

Dell EMC Unity: Replication Technologies

A Detailed Review

Abstract

This white paper explains the replication solutions for Dell EMC™ Unity systems. This paper outlines the native and non-native options available for replicating data. It also includes information about managing replication, and the benefits that replication provides. Native file synchronous replication is covered separately in the Dell EMC Unity: MetroSync white paper which can be found on Dell EMC Online Support.

June 2021

Revisions

Date	Description
May 2016	Initial release – Unity OE 4.0
December 2016	Updated for Unity OE 4.1
July 2017	Updated for Unity OE 4.2
March 2018	Updated for Unity OE 4.3
August 2018	Updated for Unity OE 4.4
January 2019	Updated for Unity OE 4.5
June 2019	Updated for Unity OE 5.0
January 2021	Metro node updates
June 2021	Template and format updates. Updated for Unity OE 5.1.

Acknowledgments

Author: Ryan Poulin

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © 2016-2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [6/21/2021] [Technical White Paper] [H15088.7]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	6
Audience	6
1 Introduction.....	7
1.1 Terminology	7
2 Native synchronous block replication	10
2.1 Licensing.....	10
2.2 Theory of operation	10
2.2.1 Synchronous replication interfaces.....	10
2.2.2 Synchronous Replication Management Ports	11
2.2.3 Replication Connections.....	12
2.2.4 Replication Sessions	12
2.2.5 Storage Resources	13
2.2.6 Replication Roles.....	13
2.3 Replication Operations	14
2.3.1 Pause and Resume	14
2.3.2 Failover and Failback	15
2.3.3 Delete	15
2.4 Data Protection Mechanisms.....	16
2.4.1 Fracture Log	16
2.4.2 Write Intent Log	16
2.5 Supported Replication Configurations.....	17
2.6 Unisphere Management	18
2.6.1 Configuring Replication	18
2.6.2 Synchronous Replication Management Ports	18
2.6.3 Creating a Replication Connection.....	20
2.6.4 Creating a Replication Session	21
2.6.5 Viewing the Replication Sessions.....	26
2.7 System Maximums	27
3 Native Asynchronous Replication.....	28
3.1 Licensing.....	28
3.2 Theory of Operation.....	28
3.2.1 Replication Interfaces	28

3.2.2	Replication Connections	30
3.2.3	Replication Sessions	35
3.2.4	Asynchronous Replication Internal Snapshots	36
3.2.5	Storage Resources	37
3.2.6	Replication Roles	38
3.3	Replication Operations	39
3.3.1	Pause and Resume	39
3.3.2	Sync	40
3.3.3	Failover with Sync, Failover, and Failback	40
3.3.4	Delete	41
3.3.5	Group Operations	42
3.4	Supported Replication Configurations	42
3.4.1	Local Replication	42
3.4.2	Remote Replication	43
3.4.3	Advanced File Remote Replication	45
3.4.4	MetroSync	49
3.5	Snapshot Replication	49
3.6	Unisphere Management	50
3.6.1	Configuring Replication	51
3.6.2	Creating Replication Interfaces	51
3.6.3	Creating a Replication Connection	52
3.6.4	Creating a Replication Session	54
3.6.5	Restore from a replicated snapshot	64
3.6.6	Override network address for file replication	64
3.6.7	Enable NDMP on NAS Server	65
3.6.8	Viewing the Replication Sessions	66
3.7	System Maximums	67
4	Proxy NAS Servers	68
5	Interoperability	69
5.1	Snapshots	69
5.1.1	Create replication session	69
5.1.2	Failback/Resume after unplanned failover	73
5.2	File Import	74
5.3	Data Reduction	74
5.4	Dynamic Pools	75
6	Upgrades	76

Table of contents

7	RecoverPoint with Dell EMC Unity	77
8	Metro node	78
9	Conclusion	79
A	Technical support and resources	80
A.1	Related resources.....	80
B	Appendix: Replication Maximums	81
C	Appendix: Replication Support Across Platforms.....	82

Executive summary

Data is often one of the most valuable assets to an organization. It is being accessed constantly by the various applications, users, and sometimes directly by customers of the organization. This makes data a crucial part of the day to day operations of an organization. Outages can happen at any time, and can be restricted to a single system, or an entire data center or location. Whether they are planned outages, such as regular maintenance, or unplanned such as a power outage, ensuring that data critical to the organization is available at all times is a top priority. A business continuity plan for critical data is suggested to prevent costly outages. To protect against the different scenarios, an organization should plan and implement a data protection strategy. A data replication solution can ensure business continuity, high availability, and data protection. Dell EMC Unity™ provides native and non-native solutions that will help you protect your data and meet the goals of your business for both data availability and protection.

Dell EMC Unity systems provide synchronous and asynchronous replication solutions which allow you to replicate data locally within the same system, or to other systems, whether they are located at the same site or a remote facility. Having remote replicas of data protects you against outages on the main system and allows you to recover quickly and easily on a destination system with minimal to no data loss depending on the replication method selected.

This white paper describes the following replication technologies for Dell EMC Unity:

- Native Synchronous Block Replication
- Native Asynchronous Block and File Replication
- Dell EMC® RecoverPoint®

Dell EMC Unity also supports native file synchronous replication. For information about native file synchronous replication and MetroSync Manager, reference the *Dell EMC Unity: MetroSync* white paper on Dell EMC Online Support.

Synchronous Block Replication, Asynchronous Block and File Replication, and Manual Replication can easily be configured and managed in Unisphere, Unisphere CLI, or REST API. Unisphere is a simple and intuitive HTML5 web-based interface which not only allows users to configure and manage their replication setup, but also provides a visual representation of the configuration.

Dell EMC RecoverPoint is an appliance-based product providing an alternative solution for Block Replication for Dell EMC Unity systems. Configuring RecoverPoint protection is done through the intuitive Unisphere for RecoverPoint interface. RecoverPoint allows for the recovery of data for any point-in-time and replicate to other supported Dell EMC storage systems.

Audience

This white paper is intended for Dell EMC customers, partners, and employees who are considering the use of Native Synchronous Block Replication, Native Asynchronous Block and File Replication, or RecoverPoint for the Dell EMC Unity family of storage systems. It assumes familiarity with Dell EMC Unity and Dell EMC's management software.

1 Introduction

To protect against outages which can interrupt data availability, it is crucial to have a redundant copy of data. To protect against a storage system outage, you can use replication to create a copy of data on a remote system. Replication is a software feature which synchronizes data to a remote system within the same site or a different location. Replicating data helps to provide data redundancy, and safeguards against storage system failures at the main production site. Having a remote Disaster Recovery (DR) site protects against system and site-wide outages and provides a remote location to resume production and minimize downtime due to a disaster. For Dell EMC Unity, the platform offers many replication solutions to easily meet disaster recovery needs in customer environments.

Synchronous replication is a data protection solution which ensures each block of data written to a storage resource is first saved locally and to a remote image before the write is acknowledged to the host. This ensures that in the event of a disaster, there is zero data loss. In synchronous replication solutions, there are also trade-offs. As each write needs to be saved locally and remotely, added response time occurs during each transaction. This response time increases as distance increases between remote images. Synchronous replication has a distance limitation based on latency between systems. This limitation is generally 60 miles or 100 km between sites. It is recommended to ensure the latency of the link between the local and remote system is less than 10 milliseconds.

Asynchronous replication is primarily used to replicate data over long distances, but also can be used to replicate storage resources between Pools within the same system. Asynchronous replication does not impact host I/O latency as host writes are acknowledged once they are saved to the local storage resource. Because write operations are not immediately replicated to a destination resource, all writes are tracked on the source. This data will be replicated during the next synchronization. Asynchronous replication introduces the concept of a Recovery Point Objective (RPO). RPO is the acceptable amount of data, measured in units of time, which may be lost due to a failure. This delta of time also affects the amount of data which needs to be replicated during the next synchronization, and the amount of potential data loss if a disaster scenario were to occur.

Dell EMC Unity's synchronous and asynchronous replication features can easily be configured using Unisphere, Unisphere CLI, or REST API. Dell EMC Unity also supports manual replication, which does not automatically update a destination image with changes on the source. Manual replication will be discussed with asynchronous replication. RecoverPoint also supports block replication for Dell EMC Unity and will be discussed later in this paper. RecoverPoint uses physical or virtual appliances to replicate data between systems and is configured using the Unisphere for RecoverPoint user interface.

The Dell EMC Unity File synchronous replication feature and MetroSync Manager is covered in the *Dell EMC Unity: MetroSync* white paper, which can be found on Dell EMC Online Support.

1.1 Terminology

Asynchronous Replication: A replication method which allows you to replicate data over long distances and maintain a replica at a destination site. Updates to the destination image can be issued manually, or automatically based on a customizable Recovery Point Objective (RPO).

Bandwidth: The amount of data, represented in MB/s, which can be transferred in a given period of time.

Common Base: A pair of Unified Snapshots taken on a replication source and destination storage resource which have the same point-in-time image.

Consistency Group (Dell EMC Unity): A storage instance which contains one or more LUNs within a storage system. Consistency groups help organize the storage allocated for a particular host or hosts.

Consistency Group (RecoverPoint): In RecoverPoint, a consistency group is a user-defined group of LUNs to be replicated. A consistency group needs to be created in RecoverPoint as a part of the workflow to configure replication. This consistency group containing Dell EMC Unity LUNs can be replicated locally or remotely. RecoverPoint consistency groups are synonymous with Dell EMC Unity consistency groups.

Destination Storage Resource: A storage resource that is used for disaster recovery in a replication session. This is also known as a target image.

Fracture Log: A bitmap held in persistent memory on each storage processor which indicates which physical areas of a source resource participating in a synchronous replication session have been updated since communication was interrupted (fracture).

Internal Snapshot (Replication Snapshot): Unified Snapshots created by the system which are part of an asynchronous replication session. These snapshots can be viewed in Unisphere, but user operations are not permitted. Each asynchronous replication session uses two internal snapshots taken on the source and destination storage resources.

Journal Volume: In RecoverPoint, a journal volume is a LUN designated to hold data associated with previous points-in-time. The journal is used to allow RecoverPoint to roll back data to any point-in-time. Journal volumes must be configured for each copy of a consistency group, including the production copy.

LUN: A block-based storage resource which a user provisions. It represents a SCSI logical unit.

RecoverPoint: An appliance-based disaster recovery solution that replicates synchronously or asynchronously and enables recovery to any point-in-time.

RecoverPoint Appliance (RPA): An industry-standard server platform that runs RecoverPoint software and manages all aspects of data protection for a consistency group. RPAs are clustered at each site in a RecoverPoint system for high availability and load balancing. Virtual machines running RecoverPoint software, or vRPAs, are also supported as an alternative to physical appliances.

Recovery Point Objective (RPO): The acceptable amount of data, measured in units of time, which may be lost due to a failure. For example, if a storage resource has an RPO of 1 hour, any data written to the storage resource within the most recent hour may be lost when the replication session is failed over to the destination storage resource.

Recovery Time Objective (RTO): The duration of time in which a business process must be restored after a disaster. For example, an RTO of 1 hour requires access to the data be restored within 1 hour after a disaster is encountered.

Replication Session: A relationship configured between two storage resources of the same type, on the same or different systems, to automatically or manually synchronize the data from one resource to another.

Snapshot: A snapshot, also called a Unified Snapshot, is a point-in-time view of a storage resource. When a snapshot is taken, the snapshot is an exact copy of the source storage resource and shares all blocks of data with it. As data changes on the source, new blocks are allocated and written to. Unified Snapshot technology can be used to take a snapshot of a Block or File storage resource.

Storage Resource: The top-level object a user can provision, associated with a specific quantity of storage. All host access and data protection activities are performed at this level. In this document, storage resource

refers specifically to those which support replication: LUNs, consistency groups, NAS Servers, file systems, and VMware VMFS and NFS Datastores.

Synchronous Replication: A replication mode in which the host initiates a write to the system at a local site and the data must be successfully stored in both local and remote sites before an acknowledgment is sent back to the host.

Thin Clone: A thin clone is a Read/Write copy of a Thin Block storage resource (LUN, consistency group, or VMware VMFS datastore) that shares blocks with the parent resource.

Unisphere: A web-based Dell EMC management interface for creating storage resources and configuring and scheduling protection of stored data on a Dell EMC Unity system. Unisphere can be used for all management of Dell EMC Unity native replication.

Unisphere for RecoverPoint: A web-based interface for managing RecoverPoint replication. It serves as a single pane of glass for replicating storage resources of multiple storage systems configured to use RecoverPoint. Consistency groups are created, replicated, and recovered through this interface.

User Snapshot: A snapshot created manually by the user or by a schedule. This is different than an internal snapshot, which is taken automatically by the system with asynchronous replication.

Write Intent Log: A record stored in persistent memory on the Dell EMC Unity storage system which tracks in-flight writes between systems participating in synchronous replication.

2 Native synchronous block replication

This section will discuss the native synchronous block replication feature found on Dell EMC Unity systems. This feature allows users to create synchronous replication sessions for Block storage resources between supported systems. All configuration and management operations in this section will be demonstrated in Unisphere, though Unisphere CLI and REST API may also be used. The following sections of this paper will discuss:

- The licensing requirements for the native synchronous block replication feature
- How the native synchronous replication feature works
- The configurations supported for synchronous replication
- Unisphere configuration and management

Synchronous block replication using RecoverPoint will be discussed in the *RecoverPoint with Dell EMC Unity* section.

2.1 Licensing

For Dell EMC Unity, synchronous block replication is only supported on physical deployments. The supported systems include the All Flash models, and the Hybrid models. For Dell EMC Unity All Flash and Hybrid systems, Synchronous Block Replication license comes included at no additional cost.

2.2 Theory of operation

Native synchronous block replication is consisted of a number of physical and software components. Each of these components will be discussed in the following sections. To fully configure synchronous replication, you must complete each of following:

- Configure FC connectivity for synchronous replication data transfer
- Configure network interfaces on the synchronous replication management ports for replication command communication
- Configure a replication connection between systems
- Configure replication sessions between systems

Also discussed in this section will be the various storage resources which support synchronous replication, and the different replication modes and roles for synchronous replication. The following sections detail the functions, requirements, and interactions of these components. The configuration and management of these components in Unisphere will be discussed in the *Unisphere Management* subsection.

2.2.1 Synchronous replication interfaces

Native synchronous block replication is performed over the first Fibre Channel (FC) port configured on the system. Older model Dell EMC Unity physical systems include two onboard Converged Network Adapter (CNA) ports which can be configured with either 16 Gb FC SFPs, 10 Gb Optical SFPs, or be used with Twinax cables. When configured as Fibre Channel, the lowest numbered port will be used for the synchronous replication interface on each SP. When the CNA is not configured for Fibre Channel, a 16 Gb Fibre Channel I/O module must be purchased, installed, and persisted on each of the SPs of the system. Dell EMC Unity XT systems support both a 16 Gb and 32 Gb Fibre Channel I/O module. Support for the 32 Gb Fibre Channel I/O module was added in Dell EMC Unity OE version 5.1. In any model type, the lowest numbered FC port on the system will be used as the synchronous block replication interface.

Starting with Dell EMC Unity OE version 4.1, the Replication Capability field is displayed for the Fibre Channel ports. To view, go to Unisphere > System > System View > Enclosures, then select a Fibre Channel port to review. Figure 1 is an example of the Enclosures tab on a Dell EMC Unity 480 system, with the SP A I/O Module 0 FC Port 0 selected. The Replication Capability can be seen from the quick properties of the port.

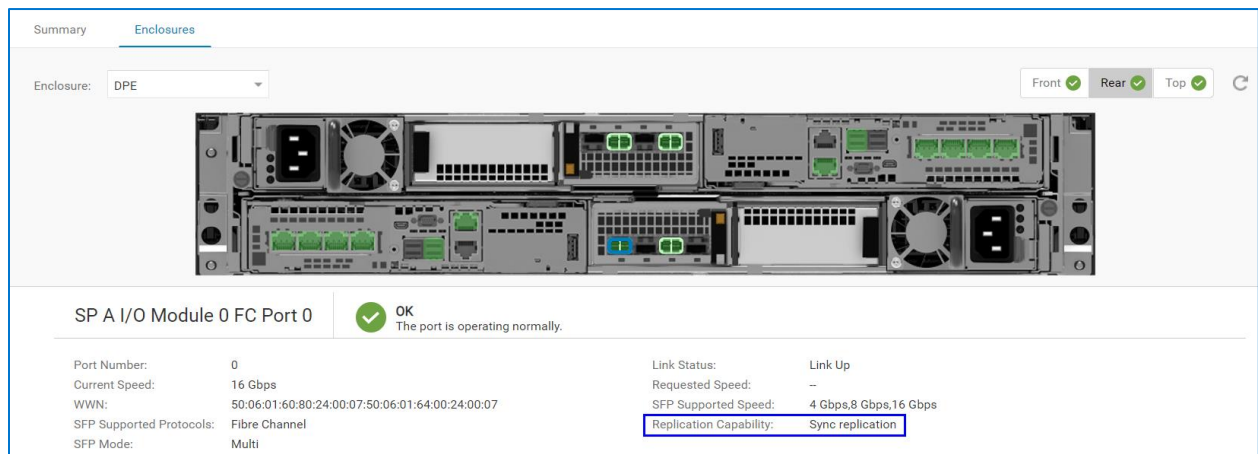


Figure 1 Replication Capability

With Unisphere CLI, you can also easily locate the synchronous block replication port by running the `/net/port/fc show -detail` command. This command will return information for Fibre Channel ports of the system. A sample output of this command with only the relevant information displayed is shown below:

```
C:\>uemcli -d ip_address -u username -p password /net/port/fc show -detail
1:   ID                               = spa_iom_0_fc0
      Replication capability          = Sync replication
4:   ID                               = spb_iom_0_fc0
      Replication capability          = Sync replication
```

In the above example, the first Fibre Channel port on the system, labeled as I/O Module 0 FC Port 0 on both Storage Processors (SP) are used for synchronous block replication. These ports can be directly connected to the synchronous ports on the destination system or zoned to them using a switch. For zoning, the proper configuration is to zone the SPA replication port for the local and remote systems together, and in a different zone, zone the SPB replication ports together. It is recommended to dedicate the FC replication interfaces strictly for replication purposes, but it is not required. Hosts can also use these ports if required. For redundancy, it is best to configure SPA's replication connection over one FC network, and SPB's over another where possible.

2.2.2 Synchronous Replication Management Ports

The synchronous replication management port is a virtual port used by synchronous replication to pass management commands between systems. A synchronous replication management port is created on each SP and is configured on the management port of a storage processor. This port must be able to communicate with the destination system's synchronous replication management port interfaces and does not support VLAN tagging. This port only needs to be configured when the system will participate in synchronous replication with a remote system. Synchronous replication management ports must be configured on the source and destination systems participating in synchronous replication.

2.2.3 Replication Connections

When configuring replication, a trusted link must be created between two systems before any replication sessions can be made. A replication connection is a logical link that is created between systems that will participate in remote replication. The replication connection establishes a link for management, and the data path between a pair of systems. After the replication interfaces and the synchronous replication management ports are configured, you must create the replication connection between the systems. Once configured, all synchronous replication sessions will use the replication connections to transport data to the remote system.

Figure 2 below shows an example configuration of a replication connection between two systems. In this example, the Dell EMC Unity XT systems are labeled “Production System” and “DR System”. All synchronous replication will occur between the Production System and the DR System. Each system is also configured with 32 Gb FC I/O Modules in slot 0, so I/O Module 0 port 0 on each SP is the synchronous replication port. After cabling each system to the network, configuring the synchronous replication management ports, and zoning the FC ports together, the replication connection can be established. All management related traffic will cross the LAN/WAN connections, and all data will be replicated across the FC Network synchronously. Once a replication connection is configured on one of the systems participating in replication, it is automatically created on the other system.

Verify and Update is a single operation that is used to update the replication connection information about the system it is issued on. This operation is performed on the replication connection itself, as opposed to an individual replication session. Verify and Update can be used to test a replication connection to a remote system or update the replication information if changes to the system have been made. Verify and Update should be issued to reestablish the replication connection to a remote system after an outage or when a configuration change to a replication interface on the system has been made. Another use case for using Verify and Update is if a Synchronous Replication Management Port IP Address has been changed on the system.

Note: A system can only be synchronously connected to one remote system at a time. FC connections to multiple Dell EMC Unity systems will cause the remote connection verification to fail.

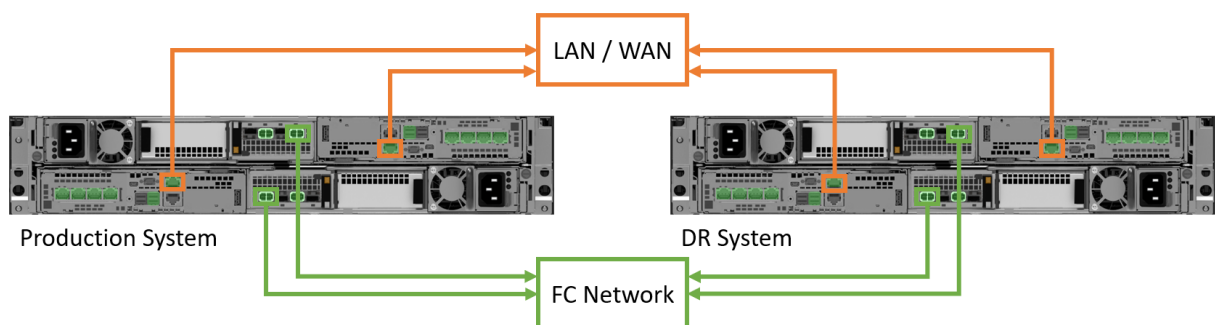


Figure 2 Native synchronous block replication connection example

2.2.4 Replication Sessions

A replication session uses a configured replication connection to transfer data from a source to a destination storage resource. If a replication session is created using Unisphere, a storage resource of the same size and type is created on the destination system and a full synchronization occurs between the source and destination storage resource. If replication is configured when creating a new resource, the synchronization operation is quick as no data needs to be copied across the link. When configuring synchronous block replication, snapshots and snapshot schedules are not replicated to the DR system.

The following steps outline a write operation to a storage resource with a synchronous replication session configured. In this example, assume the initial synchronization is complete.

1. A write I/O is sent to a storage resource on the production system.
2. System cache on the production system accepts the write I/O.
3. The production system replicates the data to the DR System.
4. The DR system accepts the data into system cache.
5. The DR system responds to the production system and acknowledges the write has been saved.
6. The production system acknowledges the host that the data has been accepted and saved on the system.

2.2.5 Storage Resources

When configuring replication, the source and destination storage resource must be of the same type. In Dell EMC Unity, native synchronous replication is supported on the following storage resources:

- LUNs
- Consistency groups
- VMware VMFS datastores
- Thin clones
- File systems and VMware NFS datastores
 - Covered separately in the *Dell EMC Unity: MetroSync* white paper on Dell EMC Online Support

On the system, replication for LUNs, thin clones, and VMware VMFS datastores function identically to each other. When configuring replication on a LUN, a single replication session is created, and the source and destination storage resources will be the same size and type. When configuring replication session on a thin clone, the destination storage resource will be a regular storage resource and not a clone. While synchronous replication is configured, you cannot extend the storage resource. Other options, such as the LUN's name or tiering policy may be configured differently between systems.

In Dell EMC Unity, a consistency group is a storage instance which contains one or more LUNs within a storage system. Consistency groups help organize the storage allocated for a particular host or hosts. consistency groups are treated as a single entity when they are replicated, meaning that a single replication session is created for the entire consistency group no matter how many LUNs it contains. When pausing or resuming replication on a consistency group, the entire group is affected by the replication operation.

2.2.6 Replication Roles

Two storage resources are required for replication:

- A source storage resource which will be replicated
- A destination storage resource which is copied to from the source

When a replication session is created in Unisphere, the destination storage resource is automatically created with the session. In Unisphere, you may not choose a previously provisioned storage resource as a replication destination when configuring a session. Upon creation, the destination resource is marked as a destination image. This restriction blocks Read/Write access on the destination storage resource. To view the data contained in the destination storage resource, you may take a snapshot of the resource and provide host access to it.

In Unisphere, you can easily determine which storage resources are replicated and which are destination images from any of the storage resource pages, such as the LUNs tab on the Block page. While in this tab,

select the gear icon, and hover over the Columns option to view the available columns that can be viewed. Select the check boxes for “Replication Type” and “Restricted Replication Access”. Replication Type displays what type of replication the storage resource is participating in, whether it is None, Remote, or Local. The Restricted Replication Access column will display “Yes” if the storage resource is labeled as a replication destination resource, or “No” if it is not. If a replication session is deleted, the destination storage resource will still be labeled as a replication destination image. The replication destination label must be edited manually over Unisphere CLI before the resource is allowed to receive I/O. For example, to remove the replication destination setting from a LUN, use the `uemcli /stor/prov/luns/lun -id <value> set - replDest no` command. For more information about removing the replication destination setting on other storage resource types, consult the *Dell EMC Unity: Unisphere Command Line Interface User Guide* on Dell EMC Online Support.

In Dell EMC Unity OE versions 5.1 and later, Source, Destination, and All filtering buttons on the replication sessions page and various storage resource pages help the user easily identify replication source and destination resources/sessions without adding columns to the view. When All is selected, all resources/sessions on the current page are displayed. When Source is selected on a resource page, all resources that are the source of a replication session are displayed. Resources that are not replicated are also shown when Source is selected. When Source is selected on the replication sessions page, only replication session originating on the system are shown. When Destination is selected on a resource page, only resources that are the destination images of a replication session are shown. While on the sessions page, Destination will only the sessions replicating to the current system. Also, sessions that are part of local replication are displayed regardless of which view is selected.

2.3 Replication Operations

Once a replication session has been created, several operations are available to manipulate the replication session as needed. Not all operations are available at all times, as some depend on the session being in a particular state. Also, certain operations will perform differently depending on which system they are issued on, source or destination. Only one replication operation can be issued and running per session at a particular point in time.

2.3.1 Pause and Resume

Pause and resume are used to stop and start replication between the systems for a particular replication session. In Unisphere, the pause operation can only be issued from the source system. While a session is paused, all write I/O to a storage resource is only saved on the source system before being acknowledged to the host. All I/O is tracked when the replication session is paused, and when the session is resumed, all blocks not already copied to the destination will be done at this time. During this time, the state of the replication session will be “Syncing”. Once all I/O is copied to the destination, the state will change to “In Sync”.

