  NOTE: This qualifier is used for asynchronous replications only. Only the user snapshots that are created after the common base snapshots were established and are in the Idle or Pending state are replicated. All eligible snapshots are replicated with 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: <ul style="list-style-type: none"> • The snapshot is created by either the user or a snapshot schedule. • The snapshot is read-only (file resource snapshot must be a checkpoint snapshot; block resource snapshot must not be attached). • The snapshot is not currently undergoing deletion.
<code>-allowAsyncSnapReplication</code>	Indicate whether snapshot replication is allowed in an asynchronous session. Valid values are: <ul style="list-style-type: none"> • yes • no  NOTE: This qualifier is used for file system asynchronous replications only.
<code>-enableUserSnapCommonBase</code>	Designate the user snapshot as the common base for Unity systems running OE version 5.1. If this qualifier is absent, the system does not treat the user snapshot as the common base.  NOTE: This qualifier must be set when the replication session is configured.
<code>-srcCommonBaseSnap</code>	Specify the user snapshot identifier that is used as the common base on the source resource for Unity systems running OE version 5.1.
<code>-dstCommonBaseSnap</code>	Specify the user snapshot identifier that is used as the common base on the destination resource for Unity systems running OE version 5.1.
<code>-autosearch</code>	Enable the system to perform an automatic search to find a user snapshot that can be used as a common base for Unity systems running OE version 5.1. If this qualifier is not included, the <code>-autosearch</code> function is disabled.  NOTE: This qualifier must be set when the replication session is configured if you want to avoid a full copy after a failover.
<code>-skipSnapSignatureCheck</code>	Indicate whether you want to skip the user snap signature check when using the user snapshot as the common base for Unity systems running OE version 5.1. This qualifier is valid only when specifying the source and destination user snapshots as the common base.  NOTE: Check if the user snapshots are qualified as a common base.
<code>-cascadeReplicatedSnap</code>	Indicates whether you want to automatically replicate snapshots downstream when the snapshot is replicated from the inbound replication session. The valid values are as follows: <ul style="list-style-type: none"> • yes • no (default)

Qualifier	Description
-overwriteDestination	Indicate whether to overwrite dirty data on the destination; it takes more time to complete full data transfer. If this qualifier is not specified, the destination might contain different data than the source if the destination has corrupted data.

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 REPl
-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_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_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              =
81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000
      Name              = MyRep1
      Session type      = block
```

```

Synchronization type = manual
RPO                  =
Sync State           = idle
Health state         = OK (5)
Operational status   = Idle (0x8404)
Time of last sync    =
Sync status          =
Element pairs        = sv_4 => sv_4

```

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
42949672965_FCNCH0972DDD67_0000_42949672967_FCNCH0972DDD67_0000
      Name                               = rep_sess_sv_1_sv_2_local
      Session type                       = lun
      Synchronization type               = manual
      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             = not replicated
      Daily snapshot allow auto-delete    = not replicated
      Allow Async Snap Replication        = N/A
      Cascade Replicated Snap             = no

```

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 [-cascadeReplicatedSnap {yes | no}] | 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. <i>i</i> NOTE: 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. <i>i</i> NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPAInterface	Type the ID of the interface on the SPA of the destination system for the replication. <i>i</i> NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPBInterface	Type the ID of the interface on the SPB of the destination system for the replication. <i>i</i> NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-syncType	This parameter indicates whether to perform the first replication sync (a full sync) automatically. The following example shows the manual option used for the first replication sync: <code>-syncType {manual auto [-rpo <value>]}</code> Valid values are: <ul style="list-style-type: none"> <code>manual</code>—Data is replicated when manually initiated. <code>auto [-rpo <value>]</code>—Data is replicated automatically in order to maintain the wanted Recovery Point Objective (RPO). <i>i</i> NOTE: This qualifier is used for asynchronous replications only. As the RPO increases, the potential for data loss also increases, and the amount of required protection space. Lowering the RPO reduces the amount of potential data loss, but also increases network traffic and may negatively impact performance. The default RPO is one hour.
-rpo	For automatic synchronization, type the time interval for when the synchronization will run. Use the following format: <HH>[:MM] Where: <ul style="list-style-type: none"> <code>HH</code>—Type the number of hours. Range is 00-24 hours (1 day). <code>MM</code>—Type the number of minutes, in 5 minute increments. Range is 05 to 55.

