

Dell PowerStore: Replication Technologies

May 2024

H18153.10

White Paper

Abstract

This white paper explains the replication technologies for the Dell PowerStore platform. It outlines the native and non-native options available for replicating data and describes managing replication and its benefits.

Copyright

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.

Copyright © 2020-2024 Dell Inc. or its subsidiaries. All Rights Reserved. Published in the USA May 2024 H18153.10.

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

Contents

Executive summary	4
Introduction	6
Remote system configuration	12
Native replication for block and file	29
Asynchronous replication for vVol based VMs	55
Metro Volume	58
System limits	58
Integration with PowerStore	58
Conclusion	61
Appendix A: Replication support across platforms	62
References	63

Executive summary

Overview

Dell PowerStore provides native and non-native solutions to protect data and to help organizations meet business goals for both data availability and protection. PowerStore native replication solutions can replicate data to other systems, whether they are at the same site or a remote facility. Having remote copies of data protects against outages on the main system. Data protection features in PowerStore also enable quick recovery on a destination system with minimal to no data loss, depending on the replication method selected.

Native replication and metro can be configured and managed in PowerStore Manager, PowerStore CLI, or REST API. PowerStore Manager is an intuitive HTML5-based interface that allows users to configure and manage their replication setup and provides a visual representation of the configuration.

This white paper describes the following technologies to replicate data for PowerStore:

- Native asynchronous replication for
 - Block
 - File
 - vVol based VMs
- Native synchronous replication for
 - Block (Active-Passive)
 - File
- Dell RecoverPoint for Virtual Machines

Additional white papers cover data protection and replication technologies in further detail.

Metro Volume

Native synchronous replication that provides active-active access is available with Metro Volumes. The [Dell PowerStore: Metro Volume](#) white paper describes this feature in detail.

Snapshots and thin clones

The native integration for snapshot backup can back up volumes and volume groups directly to a PowerProtect DD series appliance. This feature eliminates the requirement of a backup host because all backup traffic is offloaded to the storage appliances. More details are available in the white paper [Dell PowerStore: Snapshots and Thin Clones](#).

Dell RecoverPoint for Virtual Machines

Dell RecoverPoint for Virtual Machines is a virtual appliance that offers an alternative solution for VM replication for PowerStore. RecoverPoint is configured for VM protection through the intuitive Dell Unisphere Manager for RecoverPoint interface. Due to its agnostic nature, RecoverPoint for Virtual Machines enables recovering VM data for any point in time and replicating the data towards many other storage systems. More information is available in the KB article [RecoverPoint for VMs: How to protect Virtual Machines using shared VMDK or RDM disk\(s\)](#)

Dell metro node

Dell metro node offers continuous application data availability and transparent data mobility for block storage. Metro node is placed into the data path between hosts and creates a flexible storage architecture. Solution white papers that describe how to integrate metro node with PowerStore are available on the [Dell Technologies Info Hub](#).

Audience

This white paper is intended for Dell Technologies customers, partners, and employees who are considering using PowerStore native replication or RecoverPoint for Virtual Machines for PowerStore. The document assumes familiarity with the PowerStore system and management software.

Revisions

Date	Part number/ revision	Description
April 2020	H18153	Initial release: PowerStoreOS 1.0
May 2020	H18153.1	Minor updates
August 2020	H18153.2	Minor updates
January 2021	H18153.3	Metro node updates
April 2021	H18153.4	PowerStoreOS 2.0 updates including failover test
November 2021	H18153.5	Template update
July 2022	H18153.6	PowerStoreOS 3.0 updates
October 2022	H18153.7	PowerStoreOS 3.2 updates
May 2023	H18153.8	PowerStoreOS 3.5 updates: <ul style="list-style-type: none"> Added information about secure snapshots Updated metro node content
October 2023	H18153.9	PowerStoreOS 3.6 updates: <ul style="list-style-type: none"> Added reference for Metro Witness support
May 2024	H18153.10	PowerStoreOS 4.0 updates: <ul style="list-style-type: none"> Added replication network enhancements Added synchronous replication for file and block Removed references to PowerStore X

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Authors: Robert Weilhammer, Ethan Stokes, Wei Chen

Note: For links to other documentation for this topic, see the [PowerStore Info Hub](#).