Pausing a replication session may be done for a number of reasons. Some reasons include the need to power off the source or destination system for planned maintenance, a configuration change on the network between systems, or interface changes on either system. Another reason may be to physically move a system from one data center or site to another. In certain circumstances, configuring replication and synchronizing the data between systems may be done within the same site, then the destination system is later moved to its final destination.

2.3.2 Failover and Failback

When issuing the failover operation, the current destination storage resource becomes the production resource. Before issuing the failover operation, it is highly suggested to quiesce I/O to the source image first. Once the failover operation is initiated the destination resource will be available for Read/Write operations and the original source resource will no longer be available for reads or writes. If host access is configured on the destination resource, hosts can access the data at this time. The effects of issuing the failover operation will depend on which system the failover was initiated from.

A failover issued from the production system is also referred to as a planned failover. The destination resource will become Read/Write available, and the direction for replicating data will switch. When this occurs, the original destination system will start replicating all new writes it receives for the storage resource to the original source system for the replication session. Issuing a failover from the source is suggested when testing a site failover to ensure the DR configuration is working properly. To return the session to its original state, first quiesce I/O to the original destination storage resource. Then, issue a failover from the original destination system. This will make the original storage resource the production image and start replication in the original direction.

If a failover operation is issued from a destination system, an unplanned failover or disaster situation is initiated. An unplanned failover assumes a disaster has occurred on the production system, and the destination image is made Read/Write available. The replication session will also pause and not automatically switch the direction for replication. The replication session is left in this state until the user issues another replication operation.

Starting with Dell EMC Unity OE version 5.0, an unplanned failover operation can be initiated even when the replication is in a "Paused" state. Previously, unplanned failover operations were not allowed while the session was in a "Paused" state. Any changes made on the source system while the session is paused are not replicated to the destination.

When the original production system becomes available, the user has the option of issuing the Resume operation on the session from the original destination to replicate data to the original source which will result in a full sync. Then, at a later time, they can issue the Failover operation from the destination system to return to the original state for replication, from the production system to the destination system. The user also has the option of issuing the Failback operation from the destination system when the source becomes available. The Failback operation will do a full synchronization of the storage resource on the source system with the data on the destination resource, block access to the destination storage resource, and make the original production resource Read/Write. Replication will also then be resumed from the production system to the destination system for the replication session. It is suggested to quiesce all I/O before issuing the Failback operation.

2.3.3 Delete

Deleting a replication session can be issued on the source system or destination system, but it is recommended that the operation be issued on the source when the source is available. When there are no issues in the configuration and a Delete operation is issued on the source system, the replication session will be deleted from both the source and destination system. The destination storage resource is not automatically deleted when the replication session is deleted. If the delete operation is issued while the destination system cannot be reached, the session will need to be deleted from the destination system manually. If the delete operation is issued from the destination system, the source session is left configured and must be deleted manually. Once a replication session is deleted, a full sync will need to occur if replication is reconfigured.

A delete operation can also be issued for a replication connection. A replication connection can only be deleted after all configured replication sessions using the connection have been deleted.

2.4 Data Protection Mechanisms

Synchronous replication has mechanisms to resync data differences between the source and/or destination resources in the event of replication disruption thereby preventing the need for full synchronization. The fracture log protects primarily against loss of communication with the destination resource. The write intent log protects primarily against interruptions to the source resource. Both of these structures exist to enable partial synchronizations in the event of interruptions to the source or destination resources.

2.4.1 Fracture Log

The fracture log is a bitmap held in the memory of the storage processor that owns the source resource. It indicates which physical areas of the source have been updated since communication was interrupted with the destination.

The fracture log is automatically invoked when the destination resource of a replication session is lost for any reason and becomes out of sync. The replication session is out of sync (no longer replicating) if the destination is not available, or it can be administratively paused through Unisphere or UEMCLI. Dell EMC Unity sets a replication session as out of sync if an outstanding I/O to the destination is not acknowledged within 25 seconds. While in a state of out of sync, the source pings the destination every 20 seconds to determine if communication has been restored.

The fracture log tracks changes on the source resource for as long as the destination resource is unreachable. It is a bitmap that represents areas of the source resource with regions called extents. The amount of data represented by an extent depends on the size of the data resource. Since the fracture log is a bitmap and tracks changed areas of the source resource, it is not possible to run out of fracture log capacity.

When the destination resource returns to service, it must be synchronized with the source. This is accomplished by reading those areas of the source addressed by the fracture log and writing them to the destination resource. This activity occurs in parallel with any writes coming into the source and replicated to the destination. Bits in the fracture log are cleared once the area of the source marked by an extent is copied to the destination. This ability to perform a partial synchronization can result in significant time savings. It may be necessary, depending on the length of the outage and the amount of write activity, to resynchronize the entire dataset.

By default, the fracture log is stored in memory. Therefore, it would be possible for a full resynchronization to be required if a destination resource is out of sync and an interruption in service occurs on the source SP. To protect against such scenarios, the write intent log is used.

2.4.2 Write Intent Log

The write intent log is a record stored in persistent memory (disk) on the storage system on which the source resource resides. During normal operation, the write intent log tracks in-flight writes to both the source and destination resources in a sync replication relationship. Much like the fracture log, the write intent log is a bitmap composed of extents indicating where data is written. The write intent log is always active, but the fracture log is only enabled when the replication session is out of sync.

When in use, Dell EMC Unity makes an entry in the write intent log of its intent to update the source and destination resources at a particular location, then proceeds with the attempted update. After both images

respond that data has been written (that is written to write cache), the system clears previous write intent log entries. For performance reasons, the write intent log is not cleared immediately following the acknowledgment from the source and destination resources. It will be cleared while subsequent write intent log operations are performed.

In a recovery situation, the write intent log can be used to determine which extents must be synchronized from the source storage system to the destination system. For instance, if a single SP becomes unavailable (for example during a reboot or failure), there may be in-flight writes that were sent to the destination, but not acknowledged before the outage. These writes will remain marked in the write intent log. Then server software trespasses the resource to the peer SP. The remaining SP directly accesses the unavailable SP's write intent log and recovers the recent modification history. The SP then resends the data marked by the extents in the write intent log. This allows for recovery using only a partial resynchronization, rather than a full resynchronization because it ensures that any writes in process at the time of the failure are acknowledged by the destination resource. If the entire array becomes unavailable, then the write intent log is used to facilitate a partial resynchronization from source to destination once the source array is recovered.

2.5 Supported Replication Configurations

Dell EMC Unity's Native Synchronous Block Replication feature is supported in two different topologies, and deployment models will vary depending on the requirements of the configuration. The following synchronous replication topologies outline the system configuration options supported in Dell EMC Unity.

- One-Directional
 - A single source system replicating to a single destination system
- Bi-Directional
 - A two-system topology in which each system acts as a replication destination for the peer's production data

The following figure, Figure 3, is a graphical view of the supported topologies listed above. Note the figure uses LUNs as an example. For synchronous replication, the source and destination system must be a physical system, but the models used can vary depending on the configuration requirements. Below are examples of the two topologies that can be used, either One-Directional or Bi-Directional replication. The source and destination system can either be the same model type, either All Flash or Hybrid, or the source and destination can be a mix of different model All Flash and Hybrid systems.



Figure 3 Supported synchronous block replication topologies

One-Directional replication is typically deployed when only one of the systems will be used for production I/O. The second system is a replication target for all production data and sits idle within the same data center or a remote location. If the need arises, the DR system can be placed into production and host production I/O. In this scenario, mirroring the production system's configuration, including the number of drives and Pool layout, on the DR system is suggested, as each system would then have the same performance potential.

The Bi-Directional replication topology is typically used when production I/O needs to be spread across multiple systems or locations. The systems may exist within a single data center or in different, remote locations. When using this replication topology, production I/O from each system is mirrored to the peer system. If there is an outage, one of the systems can be promoted as the primary production system, and all production I/O can be sent to it. Once the outage is complete, the replication configuration can be changed back to its original configuration. This replication topology ensures both systems are in use by production I/O at all times.

On Dell EMC Unity, a block resource can only be synchronously replicated to one single storage resource on a remote system. This means that replicating a single storage resource to multiple destination resources is not supported. Therefore, only one replication connection between a source and destination pair of Dell EMC Unity systems can be used for synchronous replication. Cascading replication, when a destination storage resource is also replicated to another destination resource, is also not supported for synchronous block replication.

2.6 Unisphere Management

Creating and managing replication in Unisphere is easy and intuitive. All replication operations, including configuring of synchronous replication network interfaces, replication connections, and replication sessions can be performed in the Unisphere GUI. With the help of easy to follow wizards, replication can be configured quickly by IT generalists or advanced users alike. Replication can also be configured using Unisphere CLI, or REST API. For more information about configuring and managing replication using Unisphere CLI, refer to the *Unisphere Command Line Interface Guide* found on Dell EMC Online Support. For more information about REST API, refer to the *Unisphere Management REST API Programmer's Guide* found on Dell EMC Online Support.

2.6.1 Configuring Replication

When configuring synchronous block replication, the following components need to be configured and were previously discussed in this paper:

1. FC connectivity between the synchronous replication interfaces on the local and remote system
2. Synchronous Replication Management Port Interfaces
3. A Replication Connection
4. A Replication Session

The first step above is completed by directly connecting the synchronous replication FC ports of two systems together or zoning the ports together over the fabric as outlined in the Synchronous Replication Interfaces section.

The following sections will outline the remaining steps needed to configure remote replication in Unisphere. Each of the following operations are completed from a particular page in Unisphere. Each page will be discussed in detail below. For more information about using Unisphere to configure and manage replication, refer to the *Unisphere Online Help*.

2.6.2 Synchronous Replication Management Ports

To create the synchronous replication management ports, you must first go to the **Interfaces** page, which is found under **Protection & Mobility** in the left pane in Unisphere. For releases prior to Dell EMC Unity OE version 4.1, the Interfaces can be found as a tab in the **Replication** page, under **Data Protection**. Figure 4 below shows the **Interfaces** page. From this page, you have the option to create or delete mobility interfaces,

refresh the current page, or edit a configured mobility interface. In the example below, no mobility interfaces have been created. Only one pair of synchronous replication management ports needs to be created on the source and destination system. Mobility Interfaces may be shared network interfaces for import and replication related data or management traffic using the virtual management port. In this section, the mobility interfaces will be referenced as replication interfaces.




Figure 4 Interfaces page

To create a replication interface, click the **Create Interface** button, shown as a + sign in the **Interfaces** page. Once selected, the **Create Interface** wizard appears, which is shown in Figure 5. For synchronous block replication, you must create the Synchronous Replication Management Ports on the system, which are used for management connectivity between the source and destination system. In the **Ethernet Port** drop-down list, select **Sync Replication Management Port**. The status of the ports will be shown in parentheses. Both system management ports, also used for Unisphere management, will need to be cabled to the network for these links to work properly.

Next, configure an IP address for SPA and SPB to be used for synchronous replication management. Note that VLAN tagging is not supported on these interfaces. Once done, click **OK**.

Figure 5 Create Interface window

Once the synchronous replication management ports are created, they will be shown on the Interfaces page as shown in Figure 6. From here the status of each port is shown, currently operating normally as denoted by the green circle checkmark. To delete a synchronous replication management port, select the port and click the **Delete Interface** icon, which is shown as a trash can on this tab. To edit a replication management port, to change the IP address for example, select the **Edit** icon, which is shown as a pencil icon on this tab.

Interfaces						
+ [trash] [refresh] [edit] 2 items [filter] [settings] [download]						
<input type="checkbox"/>	!	Ethernet Port	↑ SP	IP Address	Subnet Mask / Prefix Length	Gateway
<input type="checkbox"/>	✓	Sync Replication Management Port	SP A	[IP Address]	[Subnet Mask / Prefix Length]	[Gateway]
<input type="checkbox"/>	✓	Sync Replication Management Port	SP B	[IP Address]	[Subnet Mask / Prefix Length]	[Gateway]

Figure 6 Synchronous replication management interfaces

2.6.3 Creating a Replication Connection

The next step in configuring synchronous block replication is to create a Replication Connection with another system. This configures a private replication connection for the pair of systems. To create the Replication Connection, you must first go to the **Replication** page, which is found under **Protection & Mobility** in the left navigation pane in Unisphere. The Replication page has the Sessions and Connections tabs. From the **Replication** page, select the **Connections** tab. Figure 7 below shows an example of the **Connections** tab, without any replication connections configured. From this tab, you have the option to create or delete a replication connection, refresh the current page, or edit a configured replication connection. From this tab, you can also select **Verify and Update**, which will verify the selected replication connection still exists with the remote system and update the connection details if any changes were made.

Sessions <u>Connections</u>									
+ [trash] [refresh] [edit] Verify and Update 0 items [filter] [settings] [download]									
<input type="checkbox"/>	!	Remote System Name ↓	System Type	Mode	Management IP Address	Local Interfaces		Remote Interfaces	
						SPA Interfaces	SPB Interfaces	SPA Interfaces	SPB Interfaces
You do not have any Connections configured. Click the <input type="button" value="+"/> button above to create one.									

Figure 7 Replication Connections page

To create a Replication Connection, select the **Create Replication Connection** icon, which is displayed as a + sign on this tab. The **Create Replication Connection** window appears, as shown in Figure 8. In this window, you must specify the Remote System's **Management IP Address**, which is the IP used to access Unisphere, and the Unisphere **User Name** and **Password**. Also, in this window you must enter the **Password** used to log in to Unisphere on the system you are configuring the replication connection on. Lastly, you must select the **Connection Mode** that will be used between the systems. In the drop-down list, you have the option to choose **Asynchronous**, **Synchronous**, or **Both**. When configuring synchronous replication between two systems, select **Synchronous**. If asynchronous replication will also be used for replicating data to other storage resources, select **Both**, as both synchronous and asynchronous replication will be used. After entering the required information, click **OK**. Dell EMC Unity Asynchronous and Manual Replication is discussed in the *Native Asynchronous Replication* section.

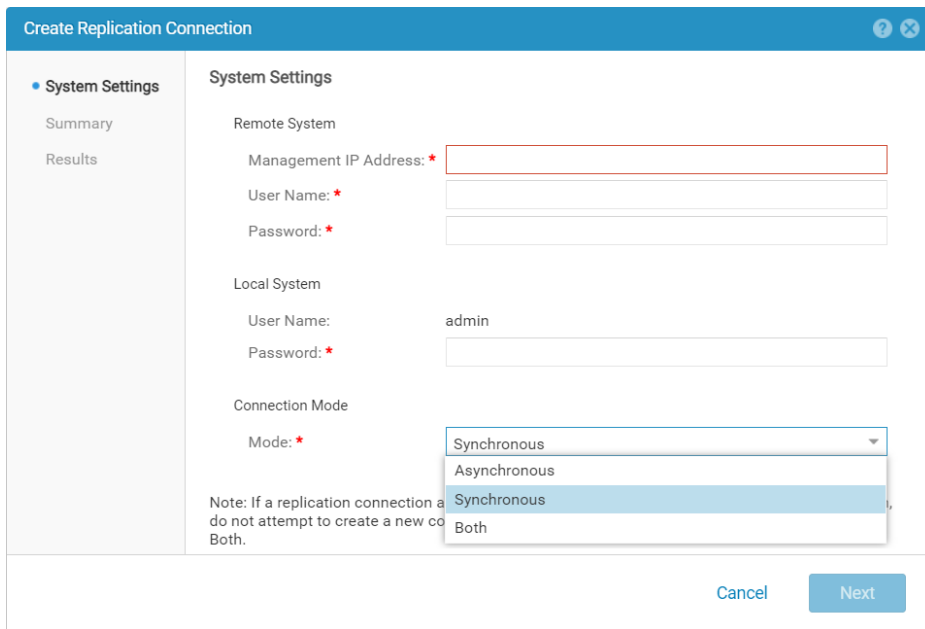


Figure 8 Create replication connection window

After selecting **OK**, a Job is generated to create the replication connection. The job has multiple steps, which includes registering the remote and local system with its peer, refreshing the connection on both systems, and validating the connections on the local system. Once the job completes, the replication connection will be shown on the source and destination system.

Figure 9 shows the Replication page **Connections** tab once the replication connection is created. In this example **DRSystem** is the name of the remote system the connection was configured with. The System Type, replication Mode, Management IP Address, and Local and Remote replication Interfaces will also be displayed. The Remote Interfaces displayed for synchronous block replication will be the IP Addresses configured for the synchronous replication management ports on the remote system. The **Verify and Update** button is also shown and is used to update the replication information for the selected connection.

Sessions		Connections					
		Remote System Name	System Type	Mode	Management IP Address	Local Interfaces	
						SPA Interfaces	SPB Interfaces
<input type="checkbox"/>	!						
<input checked="" type="checkbox"/>	✓	DRSystem	Unity 880	Synchronous Replication	192.168.1.100	192.168.1.100	192.168.1.100

Figure 9 Replication connection established

2.6.4 Creating a Replication Session

To create a replication session, you must enable replication on the storage resource you want to replicate. This can be done at time of creation, or on an existing storage resource. Figure 10 shows an example of creating a new LUN and enabling Synchronous remote replication on it. From the **Replication** step, check the checkbox in front of **Enable Replication** to configure replication. For Synchronous replication, select **Synchronous** from the **Replication Mode** drop-down box. When Synchronous is chosen as the Replication Mode, the **Replicate To** box will automatically be populated with the name of the system with which a synchronous Replication Connection has been configured. After completing this step, click **Next** and continue with the LUN creation. For thin clones, consistency groups and VMware VMFS datastores, a Replication step

exists in the creation wizards for each with the exact look and configuration options as the one shown in Figure 10.

Figure 10 Create LUNs wizard – Replication step

When the **Destination Configuration** box is selected in the Replication step, the remote storage resource's configuration can be customized. Figure 11 below shows an example of this window.

Figure 11 Destination LUN Configuration window

After synchronous replication is configured on a new storage resource, you can view information about replication from the resource's properties window. From Unisphere, select the storage resource in question and click Edit or double-click the name of the storage resource. From the properties window, view the **Replication** tab. An example of this tab is shown in Figure 12. On this tab, you can view the following information:

- The replication **Session Name**
- The replication **Mode**
- The **Local Role** of the storage resource, which can either be listed as the **Source** or **Destination**
- The Sync State

Also shown is a pictorial representation of the replication session. The picture shows which storage resource is available for I/O, which direction the data is replicating in and its current state, and the destination storage resource and the system name, IP Address, and the destination LUN name. As the state of the replication session changes, this figure will update to reflect the new state.

Also shown on the Replication tab are buttons for each Replication Operation. This tab is also used to display asynchronous replication operations, so all replication operations for both synchronous and asynchronous replication are displayed. Not all replication operations are supported on each mode of replication, so only operations supported on the current replication mode will be selectable. Also, only certain operations are available depending on what the current state of the replication session is in, so only these options are available for selection. In Figure 12, the current session is Active and replicating for the Production System to the DR System. Available replication operations include **Delete**, **Pause**, and **Failover**.

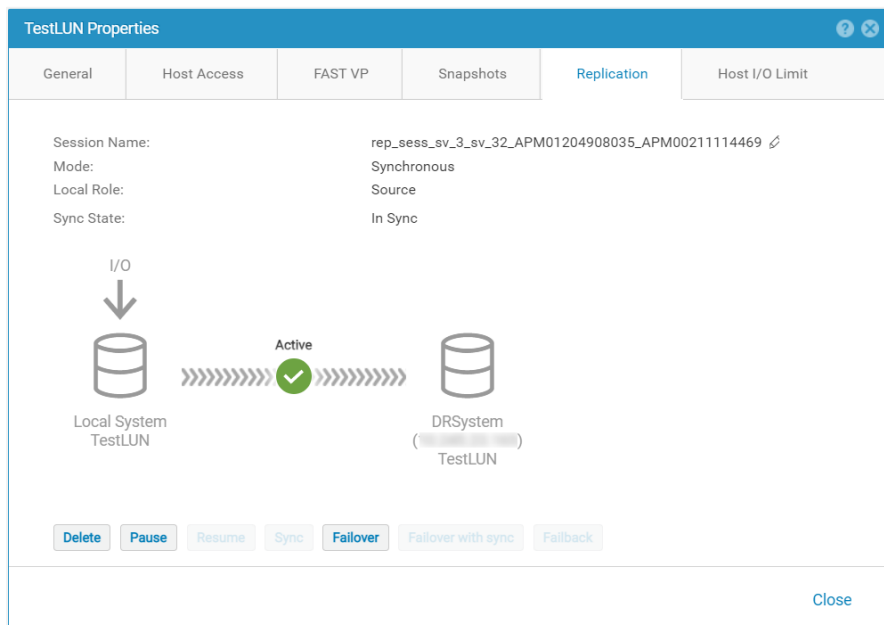


Figure 12 LUN properties – Replication tab

Replication can also be configured on an existing storage resource. When replication is not configured on the storage resource, viewing the **Replication** tab in the storage resource’s properties window will show what is displayed in Figure 13. To configure Replication on the storage resource, select **Configure Replication**.

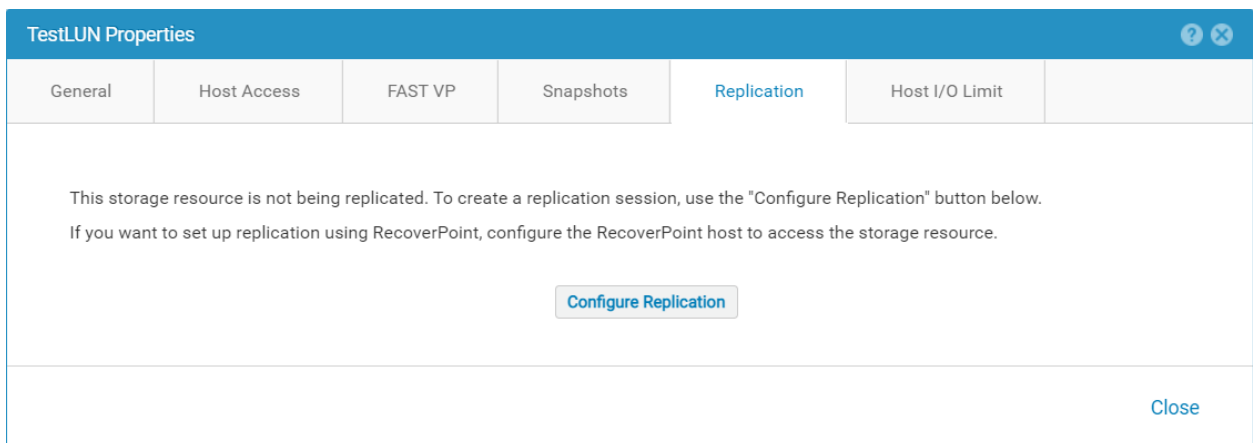


Figure 13 Replication not configured

After selecting **Configure Replication**, the **Create a Session** wizard is launched. An example of this wizard is shown in Figure 14. On the **Replication Settings** step, you need to customize which **Replication Mode** either Asynchronous, Manual, or Synchronous replication will be used. For Synchronous Replication, select Synchronous. Next, select the destination system by selecting the correct system in the **Replicate To** drop down box. After the previous selections have been made, click **Next**.

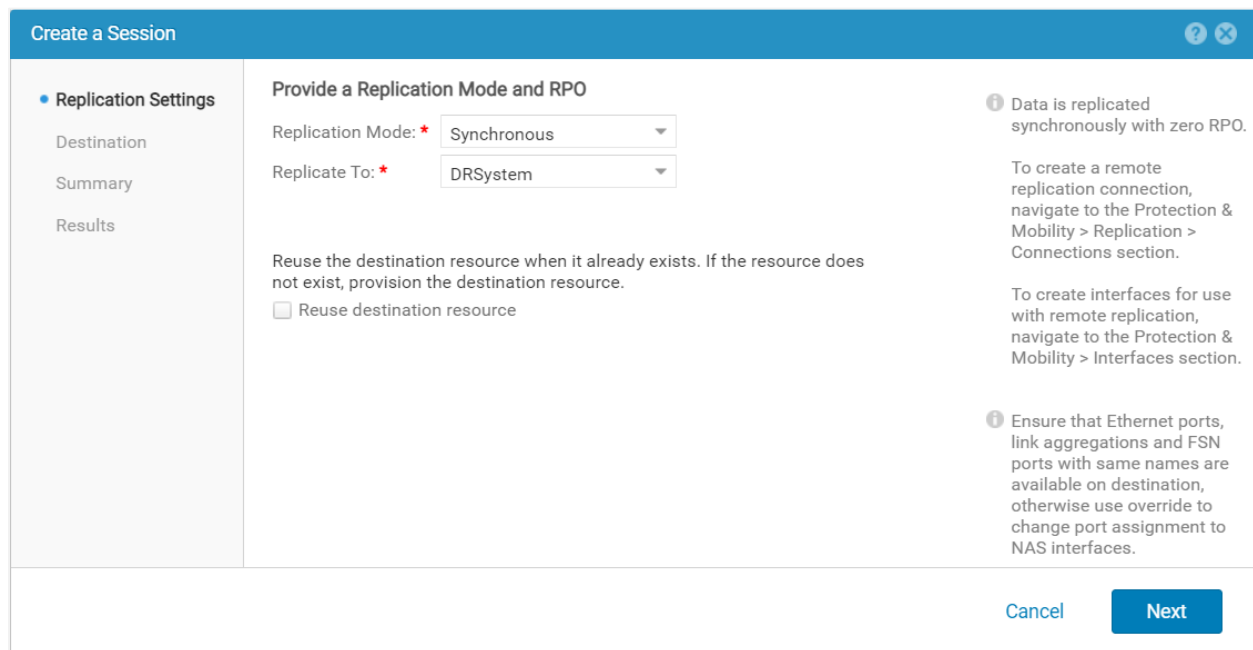


Figure 14 Create a Session

The **Destination** step of the **Create a Session** wizard is now shown. An example of this step is shown in Figure 15. From here, you can customize the storage resource's **Name** that will be displayed on the destination system, the **Pool** it will use, and the **Tiering Policy** on the destination system's Pool. For existing storage resources, you can customize this exact same information. Similarly, when creating a replication session on an existing consistency group, you can customize this same information for each LUN contained within the consistency group. After editing the available information, click **Next**.