Qualifier	Description
	<p> NOTE: For synchronous replication, specify the value 0. The value cannot be reset from zero to nonzero or from nonzero to zero.</p>
<code>-replicateHourlySnaps</code>	<p>Specify whether to mark hourly scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> • yes • no <p> NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-replicateDailySnaps</code>	<p>Specify whether to mark daily scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> • yes • no <p> NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-keepSameAsSource</code>	<p>Specify whether to use the same retention policy (expiration time and auto-delete) of the source for the destination. This option propagates changes that are 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> <p> NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-keepFor</code>	<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> <p><code><value><qualifier></code></p> <p>Where:</p> <ul style="list-style-type: none"> • <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. • <i>qualifier</i>—A value qualifier. The valid values are: <ul style="list-style-type: none"> ◦ h (hours) ◦ d (days) <p> NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-allowAutoDelete</code>	<p> NOTE: Only valid when <code>-replicateHourlySnaps yes</code> or <code>-replicateDailySnaps yes</code>.</p> <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"> • yes • no <p> NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-allowAsyncSnapReplication</code>	<p>Indicates whether to allow snapshot replication in asynchronous session. Valid values are:</p> <ul style="list-style-type: none"> • yes • no
<code>-cascadeReplicatedSnap</code>	<p>Indicates whether to replicate snapshots in the cascade replication session. This qualifier can only be modified when the session is active. Valid values are:</p> <ul style="list-style-type: none"> • yes • no (default)

Example

The following command changes the source interface and destination interface for replication session 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000:

```
uemcli /prot/rep/session -id
81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_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_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000
Operation completed successfully.
```

Pause replication sessions

Pause a replication session or sessions at the remote system level.

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 each replication session to be paused.

Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.

Example

The following command pauses replication session
81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000:

```
uemcli /prot/rep/session -id
81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_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 or sessions at the remote level.








Format

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


Object qualifiers

Qualifier	Description
-id	Type the ID of the replication session that you want to resume.
-name	Type the name of each replication session that you want to resume.

Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.
-allowFullCopy	<p>This qualifier specifies that a full synchronization can occur if there is no common base snapshot when an asynchronous replication session is resumed.</p> <p>The -allowFullCopy option is not supported when a synchronous replication session resumes after a pause operation. The full copy is performed with the same behavior with -allowFullCopy specified for synchronous file replication if the session has no common base snapshots, and neither -allowFullCopy nor -forceFullCopy is specified.</p> <p>If the source system has snapshots that were created before you upgraded your system to Unity 5.1.2 and those snapshots are used as common base snapshots, those snapshots trigger a full synchronization. Beginning with Unity 5.1.2, common base snapshots have signatures that allow delta synchronization to occur.</p> <p> NOTE: This qualifier does not support synchronous block replications.</p>
-forceFullCopy	<p>This qualifier specifies that a full synchronization occurs when an asynchronous replication session is resumed.</p> <p>The --forceFullCopy option is not supported when a synchronous replication session resumes after a pause operation.</p> <p> NOTE: This qualifier does not support synchronous block replications.</p>
-srcSPAInterface	<p>Type the ID of the interface on the SPA of the source system for the replication.</p> <p> NOTE: 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> <p> NOTE: 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	<p>Type the ID of the interface on the SPB of the source system for the replication.</p> <p> NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
-dstSPBInterface	<p>Type the ID of the interface on the SPB of the destination system for the replication.</p> <p> NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.</p>
-forceSyncData	<p>This qualifier forces a data transfer from the source system to the remote system. This transfer occurs even if the remote system has data that is not replicated from the source system.</p> <p> NOTE: This qualifier is used for asynchronous replications only.</p>

Example

The following command resumes the replication session 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000:

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

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

```
ID = 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_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
-id	Type the ID of the replication session to synchronize.
-name	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_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session -id 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_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](#) 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
-id	Type the ID of the replication session to fail over.
-name	Type the name of the replication session to fail over.

Action qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.
-sync	<p>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:</p> <ul style="list-style-type: none">• yes—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>.• no—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>. <p> NOTE: 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>
-force	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_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /prot/rep/session -id 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_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 resumes and you may connect your hosts to the source storage resource.

 **NOTE:** Ensure that hosts do not write to the destination storage resource, which becomes read-only.




Format

```
/prot/rep/session {-id <value> | -name <value>} failback [-async] [-forceFullCopy] [-allowFullCopy] [-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.
-allowFullCopy	<p>This qualifier specifies that a full synchronization can occur after a failback if there is no common base snapshot when an asynchronous replication session is resumed.</p> <p>The -allowFullCopy option is not supported when a synchronous replication session resumes after a pause operation. The full copy is performed with the same behavior with -allowFullCopy specified for synchronous file replication if the session has no common base snapshots, and neither -allowFullCopy nor --forceFullCopy is specified.</p> <p>If the source system has snapshots that were created before you upgraded your system to Unity 5.1.2 and those snapshots are used as common base snapshots, those snapshots trigger a full synchronization. Beginning with Unity 5.1.2, common base snapshots have signatures that allow delta synchronizations to occur.</p> <p> NOTE: This qualifier does not support synchronous block replications.</p>
-forceFullCopy	<p>This qualifier specifies that a full synchronization occurs when an asynchronous replication session is resumed.</p> <p>The -forceFullCopy option is not supported when a synchronous replication session resumes after a pause operation.</p> <p> NOTE: This qualifier does not support synchronous block replications.</p>
-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 fails if the remote system has data that are not replicated from the source system. Valid values are:</p> <ul style="list-style-type: none">• force – Force data transfer from the source system to the remote system, even if the remote system has out-of-sync data.• ignore – Skip data transfer from the source system to the remote system and trigger the replication session reverse from the remote system to the source system. <p> NOTE: This qualifier is used for asynchronous replications only.</p>

Example

The following command performs a fail back of replication session 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_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 [-dstSys <value>]
```

Object qualifiers

Qualifier	Description
-id	Identifies the NAS server synchronous replication session.
-name	Identifies the NAS server synchronous replication session by name.
-dstSys	Identifies the remote system which is selected as the NAS server asynchronous replication session destination system for preservation. If this qualifier is not specified, all async-type remote systems are checked, and the first one that has a matching asynchronous replication session (having the same NAS server ID as the synchronous replication session) is selected for preservation.