Introduction

Business continuity planning

Data is one of the most valuable assets to an organization. Because users and their customers access data constantly, directly and indirectly using various applications, data is a crucial part of day-to-day operations. Outages can occur at any time and can be restricted to a single system or to an entire data center or location. Whether they are planned outages such as regular maintenance, or unplanned events such as a power outage, it is a top priority to ensure that critical data is always available.

A business continuity plan for critical data can prevent these costly outages. To protect against different scenarios, an organization should plan and implement a data-protection strategy that includes a data-replication solution.

Asynchronous replication can be used to protect against a storage-system outage by creating a copy of data to a remote system. Replication is a software feature that 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. It also provides a remote location that can resume production and minimize downtime due to a disaster. The PowerStore platform offers many data-protection solutions that can meet disaster recovery needs in various environments.

Asynchronous replication is primarily used to replicate data over long distances, but it can be used to replicate to systems within the same location also. The asynchronous replication for PowerStore is designed to have minimal impact on host I/O latency. Host writes are acknowledged when they are saved to the local storage resource, and no additional writes are needed for change tracking. Because write operations are not immediately replicated to a destination resource, all writes are tracked on the source. This data is replicated during the next synchronization. With protection policies, asynchronous replication uses the concept of a recovery point objective (RPO). The RPO is the acceptable amount of data, measured in units of time, that can be lost due to an outage. This delta of time affects the amount of data that must be replicated during the next synchronization. It also reflects the amount of potential data loss in a disaster scenario. PowerStore asynchronous replication features can be configured using PowerStore Manager, PowerStore CLI, or REST API. RecoverPoint for Virtual Machines supports VM replication for PowerStore and is configured using the Unisphere Manager for RecoverPoint user interface.

Metro and synchronous replication focus on zero RPO for the replicated data. PowerStore's built-in synchronous replication technology ensures that all data is identical on the local and remote storage resources. Compared to asynchronous replication, a volume with active synchronous replication or metro configuration acknowledges the host writes after data is saved on the remote storage resources. This has some impact on performance due to the additional latency from the replicated writes.

PowerStore overview

PowerStore achieves new levels of operational simplicity and agility. It uses a container-based microservices architecture, advanced storage technologies, and integrated machine learning to unlock the power of your data. PowerStore is a versatile platform with

a performance-centric design that delivers multidimensional scale, always-on data reduction, and support for next-generation media.

PowerStore brings the simplicity of public cloud to on-premises infrastructure, streamlining operations with an integrated machine-learning engine and seamless automation. It also offers predictive analytics to easily monitor, analyze, and troubleshoot the environment. PowerStore is highly adaptable, providing the flexibility to host specialized workloads directly on the appliance and modernize infrastructure without disruption. It also offers investment protection through flexible payment solutions and data-in-place upgrades.

Terminology

The following table provides definitions for some of the terms that are used in this document.

Table 1. Terminology

Term	Definition
ALUA	Asynchronous Logical Unit Access. PowerStore uses implicit ALUA which allows PowerStore to provide a recommended active optimized path to a storage resource for the hosts.
Asynchronous replication	A replication method that allows replicating data over long distances and maintaining a replica at a destination site. Updates to the destination image can be issued manually, or automatically based on a customizable RPO.
Bandwidth	The amount of data, represented in MB/s, which can be transferred in a given period.
Common base	A pair of snapshots that are taken on a replication source and destination storage resource that have the same point-in-time image.
Destination storage resource	A storage resource that is used for disaster recovery in a replication session. This term is also known as a target image.
Internal snapshot (replication snapshot)	The system creates unified snapshots and is part of an asynchronous replication session. These snapshots are only visible in the PowerStore CLI or PowerStore REST API, and manual modification is not possible. Each asynchronous replication session uses up to two internal snapshots that are taken on the source and destination storage resources. Each session also takes up one read/write snapshot on destination storage system. The last successful internal read-only (RO) snapshots for source and destination storage resources and are used as a common base.
PowerStore Manager	A web-based management interface for creating storage resources and configuring and scheduling protection of stored data on PowerStore. PowerStore Manager can be used for all management of PowerStore native replication.
PowerStore CLI	A tool that can be installed on an operating system to manage a PowerStore system.
RecoverPoint for Virtual Machines	Protects virtual machines (VMs) in a VMware environment with VM-level granularity and provides local or remote replication for any point-in-time recovery. This feature is integrated with VMware vCenter and has integrated orchestration and automation capabilities.