Create a Session

✓ Replication Settings

• **Destination**

Summary

Results

Provide destination storage resource info

Name: * TestLUN

Destination System: DRSystem

Pool: * Pool 3 (Traditional, Multi-tier Pool, 9.4 TB free)

Size(GB): 200.0

Thin

Data Reduction

Tiering Policy: Start High Then Auto-Ti

Cancel Back Next

Figure 15 Create a Session – Destination step

The **Summary** step is now shown, and an example of this screen is displayed in Figure 16. Here you can see a summary of the settings that will be used to create replication. If anything is incorrect, you can select **Back** to correct the wanted setting. To create the replication session, click **Finish**. A Summary step will also be shown when creating a replication session on other supported storage resources.

Create a Session

✓ Replication Settings

✓ Destination

• **Summary**

Results

Summary

Settings configurations

Replication Mode: Synchronous

Destination System: DRSystem ()

Reuse destination resource: No

Destination configurations

Name: TestLUN

Pool: Pool 3

Size (GB): 200.0

Thin: Yes

Tiering Policy: Start High Then Auto-Tier

Cancel Back Finish

Figure 16 Create a Session – Summary step

The **Results** step is the last step in the Create a Session wizard. This step shows the Overall Status of each of the jobs to create the replication session. The steps to create the replication session include creating the storage resource on the destination system, allowing the remote storage resource to finish the creation process, and lastly creating the replication session. You can either wait for the Overall status to say 100% Complete or close the window at any time by clicking Close. Closing this window will not impact the creation process since it is a background job in Unisphere.

Figure 17 shows the **Replication** tab in the LUN properties window after enabling replication. After enabling replication on an existing device, a full synchronization needs to be completed. During this process, the **Sync State** shows **Syncing**. This operation may take time and will be based on the amount of data that is needed to be copied to the remote system, and the available bandwidth of the link between the systems.

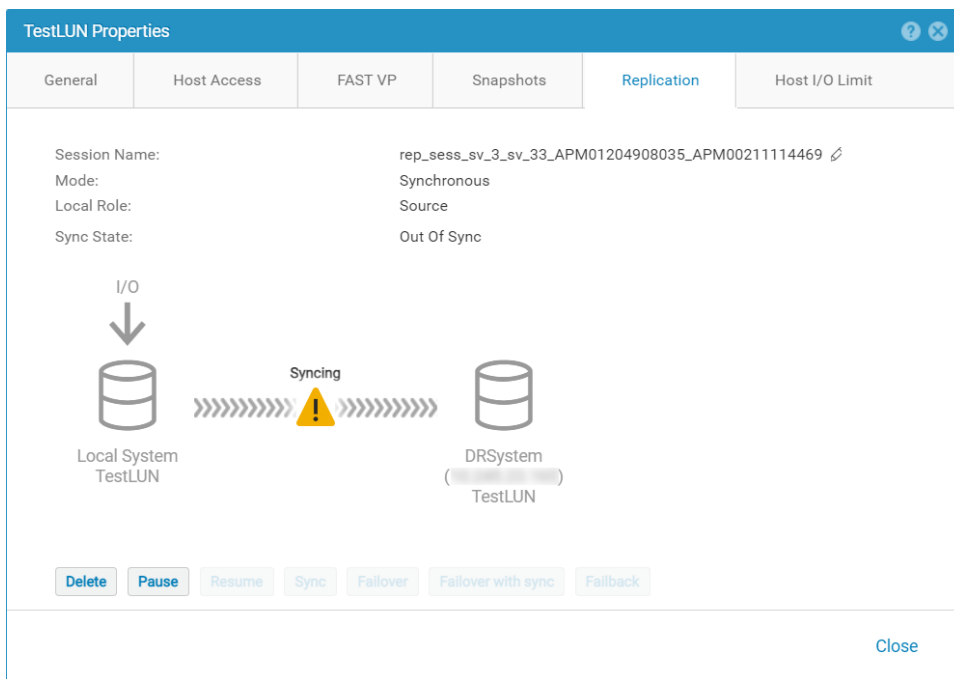


Figure 17 Sync State – Syncing

Figure 18 shows the Replication tab after the full synchronization. Notice that the **Sync State** now shows **In Sync**.

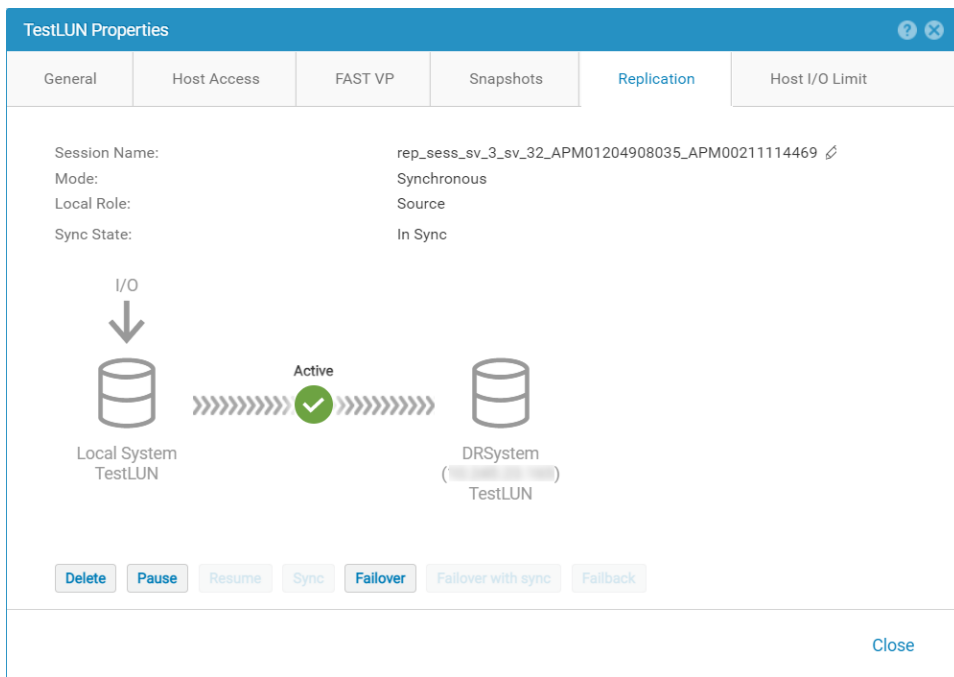


Figure 18 Sync State – In Sync

2.6.5 Viewing the Replication Sessions

All replication sessions on the system can be viewed from the **Sessions** tab within the Replication page. This includes synchronous, asynchronous, and manual replication sessions. To browse this page, click **Replication** under **Protection & Mobility** in Unisphere. For releases prior to Dell EMC Unity OE version 4.1,

the **Replication** page can be found under **Data Protection**. Figure 19 shows an example of the **Sessions** tab with multiple replication sessions created on the system. In this example, multiple LUNs, consistency groups, NAS Servers, and file systems are all being replicated. From this window, you can easily see information regarding each session. The following is a list of information displayed on this screen:

- The Replication Session **Name**
- The current **State**
- The **Source**, which includes the source system and the source storage resource
- The Resource Type
- The Replication Mode
- The **Destination**, which includes the destination system name and the destination storage resource
- The Replication Session **ID**

Name	State	Source		Resource Type	Replication Mode	Destination		id
		System	Resource			System	Resource	
RemoteRep_LUN1_to_...	Auto Sync Configured	Local System	LUN 1	LUN	Asynchronous	APM01204701525	LUN 1 - Dest - 0229	42949672965_APM00...
RemoteRep_LUN3_to_...	Auto Sync Configured	Local System	LUN 3	LUN	Asynchronous	APM01204701525	LUN 3 - Dest - 0229	42949672967_APM00...
rep_ses_sv_3_sv_33_...	Active	Local System	TestLUN	LUN	Synchronous	DRSystem	TestLUN	42949672967_APM01...
LocalRep_LUN5_to_Poo...	Auto Sync Configured	Local System	LUN 5	LUN	Asynchronous	Local System	LUN 5 - Dest - 0229	42949672969_APM00...
RemoteRep_CG1_to_1_...	Auto Sync Configured	Local System	CG1	Consistency Group	Asynchronous	APM01204701525	CG1 - Dest - 0229	81604378625_APM00...
LocalRep_CG2_to_Poo...	Auto Sync Configured	Local System	CG2	Consistency Group	Asynchronous	Local System	CG2 - Dest - 0229	81604378627_APM00...
RemoteRep_NASServe...	Auto Sync Configured	Local System	NAS_Server_2	NAS Server	Asynchronous	APM01204701525	DR_NAS_Server_2_...	103079215108_APM0...
LocalRep_NASServer3...	Auto Sync Configured	Local System	NAS_Server_3	NAS Server	Asynchronous	Local System	DR_NAS_Server_3_...	103079215109_APM0...
RemoteRep_NASServe...	Auto Sync Configured	Local System	NAS_Server_5	NAS Server	Asynchronous	APM01204701525	DR_NAS_Server_5_...	103079215111_APM0...
RemoteRep_FileSyste...	Auto Sync Configured	Local System	NS2_File_System_...	File System	Asynchronous	APM01204701525	NS2_File_System_1_...	171798691855_APM0...
RemoteRep_FileSyste...	Auto Sync Configured	Local System	NS2_File_System_...	File System	Asynchronous	APM01204701525	NS2_File_System_2_...	171798691856_APM0...

Figure 19 Replication Sessions

From the **Sessions** tab, you can also issue replication operations on the sessions. After selecting the checkbox for a replication session in the list, select **More Actions** to view the replication operations available for that session in its current state. In Figure 20 below, you can see that only **Pause** and **Failover** are valid options based on the selected resource's replication session state.

Name	State	Source		Resource Type	Replication Mode	Destination		id
		System	Resource			System	Resource	
RemoteRep_LUN1_to_...	Auto Sync Configured	Local System	LUN 1	LUN	Asynchronous	APM01204701525	LUN 1 - Dest - 0229	42949672965_APM00...
RemoteRep_LUN3_to_...	Auto Sync Configured	Local System	LUN 3	LUN	Asynchronous	APM01204701525	LUN 3 - Dest - 0229	42949672967_APM00...
rep_ses_sv_3_sv_33_...	Active	Local System	TestLUN	LUN	Synchronous	DRSystem	TestLUN	42949672967_APM01...
LocalRep_LUN5_to_Poo...	Auto Sync Configured	Local System	LUN 5	LUN	Asynchronous	Local System	LUN 5 - Dest - 0229	42949672969_APM00...
RemoteRep_CG1_to_1_...	Auto Sync Configured	Local System	CG1	Consistency Group	Asynchronous	APM01204701525	CG1 - Dest - 0229	81604378625_APM00...
LocalRep_CG2_to_Poo...	Auto Sync Configured	Local System	CG2	Consistency Group	Asynchronous	Local System	CG2 - Dest - 0229	81604378627_APM00...
RemoteRep_NASServe...	Auto Sync Configured	Local System	NAS_Server_2	NAS Server	Asynchronous	APM01204701525	DR_NAS_Server_2_...	103079215108_APM0...
LocalRep_NASServer3...	Auto Sync Configured	Local System	NAS_Server_3	NAS Server	Asynchronous	Local System	DR_NAS_Server_3_...	103079215109_APM0...
RemoteRep_NASServe...	Auto Sync Configured	Local System	NAS_Server_5	NAS Server	Asynchronous	APM01204701525	DR_NAS_Server_5_...	103079215111_APM0...
RemoteRep_FileSyste...	Auto Sync Configured	Local System	NS2_File_System_...	File System	Asynchronous	APM01204701525	NS2_File_System_1_...	171798691855_APM0...
RemoteRep_FileSyste...	Auto Sync Configured	Local System	NS2_File_System_...	File System	Asynchronous	APM01204701525	NS2_File_System_2_...	171798691856_APM0...

Figure 20 Replication Actions

2.7 System Maximums

See *Appendix A: Replication Maximums* for information about Dell EMC Unity replication maximums.

3 Native Asynchronous Replication

This section will discuss the Native Asynchronous Replication feature found on Dell EMC Unity systems. This feature allows users to create replication sessions for block and file storage resources between Pools within a system or between supported systems. All configuration and management operations in this section will be illustrated in Unisphere, though Unisphere CLI and REST API may also be used. The following sections of this paper will discuss:

- The licensing requirements for native asynchronous replication feature
- How native asynchronous replication feature works
- The configurations supported for asynchronous replication
- Unisphere configuration and management

3.1 Licensing

Asynchronous replication is supported on all Dell EMC Unity systems. The supported systems include the All Flash models, the Hybrid models, and the Dell EMC UnityVSA. For all systems, Asynchronous Replication comes included at no additional cost.

3.2 Theory of Operation

Native asynchronous replication is consisted of several physical and software components. Each of these components will be discussed in the following sections. When creating asynchronous replication between Pools within a system, only a replication session needs to be created. To fully configure asynchronous replication between systems, you must do each of following:

- Configure connectivity for the replication interfaces used by asynchronous replication
- Configure the replication connection between systems
- Configure replication sessions between systems

This section will also discuss the various block and file resources which support asynchronous replication. The replication modes and roles will also be discussed. The sections below outline the different functions, requirements, and how these components interact with each other. Configuration and management of these components will be completed using Unisphere, which will be discussed in the *Unisphere Management* section.

3.2.1 Replication Interfaces

Replication interfaces are used to transport data to a destination system for remote replication sessions and must be defined for replication to work. Asynchronous replication between Pools within a system does not require replication interfaces to be configured. Asynchronous remote replication requires a replication interface be created on each SP on the system. For asynchronous replication, replication is performed over ethernet interfaces configured on the system.

Any front-end ethernet port on Dell EMC Unity physical systems can be used for replication, such as onboard 10 GbE BaseT ports, CNA ports with ethernet personality, or ethernet based I/O modules. Supported front-end port configurations and capabilities varies by system model. Replication interfaces can also be created on link aggregated ports for high availability, increased maximum throughput, and load balancing of replication traffic across physical ports in the aggregation. It is suggested to only configure replication interfaces on ports of the same type and speed. It is also recommended to dedicate the replication interfaces strictly for

replication purposes, but it is not required, and hosts can also use these ports. In Dell EMC UnityVSA systems, any I/O port can be used for asynchronous replication.

In OE versions prior to the 5.1 release, in order to have a successful asynchronous replication connection, all asynchronous replication ports on a source system must be able to see and/or route to all asynchronous replication ports on the destination system, and the other way around. This also means that if more than two systems were part of the replication topology, all interfaces on all systems could communicate with each other. This requirement limited the possible replication topologies that could be supported on the system.

In Dell EMC OE version 5.1 the asynchronous replication interface pairing requirement that all asynchronous replication interfaces must be able to connect to each other has been removed. In OE 5.1 and later, the replication connection requires that replication interfaces on each SP of the source system be able to see and/or route to replication interfaces on both SPs of the destination system, and vice versa. When the remote connection is being created, only valid paths between the systems will be used for the connection. This feature allows for more advanced replication topologies to be created. Management connectivity is still required between the local system and all remote systems.

With the change in the OE 5.1 release and later, replication connections to multiple remote systems can be created over different physical networks. This allows a common system to remotely connect to peer systems that are physically isolated from each other. In Figure 21 below, **System 2** has replication interfaces created on physical ports that spread across **Network A** and **Network B**. Interfaces connecting to **Network A** are used to connect to **System 1**, while interfaces created on **Network B** ports only connect to **System 3**. In this example, **System 1** and **System 3** are on different data networks and cannot reach each other. In this topology, the management port on **System 2** is still required to have access to both systems at the same time.




Figure 21 Physical network isolation example

Another feature in OE version 5.1 and later is the ability to create replication interfaces on different VLANs. Replication interfaces created with different VLAN IDs can be used for connectivity to different systems. As shown in Figure 22, multiple replication interfaces can be created on the same physical port but can be tagged with different VLAN IDs. In this example **System 2** makes a remote connection with **System 1** over **VLAN 1000**, while **System 2** is using **VLAN 1100** to connect to **System 3**. In this configuration, **System 1** and **System 3** would not be able to connect to each other as they are tagged with different VLANs. Duplicate IP addresses are not allowed in this configuration. Topologies including a mix of physical network isolation and VLAN isolation is also supported. As mentioned previously, management connectivity is still required between the local system and all remote systems.

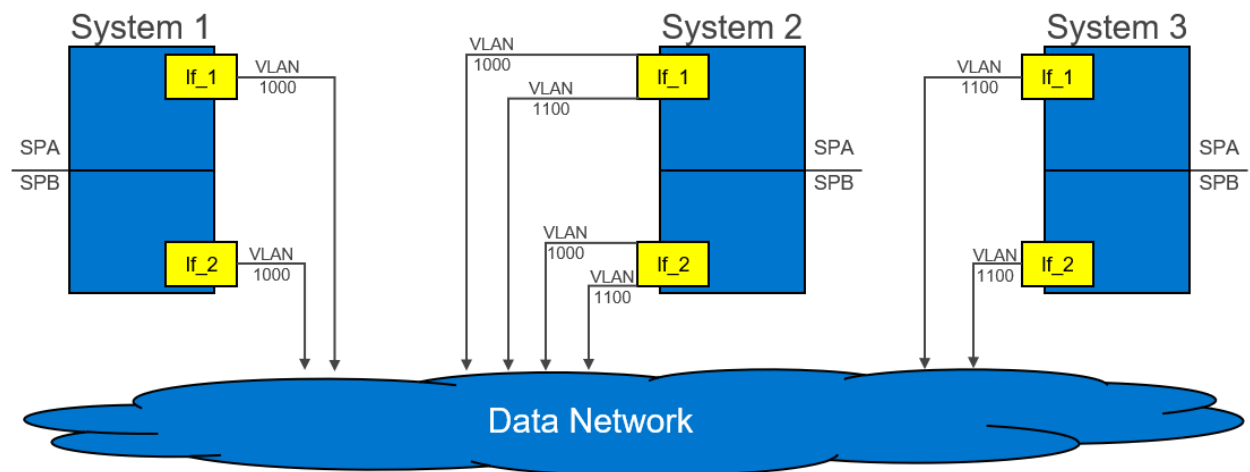


Figure 22 VLAN isolation example

To support the new replication interface topologies supported in the OE 5.1 release, all systems in the configuration must be running Dell EMC OE version 5.1 or later. If a system is not running version 5.1 or later, the remote connections within the topology will not be able to validate as the previous verification rules will apply. When changes to the replication configuration occur, the remote connections need to be re-validated to pick up any changes. Changes include adding or removing a replication interface or changing the settings within an existing replication interface.

3.2.2 Replication Connections

Before a remote replication session can be created, a trusted link must be created between systems participating in replication. A replication connection establishes this logical link between remote systems. This trusted link is used for replication management operations and the data path between the pair of remote systems. After the replication interfaces are created, you must create the replication connection between the systems. Once created, all asynchronous remote replication sessions will use the replication connections to transport data to the remote system.

Figure 23 below shows an example configuration of an asynchronous replication connection between two physical systems. In the figure, the source of the replication session is labeled “Production System”, and the destination is labeled “DR System”. For each of these example systems, the Dell EMC Unity 480 system has 10 GbE ports configured as replication interfaces. The green lines in the figure show the connections for the 10GbE replication interfaces, and the orange lines are the system management connections. Once the replication interfaces are created and cabled to the network on both systems, the replication connection between the systems can be made. Once a replication connection is configured on one of the systems participating in replication, it is automatically created on the peer system.

Verify and Update is a single operation that is used to update the replication connection information about the system it is issued on. This operation is performed on the replication connection itself, as opposed to an individual replication session. Verify and Update can be used to test a replication connection to a remote system or update the replication information if changes to the system have been made. Verify and Update should be issued to reestablish the replication connection to a remote system after an outage. A use case for using Verify and Update is if an Asynchronous Replication Interface has been added or an IP Address has been changed on the system.




Figure 23 Native asynchronous replication example

3.2.2.1 Asynchronous replication throttling

In Dell EMC OE version 5.1 and later asynchronous replication traffic can be throttled to reduce the rate at which data is replicated to a destination system. Asynchronous replication throttling is configured at the replication connection level, which allows each remote system connection to be controlled independently of each other. Also, only outbound replication traffic to a remote system is throttled. This not only allows different throttles to and from a remote system over the replication connections, but also allows replication traffic to be throttled from a system running OE 5.1 and later to a system running an earlier release. This feature can be managed using Unisphere, Unisphere CLI, and REST API.

Asynchronous replication throttling is controlled by bandwidth schedules configured directly on a remote system connection. Bandwidth schedules can be configured when creating a replication connection to a new system or added and modified at any time on an existing connection. By default, no bandwidth schedules exist. Figure 24 below shows the **Bandwidth Schedules** step within the **Create Replication Connection** window. Here, one or more bandwidth schedules can be added when creating the replication connection. Available options on this screen include **Add**, **Delete**, **View/Edit**, **Move up**, and **Move down**.

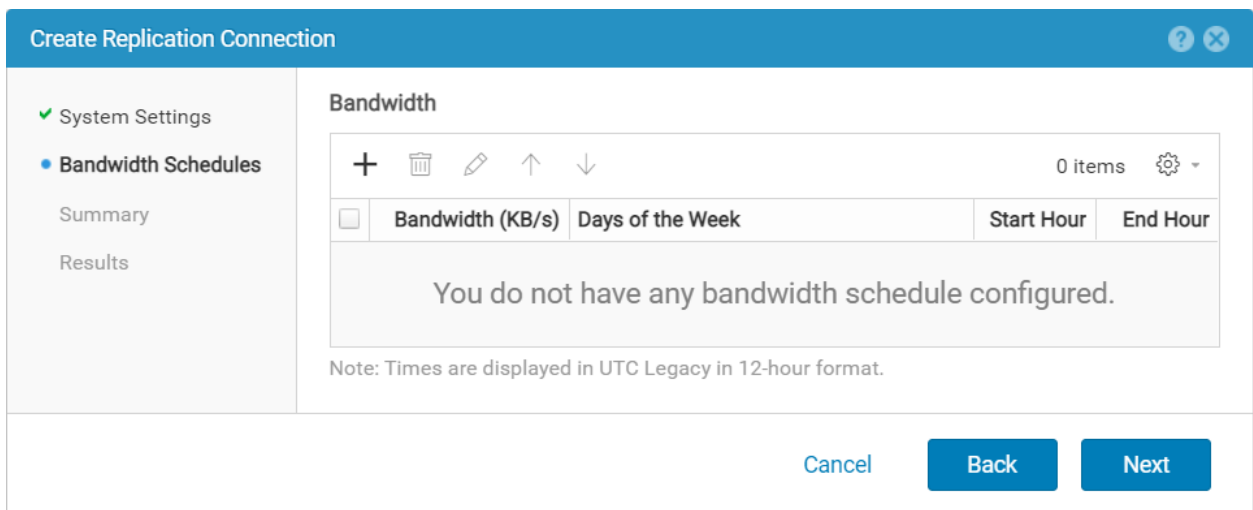


Figure 24 Create Replication Connection – Bandwidth Schedules step

When creating a bandwidth schedule, the following information can be customized:

- **Maximum Bandwidth:** Sets the maximum bandwidth in KB/sec for the bandwidth schedule.
 - When this value is left blank (default), the replication connection will use the available bandwidth of the connection for replication traffic.
 - When this value is set to 0 KB/sec, no replication traffic will be allowed to the destination system. This may impact the RPO of the replication sessions if 0 KB/sec is defined for long periods of time.

- When the maximum bandwidth is set to a custom value, throttling based on the maximum bandwidth value only occurs if the value of the throttle is configured lower than the speed of the network between the systems.
- **Days of the Week:** Sets the days of the week the bandwidth schedule will run.
 - When no checkboxes are selected (default), the bandwidth schedule will run on all days of the week.
- **Hours of the day:** Sets the Start and End Hour for when the bandwidth schedule will run.
 - When the Start and End Hour values are left blank, the schedule will run for all hours of the day.

Figure 25 below shows an example of the **Add Bandwidth Schedule** window. In this window the **Maximum Bandwidth** and **Hours of the day** have been customized. As the **Days of the Week** checkboxes are not set, this schedule will run on all days once created. When configuring the **Start Hour** and the **End Hour**, the timing entered directly depends on the Schedule Time Zone set on the system. As shown in this example, when the system's Schedule Time Zone setting is configured as UTC Legacy (default), the drop-down options must be configured with the UTC Legacy offset. Based on the selections, the timing is also displayed taking into account the current time zone offset of the computer Unisphere is being accessed from. This schedule will run as configured between 8:00 AM (UTC-04:00) and 5:00 PM (UTC-04:00). The Schedule Time Zone setting is discussed later in this section.

Figure 25 Add Bandwidth Schedule

Multiple bandwidth schedules can be created and can overlap. This allows replication traffic to be controlled differently depending on the day of the week and time of day. When multiple bandwidth schedules are present on the connection, the first schedule matching the current day and time will be enforced. The Move up and Move down arrows can be used to reorder the bandwidth schedules list as needed. If one schedule limits the replication traffic during a specific time of the day, and another is a limit for all days and times, the specific schedule should be listed first. Duplicate schedules can be created, but only the first occurrence of the schedule will be enforced. If duplicate schedules exist on the connection, a warning is displayed and a hyperlink to select the duplicates is presented so they can easily be deleted.

Figure 26 below shows an example configuration on the **Bandwidth Schedules** tab within the properties of the replication connection. In this window all bandwidth schedules for the replication connection are shown, along with the **Current bandwidth** limit, which lists the current limit in KB/sec. A note about the display times for the window is also shown.

In this example, on Saturdays and Sundays the replication connection uses all available bandwidth of the link as no limit is defined. On Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays, limits are created for

various times during the day. In this example the 7:00 AM – 9:00 AM and 4:00 PM – 7:00 PM schedules are listed and do not overlap. Both schedules will be enforced on the replication connection. The fourth schedule on the list is configured for 0 KB/sec between 7:00 AM and 7:00 PM. This schedule restricts the replication sessions from running on the connection. As this schedule overlaps with day and times defined earlier in the list, this schedule will only run during times that are not previously configured (9:00 AM to 4:00 PM). The last schedule in this example is set to run at all times on all days and enforces a 204800 KB/sec limit. As this is the last schedule in the list, all other schedules would run before this schedule is encountered. If this schedule was listed first, then no other schedules would run for the remote connection.