Example

The following command preserves asynchronous replication sessions for 81604378625_FCNCH097274B3A_0000_81604378627_FCNCH097274B37_0000:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /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.
```

Enable asynchronous snapshot replication after a replication session is created

You can enable asynchronous snapshot replication in an existing asynchronous replication session when that session is in the Auto Configured or Idle state.

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 {no}]
```

Object qualifiers

Qualifier	Description
-id	Type the ID of the replication session.
-name	Type the name of the replication session to change.
-srcSPAInterface	Type the ID of the interface on the SPA of the source system for the replication. i NOTE: 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. i NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPAInterface	Type the ID of the interface on the SPA of the destination system for the replication. i NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPBInterface	Type the ID of the interface on the SPB of the destination system for the replication. i NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-syncType	Specify how the source and destination synchronize. Valid values are: <ul style="list-style-type: none">• manual—Data is replicated when manually initiated.• auto—Data is replicated automatically in order to maintain the wanted Recovery Point Objective (RPO). i NOTE: This qualifier is used for asynchronous replications only. As the RPO increases, the potential for data loss also increases, and the amount of required protection space. Lowering the RPO reduces the amount of potential data loss, but also increases network traffic and may negatively impact performance. The default RPO is one hour.
-rpo	Type the time interval for when the synchronization runs. Use the following format: <HH> [:MM] where: <ul style="list-style-type: none">• HH—Type the number of hours. Range is 00-24 hours (1 day).• MM—Type the number of minutes, in 5 minute increments. Range is 05–55. For synchronous replication, specify the value 0. Once the value is set, it cannot be reset from zero to nonzero or from nonzero to zero.
-replicateHourlySnaps	Specify whether to mark hourly scheduled snapshots for replication. Valid values are: <ul style="list-style-type: none">• Yes• No

Qualifier	Description
	<p>You can keep hourly snapshots for the same period as specified for the source or set a specific retention time.</p> <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-keepSameAsSource</code>	<p>Specify whether to use the same retention policy (expiration time and auto-delete) of the source for the destination. This option propagates changes that are 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> <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-keepFor</code>	<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: <code><value><qualifier></code></p> <p>Where:</p> <ul style="list-style-type: none"> <code>value</code>—An integer value. If the qualifier is h (hours), the valid range is from 1 to 42840. If the qualifier is d (days), the valid range is 1–1785. <code>qualifier</code>—A value qualifier. The valid values are: <ul style="list-style-type: none"> h (hours) d (days) <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-allowAutoDelete</code>	<p>NOTE: Only valid when <code>-replicateHourlySnaps yes</code> or <code>-replicateDailySnaps yes</code>.</p> <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"> Yes No <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-replicateDailySnaps</code>	<p>Specify whether to mark daily scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> Yes No <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-allowAsyncSnapReplication</code>	<p>Indicate whether snapshot replication is allowed in an asynchronous session. Valid values are:</p> <ul style="list-style-type: none"> Yes No <p>NOTE: This qualifier is used for file system asynchronous replications only.</p>

Disable asynchronous snapshot replication after a replication session is created

You can disable asynchronous snapshot replication in an existing asynchronous replication session when that session is in the Auto Configured or Idle state.