Term	Definition
Recovery point objective (RPO)	Acceptable amount of data, which is measured in units of time, that may be lost due to a failure. For example, if a storage resource has a one-hour RPO, data that is written to the storage resource within the last hour may be lost when the replication session is failed over to the destination storage resource.
Recovery time objective (RTO)	Duration of time in which a business process must be restored after a disaster. For example, an RTO of one hour requires restoring data access within one hour after a disaster occurs.
Remote systems	Relationship that is configured between two PowerStore systems.
Replication session	A relationship that is configured between two storage resources of the same type on different systems, and automatically synchronizes data from one resource to another.
Snapshot	Also called a unified snapshot, a snapshot is a point-in-time view of a storage resource. When a snapshot is taken, it creates 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	A top-level object that a user can provision, which is associated with a specific quantity of storage. All host access and data-protection activities are performed at this level. In this document, storage resources refer to resources that support replication such as volumes, volume groups, and thin clones.
Synchronous Replication	<p>A replication method to keep the source and destination of a replication system consistent (zero RPO). PowerStore supports active-active synchronous replication with metro active-active replication and active-passive synchronous replication:</p> <p>Active-active – bi-directional synchronization where either mirrored volume can accept host IO (metro). Supports transparent failover.</p> <p>Active-passive – unidirectional synchronous replication where only the source volume accepts host IO. An external triggered failover (such as by cluster software) is required to enable the volume on destination and for re-protection.</p> <p>In both configurations, data is only acknowledged to the hosts when data has been acknowledged by the active site and the replication destination.</p>
Thin clone	A read/write copy of a volume, volume group, file system, NAS server, or snapshot that shares blocks with the parent resource.
Unisphere Manager 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 that are configured to use RecoverPoint. Consistency groups are created, replicated, and recovered through this interface.

Term	Definition
User snapshot	A snapshot that is created manually by the user or by a protection policy with an associated snapshot rule. This snapshot type is different than an internal snapshot, which is taken automatically by the system with asynchronous replication.
Volume	A block-based storage resource that a user provisions. It represents a SCSI logical unit.
Volume group	A storage instance that contains one or more volumes within a storage system. Volume groups can be configured with write-order consistency and help organize the storage that is allocated for particular hosts.

Replication methods

There are multiple replication approaches, but two methods are highly recognized in the storage industry: asynchronous and synchronous. PowerStore supports asynchronous replication of block volumes, vVols, and file. It also supports synchronous replication (active-passive) of file and block volumes. For Metro Volumes, PowerStore uses a Symmetric Active/Active architecture. This means that either volume can be a synchronous replication source and either volume can be a synchronous replication destination. Synchronous replication happens bidirectionally between clusters that support Metro Volumes.

Synchronous replication

Synchronous replication, introduced with PowerStoreOS 4.0, guarantees data consistency (zero data loss) between the replication source and destination volumes during normal operation. This data consistency is achieved by ensuring write I/O commitments at the replication source and destination before a successful write acknowledgment is sent back to the host and the requesting application. In a metro configuration synchronous replication provides a blend of data consistency and high availability with uniform storage presentation. Traditional synchronous replication requires a triggered failover to enable the DR storage resource.

With synchronous replication, any source of latency that impacts the source or destination volume, or the replication link in-between, adversely impacts applications in terms of latency (slowness) and availability. This applies to any kind of synchronous replication including Metro Volumes which is built on top of synchronous replication. For this reason, appropriate performance sizing is paramount for the source and destination storage, as well as the replication bandwidth and any other upstream infrastructure on which the storage depends.

Figure 1 demonstrates the write I/O sequence of synchronous replication:

1. The application or server sends a write request to the source volume.
2. The write I/O is mirrored to the destination volume.
3. The mirrored write I/O is committed to the destination volume.
4. The write commit at the destination volume is acknowledged back to the source volume.
5. The write I/O is committed to the source volume.