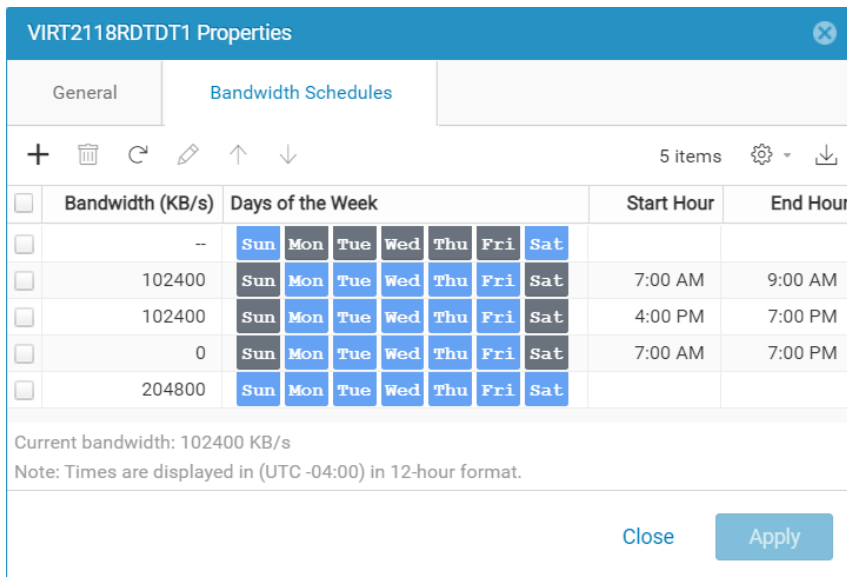


Figure 26 Replication Connection properties – Bandwidth Schedules tab

When a bandwidth schedule with a KB/sec limit above zero is defined, all active replication sessions on the replication connection will share the bandwidth allowance. Using the 102400 KB/sec limit as an example, all sessions actively replicating to the remote system will share the 102400 KB/sec limit. As replication session updates end and begin, the bandwidth allowance for each active session is adjusted. Bandwidth schedules are enforced at the replication session level not at the replication interface level and are in no way impacted by the number of replication interfaces used by the remote connection on the system.

3.2.2.2 Schedule Time Zone

Due to seasonal time changes in certain regions of the world, the bandwidth schedules can become out of sync with the local time. For example, a schedule starting at 8:00AM local time in one part of the year may start at 9:00AM in another part of the year. To avoid this issue, a user must manually validate and correct the settings of the bandwidth schedules when a seasonal time change has occurred.

In Dell EMC Unity OE version 5.1 and later the Schedule Time Zone option can be set to correct bandwidth schedule timing. This feature automatically adjusts the timing of bandwidth schedules with the seasonal time changes of the assigned time zone. The Schedule Time Zone option not only applies to asynchronous replication throttling, but also snapshot schedules. This option can be found in Settings > Management > Schedule Time Zone, as shown in Figure 27. A link to this page is also available within the Create Schedule page, as shown in Figure 25.

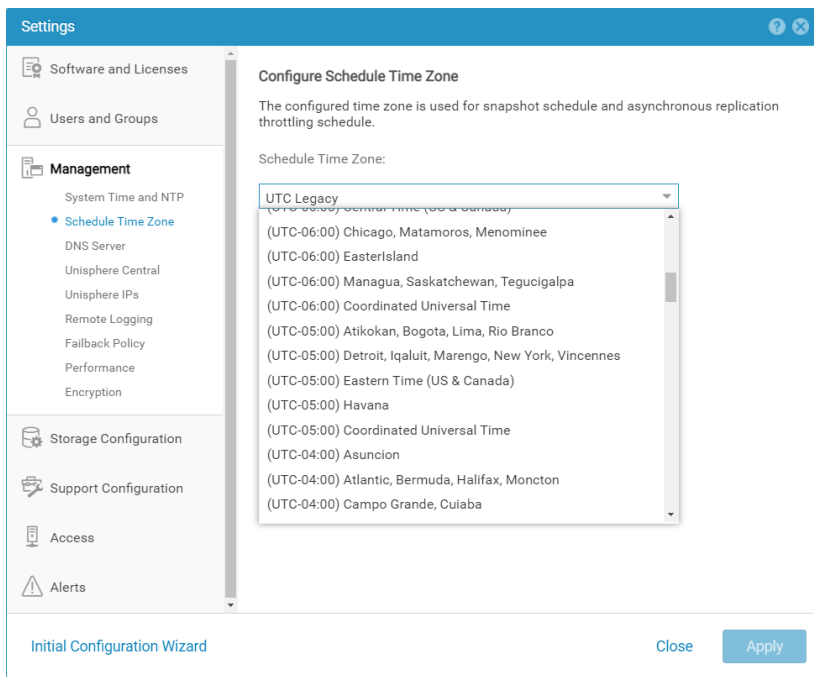


Figure 27 Configure Schedule Time Zone

After installing or upgrading to the 5.1 release or later, the system is set to UTC Legacy setting by default. UTC Legacy keeps the default handling and does not adjust schedules due to seasonal time changes. The user has the option to choose from over 100 times zones, covering over 300 cities around the world. The time zone of a system can be changed at any time, which may be required if the system is relocated from one site to another.

After changing the time zone, it is recommended to verify the timing is correct for bandwidth schedules that are in use on the system. When the Schedule Time Zone option is changed, the timing of schedules is not updated to reflect the new time zone. This may cause bandwidth throttle to run at the incorrect time. If the timing is no longer correct for a bandwidth schedule, it should be modified to reflect the required timing.

In the following example, the Schedule Time Zone is set to (UTC-05:00) Eastern Time (US & Canada). The Add Bandwidth Schedule window displays the time zone setting on the system. When configuring the Start Hour and End Hour, the timing will be based on the time zone of the system. At the timing of this update, daylight saving time / summer time is in effect. The Add Bandwidth Schedule window also displays the schedule time based on the current time zone offset of the computer Unisphere is being accessed from. In this case the current time is based on UTC-4:00. Both times are displayed to confirm proper timing.

Figure 28 Add Bandwidth Schedule

3.2.3 Replication Sessions

A replication session is used as a logical link between storage resources participating in replication. When utilizing asynchronous remote replication, a configured replication connection is used to transfer data to the remote system. A replication connection is not required when asynchronously replicating data between Pools within a system. When a replication session is created in Unisphere, a storage resource of the same size and type is created on the destination Pool, whether it is locally within the same system or on a destination system. Snapshot schedules are not replicated on the destination storage resource.

Asynchronous replication image synchronizations are triggered by a user defined Recovery Point Objective (RPO) or at any time manually by the user. The following characteristics define asynchronous replication:

- Writes to a storage resource are saved to the source storage resource and acknowledged to the host before being replicated to the destination storage resource.
 - When a host writes to a storage resource, the write is saved and tracked by the replication session. This data will be replicated at a later time.
- A user defined RPO is utilized to define the maximum amount of time between scheduled synchronizations.
 - During the time between synchronizations, new data is only saved on the source storage resource. The RPO is the maximum amount of data the user is willing to lose in the event of a disaster, measured in time. The RPO determines how often synchronizations occur at a minimum.

Manual replication operates the same as asynchronous replication, except all replication updates must be started manually by the user as opposed to using an RPO. Updates occur when the sync operation is initiated on the session, which is described later. When writing to a storage resource configured with a manual replication session, the data is stored locally and acknowledged and only replicated when a sync is issued.

When an asynchronous replication session is created, a full synchronization of the source and destination storage resource may be required. If replication is configured when a new resource is being created, the synchronization is quick as no data needs to be copied to the destination storage resource. If replication is added to an existing storage resource, a full synchronization may be required between the source and destination storage resource if no common base snapshot exists. In Dell EMC Unity OE 5.1 and later, a replicated snapshot may be used to avoid a full synchronization under certain circumstances. This feature is discussed in the Snapshots section of this paper. Writes occurring during the initial synchronization period are

not copied to the destination storage resource at this time, but rather tracked for a later synchronization. Once the initial synchronization is complete, a common base is established between the source storage resource and the destination. When creating a manual replication session, an initial synchronization does not automatically start. To synchronize the local and remote images, a manual sync needs to be initiated. Host write operations which occur after the initial synchronization are acknowledged with the host normally, and no data is replicated to the destination until the next sync. Manually at a later time, or at the RPO interval for asynchronous replication, all changes made to the source storage resource since the last synchronization will be replicated to the destination. A new common base is then established. If a failure is encountered on the source, all data not copied to the destination will be lost as the changes have not been copied to the destination.

3.2.4 Asynchronous Replication Internal Snapshots

Asynchronous replication in Dell EMC Unity leverages Unified Snapshots to maintain the common base images explained in the previous section. Figure 29 below shows how Unified Snapshots are used with asynchronous and manual replication. This section assumes no common base snapshots exist at time of session creation.

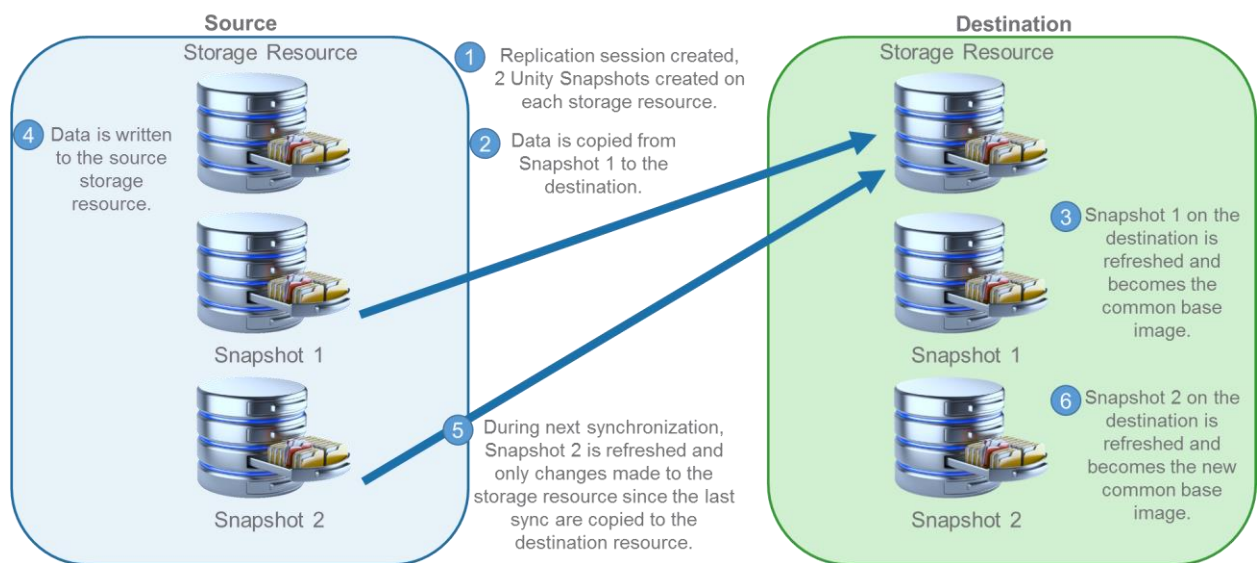


Figure 29 Asynchronous and Manual Replication Theory

Dell EMC Unity Asynchronous Replication operates in the following way:

1. When a replication session is created on a storage resource, two internal Unified Snapshots are created on the source and destination storage resources. At time of creation, Snapshot 1 and Snapshot 2 on the source have the same content as the source storage resource.
2. Data is then replicated from Snapshot 1 to the newly created destination storage resource. This is the initial synchronization of the source and destination storage resources and is a full copy of all the data.
3. Once the initial synchronization is complete, Snapshot 1 is refreshed on the destination storage resource. Snapshot 1 on the source and destination storage resource contains the same information and represent the point-in-time in which the synchronization started. Snapshot 1 on each system is now a common base for the replication session.
4. Over time the host application writes new data to the source storage resource.

5. During the next update, either manually started or by the RPO with asynchronous replication, Snapshot 2 on the source storage resource is refreshed to reflect the current point-in-time view of the source storage resource. All changes since the last update of the destination are copied to the destination storage resource.
6. After the incremental copy is complete, Snapshot 2 on the destination storage resource is refreshed to reflect the current information located in the destination storage resource. Snapshot 2 on the source and destination contains the latest information and are the latest common base image for the replication session.

Each time the RPO is reached or a manual update is started, the common base image will alternate between Snapshot 1 and Snapshot 2.

Snapshots used for asynchronous replication behave the same as user Unified Snapshots and are based on redirect on write technology. Space needed to preserve the point-in-time snapshot is allocated from the same Pool as the source storage resource. Although user snapshots and replication snapshots share the same technology, replication snapshots have restrictions on their usage. Replication snapshots can be viewed in Unisphere, but user operations, such as restore, or mount operations are not allowed. Snapshots allocated for replication purposes do not count towards user snapshot maximums.

3.2.5 Storage Resources

When configuring asynchronous replication, the source and destination storage resource must be of the same type. In Dell EMC Unity, native asynchronous replication is supported on the following storage resources:

- LUNs
- Thin clones
- Consistency groups
- VMware VMFS Datastores
- File systems
- NAS Servers
- VMware NFS Datastores

In Dell EMC Unity, asynchronous replication for LUNs, thin clones, and VMware VMFS Datastores function the same. When configuring asynchronous replication on a LUN in Unisphere, a single replication session is created, and the destination storage resource will be created the same size and type as the source storage resource. When configuring replication session on a thin clone, the destination storage resource will be a regular storage resource and not a clone. While replication is configured, you can extend the size of LUNs, thin clones, and VMware VMFS Datastores, and the changes will be reflected on the destination storage resource after the next sync. Options such as the LUN's, thin clone's, or VMware VMFS Datastore's name or tiering policy may be configured differently between systems.

In Dell EMC Unity, a consistency group is a storage resource which contains one or more LUNs within a storage system. consistency groups help organize storage resources allocated for a particular host or hosts. Consistency groups are treated as a single entity when they are replicated, meaning that a single replication session is created for the entire consistency group no matter how many LUNs it contains. When replication is configured in Unisphere for a consistency group, the destination storage resource and its contents are created automatically. While a consistency group is part of an asynchronous replication session, LUNs within the consistency group can be expanded. All changes to LUNs within a consistency group will be reflected on the destination image after the next completed synchronization. LUNs cannot be added or removed while replication is configured. When pausing or resuming replication on a consistency group, the entire group is affected by the replication operation.

When replicating existing file systems or VMware NFS Datastores with asynchronous replication, you must first configure replication for the NAS Server it is mounted on. If replication is configured in Unisphere for the NAS Server, all file systems and NFS Datastores on the NAS Server will also be replicated to the destination. Replication sessions for resources which do not require replication can be deleted later. All replication sessions automatically configured when the NAS Server is replicated will have the same RPO. The RPO for the individual replication sessions can be changed later. While replicated, all size changes to the file systems and NFS Datastores will be reflected on the destination after the next synchronization. If a different IP address for NAS Server access needs to be specified on the destination NAS Server, you may specify an Override Address by editing a NAS Server Network Interface.

3.2.6 Replication Roles

For asynchronous replication to operate, two storage resources are required for replication:

- A source storage resource which will be replicated
- A destination storage resource to which the source is copied

When a replication session is created in Unisphere, the destination storage resource is automatically created. In code versions prior to the OE 5.1 release, you may not choose a previously provisioned storage resource as a replication destination when configuring a session in Unisphere. In Dell EMC OE version 5.1 and later, Unisphere includes an option to reuse a destination resource if one exists. This feature is discussed in detail in the Snapshots section. Upon creation, the destination resource is marked as a destination image. This restriction blocks Read/Write access on the destination storage resource. To view the data contained in the destination storage resource, you may take a snapshot of the resource and provide host access to the snapshot.

In Unisphere, you can easily determine which storage resources are replicated and which are destination images from any of the storage resource pages, such as the LUNs tab on the Block page. While in this tab, select the gear icon, and hover over the Columns option to view the available columns that can be viewed. Select the check boxes for “Replication Type” and “Restricted Replication Access”. Replication Type displays what type of replication the storage resource is participating in, whether it is None, Remote, or Local. The Restricted Replication Access column will display “Yes” if the storage resource is labeled as a replication destination resource, or “No” if it is not. If a replication session is deleted, the destination storage resource will still be labeled as a replication destination image. The replication destination label must be edited manually over Unisphere CLI before the resource is allowed to receive I/O. For example, to remove the replication destination setting from a LUN, use the `uemcli /stor/prov/luns/lun -id <value> set - replDest no` command. For more information about removing the replication destination setting on other storage resource types, consult the *Dell EMC Unity Unisphere Command Line Interface User Guide*.

In Dell EMC Unity OE versions 5.1 and later, Source, Destination, and All filtering buttons on the replication sessions page and various storage resource pages help the user easily identify replication source and destination resources/sessions without adding columns to the view. When All is selected, all resources/sessions on the current page are displayed. When Source is selected on a resource page, all resources that are the source of a replication session are displayed. Resources that are not replicated are also shown when Source is selected. When Source is selected on the replication sessions page, only replication session originating on the system are shown. When Destination is selected on a resource page, only resources that are the destination images of a replication session are shown. While on the sessions page, Destination will only show the sessions replicating to the current system. Also, sessions that are part of local replication are displayed regardless of which view is selected.

3.3 Replication Operations

Several operations are available to manipulate replication sessions as needed. Not all operations are available at all times, as some depend on the session being in a particular state. Also, certain operations will perform differently depending on which system they are issued on, source or destination. Only one replication operation can be issued and running at a particular point in time.

3.3.1 Pause and Resume

Pause and resume are used to stop and start replication between the resources for a particular replication session. In Unisphere, the pause operation can only be issued from the source system. If the session is paused while a sync is in progress, all incremental changes on the destination are discarded and the destination is reverted back to the common base. All I/O is tracked when the replication session is paused. When the session is resumed, replication will resume and synchronizations to the destination storage resource will be based on the current RPO if asynchronous replication is being used. When a manual replication session is resumed, no replication updates occur until a manual sync is issued. While an asynchronous and manual replication session is paused, no sync updates can be performed on the session.

The resume option is also available after a failover operation. If a resume operation is initiated, replication begins running in the reverse direction. After a failover, it is possible for changes to be made to the destination resource. When running the resume operation, a checkbox is available to resync and overwrite any data written to the remote system. This checkbox may be necessary in situations where there are different changes on each system. The administrator must acknowledge they want replication to overwrite any changes made to the other system before the resume operation can complete successfully, as shown in Figure 30.

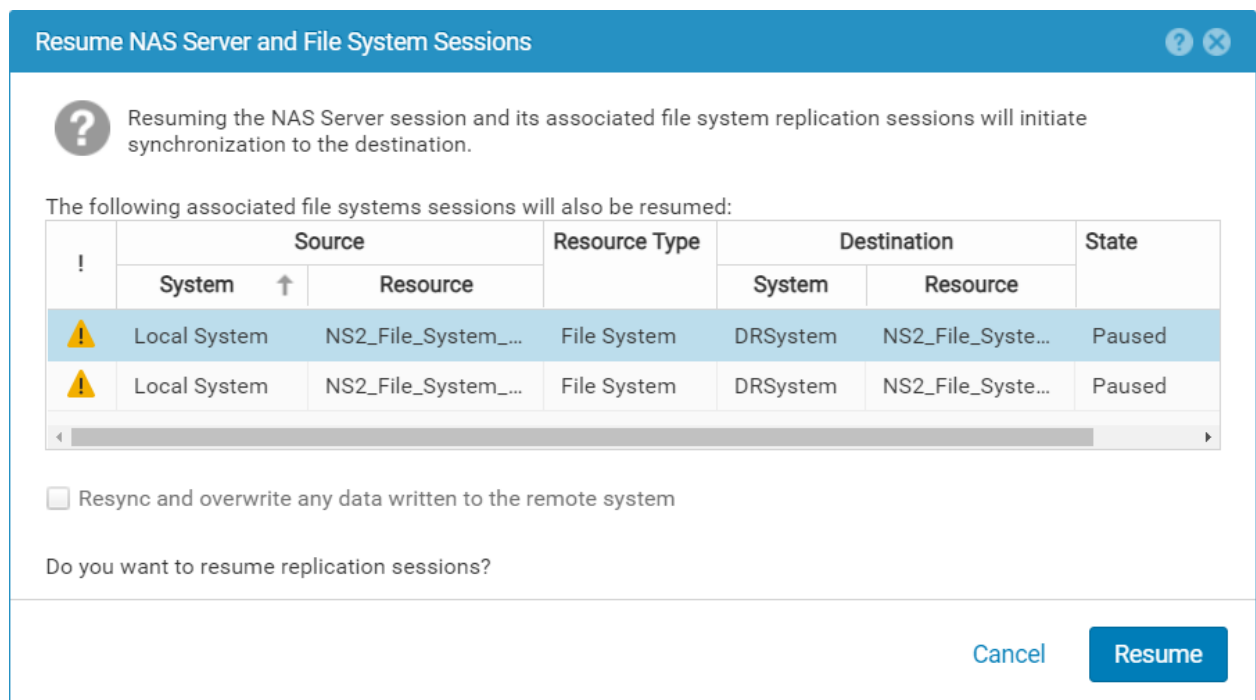


Figure 30 Resume with Resync

Pausing a replication session may be done for several reasons to stop updates from occurring to the destination. Some reasons include the need to power off the source or destination system for planned maintenance, a configuration change on the network between systems, or interface changes on either

system. Another reason may be to physically move a system from one data center or site to another. In certain circumstances, configuring replication and synchronizing the data between systems may be done within the same site, then the destination system is later moved to its final destination. Depending on the network speeds between sites, this may save time and avoid lengthy initial synchronizations.

3.3.2 Sync

With asynchronous replication, updates to a destination storage resource happen at a set interval based on the defined RPO. At any point in time when replication is active, and an update is not already occurring, a sync operation can be issued to synchronize the latest changes to the destination resource. The sync operation is also used to update a remote image when manual replication is configured. After the sync operation is selected, all data changed since the last update will be copied to the destination storage resource. Issuing a manual sync operation also updates a destination image's size when the source image size has been changed.

3.3.3 Failover with Sync, Failover, and Failback

When using asynchronous replication, the failover with sync option allows you to failover the session after first completing a last synchronization between the images. The failover with sync option is used in planned failover events. The failover with sync option is only available on the source storage resource when the replication session is operating normally and will result in a short period of data being unavailable during the failover operation. Before issuing the failover with sync operation, it is suggested to issue a manual sync first, to reduce the amount of data to copy during the failover with sync. When issuing the failover with sync, it is highly suggested to quiesce I/O to the source image first. After the synchronization completes, the destination storage resource is available for production I/O and the original source will no longer allow Read/Write I/O. If host access is configured on the destination resource, hosts can access the data at this time. Replication will not be resumed in either direction when the failover with sync operation is used.

The failover option is only available on the destination of the replication session. When issuing a failover operation on a replication session from the destination system, an unplanned failover is initiated, and a final synchronization of the data from the source storage resource is not completed. An unplanned failover assumes a disaster has occurred on the production system, and the destination image is made Read/Write available. When failover is selected on a destination resource of a replication session, Read/Write access is removed from the original source if the source is available to receive management commands. The replication session will also pause and not automatically switch the direction for replication. The replication session is left in this state until the user issues another replication operation. If I/O occurs to the original destination resource while in this state, the data must be replicated to the original source when the source becomes available.

Starting with Dell EMC Unity OE version 5.0, an unplanned failover operation can be initiated even when the replication is in a "Paused" state. Previously, unplanned failover operations were not allowed while the session was in a "Paused" state. Any changes made on the source system while the session is paused are not replicated to the destination.

When the failover with sync or failover option is used, the failback option becomes available. Failback restores the replication session to the state before failover with sync or failover was issued. After a failover, it is possible for changes to be made to the destination resource. When running a failback operation, the following options are available:

- **Keep local data changes by updating the remote resource:** Preserves the changes that were made on the destination site by replicating them back to the original source

- **Keep remote data by discarding all local data changes:** Discards any changes made on the destination site after the failover by replicating in the original direction and overwriting any changed data

These options may be necessary in situations where there are different changes on each system. The administrator must acknowledge which changes they want to keep before the failback operation can complete successfully.

Alternatively, running a resume, failover with sync, and another resume from the other system results in the same behavior as the first option. Once the failover with sync operation completes, the original configuration and replication direction is applied. Note that once a resume operation is initiated, the failback option is no longer available since replication is already running in the reverse direction. The failback options are shown in Figure 31.

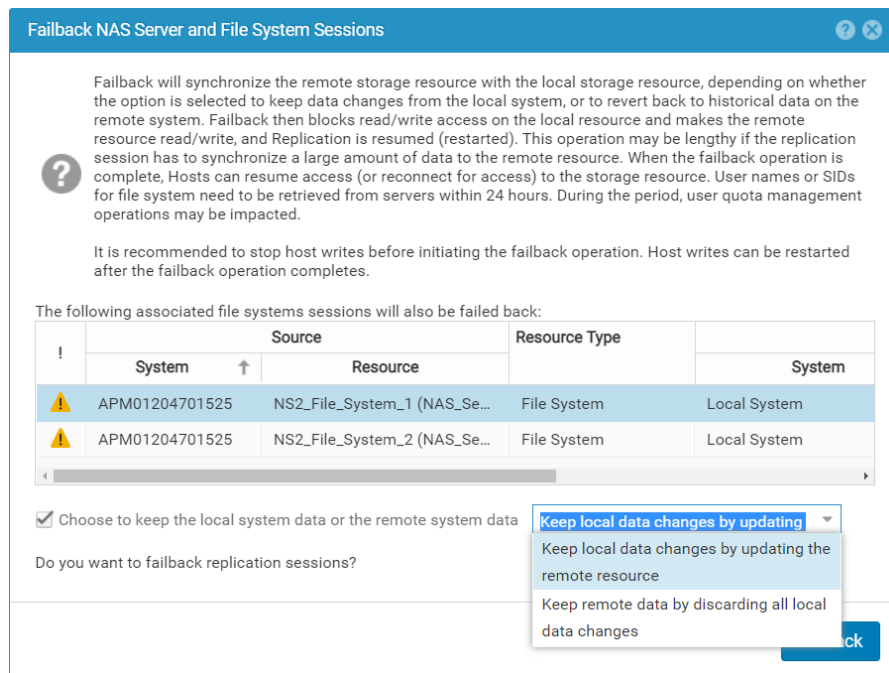


Figure 31 Failback options

When the original source resource becomes available, the user has the option of issuing the Resume operation to replicate data to the original source. Clicking resume will reverse the direction of replication and resume updates based on the user RPO from the original destination resource to the original source. A manual sync must be issued when using manual replication. To return replication to the original replication configuration, the failover with sync operation can be issued on the original destination. A final synchronization of the data is performed, and the replication will be failed over to the original source. It is suggested to quiesce all I/O before issuing the failover with sync operation. The original source will become Read/Write available and host access will be removed from the destination resource. You must issue Resume on the source to resume replication from the source to destination.