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 {no}]
```

Object qualifiers

Qualifier	Description
-id	Type the ID of the replication session.
-name	Type the name of the replication session to change.
-srcSPAInterface	Type the ID of the interface on the SPA of the source system for the replication. <i>NOTE:</i> 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. <i>NOTE:</i> This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPAInterface	Type the ID of the interface on the SPA of the destination system for the replication. <i>NOTE:</i> This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPBInterface	Type the ID of the interface on the SPB of the destination system for the replication. <i>NOTE:</i> This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-syncType	This parameter indicates whether to perform the first replication sync (a full sync) automatically. The following example shows the manual option used for the first replication sync: <code>-syncType {manual [-autoInitiate {yes no}] auto [-rpo <value>]}</code> Valid values are: <ul style="list-style-type: none"> manual—Data is replicated when manually initiated. auto [-rpo <value>]—Data is replicated automatically in order to maintain the wanted Recovery Point Objective (RPO). <i>NOTE:</i> This qualifier is used for asynchronous replications only. As the RPO increases, the potential for data loss also increases, and the amount of required protection space. Lowering the RPO reduces the amount of potential data loss, but also increases network traffic and may negatively impact performance. The default RPO is one hour.
-rpo	Type the time interval for when the synchronization runs. Use the following format: <code><HH> [:MM]</code> where: <ul style="list-style-type: none"> HH—Type the number of hours. Range is 00-24 hours (1 day). MM—Type the number of minutes, in 5 minute increments. Range is 05-55. For synchronous replication, specify the value 0. Once the value is set, it cannot be reset from zero to nonzero or from nonzero to zero.
-replicateHourlySnaps	Specify whether to mark hourly scheduled snapshots for replication. Valid values are: <ul style="list-style-type: none"> Yes No You can keep hourly snapshots for the same period as specified for the source or set a specific retention time. <i>NOTE:</i> This qualifier is used for asynchronous replications only.

Qualifier	Description
<code>-keepSameAsSource</code>	Specify whether to use the same retention policy (expiration time and auto-delete) of the source for the destination. This option propagates changes that are 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. i NOTE: This qualifier is used for asynchronous replications only.
<code>-keepFor</code>	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: <code><value><qualifier></code> Where: <ul style="list-style-type: none"> <code>value</code>—An integer value. If the qualifier is h (hours), the valid range is from 1 to 42840. If the qualifier is d (days), the valid range is 1–1785. <code>qualifier</code>—A value qualifier. The valid values are: <ul style="list-style-type: none"> h (hours) d (days) i NOTE: This qualifier is used for asynchronous replications only.
<code>-allowAutoDelete</code>	i NOTE: Only valid when <code>-replicateHourlySnaps yes</code> or <code>-replicateDailySnaps yes</code> . Specify whether auto delete is allowed on the replicated copy of this snapshot or snapshot set. Valid values are: <ul style="list-style-type: none"> Yes No i NOTE: This qualifier is used for asynchronous replications only.
<code>-replicateDailySnaps</code>	Specify whether to mark daily scheduled snapshots for replication. Valid values are: <ul style="list-style-type: none"> Yes No i NOTE: This qualifier is used for asynchronous replications only.
<code>-allowAsyncSnapReplication</code>	Indicate whether snapshot replication is allowed in an asynchronous session. Valid values are: <ul style="list-style-type: none"> Yes No i NOTE: This qualifier is used for file system asynchronous replications only.

Cascade a replicated snapshot in an asynchronous replication session

In cascade topologies, user snapshots are synchronized in asynchronous replication sessions. However, user snapshots are not automatically replicated downstream. You must manually enable downstream replication.

Format

```
/prot/rep/session {-id <value> | -name <value>} set [-async] ... [-allowAsyncSnapReplication {yes [-cascadeReplicatedSnap {yes | no}] | no}] [-enableUserSnapCommonBase {-autoSearch | -srcCommonBaseSnap <value> -dstCommonBaseSnap <value> [ -skipSnapSignatureCheck
```

Object qualifiers

Qualifier	Description
-allowAsyncSnapReplication	Indicate whether snapshot replication is allowed in an asynchronous session. Valid values are: <ul style="list-style-type: none">• Yes• No NOTE: This qualifier is used for file system asynchronous replications only.
-cascadeReplicatedSnap	Indicate whether to replicate snapshots in the cascade replication session. Valid values are: <ul style="list-style-type: none">• Yes• No NOTE: This qualifier is used for file system asynchronous replications only.
-enableUserSnapCommonBase	Designate the user snapshot as the common base for Unity systems running OE version 5.1. If this qualifier is absent, the system does not treat the user snapshot as the common base. NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.
-autosearch	Enable the system to perform an automatic search to find the user snapshot was designated as the common base snapshot for Unity systems running OE version 5.1. If this qualifier is not included, the -autosearch function is disabled. NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.
-srcCommonBaseSnap	Specify the user snapshot identifier that is used as the common base on the source resource for Unity systems running OE version 5.1. NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.
-dstCommonBaseSnap	Specify the user snapshot identifier that is used as the common base on the destination resource for Unity systems running OE version 5.1. NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.
-skipSnapSignatureCheck	Indicate whether you want to skip the user snap signature check when using the user snapshot as the common base for Unity systems running OE version 5.1. This qualifier is valid only when specifying the source and destination user snapshots as the common base. NOTE: Confirm that the user snapshots are qualified to use as a common base.

Provision an asynchronous replication session to autoreplicate user snapshots





When you create an asynchronous replication session, you can provision that session to autoreplicate scheduled user snapshots.

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
[-cascadeReplicatedSnap {yes | no}} | no}} [-enableUserSnapCommonBase { -autoSearch |
-srcCommonBaseSnap <value> -dstCommonBaseSnap <value> [ -skipSnapSignatureCheck ] } ] [ -
overwriteDestination ]
```