6. The write acknowledgment is sent to the application or server.

This process is repeated for each write I/O requested by the application or server.

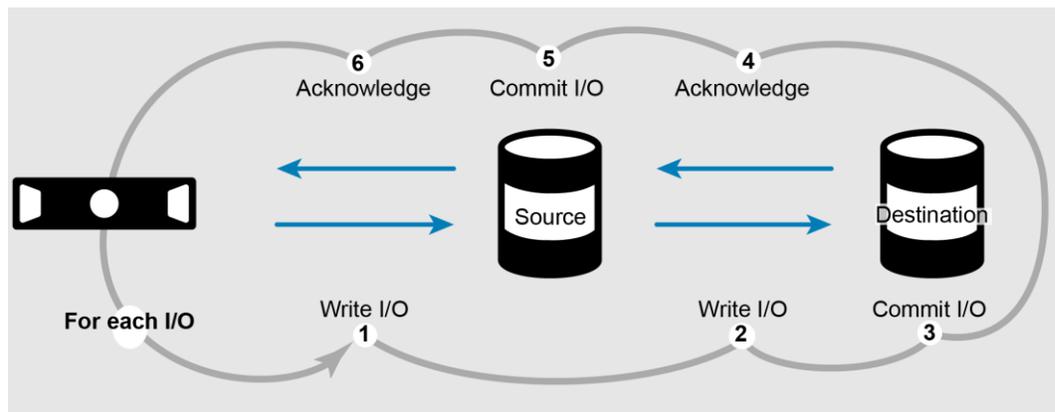


Figure 1. Block storage synchronous replication write I/O sequence

Note: PowerStore supports synchronous replication as part of the Metro Volume feature introduced in PowerStoreOS 3.0. For more information, see the [Dell PowerStore: Metro Volume white paper](#).

Asynchronous replication

Asynchronous replication accomplishes similar data protection goals in that data is replicated from source storage to destination storage in a unidirectional way. However, the manner and frequency with which the data is replicated differs from synchronous replication. With asynchronous replication, instead of committing a write at both replication source and destination simultaneously, the write is committed only at the source and an acknowledgment is immediately sent to the host and application. The accumulated committed writes at the source volume are replicated to the destination volume in one batch at scheduled intervals and committed to the destination volume.

PowerStore replication rules combine to form a protection policy (applied granularly per volume or NAS server) that dictates the asynchronous replication intervals and RPO for the volume. A PowerStore protection policy can include a replication rule and snapshot rules at the same time for local and remote protection of storage resources.

Starting with PowerStoreOS 3.5, secure snapshots are available for volumes and volume groups. When the secure snapshot setting is enabled, the snapshot is protected from deletion until the retention period expires. If both the source and destination systems are running PowerStoreOS 3.5 or later, snapshots and snapshot rules that have secure snapshots enabled are replicated as secure to the destination system. However, if the destination system is running an earlier version of PowerStoreOS, snapshots and snapshot rules are replicated as regular snapshots and rules without the secure option. In a failover scenario, where the primary system is running a code earlier than the PowerStoreOS 3.5 release, the system takes regular snapshots instead. If secure snapshots are used, having both systems running PowerStoreOS 3.5 or later is recommended for compatibility purposes.

Volumes can adhere to their own independent replication schedule, or they can share a replication schedule with other volumes that leverage the same protection policy. Because

asynchronously replicated transactions are not required to wait for write committals at the replica destination volume, the replication link and/or destination storage will not contribute to application or transaction latency at the source volume.

Figure 2 demonstrates the write I/O pattern sequence for asynchronous replication:

1. The application or server sends a write request to the source volume.
2. The write I/O is committed to the source volume.
3. The write acknowledgment is sent to the application or server.

Steps 1 through 3 are repeated for each write I/O requested by the application or server.

4. Periodically, a batch of write I/Os that have already been committed to the source volume are transferred to the destination volume.
5. The write I/Os are committed to the destination volume.
6. A batch acknowledgment is sent to the source.