3.3.4 Delete

Deleting a replication session can be issued on the source system or destination system, but it is recommended that the operation be issued on the source when the source is available. When there are no issues in the configuration and a delete operation is issued on the source system, the replication session will

be deleted from both the source and destination system. The destination storage resource is not automatically deleted when the replication session is deleted. If the delete operation is issued while the destination system cannot be reached, the session will need to be deleted from the destination system manually. If the delete operation is issued from the destination system, the source session is left configured and must be deleted manually. Once a replication session is deleted, a full sync may need to occur if replication is reconfigured. In Dell EMC Unity OE 5.1 and later, a replicated snapshot may be used to avoid a full synchronization under certain circumstances. This feature is discussed in the Snapshots section of this paper.

A delete operation can also be issued for a replication connection. A replication connection can only be deleted after all configured replication sessions using the connection have been deleted.

3.3.5 Group Operations

Dell EMC Unity OE version 4.2 introduces the feature group operations at a NAS Server level. Group operations allow for the replication operation to be propagated from the NAS Server to the file systems automatically and works with local and remote asynchronous replication. The group operations are available for the following operations: Failover, Failover with Sync, Failback, Pause, and Resume. These operations remain available at the file system level. All other operations (Create, Sync, Delete, and Modify) remain individual commands. Before issuing a failover with sync operation, it is suggested to issue a manual sync first, to reduce the amount of data to copy during the failover with sync.

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 performed simultaneously at both sides of a replication session can cause the group replication session to enter an unhealthy state. Although a group asynchronous 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. Group operations skips file system replication sessions that are in a paused, error, or non-replicated state. Figure 65 in the Viewing the Replication Sessions section provides an example for the group Failover operation.

3.4 Supported Replication Configurations

Dell EMC Unity's Native Asynchronous Replication features allow supported storage resources to be replicated locally within the same system, or remotely between systems. The following sections outline the supported configurations for asynchronous replication. For more information about which systems are supported for asynchronous replication, please review Appendix C: Replication Support Across Platforms.

3.4.1 Local Replication

Dell EMC Unity's Native Asynchronous Replication feature allows Block and File resources to be replicated locally within the same system. When replicating file systems or VMware NFS datastores, the NAS Server must also be replicated. When configuring local replication, the source and destination storage resources cannot exist within the same Pool. By replicating to a different Pool, a storage resource is protected against the unlikely event the source Pool encounters a data unavailable situation. All asynchronous replication operations are supported when local replication is configured. Replication Connections and Interfaces are not required when local replication is configured. Figure 32 below shows an example of local replication.



Figure 32 Local asynchronous replication

3.4.2 Remote Replication

Dell EMC Unity's Native Asynchronous Replication feature is supported in many different topologies, and deployment models will vary depending on the requirements of the configuration. At a system level, the following configurations are supported:

System Level

- One-Directional
 - A single source system replicating to a single destination system
- Bi-Directional
 - A two-system topology in which each system acts as a replication destination for the peer's production resources
- One-to-Many
 - A system topology in which a single system replicates multiple resources, each to a different remote system
- Many-to-One
 - A system topology in which multiple systems replicate their respective resources to a single system

Figure 33 is a graphical view of the supported topologies listed above. Note the figure uses LUNs to represent the storage resource, but file resources are also supported in these topologies. In all topologies mentioned, a Dell EMC Unity All Flash system, a Dell EMC Unity Hybrid system, or a Dell EMC UnityVSA system can be used for any system in the configuration. For a list of other Dell EMC storage systems supporting asynchronous replication to and from Dell EMC Unity, please see Appendix C. Replication Interfaces are required to be configured on each system participating in remote replication. A Replication Connection also needs to be configured between each system pair to allow replication sessions to be configured. Asynchronous replication allows for many different deployment models to meet the needs of an organization.

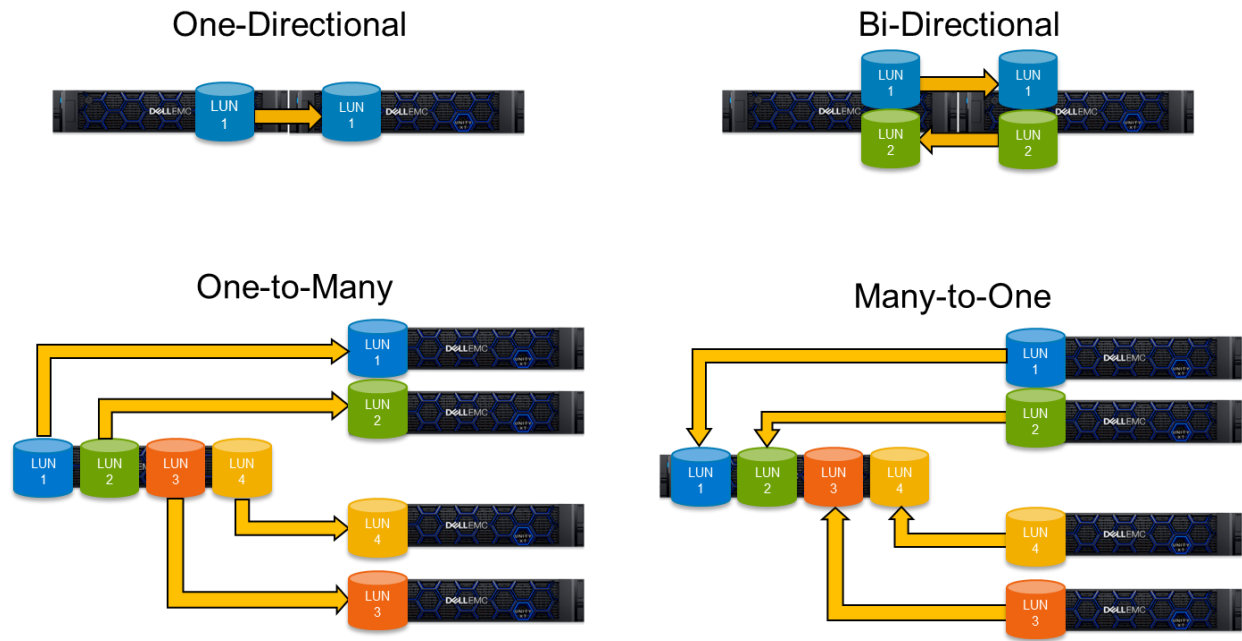


Figure 33 System level asynchronous replication topologies

One-Directional replication is typically deployed when only one of the systems will be used for production I/O. The second system is a replication target for all production data and sits idle within the same data center or a remote location. If the need arises, the DR system can be placed into production and host production I/O. In this scenario, mirroring the production system's configuration on the DR system is suggested, as each system would then have the same performance potential. For physical systems this would mean mirroring the drive configurations and Pool layout, while on Dell EMC UnityVSA systems this would mean configuring similar Virtual Drives and Pools.

The Bi-Directional replication topology is typically used when production I/O needs to be spread across multiple systems or locations. The systems may exist within a single data center or in different, remote locations. When using this replication topology, production I/O from each system is replicated to the peer system. If there is an outage, one of the systems can be promoted as the primary production system, and all production I/O can be sent to it. Once the outage is addressed, the replication configuration can be changed back to its original configuration. This replication topology ensures both systems are in use by production I/O at all times.

The One-to-Many replication topology is usually deployed when production exists on a single system, but replication needs to occur to multiple remote systems. This replication topology can be used to replicate data from a production system to a remote location to provide local data access to a remote team. At the remote location, Dell EMC Unity Snapshots can be used to provide host access to the local organization or test team. In this topology, any combination of Dell EMC Unity All Flash systems, Dell EMC Unity Hybrid systems, and Dell EMC UnityVSA systems can be used. The production system may be an All Flash system replicating to multiple physical All-Flash or Hybrid systems and/or Dell EMC UnityVSA systems.

The Many-to-One replication topology is deployed when multiple production systems exist and replicating to a single system to consolidate the data is required. This topology is useful when multiple production data sites exist, and data must be replicated from these sites to a single DR data center. One example of this configuration is Remote Office Branch Office (ROBO) locations. A Dell EMC UnityVSA may be deployed at each ROBO site, and all replicate back to a single All Flash or Hybrid Flash system. Utilizing Dell EMC UnityVSA at ROBO locations eliminates the need for a physical Dell EMC Unity system at each site.

For the One-to-Many and Many-to-One replication topology examples in Figure 33, One-Directional replication is depicted. One-Directional replication is not a requirement when configuring the One-to-Many and Many-to-One replication topologies. Each individual Replication Connection can be used for bi-directional replication between systems, which allows for more replication options than what is depicted.

3.4.3 Advanced File Remote Replication

Starting with Dell EMC Unity OE version 5.0, asynchronously replicated file resources can also be configured in advanced replication topologies. This allows for configurations such as fan-out and cascading replication at the granularity of a NAS Server and its associated file resources. Prior to this, these configurations were not supported since asynchronous replication was limited to a single destination. The ability to replicate and store the same dataset on multiple systems provides additional data protection and enables use cases such as content distribution. This feature does not support synchronous replication and is not available for block resources.

To use this feature, all systems in the topology must be running Dell EMC Unity OE version 5.0 or later. Replication between Dell EMC Unity OE 5.0 and 4.x is still supported in a one-directional configuration. This feature is available on both physical and virtual Dell EMC Unity systems. The following advanced topologies are supported:

File Resource Level

- One-Directional
 - A single file resource replicating to a single destination resource
 - Example: A → B
- Fan-Out
 - A single file resource replicating to up to four different destination systems
 - Example: A → B and A → C
- Cascade
 - A single file resource replicating to a second system and from there, replicating to a third system
 - Example: A → B → C
- Mixed
 - Leveraging a combination of both cascade and fan-out, or vice versa
 - Example: A → B and B → C and B → D

Fan-out replication configurations allow a file resource to be replicated to up to four different destination systems. Cascading replication configurations allow a destination file resource to be replicated again to another system. When the RPO is reached from B → C, the data that is replicated is based on the last completed sync from A. That means if there is an ongoing sync from A → B when the B → C RPO is reached, the new changes from A are not replicated until the next RPO between B → C is reached.

With cascaded configurations, you can only failover to the directly adjacent site. For example, assume a cascaded configuration from A → B → C. In this configuration, you cannot failover directly from A → C. To accomplish this, you would need to failover from A → B and then failover from B → C. Alternatively, you can failover from B → C directly if a failover with sync from A is not required.

A combination of fan-out and cascading can be configured as long as each resource does not exceed four total replication sessions. When designing replication topologies, it is important to note that a maximum of four total replication sessions can be created for each resource. For example, if you create a cascade from A → B → C, then C can only fan-out to three other systems due to the existing session from B → C. The tested limit for the number of cascaded sessions is three, but there is no hard limit.

Figure 34 shows an example of a supported advanced replication topology. An example of a cascaded topology is replicating from Hopkinton → Boston → London. An example of a fan-out topology is San Francisco → Tokyo and Mexico City. Also shown in the figure is the ability for each session in the topology to have its own custom RPO configured. The valid range for the RPO value is 5 minutes through 1440 minutes (24 hours).

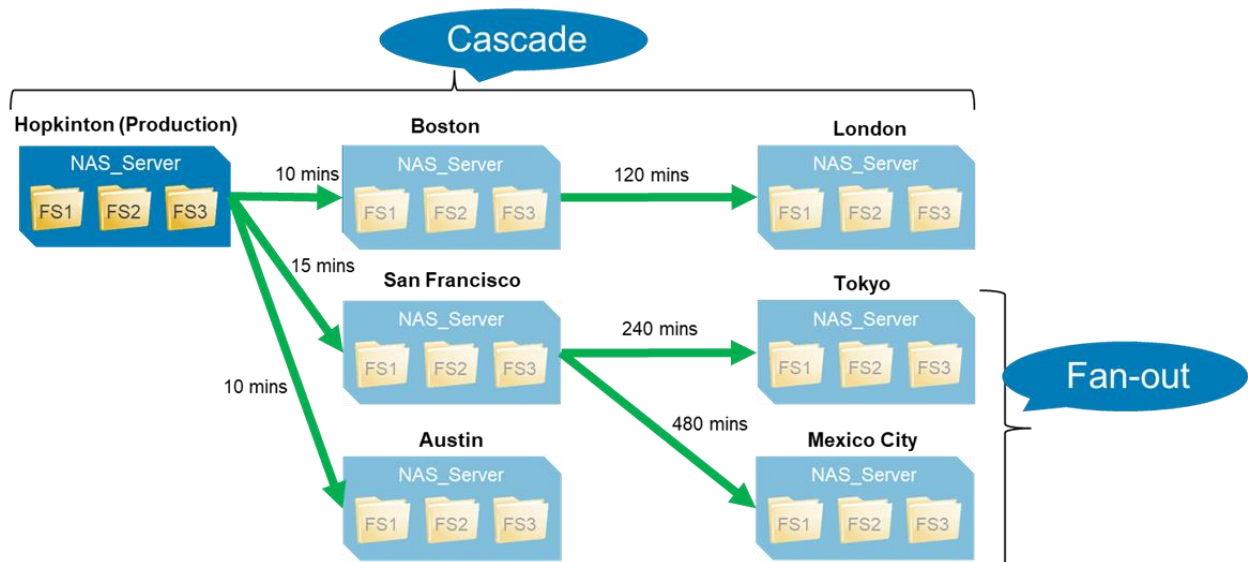


Figure 34 Advanced Replication Topologies

Advanced replication topologies can be used in conjunction with Proxy NAS Servers to enable access from any one of the destination systems. This is useful for use cases such as DR testing, test/dev, and analytics. Also, NDMP can be enabled on any one of the systems for backup operations.

Once a replication session is configured between two systems, a second replication session cannot be configured using the same two systems. A fan-out configuration requires each destination system to be unique physical or virtual systems. However, one of the sessions can be configured for local replication to a different pool within the same system. The local replication session counts towards the maximum limit of four sessions for that resource.

It is prohibited to configure a replication session to a resource that is already a replication destination for another resource. For example, assume a fan-out topology with $A \rightarrow B$ and $A \rightarrow C$. In this configuration, you cannot create a replication session from $B \rightarrow C$ as that results in multiple sessions writing to the same destination. You also cannot replicate a resource back to its original source system.

It is also prohibited to create a configuration or run operations that result in multiple systems replicating to the same destination resource. For example, assume a cascaded topology with $A \rightarrow B$ and $B \rightarrow C$. If you initiate a failover from $B \rightarrow C$, you cannot initiate a resume operation to start replicating from $C \rightarrow B$.

After advanced replication is configured, it is crucial to properly document the topology for future reference. This is very valuable during disaster scenarios where failing over the correct session is necessary to restore data access and minimize downtime. Each system is only aware of the systems that it is replicating to, so it is unable to provide an end-to-end topology view. Remember to keep this updated if any systems are added, removed, or has its role changed in the topology. For example, if you have a fan-out from $A \rightarrow B$ and $A \rightarrow C$, and a failover is initiated from $A \rightarrow B$. Afterwards, if you resume the failed over session from $B \rightarrow A$, this becomes a cascaded configuration from $B \rightarrow A \rightarrow C$.

Management of advanced replication sessions is performed at the NAS Server level and are automatically propagated to all associated file system sessions that are replicated to the same system. These operations include pause, resume, sync, failover with sync, failover, and failback. When failing over sessions in an advanced topology, it is also important to avoid failing over multiple sessions as that could result in a duplicate IP scenario. For example, you can create a fan-out configuration from A → B and A → C. Initiating failover operations on both B and C would result in both of those sites turning to production mode and cause a duplicate IP conflict. This could also happen if you initiate a failover from B → C in a cascaded configuration, since both A and C would be running in production mode.

After a failover, it is possible for changes to be made to the destination resource. For example, assume a fanout configuration from A → B and A → C, and a failover from A to B has occurred. In this case, both A and B could have different changes so a normal failback operation fails. The admin must check the box to either preserve any changes made or discard them by overwriting the data to successfully failback. The available options are:

- **Keep local data changes by updating the remote resource:** Preserves the changes that were made on B by replicating them back to A, which will also be replicated to C due to the fanout configuration
- **Keep remote data by discarding all local data changes:** Discards any changes made on B after the failover by replicating from A → B and overwriting any changed data

Using the same example as above, a normal resume operation also fails since both A and B could have different changes. To successfully resume the session, the admin must check the box to resync and overwrite any data written to the remote system. In this case, B starts replicating its changes to A and overwrites any changes that were made to A. After the resume operation completes, the topology is transformed into a cascaded configuration where B → A → C.

In each of these examples, the original configuration is restored after the failback or resume operation completes. All the replication updates leverage the internal replication common base snapshots to replicate only the changed data.

3.4.3.1 Snapshot replication with OE versions prior to the 5.1 release

With advanced replication configured, snapshot replication is also supported. For a single storage object, snapshot replication can only be enabled on one session at a time when using advanced replication. For example, if you create a cascade from A → B → C, then either the A → B or B → C session can have snapshot replication enabled, but not both at the same time, as shown in Figure 35.

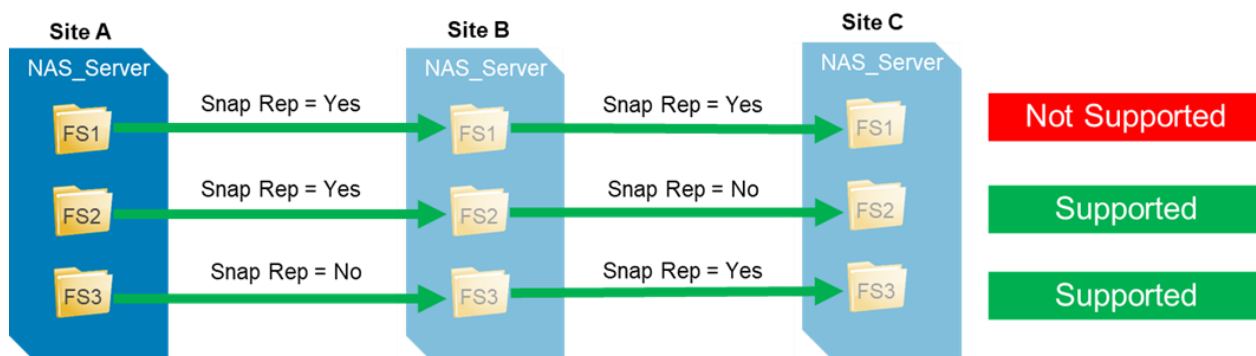


Figure 35 Snapshot Replication – Pre-OE 5.1 Configurations

In this configuration, it is still possible to replicate a snapshot from A → C, but a workaround is required. To accomplish this, follow the procedure below:

1. Enable snapshot replication from A → B
2. Replicate the snapshot from A → B
3. Disable snapshot replication between A → B
4. Enable snapshot replication between B → C
5. Replicate the snapshot from B → C

To enable or disable the snapshot replication feature, the replication session must first be paused and then the operation must be performed on the source system. This operation can be scripted using UEMCLI or REST API if this needs to be done often.

When configuring a cascaded replication topology that includes four or more systems, it may be possible to enable snapshot replication on multiple sessions simultaneously. This can be accomplished by enabling snapshot replication between A → B, leaving snapshot replication disabled between B → C, and also enabling snapshot replication between C → D. In this configuration, the replicated snapshots for the separate sessions are not consistent with each other.

With a mixed topology, it may be possible to have snapshot replication enabled on multiple sessions. For example, you can create a fan-out from A → B and A → C along with a cascade from C → D. With this configuration, you can enable snapshot replication on both A → B and from C → D. These configurations are allowed because from each site's point of view, there is still only one session that has snapshot replication enabled. An example of this configuration is shown in Figure 36.

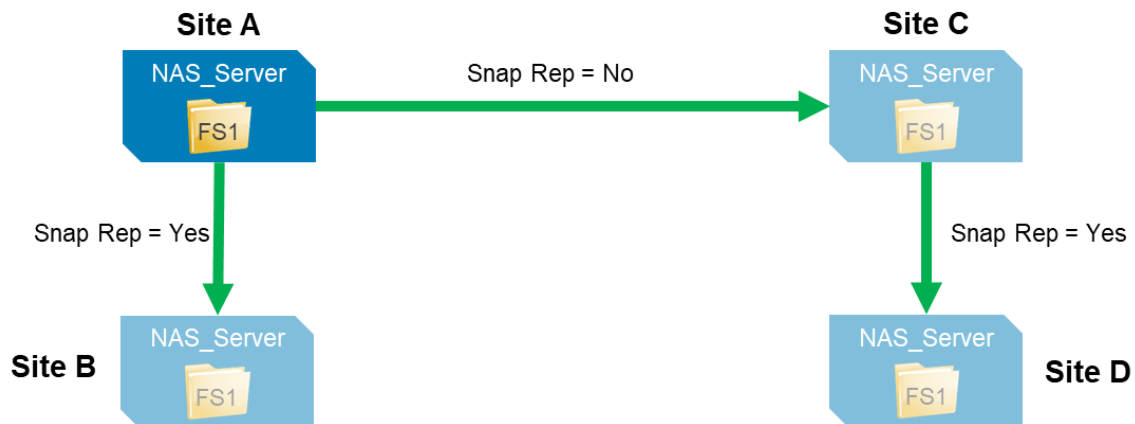


Figure 36 Dual snapshot replication configuration – Pre-5.1 OE versions

3.4.3.2 Snapshot replication with OE versions 5.1 and later

In Dell EMC Unity OE version 5.1 and later, snapshot replication is supported on all replication sessions within an advanced File replication environment. Advanced File replication topologies include fan-out and cascade topologies. To support snapshot replication, all systems within the advanced File replication topology must be running OE version 5.1 or later. By default, snapshot replication is not enabled on a replication session, but can be enable at time of session creation or modified at any time using Unisphere, Unisphere CLI, or REST API.

In Unisphere, the Cascade replicated snapshots option has been added to allow for snapshot replication across all replication session in an advanced File remote replication topology. This setting allows for snapshots to be replication from site A → B and the same snapshots replicated from site B → C. Using Figure

34 as an example, enabling the Cascade replicated snapshots option on replication sessions from Boston to London allows snapshots created and replicated from Hopkinton to Boston to also replicate to London. This setting is enabled and disabled on a per replication session basis.

3.4.3.3 Advanced File Remote Replication and MetroSync

Advanced replication configurations are also supported with MetroSync's ability to asynchronously replicate to a third site. From the third site, cascaded and fan-out topologies can be configured. Since this feature is not supported for synchronous replication, you cannot cascade or fan-out from the synchronously replicated systems. Figure 37 shows a supported advanced replication topology with MetroSync.

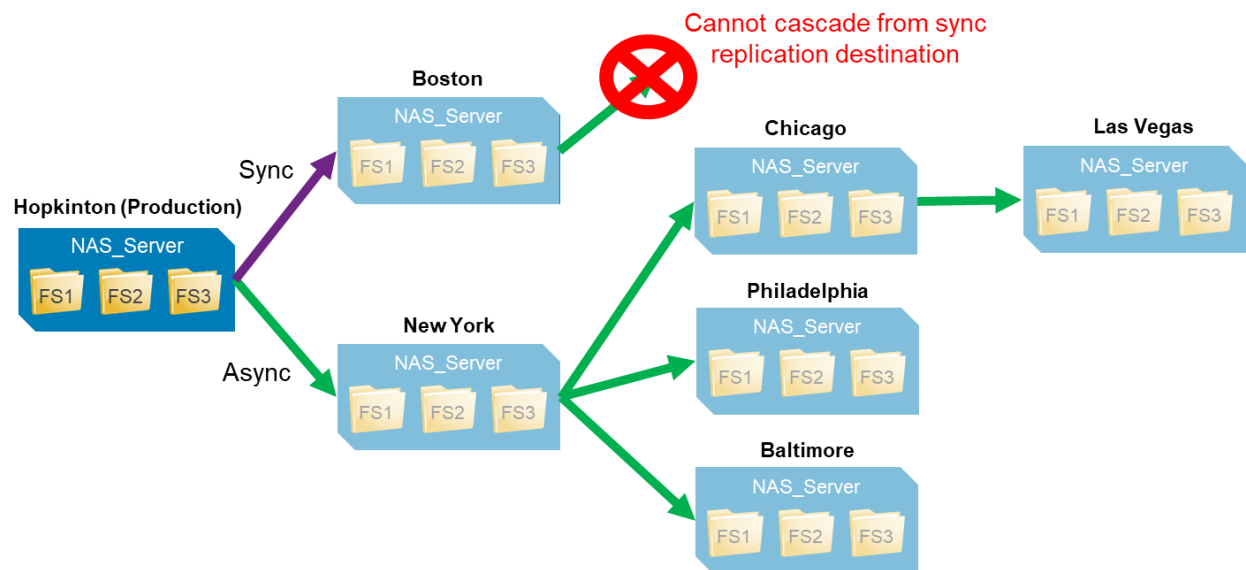


Figure 37 MetroSync with advanced replication

3.4.4 MetroSync

In OE version 4.4, Dell EMC Unity supports the ability to use the MetroSync feature, also known as native file synchronous replication, which allows a file resource to be replicated synchronously to one destination system and asynchronously to another destination system simultaneously. For more information, see the *Dell EMC Unity: MetroSync* white paper on Dell EMC Online Support.

3.5 Snapshot Replication

With Dell EMC Unity OE version 4.2, Asynchronous Replication supports the replication of Read-Only file and block snapshots locally or to a remote site along with the storage resource's data. Both scheduled snapshots and user created snapshots can be replicated. Snapshot Replication can be enabled on all resources that support Asynchronous Replication, including LUNs, thin clones, consistency groups, file systems, and VMware NFS and VMFS datastores. Snapshot Replication for scheduled snapshots can be enabled or disabled at any time in the lifetime of the replication session. The scheduled and user snapshots can be replicated with a remote retention policy that is longer than that of the source resource. For example, the retention policy for the source snapshots might be set for 2 hours while the retention policy for the replicated snapshot might be set for 1 week.