Object qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.
-srcRes	Type the ID of the storage resource to use as the source.
-dstType	Specify the type of destination. Valid values are: <ul style="list-style-type: none"> local—Maintain a full copy of the storage resource on the local system. This option has advantages over snapshots in that a full copy, not just a copy of changes, is retained. remote—Maintain a full copy of the storage resource on a remote system by transferring the data over the network. If there is catastrophic data loss, remote replication is often used to ensure that a copy is available at a remote site. For example, a catastrophic loss might be the result of a natural disaster at the local site.
-dstSys	For remote replication, type the ID of the destination system. View settings for remote storage systems explains how to view the IDs of the remote system configuration on the local system.
-dstRes	Type the ID of the destination storage resource.
-srcSPAInterface	Type the ID of the interface on the SPA of the source system for the replication.  NOTE: 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.  NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPAInterface	Type the ID of the interface on the SPA of the destination system for the replication.  NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-dstSPBInterface	Type the ID of the interface on the SPB of the destination system for the replication.  NOTE: This qualifier is used for asynchronous replications on remote systems only. If the qualifier is not specified, the system identifies the interface automatically.
-name	Type the name of the replication session to change.
-elementPairs	A list of mappings between source and destination including all storage elements. This applies to consistency groups only. Use the following format: source_id_1:dest_id_1,source_id_2:dest_id_2 For consistency group and VMware vStorage VMFS datastore replications, this qualifier specifies the LUN element pairs within the replication.
-syncType	This parameter indicates whether to perform the first replication sync (a full sync) automatically. The following example shows the manual option used for the first replication sync: <code>-syncType {manual [-autoInitiate {yes no}} auto [-rpo <value>}</code>

Qualifier	Description
	<p>Valid values are:</p> <ul style="list-style-type: none"> • <code>manual</code>—Data is replicated when manually initiated. • <code>-autoInitiate</code>—The values are for <code>-autoInitiate</code> are as follows: <ul style="list-style-type: none"> ◦ <code>yes</code> enables automatic synchronization. ◦ <code>no</code> disables automatic synchronization. • <code>auto [-rpo <value>]</code>—Data is replicated automatically in order to maintain the wanted Recovery Point Objective (RPO). <p>NOTE: This qualifier is used for asynchronous replications only.</p> <p>As the RPO increases, the potential for data loss also increases, and the amount of required protection space. Lowering the RPO reduces the amount of potential data loss, but also increases network traffic and may negatively impact performance. The default RPO is one hour.</p>
<code>-replicateHourlySnaps</code>	<p>Specify whether to mark hourly scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> • <code>Yes</code> • <code>No</code> <p>You can keep hourly snapshots for the same period as specified for the source or set a specific retention time.</p> <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-keepSameAsSource</code>	<p>Specify whether to use the same retention policy (expiration time and autodelete) of the source for the destination. This option propagates changes that are 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> <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-keepFor</code>	<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: <code><value><qualifier></code></p> <p>Where:</p> <ul style="list-style-type: none"> • <code>value</code>—An integer value. If the qualifier is h (hours), the valid range is from 1 to 42840. If the qualifier is d (days), the valid range is 1–1785. • <code>qualifier</code>—A value qualifier. The valid values are: <ul style="list-style-type: none"> ◦ h (hours) ◦ d (days) <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-allowAutoDelete</code>	<p>NOTE: Only valid when <code>-replicateHourlySnaps yes</code> or <code>-replicateDailySnaps yes</code>.</p> <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"> • <code>Yes</code> • <code>No</code> <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-replicateDailySnaps</code>	<p>Specify whether to mark daily scheduled snapshots for replication. Valid values are:</p> <ul style="list-style-type: none"> • <code>Yes</code> • <code>No</code> <p>NOTE: This qualifier is used for asynchronous replications only.</p>
<code>-replicateExistingSnaps</code>	<p>Indicate whether to replicate snapshots that exist on the source resource. This qualifier is a one-time option available during session creation that replicates snapshots existing on the source at that moment in time.</p>

Qualifier	Description
	<p> NOTE: This qualifier is used for asynchronous replications only.</p> <p>Only the user snapshots that are created after the common base snapshots were established and are in the Idle or Pending state are replicated.</p> <p>All eligible snapshots are replicated with the source retention policy that is applied for the destination retention policy. For a snapshot to be eligible for this option, it must meet these three criteria:</p> <ul style="list-style-type: none"> • The snapshot is created by either the user or a snapshot schedule. • The snapshot is read-only (file resource snapshot must be a checkpoint snapshot; block resource snapshot must not be attached). • The snapshot is not currently undergoing deletion.
-allowAsyncSnapReplication	<p>Indicate whether snapshot replication is allowed in an asynchronous session. Valid values are:</p> <ul style="list-style-type: none"> • Yes • No <p> NOTE: This qualifier is used for file system asynchronous replications only.</p>
-cascadeReplicatedSnap	<p>Indicate whether to replicate snapshots in the cascade replication session. Valid values are:</p> <ul style="list-style-type: none"> • Yes • No <p> NOTE: This qualifier is used for file system asynchronous replications only.</p>
-rpo	<p>Type the time interval for when the synchronization runs. Use the following format: <HH> [:MM]</p> <p>where:</p> <ul style="list-style-type: none"> • HH—Type the number of hours. Range is 00-24 hours (1 day). • MM—Type the number of minutes, in 5 minute increments. Range is 05-55. <p>For synchronous replication, specify the value 0. Once the value is set, it cannot be reset from zero to nonzero or from nonzero to zero.</p>
-enableUserSnapCommonBase	<p>Designate the user snapshot as the common base for Unity systems running OE version 5.1 and later. If this qualifier is absent, the system does not treat the user snapshot as the common base.</p> <p> NOTE: This qualifier must be set when the replication session is configured.</p>
-autoSearch	<p>Enable the system to perform an automatic search to find a user snapshot that can be used as a common base for Unity systems running OE version 5.1 and later. If this qualifier is not included, the -autosearch function is disabled.</p> <p> NOTE: This qualifier must be set when the replication session is configured if you want to avoid a full copy after a failover.</p>
-srcCommonBaseSnap	<p>Specify the user snapshot identifier that is used as the common base on the source resource for Unity systems running OE version 5.1 and later.</p>
-dstCommonBaseSnap	<p>Specify the user snapshot identifier that is used as the common base on the destination resource for Unity systems running OE version 5.1 and later.</p>
-skipSnapSignatureCheck	<p>Indicate whether you want to skip the user snap signature check when using the user snapshot as the common base for Unity systems running OE version 5.1 and later. This qualifier is valid only when specifying the source and destination user snapshots as the common base.</p>
-overwriteDestination	<p>Indicate whether to overwrite dirty data on the destination; it takes more time to complete a full data transfer. If this qualifier is not specified, the destination might contain different data than the source if the destination has corrupted data.</p>