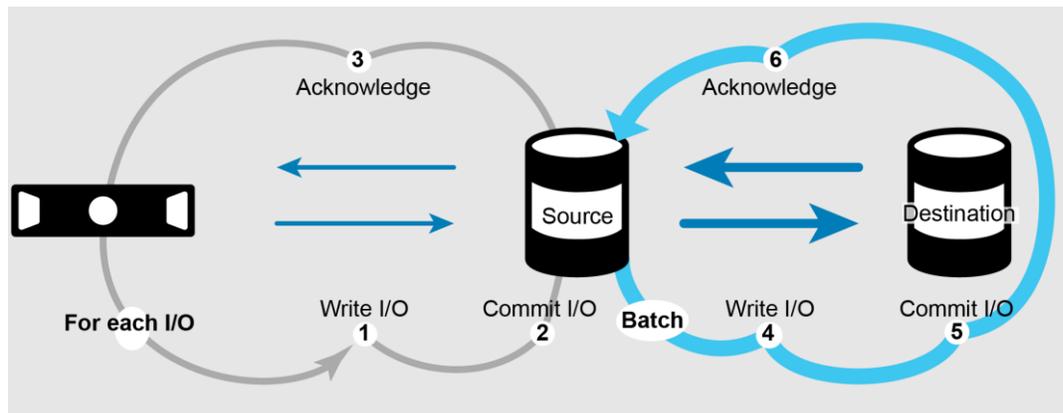


Figure 2. Block storage asynchronous replication write I/O sequence

PowerStore replication overview

Table 2 provides an overview of the native replication types that can be configured to replicate data between two PowerStore clusters.

Table 2. PowerStore replication overview

	Asynchronous	Synchronous	Metro
Type	Block and File	Block and File	Block only
Storage Resources	Volumes, Volume Groups, Thin Clones, NAS Servers, vVols	Volumes, Volume Groups, Thin Clones, NAS Servers	Volumes, Volume Groups
Replication	Asynchronous	Synchronous	Synchronous
Target RPO	Fixed values: 5min – 24h	0	0
Host access	Active – Passive Requires a failover	Active – Passive Requires a failover	Active – Active ALUA path change
Host Protocols*	SCSI, NVMe	SCSI, NVMe	SCSI

	Asynchronous	Synchronous	Metro
Block WWN / NQN	Different	Different	Same WWN on both ends
OS Support	“Any” supported OS on PowerStore	“Any” supported OS on PowerStore	ESXi, Linux, Windows
Witness	No	No	Yes
RTT	Not relevant	5ms	5ms*
Performance impact for host access	Minimal impact depending on sizing and workload	Adds additional latency (mirror round-trip-time)	Adds additional latency (mirror round-trip-time)
Snapshot replication	Replication of block snapshots on source / No Snapshot replication for File	Block snapshots nearly identical / No Snapshot replication for File	Snapshots nearly identical
Failover Test	Yes	Yes	Not applicable
Conversion Async <-> Sync	Without full sync for block resources in both directions. File replication requires full sync after reassigning a protection policy containing a replication rule with a different replication type.	Not supported	
Recovery Snap	Common base at each replication cycle	Every 30 min	Every 30 min
Inter cluster migration	Supported when no ongoing replication cycle	Supported when paused	Not supported
NDU	Paused during NDU	Active sessions continue	Active sessions continue

* Some protected applications may require lower RTT / Distance for metro configurations.

Remote system configuration

Introduction

This section describes the remote system configuration that is required for all native replication technologies supported by PowerStoreOS. A remote system configuration specifies the replication relationship between two configured PowerStore Clusters or between a PowerStore cluster and a Dell PowerProtect DD, and contains information about related network information for management and data connection. After a remote system pair is set up, the remote system configuration can be used on both participating PowerStore clusters for replication in any direction for available capabilities. The allowed capabilities for a remote system pair configuration depend on the expected and configured network latency configuration for the replication data network.

All synchronous replication types require a low latency replication data network with 5ms round trip latency or less. The replication network latency has a direct impact on the

expected host IO performance because each IO is acknowledged to the host only after the IO is committed for both the replication source and the destination.