Dell EMC Unity OE version 4.2 introduced the Replicated attribute for all snapshots. This attribute states the replication state for the snapshot. A snapshot can be in one of the following states:

- **Yes:** The snapshot has successfully transferred to the destination
- **No:** The snapshot is not marked for replication
- **Pending:** The snapshot is marked for replication but is awaiting transfer
- **Failed:** The snapshot failed to replicate

In Dell EMC Unity OE version 4.4 and later, the Replicated attribute was replaced by two separate attributes to account for the introduction of snapshot replication for file synchronous replication resources. In OE 4.4 and later, these two attributes for snapshots are called “Async Replicated” and “Sync Replicated”. The states referenced above are not affected by the introduction of the two new attributes.

Replicating a manually created snapshot from a system running Dell EMC Unity OE version 4.2 through 4.4 to a destination system running Dell EMC Unity OE version 5.x is not supported. To support this, the source system must be upgraded to Dell EMC Unity OE version 4.5 or later. Upgrading to Dell EMC Unity OE version 4.5 or later is recommended, but not required, if you want to leverage replication without snapshot replication.

For more details in regards operation that can be completed on snapshots, please review the white paper titled *Dell EMC Unity: Snapshots and Thin Clones* on Dell EMC Online Support.

To support snapshot replication, both the source and destination systems must be running Dell EMC Unity OE version 4.2 or later. Snapshots that existed prior to the creation of the replication session 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. The Unisphere Management section provides more information regarding the configuration and management of snapshot replication.

Also, with Dell EMC Unity OE 4.2 and Cloud Tiering Appliance (CTA) version 12, the block snapshots can be archived to and restored from the cloud. For more information about Cloud Tiering Appliance (CTA) and how it works with Dell EMC Unity, please review the white paper titled *Dell EMC Unity: Cloud Tiering Appliance (CTA)* on Dell EMC Online Support.

In Dell EMC Unity OE version 5.1 and later, snapshots created on a resource contain a new internal system ID for tracking purposes. This ID is not visible to a user. When a snapshot is replicated, the ID will also be replicated with the snapshot to all destinations, including across all replication sessions within an advanced File remote replication topology. To support this feature both the source and destination systems must be running OE version 5.1 or later. For local replication, the ID is supported if the system is running OE version 5.1 or later.

In OE version 5.1 and later, replicated snapshots can be used to avoid full synchronization of data in certain situations. The snapshot ID allows the system to easily determine if a snapshot on the source and destination are exactly the same. This feature is explained in more detail in section 5.1 Snapshots.

3.6 Unisphere Management

Creating and managing Replication in Unisphere is easy and intuitive. All replication operations, including configuring of replication network interfaces, replication connections, and replication sessions can be performed in the Unisphere GUI. With the help of easy to follow wizards, replication can be configured quickly by IT generalists or advanced users alike. Replication can also be configured using Unisphere CLI, or REST API. For more information about configuring and managing replication using Unisphere CLI, refer to the *Unisphere Command Line Interface Guide* found on Dell EMC Online Support. For more information about REST API, refer to the *Unisphere Management REST API Programmer's Guide* found on Dell EMC Online Support.

3.6.1 Configuring Replication

When configuring remote asynchronous replication, the following components need to be configured and were previously discussed in the paper:

1. Replication Interfaces
2. A Replication Connection
3. A Replication Session

Note: When configuring local replication between Pools on a system, Replication Interfaces and a Replication Connection do not need to be configured on the system.

The following sections will outline the remaining steps needed to configure remote replication in Unisphere. Each of the following operations is completed from a particular page in Unisphere. Each page will be discussed in detail below. For more information about using Unisphere to configure and manage replication, refer to *Unisphere Online Help*.

3.6.2 Creating Replication Interfaces

Mobility Interfaces may be shared network interfaces for import and replication-related data or management traffic using the virtual management port. In this section, the mobility interfaces will be referenced as replication interfaces. To create replication interfaces, you must first go to the **Interfaces** page, which is found under **Protection & Mobility** in the left pane in Unisphere. Creating replication interfaces is completed on the **Interfaces** page. For releases prior to Dell EMC Unity OE version 4.1, the **Interfaces** can be found as a tab in the **Replication** page, under **Data Protection** in the left pane.

Figure 38 below shows the Unisphere Interfaces page. From this page, you have the option to create or delete replication interfaces, refresh the current page, or edit a configured replication interface. In the example below, no replication interfaces have been created. One or more pairs of replication interfaces need to be created on the source and destination system to configure remote replication.




Figure 38 Interfaces page

To create a replication interface for remote replication, click the **Create Interface** button, shown as a + sign in the Interfaces page. Once selected, the **Create Interface** wizard appears, which is shown on the left in Figure 39. For remote replication, you must create replication interfaces on the system, which are used for connectivity between the source and destination system. In the **Ethernet Port** drop down list, select an available port. The status of the ports will be shown in parentheses. Both ports on each SP will need to be cabled to the network for these links to work properly. Figure 39 also shows an example of the Ethernet Port drop down choices on a physical system. All Ethernet ports are available to create a network interface. Also shown are Link Aggregation ports previously created on the system.

Next, configure an IP address for SPA and SPB to be used for the replication interfaces. Once done, click **OK**.

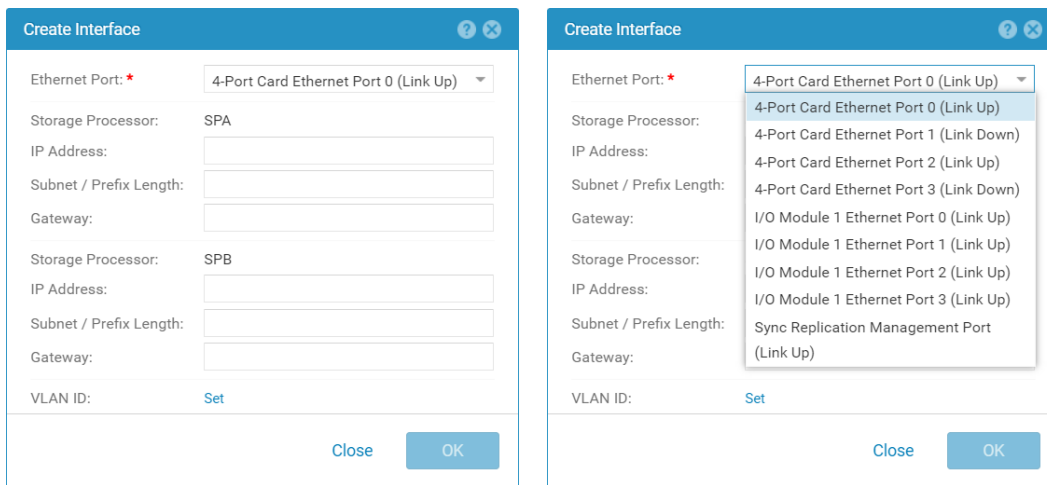


Figure 39 Create Interface window

Once the replication interfaces are created, they will be shown on the Interfaces page as shown in Figure 40. From here the status of each port is shown, currently operating normally as denoted by the green circle checkmark. To delete a replication interface, select the port and click the **Delete Interface** icon, which is shown as a trash can on this tab. To edit a replication interface, to change the IP address for example, select the **Edit** icon, which is shown as a pencil icon on this tab.

Interfaces							
+ [trash] [refresh] [edit] 6 items [filter] [settings] [download]							
<input type="checkbox"/>	!	Ethernet Port	↑ SP	IP Address	Subnet Mask / Prefix...	Gateway	VLAN ID
<input type="checkbox"/>	✓	4-Port Card Ethernet Port 0	SP A	10.10.10.5	255.255.255.0	--	1001
<input type="checkbox"/>	✓	4-Port Card Ethernet Port 0	SP B	10.10.10.6	255.255.255.0	--	1001
<input type="checkbox"/>	✓	4-Port Card Ethernet Port 0	SP A	10.10.10.7	255.255.255.0	--	1002
<input type="checkbox"/>	✓	4-Port Card Ethernet Port 0	SP B	10.10.10.8	255.255.255.0	--	1002
<input type="checkbox"/>	✓	4-Port Card Ethernet Port 1	SP A	10.10.10.9	255.255.255.0	10.10.10.10	--
<input type="checkbox"/>	✓	4-Port Card Ethernet Port 1	SP B	10.10.10.11	255.255.255.0	10.10.10.12	--


Figure 40 Configured replication interfaces

Figure 40 above is also an example of a replication configuration leveraging the asynchronous replication interface pairing enhancements introduced in Dell EMC Unity OE version 5.1. This feature is described in detail in section 3.2.1 Replication Interfaces. In this example, multiple replication interfaces using different VLAN IDs have been created on the same front-end ports. A pair of these ports on VLAN 1001 will connect to one remote system, while the ports configured with VLAN 1002 will connect to a different system. The interfaces created on 4-Port Card Ethernet Port 1 on both SPs will be used to connect to a third system over a different physical port and network.

3.6.3 Creating a Replication Connection

The next step in configuring remote replication is to create a Replication Connection with another system. This configures a private replication connection for the pair of systems. On the Replication page, multiple tabs exist, including Sessions and Connections. From the **Replication** page, select the **Connections** tab. Figure 41 below shows an example of the **Connections** tab without any replication connections configured. From this tab, you have the option to create or delete a replication connection, refresh the current page, or edit a configured replication connection. From this tab, you can also select **Verify and Update**, which will verify the


selected replication connection still exists with the remote system and update the connection details if any changes were made.



Sessions		Connections				Local Interfaces	
		Remote System Name	System Type	Mode	Management IP Address	SPA Interfaces	SPB Interfaces
You do not have any Connections configured. Click the + button above to create one.							

Figure 41 Replication Connections

To create a Replication Connection for remote replication, select the **Create Replication Connection** icon, which is displayed as a + sign on this tab. The **Create Replication Connection** window appears, as shown in Figure 42. On the **System Settings** step, you must specify the Remote System's **Management IP Address**, which is the IP used to access Unisphere, and the Unisphere **User Name** and **Password**. Also, in this window you must enter the **Password** used to log in to Unisphere on the system you are configuring the replication connection in. Lastly, you must select the **Connection Mode** that will be used between the systems. In the drop-down list, you have the option to choose **Asynchronous**, **Synchronous**, or **Both**. When configuring asynchronous replication between two systems, select **Asynchronous**. If synchronous and asynchronous replication will be used, select **Both**. After entering the required information, click **Next**. The Bandwidth Schedules step is then shown. This feature was discussed in section 3.2.2.1 Asynchronous replication throttling.



Create Replication Connection

- System Settings
- Bandwidth Schedules
- Summary
- Results

System Settings

Remote System

Management IP Address: *

User Name: *

Password: *

Local System

User Name:

Password: *

Connection Mode

Mode: *

Asynchronous ▼

- Asynchronous
- Synchronous
- Both

Note: If a replication connection already exists, do not attempt to create a new connection mode to Both.

Cancel
Next

Figure 42 Create Replication Connection window

After completing the System Settings and Bandwidth Schedules steps and clicking OK on the Summary step, a Job is created to create the replication connection. The job has multiple steps, which include registering the remote and local system with its peer, refreshing the connection on both systems, and validating the connections on the local system. Once the job completes, the replication connection will be shown on the source and destination system.

Figure 43 shows the Replication page **Connections** tab once the replication connection is created. In this example, the name of the remote system is displayed. The System Type, Management IP Address, Replication Mode, and Remote Interfaces are also displayed. The Remote Interfaces displayed will be the IPs for the synchronous replication management ports and asynchronous replication interfaces on the remote system.

		Connections							
	Remote Syst... ↑	System Type	Mode	Management IP ...	Local Interfaces		Remote Interfaces		
					SPA Interfaces	SPB Interfaces	SPA Interfaces	SPB Interfaces	
<input type="checkbox"/>	✓	DRSystem	Unity 880	Asynchronous ...					
<input type="checkbox"/>	✓	VIRT2117N85...	UnityVSA	Asynchronous ...		10.10.10.5	10.10.10.6	10.10.10.10,10...	
<input type="checkbox"/>	✓	VIRT2118RDT...	UnityVSA	Asynchronous ...		10.10.10.7	10.10.10.8	10.10.10.20,10...	

Figure 43 Replication Connections

3.6.4 Creating a Replication Session

To create a replication session, you must enable replication on the storage resource you want to replicate. This can be done at time of creation, or on an existing storage resource. Figure 44 shows an example of creating a new LUN and enabling asynchronous remote replication on it. From the **Replication** step, check the checkbox in front of **Enable Replication** to configure replication. For asynchronous replication, select **Asynchronous** from the **Replication Mode** drop-down box and customize the **RPO** as needed. Choose **Manual** to create a manual replication session. When asynchronous is chosen as the Replication Mode, the **Replicate To** box will automatically be populated with a remote system name asynchronous replication can be configured with. If you want to replicate locally with asynchronous or manual replication, choose **Local** found under **Replicate To**. The **Replicate scheduled snapshots** option is also shown and allows the replication of scheduled snapshots. The retention policy for snapshots on the destination can also be customized. After completing this step, click **Next** and continue with the creation. For LUNs, consistency groups, VMware VMFS datastores, thin clones, file systems, and VMware NFS datastores, a Replication step exists in the creation wizards for each with the similar look and configuration options as the one shown in Figure 44.

Create LUNs

- ✓ Configure
- ✓ Access
- ✓ Snapshot
- **Replication**
- Summary
- Results

Provide a Replication Mode and RPO

Enable Replication

Replication Mode: * Asynchronous

RPO: * 60 minutes

Replicate To: * DRSystem

Destination Configuration

Replicate scheduled snapshots: Default Protection

Remote Retention Policy

Hourly: Retain for 2 days (same as source)

[Customize](#)

Overwrite any data on destination block resource. Otherwise destination may contain different content than source if destination has corrupted data.

Overwrite Destination

Info Data is replicated asynchronously with a predefined RPO interval.

Selecting a longer RPO may result in a larger data loss in the case of a failure. A shorter RPO may impact performance.

To create a remote replication connection, navigate to the Protection & Mobility > Replication > Connections section.

To create interfaces for use with remote replication, navigate to the Protection & Mobility > Interfaces section.

Info If you want to set up replication using RecoverPoint, go back to the Access step and configure the RecoverPoint host to access the storage resource.

Cancel [Back](#) [Next](#)

Figure 44 Enable Asynchronous Replication

When the **Destination Configuration** box is selected in the Replication step, information about the remote storage resource's configuration will be displayed. Depending on the capabilities of the remote system, the destination Pool, and the resource type, options for Thin and Data Reduction will be displayed. Figure 45 below shows an example of this window.

Figure 45 Destination LUN Configuration window

After asynchronous replication is configured on a new storage resource, you can view information about replication from the resource's properties window. From Unisphere, select the storage resource in question and click Edit or double-click the name of the storage resource. From the properties window, view the **Replication** tab. An example of this tab for a LUN is shown in Figure 46. On this tab, you can view the following information:

- The replication **Session Name**
- The replication **Mode**
- The **Local Role** of the storage resource, which can either be listed as the **Source** or **Destination**
- The Time of Last Sync
- The Replicate Scheduled Snapshots option

Also shown is a pictorial representation of the replication session. The figure shows which storage resource is available for I/O, which direction the data is replicating in and its current state, and the destination storage resource and the system name, IP Address, and the destination LUN name. As the state of the replication session changes, this figure will update to reflect the new state.

Also shown on the Replication tab are buttons for each Replication Operation. This tab is also used to display synchronous replication operations, so all replication operations for synchronous and asynchronous replication are displayed. Not all replication operations are supported on each mode of replication, so only operations supported on the current replication mode will be selectable. Also, only certain operations are available depending on what the current state of the replication session is in, so only these options are available for selection. In Figure 46, the current session is Auto Sync Configured and replicating from the Production System (Local System) to the DRSystem. Available replication operations for the asynchronously replicated LUN include **Delete**, **Pause**, **Sync**, and **Failover with Sync**.

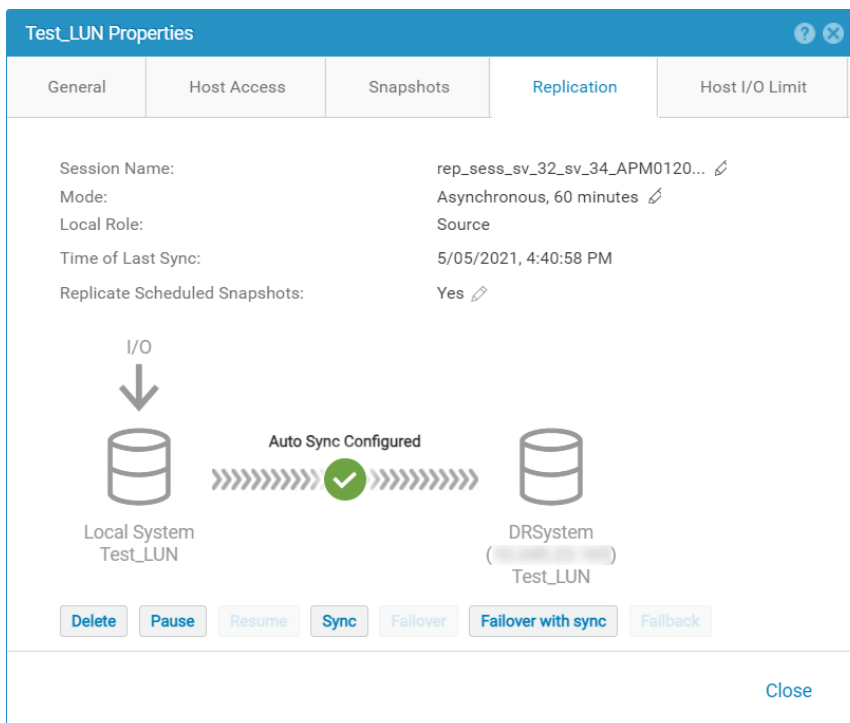


Figure 46 LUN replication configured

Replication can also be configured on an existing storage resource. When replication is not configured on the storage resource, viewing the **Replication** tab in the storage resource's properties window will show what is displayed in Figure 47. The **Replication** tab for a consistency group, VMware VMFS datastore, thin clone, file system, and VMware NFS datastore displays the same information below. To configure Replication on the storage resource, select **Configure Replication**.

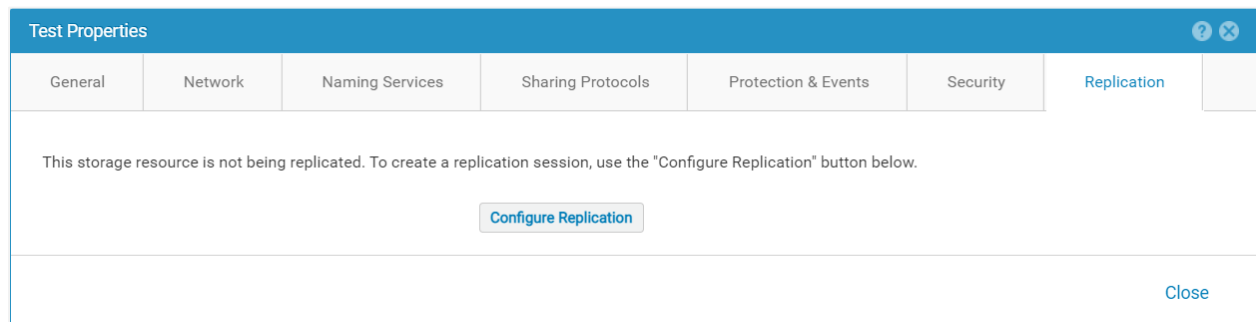


Figure 47 Replication not configured

After selecting Configure Replication, the **Create a Session** wizard is launched. An example of this wizard is shown in Figure 48. On the **Replication Settings** step, you need to customize which **Replication Mode** either Asynchronous, Manual, or Synchronous replication will be used. For Asynchronous Replication, select Asynchronous. Next, customize the wanted RPO if needed. When manual replication is selected, RPO will not be an option. Then, select the destination system by selecting the correct system in the **Replicate To** drop down box, either Local for local replication, or a listed remote system. As mentioned previously, Dell EMC Unity OE version 4.2 supports the replication of scheduled and user snapshots. The **Replicate all existing snapshots** is a one-time option, given at creation time of the replication session, that can be used to replicate all existing snapshots. The **Replicate scheduled snapshots** option can be used to replicate the scheduled snapshots of a storage resource, this option can be modified after the replication session creation.

Figure 48 shows both options already selected. The **Reuse destination resource**, **Automatically search user snap as common base**, and **Overwrite Destination** options will be discussed in detail in section 5.1 Snapshots.

Figure 48 Create a Session wizard – Replication Settings step

In the example in Figure 48, the LUN has both Hourly and Daily/Weekly snapshot schedules. To customize the remote retention policy for the snapshots that will be replicated, click the **Customize** button, which opens the **Customize Remote Retention** window. An example of this window is shown in Figure 49. This provides the options to either keep the same retention policy as the source, follow the Pool Automatic Deletion Policy, or retain the snapshots for a custom number of hours, days, or weeks. The retention policy can go up to 255 weeks. In this example, Hourly Snapshots are retained for seven Hours while the Daily/Weekly Snapshots are retained for four Weeks. After selecting the retention policy, click **OK**, and then **Next**.

Figure 49 Customize Remote Retention

The **Destination** step of the **Create a Session** wizard is now shown. An example of this step for a LUN is shown in Figure 50. From here, you can customize the storage resource's **Name** that will be displayed on the destination Pool or system, the **Pool** it will use, the **Storage Processor** if applicable, the **Tiering Policy** if applicable on the destination's Pool, and enable **Data Reduction/Advanced Deduplication** if supported on the resource and the destination. For other existing storage resources being replicated, you can customize this exact same information. When creating a replication session on an existing consistency group, you can customize this same information for each LUN contained within the consistency group. After editing the available information, click **Next**.

The screenshot shows a 'Create a Session' wizard window with a blue header. On the left is a sidebar with 'Replication Settings' (checked), 'Destination' (selected), 'Summary', and 'Results'. The main area is titled 'Provide destination storage resource info' and contains the following fields:

- Name:** * Test_LUN
- Destination System:** DRSystem
- Pool:** Pool 1 (Traditional, Extreme Performance Pool, ▾)
- Size(GB):** 225.0
- Thin
- Data Reduction
- Tiering Policy:** ▾

A warning icon (i) in the top right corner says: 'The selected pool is not tiered. The tiering policy will have no effect on the storage resource.' At the bottom right are three buttons: 'Cancel', 'Back', and 'Next'.

Figure 50 Create a Session wizard – Destination step

The **Summary** step is next in the Create a Session wizard. Here you can see a summary of the settings that will be used to create replication. If anything is incorrect, you can select **Back** to correct the wanted setting. To create the replication session, click **Finish**. A Summary step will also be shown when creating a replication session on other storage resources.

The **Results** step is the last step in the wizard. This step shows the Overall Status of each of the jobs to create the replication session. The steps to create the replication session include creating the storage resource on the destination system, allowing the remote storage resource to finish the creation process, and lastly creating the replication session. You can either wait for the Overall status to say 100% Complete or close the window at any time by clicking **Close**. Closing this window will not impact the creation process since it is a background job in Unisphere.

3.6.4.1 Advanced File Remote Replication

As mentioned previously, advanced File replication topologies allows File resources to be replicated asynchronous in fan-out and cascade topologies. Each replication session within an advanced replication topology is created separately. Before file systems or VMware NFS datastores can be replicated, the NAS Server must first be replicated to the desired destination Pool or remote system. The replication session for the NAS Server can be established when the NAS Server is created, or on an existing resource. Figure 51 shows an example of creating a new NAS Server and enabling Asynchronous remote replication on it.

From the **Replication** step, check the checkbox in front of **Enable Asynchronous Replication** to configure replication. For asynchronous replication, select **asynchronous** from the **Replication Mode** drop-down box and customize the **RPO** as needed. Choose **Manual** to create a manual replication session. When asynchronous is chosen as the Replication Mode, the **Replicate To** box will automatically be populated with a

remote system name asynchronous replication can be configured with. If you want to replicate locally with asynchronous or manual replication, choose **Local** found under **Replicate To**.

Figure 51 Enable Asynchronous Replication

When the **Destination Configuration** link is selected in the Replication step, information about the remote storage resource's configuration will be displayed. Figure 52 below shows an example of this window. After updating this window as needed, click **OK** and then **Next** to continue with the NAS Server creation.

Figure 52 Destination NAS Server Configuration window

After asynchronous replication is configured on the NAS Server, you can view information about replication from the resource's properties window. From Unisphere, select the storage resource in question and click Edit or double-click the name of the storage resource. From the properties window, view the **Replication** tab. An example of this tab for a NAS Server is shown in Figure 53. As File resources can have multiple replication session, they are all shown on this tab. In this example the NAS Server has been replicated in a fan-out configuration to three destination systems. On this tab, you can view the following information:

- The replication session **Name**
- The **Replication Mode**
- The **Destination**, including the remote **System** and **NAS Server**

- The State

TEST Properties								
General		Network		Naming Services	Sharing Protocols	Protection & Events	Security	Replication
+ [Icons] More Actions ▾							3 items	[Filter] [Settings] [Download]
[]	!	Name	Replication Mode	Destination		State		
				System	NAS Server			
[]	✓	rep_sess_nas_2_nas_2_APM0120490803...	Asynchronous	VIRT2118RDTDT1	TEST	Auto Sync Configured		
[]	✓	rep_sess_nas_2_nas_3_APM0120490803...	Asynchronous	VIRT2117N85PXE	TEST	Auto Sync Configured		
[]	✓	rep_sess_nas_2_nas_9_APM0120490803...	Asynchronous	DRSystem	TEST	Auto Sync Configured		

Figure 53 NAS Server replication configured

Replication actions, such as Pause, Resume, and Failover, can be issued from the Replication tab. After selecting a replication session in the list, click More Actions. The options available depends directly on the replication state of the resource. The properties of each replication session provides more information to the user. To view information about a replication session, double click the session name or select the session and click the View/Edit icon.