Common base snapshots

Referencing a common base snapshot enables you to avoid a full copy when performing a failback operation to the original source after an unplanned failover.


If an unplanned failover interrupts a file system replication session, the common base snapshot enables you to re-create that session.

Snapshot precheck attributes


Attribute	Description
Session ID	ID of the file synchronous replication session.
Session name	Name of the file synchronous replication session.
Source common base snapshot ID	ID of the common base snapshot on the source site.
Source common base snapshot name	Name of the common base snapshot on the source site.
Destination common base snapshot ID	ID of the common base snapshot on the destination site.
Destination common base snapshot name	Name of the common base snapshot on the destination site.

Run a precheck for existing common base snapshots

Find common base snapshots on a NAS server or in a file replication session before you perform a failback operation. Using existing common base snapshots allows you to use the avoid full copy feature and perform a delta synchronization. This type of synchronization sends only data that has changed after the failover back to the source system.

 **NOTE:** This command applies only to file synchronous replication.

Common base snapshots can be snapshots that were created according to an established schedule and replicated to the destination. Common base snapshots can also be manually created on the source and replicated to the destination system.

 **NOTE:** The precheck can only be performed with the CLI. It cannot be performed in the Unisphere UI.

Format

```
/prot/rep/session/commonbase {-session <value> | -sessionName <value>} show -detail
```

Object qualifiers

Qualifier	Description
-session	Identifies the replication session by ID number.
-sessionName	Identifies the replication session by an assigned name.
-detail	Show snapshot details.

Example

The following example shows a check for common base snapshots in a synchronous file replication session.

```
uemcli /prot/rep/session/commonbase -session  
171798691845_FCNCH097274B3A_0000_171798691846_FCNCH097274B37_0000 show -detail  
[Response]
```

```
Storage system address: 10.64.75.201
Storage system port: 443
HTTPS connection

1:      Session ID                               =
171798691845_FCNCH097274B3A_0000_171798691846_FCNCH097274B37_0000
      Session name                               = Rep1
      Source common base snapshot ID             = 171798691846
      Source common base snapshot name           = Snap1
      Destination common base snapshot ID        = 171798691848
      Destination common base snapshot name      = Snap1

ID = 171798691845_FCNCH097274B3A_0000_171798691846_FCNCH097274B37_0000
Operation completed successfully.
```

Re-create a replication session

If the source and destination systems share a common base snapshot, you can re-create a replication session.

Both synchronous and asynchronous replication sessions can be re-created. The following sessions of each type can be re-created:

- Synchronous file system sessions
- Asynchronous file system sessions
- Asynchronous block sessions

The following types of synchronous and asynchronous replication sessions cannot be re-created:







- Synchronous block sessions
- Synchronous NAS server sessions
- Asynchronous NAS server sessions

Format

```
/prot/rep/session create [-async] -srcRes <value> -dstType remote -dstSys <value>
-dstRes <value> -syncType {manual [-autoInitiate {yes | no}] | auto [-rpo <value>]]
-enableUserSnapCommonBase -autoSearch -srcCommonBaseSnap <value> -dstCommonBaseSnap <value>
-skipSnapSignatureCheck
```


Object qualifiers

Qualifier	Description
-async	Run an action in asynchronous mode.
-srcRes	Type the ID of the storage resource to use as the source.
-dstType	Specify the type of destination. Valid values are: <ul style="list-style-type: none"> • <code>local</code>—Maintain a full copy of the storage resource on the local system. This option has advantages over snapshots in that a full copy, not just a copy of changes, is retained. • <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 if there is catastrophic data loss. For example, a catastrophic loss might be caused by a natural disaster at the local site.
-dstSys	For remote replication, type the ID of the destination system. View settings for remote storage systems explains how to view the IDs of the remote system configuration on the local system.
-dstRes	Type the ID of the destination storage resource.