A pictorial representation of the replication session is shown in the properties of the session, as seen in Figure 54. On this screen the **Session Name**, **Mode**, and **Time of Last Sync** is displayed. The figure also shows which storage resource is available for I/O, which direction the data is replicating in and its current state, and the destination storage resource and the system name, IP Address, and the destination LUN name. As the state of the replication session changes, this figure will update to reflect the new state.

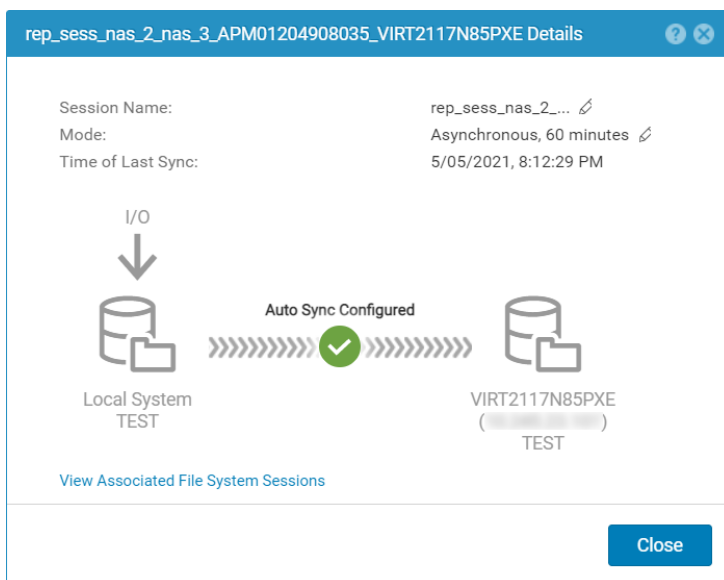


Figure 54 Replication session properties


When replication is configured on an existing NAS Server, all resources on that NAS Server are also replicated. A replication session can be deleted from a resource that should not be replicated at any time. When a new resource is added to a replicated NAS Server, the **Replication** step allows the new resource to be created with the same replication topology as the NAS Server. Figure 55 shows the **Create a File System**

wizard. In this example, the parent NAS Server is replicated to three remote systems. On the **Replication** step, replication on the new file system can also be configured. A checkbox, disabled by default, is specified for each destination. When a checkbox is selected, the replication session will be created with the same attributes as the associated NAS Server session.

Figure 55 Replication not configured

To customize a replication session to a specific destination, select the **Destination Configuration** link for the corresponding session. In the **Replication details Configuration** screen, as seen in Figure 56, many options can be customized. These options include the **Replication Mode**, either **Asynchronous** or **Manual**, the **RPO**, the destination **Pool**, the **Tiering Policy**, and the **Snapshot Replication Configuration**. These options can be independently configured on each replication session. The **Reuse destination resource** option will be discussed in detail in section 5.1 Snapshots.


Native Asynchronous Replication



The 'Replication details Configuration' dialog box is shown. It is divided into three main sections: 'Replication Configuration', 'Destination File System Configuration', and 'Snapshot Replication Configuration'. In the 'Replication Configuration' section, 'Replication Mode' is set to 'Asynchronous', 'RPO' is '60 minutes', and 'Replicate To' is 'VIRT2118RDTDT1'. The 'Destination File System Configuration' section includes a 'Name' of 'TESTFS', 'Destination System' of 'VIRT2118RDTDT1', 'Pool' of 'Pool 0 (Traditional, Extreme Performance)', and 'Size(GB)' of '100.0'. There are checkboxes for 'Reuse destination resource' and 'Data Reduction', both of which are unchecked. The 'Tiering Policy' is set to 'Start High Then Auto-Ti'. The 'Snapshot Replication Configuration' section has three checked options: 'Support Asynchronous Snap Replication', 'Cascade replicated snapshots', and 'Replicate scheduled snapshots: Default Protection'. Below these is a 'Remote Retention Policy' section with 'Hourly: Retain for 2 days (same as source)'. At the bottom, there are 'Cancel' and 'OK' buttons.

Figure 56 Create a Session

On the **Summary** step of the **Create a File System** wizard, all configuration information for the resource is displayed. As with other resources, if a correction needs to be made, the **Back** button can be used. Each replication session being created is shown on the **Summary** screen, as seen in Figure 57. The **Replication Mode**, **RPO**, and **Destination System** are shown. To view more information about a specific session, click the **More Info** link under **Session Details**.



The 'Create a File System' wizard is shown at the 'Summary' step. On the left, a list of steps is shown with 'Summary' selected. The main area displays configuration details for the file system, including 'File-level Retention: Off', 'SMB Share (Windows): Yes', and 'Name: TESTFS'. Below this are 'SMB Share's Other Settings' and 'Snapshot' settings. A 'Replication' table is displayed at the bottom of the main area.

Replication Mode	RPO	Destination System	Session Details
Asynchronous	60	VIRT2117N85PXE	More Info
Asynchronous	60	DRSystem	More Info
Asynchronous	60	VIRT2118RDTDT1	More Info

At the bottom of the dialog, there are 'Cancel', 'Back', and 'Finish' buttons.

Figure 57 Create a File System – Summary step

When **More Info** is selected, the **Replication details Configuration** window is shown. If any settings for the replication session need to be changed, simply close this window and click the **Back** button.

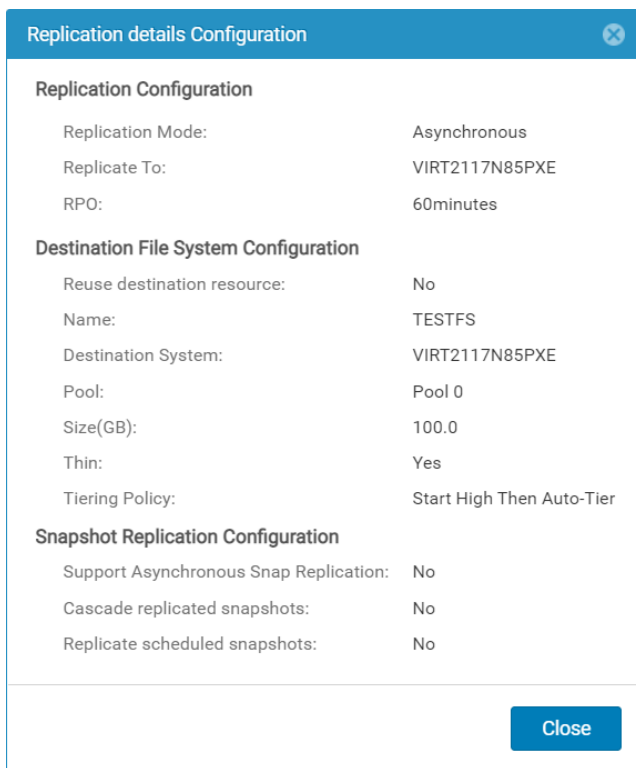


Figure 58 Replication details Configuration window

Once the resource is created, replication information can be seen on the **Replication** tab within the properties of the resource. As with NAS Servers, the replication tab for file systems and NFS datastores will display all replication sessions currently configured on the resource. An example of this screen can be seen below. Replication operations, such as Pause, Resume, and Failover can be found under **More Actions**. To view more information about a specific replication session, simply select the session and click the **View/Edit** icon.

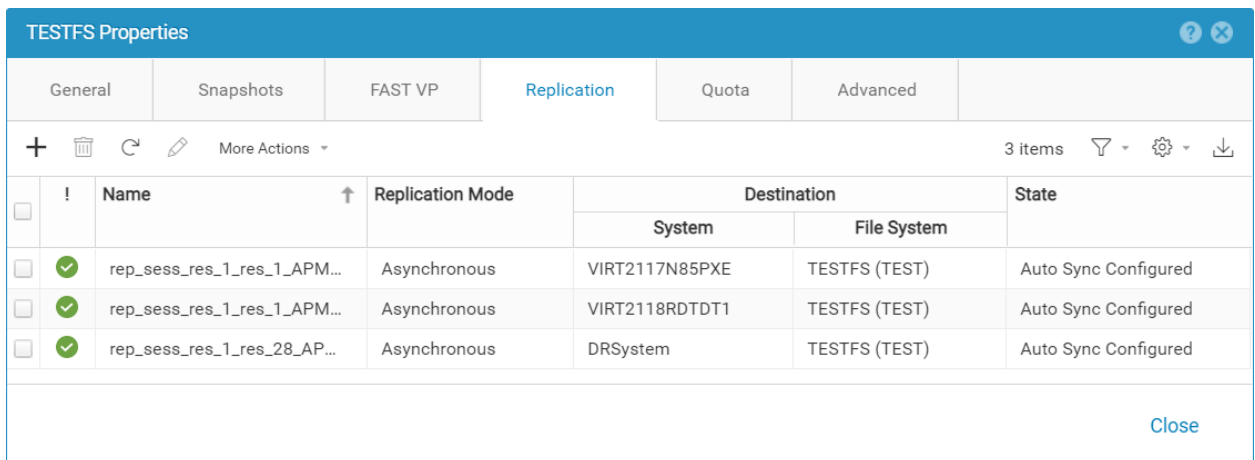


Figure 59 File system properties – Replication tab

A pictorial representation of the replication session is shown in the properties of the session, as seen in Figure 60. On this screen the **Session Name, Mode, and Time of Last Sync** is displayed. The snapshot replication settings are also displayed. The figure also shows which storage resource is available for I/O, which direction the data is replicating in and its current state, and the destination storage resource and the

system name, IP Address, and the destination LUN name. As the state of the replication session changes, this figure will update to reflect the new state.

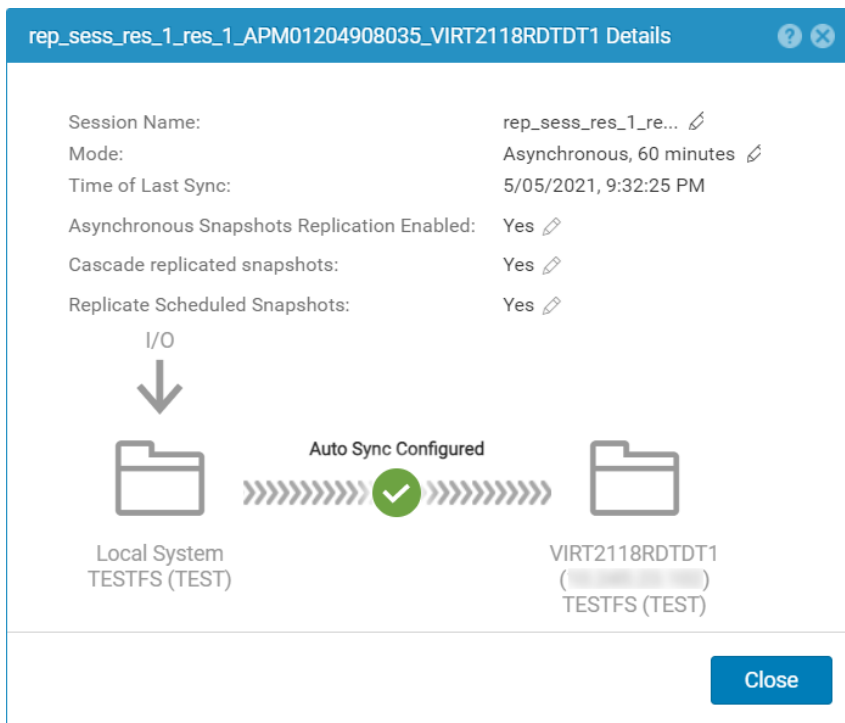


Figure 60 Replication session properties screen

3.6.5 Restore from a replicated snapshot

Once the source snapshots have been replicated to the destination, the source snapshot can be deleted from the source system to free up space. If there is ever a need to recover data that is in a replicated snapshot, the replicated data can be manually restored to the source storage resource. Taking a backup of the source resource is suggested, as this process will overwrite data within the source resource. To manually restore from a replicated snapshot, once the snapshot has been replicated, failover the storage resource's replication session, and on the destination system, restore the storage resource to the snapshot that has the data you are interested in. Once the restore is complete, failback the replication session. As result, the source storage resource will contain the data that was on the snapshot.

3.6.6 Override network address for file replication

If you have file replication configured with a destination site on a different physical network, to ensure minimal downtime during a failover, ensure that you modify the destination NAS Server properties to include an override address for the network interface. Figure 61 shows an example of the destination NAS Server with the Production Interface being overridden to a different IP Address. A different Ethernet Port can also be selected in this process.

Edit NAS Server Network Interface

Role: Production

Ethernet Port: Ethernet Port 2

IP Address: 20.20.20.20

Subnet Mask / Prefix Length: 255.255.255.0

Gateway:

VLAN ID: 1000

Override address

Ethernet Port: * Ethernet Port 2 (Link Up, Actual Speed: 10 G)

IP Address: * 30.30.30.30

Subnet Mask / Prefix Length: * 255.255.255.0

Gateway:

VLAN ID: 1003 None

Warning: This is a Production interface. The interface of this type:

- Allows CIFS, NFS and FTP access;
- Is replicated during the Replication session;
- During the Replication session, is active in the Source mode only.

This interface is NOT active now.

To override the Source network settings of this interface, please select the "Override" checkbox.

Cancel OK

Figure 61 Override address

3.6.7 Enable NDMP on NAS Server

Dell EMC Unity supports 3-way NDMP (Network Data Management Protocol) backup over Ethernet. Also, as of Dell EMC Unity OE version 4.4, 2-way NDMP backup over Fibre Channel is also supported. This support allows for data backup to NDMP backup clients which could be a tape library or virtual tape appliance. NDMP can also work in cohesion with replication for additional data protection. If replication is enabled on a NAS Server and additional data protection is needed, NDMP can be enabled through the Protection & Events tab in the properties page of a NAS Server, as seen in Figure 62. NDMP can be enabled either on the source or destination NAS Server giving configuration flexibility in running incremental NDMP backups of file systems. If you enable NDMP on a destination NAS Server, you will first need to create a Backup & DR testing network interface to be able to connect to destination file resources. Also note that enabling NDMP for the destination NAS Server is not supported in the initial OE 4.0 release but is supported in OE 4.1 and later.

TEST Properties

General Network Naming Services Sharing Protocols **Protection & Events** Security Replication

NDMP Backup

Enable NDMP [Change Password](#)

DHSM

Events Publishing

Username: * ndmp

Password: *

Confirm Password: *

Cancel Apply

Figure 62 Enable NDMP

3.6.8 Viewing the Replication Sessions

All replication sessions on the system can be viewed from the Sessions tab within the Replication page. This includes synchronous and asynchronous replication sessions. To go to this page, click Replication under Protection & Mobility in Unisphere. For releases prior to Dell EMC Unity OE version 4.1, the Replication page can be found under Data Protection. Figure 63 shows an example of the Sessions tab with multiple replication sessions created on the system. In this example, you can see replication sessions for the **TEST** NAS Server and **TESTFS** file system created in the Advanced File Remote Replication section. In this example, the NAS Server and file system are configured in an asynchronous replication fan-out topology, replicating to three destination systems. From this page, you can easily see information regarding each session. The following is a list of information displayed on this screen. More columns can be added to the view if needed.

- The **Source**, which includes the source system and the source storage resource
- The Resource Type
- The **Replication Mode**, either Synchronous, Asynchronous, or Manual
- The **Destination**, which includes the destination system name and the destination storage resource
- The current **State**
- The current **Transfer Rate (MB/Sec)**
- The Replication Session **Name**

		Source		Resource Type	Replication Mode	Destination		State	Transfer Rate (MB/Sec)	Name
System	Resource	System	Resource							
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TEST	NAS Server	Asynchronous	VIRT2118RDTDT1	TEST	Auto Sync Configured	0.0	rep_sess_nas_2_nas_2_APM01...
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TEST	NAS Server	Asynchronous	VIRT2117N85PXE	TEST	Auto Sync Configured	0.0	rep_sess_nas_2_nas_3_APM01...
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TEST	NAS Server	Asynchronous	DRSystem	TEST	Auto Sync Configured	0.0	rep_sess_nas_2_nas_9_APM01...
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TESTFS (TEST)	File System	Asynchronous	VIRT2118RDTDT1	TESTFS (TEST)	Auto Sync Configured	0.0	rep_sess_res_1_res_1_APM012...
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TESTFS (TEST)	File System	Asynchronous	VIRT2117N85PXE	TESTFS (TEST)	Auto Sync Configured	0.0	rep_sess_res_1_res_1_APM012...
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TESTFS (TEST)	File System	Asynchronous	DRSystem	TESTFS (TEST)	Auto Sync Configured	0.0	rep_sess_res_1_res_28_APM01...

Figure 63 Replication Sessions

From the Sessions tab, you can also issue replication operations on the sessions. After selecting the checkbox for a replication session in the list, select **More Actions** to view the replication operations available for that session in its current state. In Figure 64 below, you can see that only **Pause**, **Sync**, and **Failover with Sync** are valid options based on the current replication session state of the selected storage resource.

		Source		Resource Type	Replication Mode	Destination	
System	Resource	System	Resource				
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TEST	NAS Server	Asynchronous	VIRT2118RDTDT1	TEST
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TEST	NAS Server	Asynchronous	VIRT2117N85PXE	TEST
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TEST	NAS Server	Asynchronous	DRSystem	TEST
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TESTFS (TEST)	File System	Asynchronous	VIRT2118RDTDT1	TESTFS (TEST)
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TESTFS (TEST)	File System	Asynchronous	VIRT2117N85PXE	TESTFS (TEST)
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Local System	TESTFS (TEST)	File System	Asynchronous	DRSystem	TESTFS (TEST)

More Actions

- Pause
- Resume
- Sync
- Failover
- Failover with sync
- Failback
- Preserve async replication

Figure 64 Replication Sessions – More Actions

Dell EMC Unity OE version 4.2 also includes the group operations feature at a NAS Server level. The group operations are available for the following operations: Failover, Failover with Sync, Failback, Pause, and Resume. Figure 65 gives an example of performing a group Failover for a NAS Server.

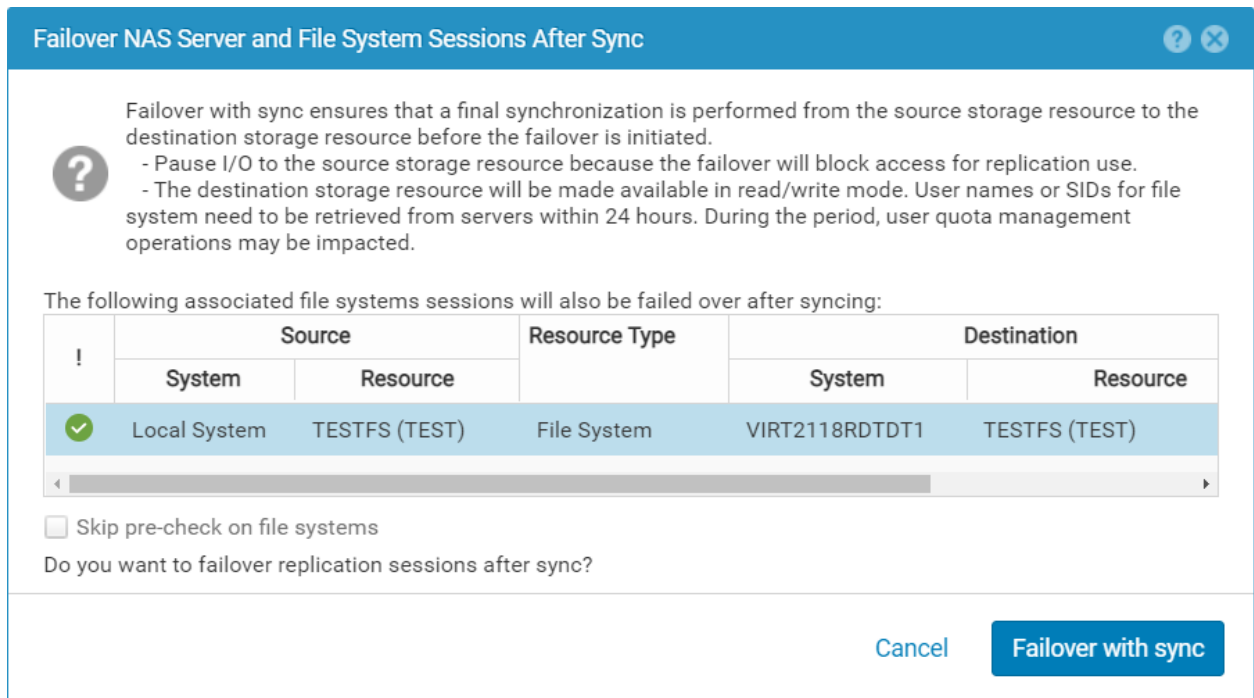


Figure 65 Failover with sync

3.7 System Maximums

See Appendix A: Replication Maximums for information about Dell EMC Unity replication maximums.

4 Proxy NAS Servers

Dell EMC Unity OE version 4.3 introduces Proxy NAS Servers, providing the ability to access files on the destination side of a replicated File resource. This feature provides Read-Only access to file system and snapshot data through SMB and NFS. There is no ability to write to the file systems or snapshots using the proxy NAS server, even if the snapshot is Read/Write.

All the NAS Servers' file systems and their snapshots are displayed when connecting to the proxy NAS Server. Due to this, the user must be part of the Local Administrators group for SMB or root for NFS. You can add users and groups to the local Administrators group of the proxy NAS server through MMC, just like a regular SMB server.

Dell EMC Unity OE version 4.5 introduces the ability to create SMB shares for writeable and Read-Only snapshots on the destination NAS Server. This feature is designed to enable DR testing without any impact to the ongoing replication session. It allows customers to confirm that an application can be brought online and write to a share hosted on the destination system. This feature works with both asynchronous and synchronous replication. This feature leverages a Proxy NAS Server and Proxy share created on the destination system to provide access to the snapshot.

In contrast to the Read-Only Proxy NAS Server feature, this feature allows any domain user to access the share and is not limited to Administrators or root. This is because each share points to a specific snapshot, as opposed to the entire contents of the NAS Server. The proxy share can be configured to point to either a Read-Only or Read/Write snapshot that exists on the destination file system. If a Read/Write snapshot is selected, then the client can write to the share.

For more information about Proxy NAS Server, please review the *Dell EMC Unity: DR Access and Testing* white paper found on Dell EMC Online Support.

5 Interoperability

Native Synchronous and Asynchronous Replication is supported between Dell EMC Unity OE versions 4.0 and later systems in either direction. If there is a specific feature related to Dell EMC Unity OE version 4.1 or later that is in use in the source system, then the replication session may not be supported and may fail to setup. In some cases, the replication session may succeed, and ignore the newer capability on the target system.

For File Replication, some features that should be considered when configuring replication include Common Event Enabler (CEE), IP Multi-Tenancy, Advanced Static Routing (ASR), Packet Reflect, and Folder Rename and Locking policies. The Advanced Static Routing (ASR) and Packet Reflect features will not prevent replication to Dell EMC Unity OE version 4.0, but the settings for ASR and Packet Reflect will not be replicated. Similarly, if remote replication is already configured for a NAS Server and the destination system is running a previous version of Dell EMC Unity OE, and the source system is upgraded to the latest Dell EMC Unity OE, new features in the NAS Server cannot be enabled on the source system. To enable new features in the NAS Server on the source system, also upgrade the destination system to the latest version of Dell EMC Unity OE. If replication between Dell EMC Unity systems is critical, consider upgrading all systems participating in the replication to the latest Dell EMC Unity OE version.

For more information about NAS Capabilities in the Dell EMC Unity systems, please review the white paper titled *Dell EMC Unity: NAS Capabilities* on Dell EMC Online Support.

5.1 Snapshots

In Dell EMC Unity OE version 5.1 and later, Read-Only snapshots created on a resource contain a new internal system ID for tracking purposes. This ID is not visible to a user. When a snapshot is replicated, the ID will also be replicated with the snapshot to all destinations, including across all replication sessions within an advanced File remote replication topology. To support this feature both the source and destination systems must be running OE version 5.1 or later. For local replication, the ID is supported if the system is running OE version 5.1 or later.

In certain circumstance, a full synchronization of data is needed to restart or continue replication. In OE version 5.1 and later, replicated snapshots can be used to avoid full synchronization of data in certain situations. The snapshot ID allows the system to easily determine if a snapshot on the source and destination are exactly the same. If the systems confirm the IDs of a pair of snapshots are the same, the snapshots can be used as a common base to start or continue replication. This can help reduce the amount of time it takes to return replication to a good state. The amount of data to be copied directly depends on the amount of data that has changed since the creation of the common base snapshots. This feature is supported on LUNs, thin clones, consistency group, file systems, and VMFS and NFS datastores for specific replication modes described below.

5.1.1 Create replication session

In situations where a replication session must be deleted and recreated, a full synchronization of data is required in OE versions prior to the 5.1 release. This is true even if no data has changed between the source and destination. If a snapshot has been created on a system running 5.1 or later and replicated to the destination running 5.1 or later, the replication session can leverage the replicated snapshot as a common base. Full copy avoidance is supported for file synchronous replication, and block and file asynchronous replication. It is highly suggested to create a new snapshot and replicate it to the destination just before deleting the replication session. Creating the snapshot with No Automatic Deletion policy is also suggested. For synchronous replication, only the latest 21 snapshots on a resource are searched for a common base,

therefore it is also suggested to delete and recreate the sessions as soon as possible. Once the replication session has been recreated and is active, the snapshot can be deleted.

In the following example, an asynchronously replicated NAS Server and file system are being transitioned to a synchronous replication configuration. Full copy avoidance can occur as long as a common base snapshot is available. When creating the replication session on the NAS Server, the **Reuse destination resource** and **Automatically search user snap as common base** options are available. Figure 66 below shows an example of these options on the **Replication Settings** step within the **Create a Session** wizard.

The screenshot shows the 'Create a Session' wizard with the 'Replication Settings' step selected. The 'Provide a Replication Mode and RPO' section contains the following fields and options:

- Replication Mode:** Synchronous (dropdown menu)
- Replicate To:** APM00211114469 (dropdown menu)
- Reuse the destination resource when it already exists. If the resource does not exist, provision the destination resource.**
 - Reuse destination resource
 - Automatically search user snap as common base

Informational text on the right side of the wizard includes:

- Data is replicated synchronously with zero RPO.
- To create a remote replication connection, navigate to the Protection & Mobility > Replication > Connections section.
- To create interfaces for use with remote replication, navigate to the Protection & Mobility > Interfaces section.
- Ensure that Ethernet ports, link aggregations and FSN ports with same names are available on destination, otherwise use override to change port assignment to NAS interfaces.

At the bottom right, there are 'Cancel' and 'Next' buttons.

Figure 66 Create a Session wizard – NAS Server example

In Figure 67 below the Destination step is shown. Enabling the **Reuse destination resource** option will automatically search the destination system and establish replication to an existing resource if the Name entered on this step matches one on the destination. In this example, the system will search for existing NAS Servers on the destination with the name **TEST**. If a NAS Server named TEST exists on any Pool within the destination system, replication will be established to the resource. If a resource with the same name is not found, the resource will be created based on the configuration on the **Destination** step. This is also true for file systems and NFS datastore on the NAS Server, and other resource types.

Create a Session

✓ Replication Settings

• Destination

Summary

Results

Configure Destination NAS Server

Name: *

Pool:

Storage Processor:

Destination System: APM00211114469

Associated File Systems/VMware NFS Datastore:

Source ... ↑	Destinatio...	Destinatio...	Thin	Data Redu...	Advanced ...
TESTFS	TESTFS	Pool 1	Yes	No	No

Cancel

Info: The table lists the recommended Pool for each File System or/and VMware NFS Datastore associated to the destination NAS server.

You may modify the Pool selection by selecting a File System/VMware NFS Datastore and clicking the details icon above the table.

Info: The selected Pool is only used for resource provision when the associated File Systems or VMware NFS Datastores do not exist on the destination NAS server.

Figure 67 Create a Session wizard – NAS Server example

When replication is configured on a NAS Server, replication is automatically configured for file systems and NFS datastores on the NAS Server. When the **Automatically search user snap as common base** option is enabled, snapshots on the source and destination file resources will be searched for a snapshot with a matching ID. If a match is found that resource will use the latest matched snapshot as a common base. Only data that has changed on the source since the snapshot was created will be copied to the destination. This comparison will occur for each file system or NFS datastore on the NAS Server. For tracking purposes, the Logs page in Unisphere will have an entry for each storage resource if a common base snapshot was found or not.

If replication needs to be deleted and recreated on a file system or NFS datastore, similar options on the **Replication Settings** page for the NAS Server are available. Figure 68 below shows the **Replication Settings** step within the **Create a Session** wizard for a file system. Here the **Reuse destination resource** and **Automatically search user snap as common base** options are available directly on the resource.

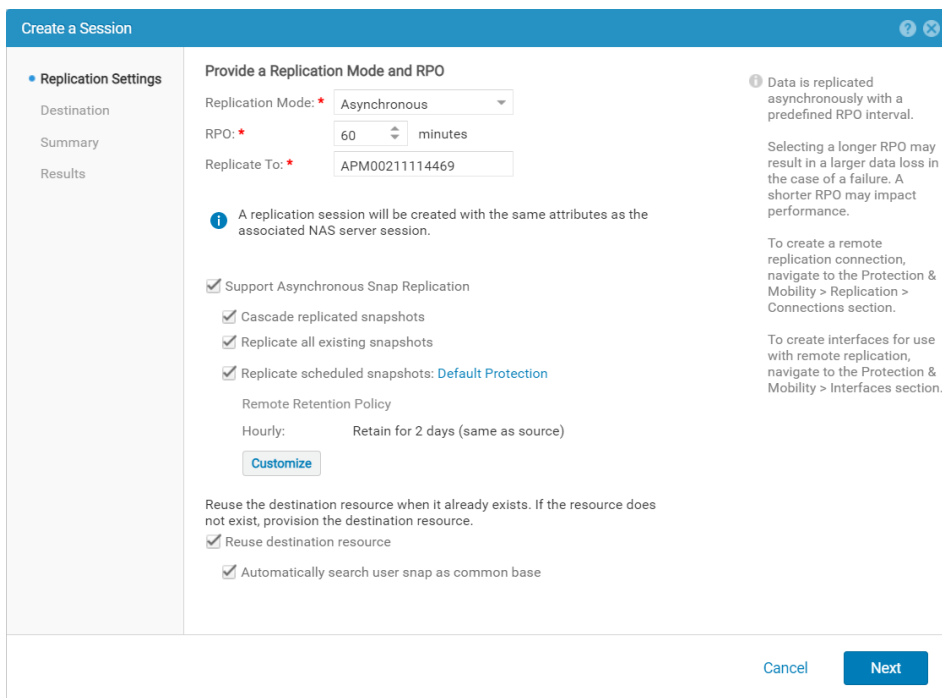


Figure 68 Create a Session wizard – Replication Settings step – file system example

For Block resources the **Reuse destination resource** and **Automatically search user snap as common base** options are available. The **Overwrite Destination** option is also available. If this checkbox is selected the entire contents of the destination resource are overwritten with data from the source. This option can be used when the source resource has been restored to a previous point in time, and the destination must be completely replaced. Figure 69 below shows an example of this window.

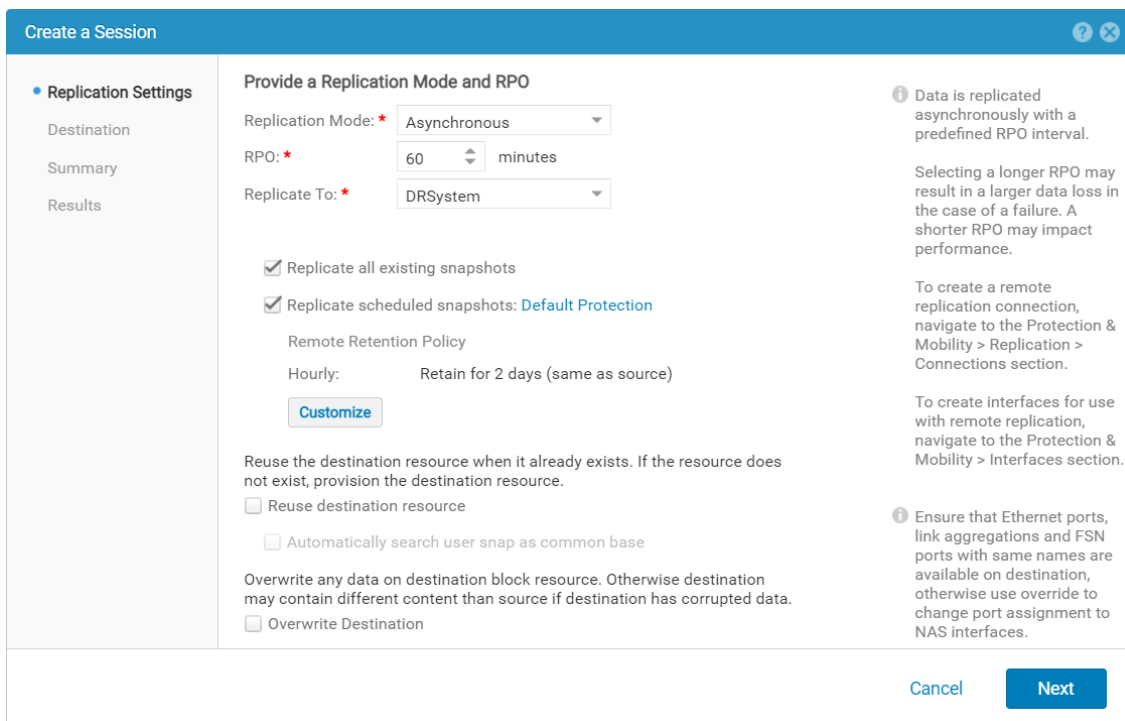


Figure 69 Create a Session wizard – Replication Settings step – LUN example

With the enhancements in the 5.1 release, changing File replication topologies with advanced File replication is possible. One example is if replication is configured in an A → B → C cascade configuration. If snapshots are created on System A and replicated to System B and System C, they can be used as common base when recreating replication sessions if required. As shown in the figure below, System B can easily be removed from the configuration as common base snapshots can be leveraged to create replication directly from System A to System C. These snapshots would reduce the amount of data that needs to synchronize between the source and destination.

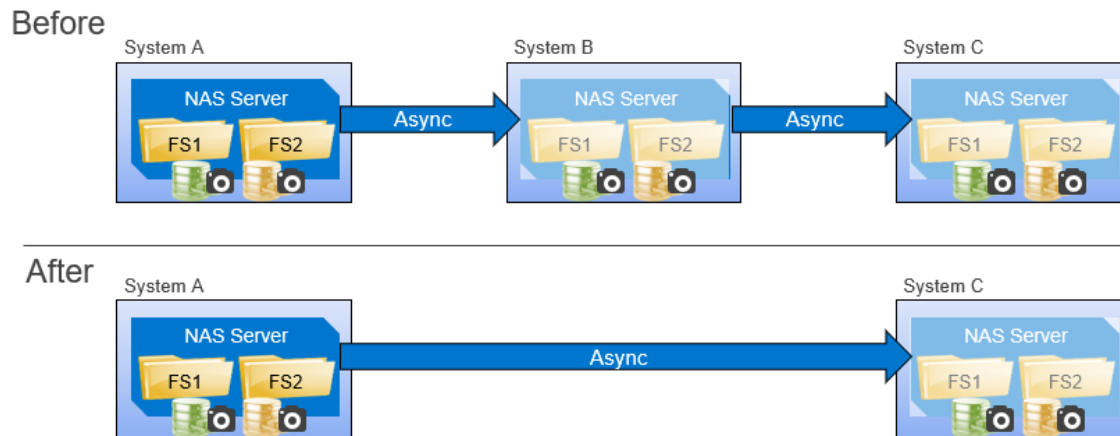


Figure 70 Create a Session wizard – Replication Settings step – LUN example

5.1.2 Failback/Resume after unplanned failover

Common base snapshots can also be used to recover from an unplanned failover event. Using common base snapshots may help reduce the amount of data that needs to be synchronized between the source and destination after the Failback or Resume option is initiated. An unplanned failover occurs when the Failover option is issued on the replication session for a destination resource. When this occurs, a disaster situation is assumed, and the destination resource becomes Read/Write. If management connectivity is still available between the source and destination systems, the source will automatically be made unavailable or Read-Only depending on the resource type. Common base snapshots and full copy avoidance is not available for block synchronous replication sessions.

Some customers require regular failover testing to ensure their disaster recovery operations are functional. When Failover is issued from the destination of a file synchronous replication session, a full synchronization of data is required when re-establishing the replication session. With the enhancements in OE version 5.1 and later, common base snapshots can help reduce the overall recovery time. It is highly suggested to create a new snapshot and replicate it to the destination on all resources involved just before issuing Failover on the replication session. As mentioned previously, only the latest 21 snapshots on a file synchronous replication session are searched for a common base.

To confirm a common base snapshot is present on a file synchronous replication session, the `uemcli /prot/rep/session/commonbase` command can be utilized. This command will search the source and destination resources and determine if a common base snapshot is present. This command should be run on the active replication session before issuing the Failover operation. If the command is run on a NAS Server, the outputs for all file systems and NFS datastores will be displayed. This command can also be run directly on a file system or NFS datastore replication session.

As shown below, common base snapshots were found for the first replication session, but not the second. The Session ID shown can be used to locate the resource if required. If the Failover operation was issued

while in this state, the first session would use snapshots to establish a common base, while the second session would require a full synchronization of data.

```
uemcli /prot/rep/session/commonbase -sessionName
rep_sess_nas_2_nas_9_APM01204908035_APM00211114469 show

1:   Session ID                               =
171798691855_APM01204908035_0000_171798692131_APM00211114469_0000
     Source common base snapshot ID           = 171798691873
     Destination common base snapshot ID      = 171798692191

2:   Session ID                               =
171798691889_APM01204908035_0000_171798692338_APM00211114469_0000
     Source common base snapshot ID           = (No common base snapshot)
     Destination common base snapshot ID      = (No common base snapshot)
```

For asynchronous block and file replication sessions, the internal system snapshots created by asynchronous replication can be used to recover after an unplanned failover event. In Dell EMC OE version 5.1 and later, replicated user snapshots can also be utilized. If for any reason the asynchronous replication snapshots are unavailable, the user created and replicated snapshots can be used to establish common base and avoid a full copy of data. This is also true for asynchronous replication sessions that are part of a MetroSync (file synchronous replication) solution.

5.2 File Import

Once a VNX VDM is imported as a NAS Server into a Dell EMC Unity system, the destination NAS Server and the file systems under it can participate in a replication session. If the import session for the source VNX VDM was committed, the destination NAS Server can be replicated to a system running Dell EMC Unity OE version 4.0 and later. For more information about File Import, please review the white paper titled *Dell EMC Unity: Migration Technologies* on Dell EMC Online Support.

5.3 Data Reduction

The following section gives information about Data Reduction on Dell EMC Unity systems. Systems with data efficiency technologies enabled running versions prior to Dell EMC Unity OE version 4.3 will see Compression instead of Data Reduction.

Storage Resources using data reduction can be replicated using any supported replication software, such as Native Synchronous Block Replication or Native Asynchronous Replication to any supported destination system. All data that is replicated, regardless if it is local replication or to a remote system, is first rehydrated, and then replicated to the destination. This method of replication ensures that all replication topologies are supported, including mixed configurations where data reduction is enabled only on one side. Replicating to systems which do not support data reduction is also supported, such as replicating to Dell EMC UnityVSA or a physical Dell EMC Unity system not running a code version which supports data reduction.

Dell EMC Unity Data Reduction can also be enabled on only the source, only the destination, or both the source and destination storage resources, depending on if the system and Pool configuration support Dell EMC Unity Data Reduction. This allows the user to fully control where to implement data reduction. One example of a supported replication configuration is when using Asynchronous Local Replication. The source storage resource may reside on an All Flash Pool and have data reduction enabled, but the destination may

be on a large capacity Hybrid Pool which does not support data reduction. Another example of a supported configuration is when replicating a storage resource from a Dell EMC UnityVSA system or a production system not using data reduction, to a storage resource with data reduction enabled on a remote system.

From Unisphere to enable data reduction on the destination resource, in the **Create a Session** wizard for the Destination step, select the checkbox in front of the **Data Reduction** option. Figure 71 gives an example on enabling data reduction on the destination for a LUN storage resource. The **Create a Session** wizard for thin clones, consistency groups, VMware VMFS datastores, file systems, and VMware NFS datastores will have the same **Data Reduction** option in the **Replication** step.

In Dell EMC Unity OE version 4.5, Advanced Deduplication was introduced for Dell EMC Unity 450F, 550F, and 650F systems and can be enabled once Data Reduction is enabled. In Dell EMC Unity OE version 5.0, Advanced Deduplication is also available on Dell EMC Unity 380(F), 480(F), 680(F), and 880(F). This additional data efficiency technology has the same restrictions as noted above in relation to replication.

For more information about Data Reduction and Advanced Deduplication, please review the white paper titled *Dell EMC Unity: Data Reduction* on Dell EMC Online Support.

The screenshot shows a 'Destination LUN Configuration' dialog box. It contains the following fields and options:

- Reuse destination resource
- Name: * TESTLUN
- Destination System: APM00211114469
- Pool: * Pool 1 (Traditional, Extreme Performance Pool, 26.5 TB free)
- Size(GB): 100.0
- Thin
- Data Reduction
- Advanced Deduplication
- Tiering Policy: Start High Then Auto-Ti

Buttons for 'Cancel' and 'OK' are located at the bottom right of the dialog.

Figure 71 Data Reduction and Advanced Deduplication

5.4 Dynamic Pools

Dell EMC Unity Native Synchronous and Asynchronous Replication is fully supported with Traditional and Dynamic Pools. When configuring local or remote replication, the source or destination Pool can be either type of Pool, either Traditional or Dynamic. This flexibility in configuration options allows systems which do not support Dynamic Pools, such as a Dell EMC UnityVSA. When configuring replication, a Type field has been added to easily determine the type of Pool that was selected as the destination. This Type field is only present on systems running Dell EMC Unity OE version 4.2 or later.

For more information about Dell EMC Unity Dynamic Pools, please review the white paper titled *Dell EMC Unity: Dynamic Pools* on Dell EMC Online Support.

For more details regarding the Interoperability of the features released in Dell EMC Unity OE version 4.1 or later, search for Interoperability for the Dell EMC Unity Family on Dell EMC Online Support.

6 Upgrades

Prior to initiating an upgrade of the Dell EMC Unity OE, it is recommended to pause all ongoing replication sessions. The Pre-Upgrade Health Check (PUHC) checks for any active replication sessions and generates a warning if any are found. Once the upgrade completes, you can resume each replication session. The pause and resume operations should be initiated from the source system for each replication session. If you are running Dell EMC Unity OE version 4.2 or later, pausing and resuming the NAS server replication session also pauses and resumes all its associated file system replication sessions automatically.

7 RecoverPoint with Dell EMC Unity

Along with the native replication options provided with physical Dell EMC Unity systems, RecoverPoint is also supported. RecoverPoint is used for disaster recovery and data loss protection, protecting organizations from site outages due to unforeseen circumstances and data loss due to corruption or human error. RecoverPoint also helps in data migration solutions, providing the ability to move data between data centers and supported systems. RecoverPoint provides a DVR-like roll back function which allows data recovery to any point-in-time. RecoverPoint uses physical or virtual RecoverPoint appliances to replicate data locally within the same Dell EMC Unity system, or to remote systems. Replication solutions are designed to ensure the integrity of the replicated data at local and remote sites. Performance is also not compromised when using RecoverPoint with Dell EMC Unity. By using proprietary software called the RecoverPoint Splitter, incoming write operations are split and sent to the local and remote storage system. This ensures there is minimal performance impact on the storage system when replicating data.

RecoverPoint can be configured to replicate Dell EMC Unity LUNs, VMware VMFS Datastores, and thin clones. Destination storage resources can be created locally within the same system, or on a remote system supported with RecoverPoint. Dell EMC Unity includes the RecoverPoint Basic license for free, allowing replication between Dell EMC Unity, VNX, and VNXe systems. Through RecoverPoint's own user interface, replication can be configured quickly and easily. For more information about which systems are supported, please view Appendix: Replication Maximums.

When configuring RecoverPoint, LUNs, VMware VMFS Datastores, and thin clones already replicated with Dell EMC Unity native replication software are not supported. Also, RecoverPoint should not be configured to use the Dell EMC Unity Synchronous Replication Port on either storage processor. The Synchronous Replication Port is discussed in the *Synchronous Replication Interfaces* section of this paper.

For more information about RecoverPoint, including RecoverPoint specific concepts and management, refer to the *RecoverPoint Administrator's Guide* found on Dell EMC Online Support.

8 Metro node

Metro node is an external hardware and software add-on feature for Unity XT for which it provides active/active synchronous replication and standard local use cases. Additionally, it also provides a solution locally with the Local mirror feature to protect data from a potential array failure. Both of these use cases provide solutions for true continuous availability with zero downtime.

Unity XT is viewed by metro node as ALUA array based on SCSI response data and therefore is required to follow the four active, four passive path connectivity rules. This rule states that both nodes of the metro node must each have four active and four passive paths to all volumes provisioned from the array. For additional information about Metro node, go to [Dell Support](#) and reference the Metro node best practices white paper.

9 Conclusion

In this paper, the various native replication solutions provided in Dell EMC Unity were discussed. Configuring a data protection solution helps guard against unforeseen situations, such as data loss or site-wide outages if they are encountered. Dell EMC Unity provides local and remote data protection solutions to help minimize the cost associated with downtime and provides easy recovery in the event of a disaster. Through the use of synchronous and asynchronous replication solutions, data protection can be configured to meet the needs of the application and organization.

Native Synchronous Block Replication is a data protection solution which replicates each write to a storage resource remotely to a peer system before the write is acknowledged with the host. In the event of a disaster, Synchronous Block Replication provides maximum protection and ensures that there is zero data loss. Dell EMC Unity Synchronous Block Replication is supported on LUNs, thin clones, consistency groups, and VMware VMFS datastores configured on Dell EMC Unity All Flash and Hybrid Flash systems. Starting in Dell EMC Unity OE version 4.4, native file synchronous replication is also supported on Dell EMC Unity. See the *Dell EMC Unity: MetroSync* white paper on Dell EMC Online Support for more information.

Native Asynchronous Block and File Replication is a data protection solution which replicates storage resources locally within the same Dell EMC Unity system, or remotely to other Dell EMC Unity systems. Asynchronous replication leverages Dell EMC Unity's Unified Snapshot technology to provide consistent point-in-time replicas which can be used in the event of a disaster. Asynchronous replication is not only supported on Dell EMC Unity All Flash and Hybrid Flash systems, but also on Dell EMC UnityVSA. When using asynchronous replication, no impact to host I/O is seen as data is not replicated as it enters the system. Asynchronous replication uses a customizable Recovery Point Objective, which automatically replicates changes in data at consistent intervals. When data needs to be replicated over long distances, asynchronous replication can meet the needs of an organization.

RecoverPoint support allows Dell EMC Unity to leverage the advanced replication features it provides. With RecoverPoint, Dell EMC Unity LUNs and VMware VMFS Datastores can be replicated locally or remotely to another supported system. With RecoverPoint's advanced functionality, such as point-in-time data recovery, Dell EMC Unity can be protected from disaster scenarios.

A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

A.1 Related resources

The following references can be found on Dell EMC Online Support:

- Dell EMC Unity: Best Practices Guide
- Dell EMC Unity: Cloud Tiering Appliance (CTA)
- Dell EMC Unity: Compression for File
- Dell EMC Unity: Data at Rest Encryption
- Dell EMC Unity: Data Integrity
- Dell EMC Unity: Data Reduction
- Dell EMC Unity: DR Access and Testing
- Dell EMC Unity: Dynamic Pools
- Dell EMC Unity: FAST Technology Overview
- Dell EMC Unity: File-Level Retention (FLR)
- Dell EMC Unity: High Availability
- Dell EMC Unity: Introduction to the Platform
- Dell EMC Unity XT: Introduction to the Platform
- Dell EMC Unity: NAS Capabilities
- Dell EMC Unity: MetroSync
- Dell EMC Unity: MetroSync and Home Directories
- Dell EMC Unity: MetroSync and VMware vSphere NFS Datastores
- Dell EMC Unity: Migration Technologies
- Dell EMC Unity: OpenStack Best Practices for Ocata Release
- Dell EMC Unity: Performance Metrics
- Dell EMC Unity: Snapshots and Thin Clones
- Dell EMC Unity: Operating Environment (OE) Overview
- Dell EMC Unity: Unisphere Overview
- Dell EMC Unity: Virtualization Integration
- Dell EMC UnityVSA
- Dell EMC Unity Cloud Edition with VMware Cloud on AWS
- Dell EMC Unity Data Reduction Analysis
- Dell EMC Unity: Migrating to Dell EMC Unity with SAN Copy
- Dell EMC Unity Storage with Microsoft Hyper-V
- Dell EMC Unity Storage with Microsoft SQL Server
- Dell EMC Unity Storage with Microsoft Exchange Server
- Dell EMC Unity Storage with VMware vSphere
- Dell EMC Unity Storage with Oracle Databases
- Dell EMC Unity 350F Storage with VMware Horizon View VDI
- Dell EMC Unity: 3,000 VMware Horizon Linked Clone VDI Users
- Dell EMC Storage with VMware Cloud Foundation
- Metro node best practices

B Appendix: Replication Maximums

The following table outlines the Dell EMC Unity replication system maximums.

	Dell EMC Unity 600/600F/650F 680/680F	Dell EMC Unity 500/500F/550F 580/580F	Dell EMC Unity 400/400F/450F 480/480F	Dell EMC Unity 300/300F/350F 380/380F	Dell EMC UnityVSA
Max Replication Sessions (Synchronous + Asynchronous)	2000	1500	1000	1000	16
Max Replication Sessions (Synchronous Replication)	2000	1000	750	500	N/A
Max Consistency Group Replication Sessions (Synchronous Replication)	128	64	64	64	N/A
Max LUNs per Replicated Consistency Group (Synchronous Replication)	32	32	32	32	N/A
Max Target Systems (Synchronous Replication)	1	1	1	1	N/A
Max Concurrent Replication Sessions (Asynchronous Replication)	256	256	256	256	8
Max Concurrent Initial Syncs (Asynchronous Replication)	32	32	32	32	4
Max LUNs per Replicated Consistency Group (Asynchronous Replication)	75	75	75	75	50
Max replicated NAS Servers (Asynchronous Replication)	256	128	128	90	4
Max Target Systems (Asynchronous Replication)	16	16	16	16	16

The table above outlines several system maximums when using synchronous and asynchronous replication. The maximum replication sessions include all replication sessions on the system, which includes both synchronous and asynchronous replication sessions, local or remote. The replication destination storage resources count towards the system maximums, even though they are not host accessible when currently a destination image. In Dell EMC Unity, only one replication connection used for synchronous replication, or synchronous and asynchronous replication, can be created. This also means that only one pair of systems can replicate synchronously to each other.

C Appendix: Replication Support Across Platforms

The following table outlines replication support across the various Dell EMC midrange platforms.

Source	Destination	Block Replication		File Replication		RecoverPoint (Block)
		Sync	Async	Sync**	Async	
Dell EMC Unity	Dell EMC Unity	✓	✓	✓	✓	✓
Dell EMC Unity	Dell EMC UnityVSA	✗	✓	✗	✓	✗
Dell EMC Unity	VNXe3200*	✗	✓	✗	✗	✓
Dell EMC Unity	VNX1/VNX2	✗	✗	✗	✗	✓
Dell EMC Unity	VNXe1600	✗	✓	✗	✗	✗
Dell EMC UnityVSA	Dell EMC UnityVSA	✗	✓	✗	✓	✗
Dell EMC UnityVSA	Dell EMC Unity	✗	✓	✗	✓	✗
Dell EMC UnityVSA	VNXe3200*	✗	✓	✗	✗	✗
Dell EMC UnityVSA	VNX1/VNX2	✗	✗	✗	✗	✗
Dell EMC UnityVSA	VNXe1600	✗	✓	✗	✗	✗
VNX1/VNX2	Dell EMC Unity	✗	✗	✗	✗	✓

*Min VNXe3200 OE v3.1.5.6801782
 **Min Dell EMC Unity OE version 4.4

Footnote:
 ✓ – Supported
 ✗ – Not Supported