Qualifier	Description
-syncType	<p>This parameter indicates whether to perform the first replication sync (a full sync) automatically. The following example shows the manual option used for the first replication sync: <code>-syncType {manual [-autoInitiate {yes no}] auto [-rpo <value>]}</code></p> <p>Valid values are:</p> <ul style="list-style-type: none"> • manual—Data is replicated when manually initiated. • -autoInitiate—The values are for -autoInitiate are as follows: <ul style="list-style-type: none"> ◦ yes enables automatic synchronization. ◦ no disables automatic synchronization. • auto [-rpo <value>]—Data is replicated automatically in order to maintain the wanted Recovery Point Objective (RPO). <p> NOTE: This qualifier is used for asynchronous replications only.</p> <p>As the RPO increases, the potential for data loss also increases, and the amount of required protection space. Lowering the RPO reduces the amount of potential data loss, but also increases network traffic and may negatively impact performance. The default RPO is one hour.</p>
-rpo	<p>Type the time interval for when the synchronization runs. Use the following format: <code><HH> [:MM]</code></p> <p>where:</p> <ul style="list-style-type: none"> • HH—Type the number of hours. Range is 00-24 hours (1 day). • MM—Type the number of minutes, in 5 minute increments. Range is 05–55. <p>For synchronous replication, specify the value 0. Once the value is set, it cannot be reset from zero to nonzero or from nonzero to zero.</p>
-enableUserSnapCommonBase	<p>Designate the user snapshot as the common base for Unity systems running OE version 5.1.</p> <p>If this qualifier is absent, the system does not treat the user snapshot as the common base.</p> <p> NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.</p>
-autosearch	<p>Enable the system to perform an automatic search to find the user snapshot was designated as the common base snapshot for Unity systems running OE version 5.1.</p> <p>If this qualifier is not included, the <code>-autosearch</code> function is disabled.</p> <p> NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.</p>
-srcCommonBaseSnap	<p>Specify the user snapshot identifier that is used as the common base on the source resource for Unity systems running OE version 5.1.</p> <p> NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.</p>
-dstCommonBaseSnap	<p>Specify the user snapshot identifier that is used as the common base on the destination resource for Unity systems running OE version 5.1.</p> <p> NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.</p>
-skipSnapSignatureCheck	<p>Indicate whether you want to skip the user snap signature check when using the user snapshot as the common base for Unity systems running OE version 5.1. This qualifier is valid only when specifying the source and destination user snapshots as the common base.</p> <p> NOTE: Confirm that the user snapshots are qualified to use as a common base.</p>

Create a replication session with a user snapshot as a common base





When you create a replication session with a user snapshot as a common base, you can re-create the session if an unplanned failover occurs. Using the common base snapshot enables you to avoid a full copy when re-creating either an asynchronous or synchronous replication session.

 **NOTE:** Using a common base user snapshot to avoid a full copy works only with file replication.

Format

```
/prot/rep/session create [-async] ... [-enableUserSnapCommonBase {-srcCommonBaseSnap <value>
-dstCommonBaseSnap <value> -skipSnapSignatureCheck
```

Object qualifiers

Qualifier	Description
-enableUserSnapCommonBase	Designate the user snapshot as the common base for Unity systems running OE version 5.1. If this qualifier is absent, the system does not treat the user snapshot as the common base.  NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.
-srcCommonBaseSnap	Specify the user snapshot identifier that is used as the common base on the source resource for Unity systems running OE version 5.1.  NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.
-dstCommonBaseSnap	Specify the user snapshot identifier that is used as the common base on the destination resource for Unity systems running OE version 5.1.  NOTE: If you want to avoid a full copy when re-creating the session, this qualifier must be set when the replication session is configured.
-skipSnapSignatureCheck	Indicate whether you want to skip the user snap signature check when using the user snapshot as the common base for Unity systems running OE version 5.1. This qualifier is valid only when specifying the source and destination user snapshots as the common base.  NOTE: Confirm that the user snapshots are qualified to use as a common base.

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
<code>-output -o</code>	Specify the output format. Value is one of the following: <ul style="list-style-type: none"><code>nvp</code> — The name-value pair (NVP) format displays output as <code>name=value</code>. Name-value pair format provides an example of the NVP format.<code>table</code> — 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 <code>-wrap</code> after the table qualifier, separated by a space, so that the values wrap. Table format provides an example of the table format.<code>csv</code> — The comma-separated values (CSV) format is similar to the table format, but the names and values are separated by commas. Comma-separated values format provides an example of the CSV format.
<code>-detail</code>	Display all attributes.
<code>-brief</code>	Display only the basic attributes (default).
<code>-filter</code>	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">• <code>mgmt</code> — Management interface.• <code>file</code> — Network interface for Windows (SMB) and Linux/UNIX (NFS) storage.• <code>iscsi</code> — iSCSI interface for iSCSI storage. Manage network interfaces 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.
Available MTU size	List of available MTU sizes.  NOTE: This can display as either a comma-separate 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.
Speed	Current link speed of the port.
Requested speed	Link speed set by the user.
Available speeds	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">• <code>Unknown (0)</code> — Status is unknown.• <code>OK (5)</code> — Port is operating normally.• <code>OK BUT (7)</code> — Lost communication, but the port is not in use.• <code>Minor failure (15)</code> — Lost communication. Check the network connection and connected cables.• <code>Major failure (20)</code> — Port has failed. Replace the SP that contains the port.
Health details	Additional health information. See Appendix A, Reference, for health information details.

Table 9. Network port attributes (continued)

Attribute	Description
Aggregated port ID	If the port is in a link aggregation, the ID of the link aggregation appears. Manage link aggregations explains how to configure link aggregations on the SP ports.
Connector type	Physical connector type. Valid values are: <ul style="list-style-type: none"> • unknown • RJ45 • LC • MiniSAS_HD • CopperPigtail • NoSeparableConnector
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	List of supported protocols of the inserted Small Form-factor Pluggable. Valid values are: <ul style="list-style-type: none"> • unknown • FibreChannel • Ethernet • SAS

View Ethernet port settings

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

 **NOTE:** The `show action command` 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	Type the maximum transmission unit packet size (in bytes) for the port: <ul style="list-style-type: none"> If an Ethernet port carries File interfaces only, the MTU size can be set to a custom value between 1280 and 9216. If an Ethernet port carries iSCSI interfaces, the allowed MTU sizes are 1500 and 9000. 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.
-speed	Type the port speed. Available speed attribute contains the list of valid values. <p>NOTE: The port will be link down temporarily during change of speed.</p>

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">Unknown (0) — Status is unknown.OK (5) — Port is operating normally.OK BUT (7) — Lost communication, but the port is not in use.Minor failure (15) — Lost communication. Check the network connection and connected cables.Major failure (20) — Port has failed. Replace the SP that contains the port.
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">unknownRJ45LCMiniSAS_HDCopperPigtailNoSeparableConnector
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: <ul style="list-style-type: none">unknownFibreChannelEthernetSAS
Replication capability	Type of replication capability. Valid values are: <ul style="list-style-type: none">Sync replicationRecoverPoint

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. Available speed attribute contains the list of valid values.  NOTE: The port will be link down temporarily during change of 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 an SP to a single logical port, allowing you to 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.

An ID is used to identify each link aggregation.

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: Ports must have the same MTU size in order to be aggregated. Linked ports must connect to the same logical switch. The switch must support and be configured to use the IEEE 802.3ad Dynamic Link Aggregation Control Protocol (LACP), which provides fault tolerance and load balancing. In this LACP mode, the system aggregates interfaces into groups that share the same speed and duplex settings. The documentation that came with your switch should provide more information about 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">• SPA• SPB
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. NOTE: This option 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:

Table 11. Link aggregation attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> Unknown (0) —Status is unknown. OK (5) —Working correctly. OK BUT (7) —Lost connection, but the link aggregation is not in use. Degraded/Warning (10) —Working and performing all functions, but the performance may not be optimum. 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. 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. Critical failure (25) —Failed and recovery may not be possible. This condition has resulted in data loss and should be remedied immediately. Non-recoverable error (30) —Completely failed and cannot be recovered.
Health details	More 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 = 1a0_SPB  
Operation completed successfully.
```

Use case for CLI

Topics:

- [Replication configuration use case](#)

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

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.

 **NOTE:** Local replication is applicable to asynchronous replication only. Local replication is not applicable to synchronous replication.

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

Prerequisites

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 a NAS server is being replicated, the destination NAS server must have a matching tenant configuration. For example, you cannot replicate a non-tenanted NAS server to a tenanted NAS server. Tenants must be created on the target system using the same UUID as on the source system.

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.

Steps

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.

5. (Optional) Create one or more bandwidth schedules to control when replication occurs, and how much bandwidth should be used during replication.

Bandwidth schedules are configured to replication connections to the remote system. All asynchronous replication sessions with active data transfer on their replication connections evenly share the bandwidth assigned for the specified days and hours.

NOTE: If no schedule is set, all available bandwidth will be used during replication.

Configure synchronous replication

About this task

Only remote replication is supported for synchronous replication.

If a NAS server is being replicated, the destination NAS server must have a matching tenant configuration. For example, you cannot replicate a non-tenanted NAS server to a tenanted NAS server. Tenants must be created on the target system using the same UUID as on the source system.

Steps

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 that 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

Prerequisites

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.

About this task

Ensure that you create replication interfaces on each SP.

Steps

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::20a8bfff: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
      SP          = spa


3:      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

Prerequisites

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**.

 **NOTE:** If a replication connection already exists and you plan to add a different mode of replication, do not attempt to create a new connection. Change the existing replication connection mode to Both.

About this task

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

Steps

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

Prerequisites

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.

Steps

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

Prerequisites

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.

NOTE: 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.

About this task

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.

Steps

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

Prerequisites

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.

About this task

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

Steps

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 MyNSRep1 -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_FCNCH097274999_0000_103079215106_FCNCH0972749A9_0000
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 MyFSRep1 -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_FCNCH097274999_0000_171798691844_FCNCH0972749A9_0000
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
```