Latency independent replication types:

- Block asynchronous
- vVol asynchronous
- File asynchronous

Only when latency is set to less than 5ms and latency dependent replication types:

- Block Metro synchronous
- Block synchronous
- File synchronous

When a remote system pair is configured with 5ms and used for synchronous replication or metro, it is not supported to change the latency configuration to a higher value.

The replication with PowerStoreOS supports up to eight remote system pairs.

Overview and prerequisites

PowerStore embedded replication has several physical and software components. Each of these components is described in the following sections for the supported types of replication. To prepare a remote system configuration, perform the following steps:

1. Verify that a storage network for replication is configured on both PowerStore clusters. Both systems can be either in the same network or in different networks with bi-directional routing. In a routed environment over a WAN connection, consider the latency requirements for the planned use case. For PowerStore running on OS 1.x-3.x, it is required to tag a storage network for replication in the ports view of an appliance. PowerStoreOS 4.0 and later uses network groups to define the network and ports for replication. Both are covered in detail in the section [Storage network IPs and replication ports](#).
2. For management traffic, ensure that network connectivity exists between the participating PowerStore clusters on the management interfaces. The management connection is critical for orchestrating replication across participating systems initiated by PowerStore. Similar to the data connection, it can be within the same network or over a routed connection.
3. For file replication, configure an additional file mobility network. The file mobility network consists of three additional IP addresses per PowerStore cluster that leverage the existing management network VLAN, gateway, and netmask. These interfaces are mapped to the 1 GbE management ports, sharing the physical port with the existing management interfaces. Each PowerStore cluster intended to support file replication must have the file mobility network configured. In PowerStoreOS 3.0, no changes to the file mobility network are supported while a file replication session is in place. PowerStoreOS 3.2 and later allows deleting and re-creating the file mobility network while having file replication in place.

The following sections outline the different functions, including requirement details, and how these components interact with each other. Use PowerStore Manager to configure and manage these components.

Port and network configuration for replication

In PowerStoreOS 1.x-3.x tagged ports for replication are used to transport data to a destination system for remote replication. By default, the system tags the bond0 port group on the 4-port card (port 0 + port 1) for replication traffic. In this configuration, the system uses the same storage network for host access to storage resources and replication data traffic. Tagged ports for remote replication can be modified in PowerStore Manager and are the same for both nodes of a PowerStore appliance. Tagging replication ports in PowerStore Manager is not related to VLAN tagging on network infrastructure. The replication is performed over Ethernet ports available on the system.

Starting with PowerStoreOS 4.0, the replication network configuration is enhanced to increase flexibility, scalability, and resiliency. Different port and network pairs can be used to declare one or multiple network groups for a remote system pair leveraged. Each remote system pair has its own replication network group configuration for the replication data traffic. Details of the configuration are covered in the section [Replication data network configuration \(PowerStoreOS 4.0 and later\)](#).

For replication, all supported Ethernet ports that can be used for a storage network can be used for replication. This includes the following:

- 4-Port Card
 - 10 GbE BaseT
 - 10/25 GbE optical
- Onboard Ports 0/1 (PowerStore 500T model only)
 - 10GbE optical
- IO-Module
 - 1-10GbE BaseT
 - 10/25GbE optical
 - 100 GbE IO-Modules

The figures in the section [Replication connection](#) show examples of minimal cabling for replication between PowerStore appliances ([Figure 3](#)). The link-aggregated ports (4-Port card Port 0 and Port 1) provide high availability, maximum throughput, and load balancing of replication traffic across physical ports in the aggregation. For a successful replication connection, all replication ports on a source system must be able to communicate with all replication ports on the destination system, and conversely. The communication could be either on a local network or in a routed network.

When planning for replication using a configuration with system bond0, consider that the ports might also be used for other traffic such as I/O for block storage host access, migration, or file. If these features are used, planning replication using dedicated interfaces is recommended. Extra ports use dynamic storage IP configuration from the range added after cluster initialization in the networking section of PowerStore Manager.

Replication connection

[Figure 3](#) shows a sample configuration of a replication connection between two physical systems. It shows cabling for a pair of PowerStore appliances using system bond0. The source of the replication session is the Production System, and the destination is the DR

System. For each of these systems, the system bond0 is used as replication ports and all management ports are connected to the same L2/L3 network.

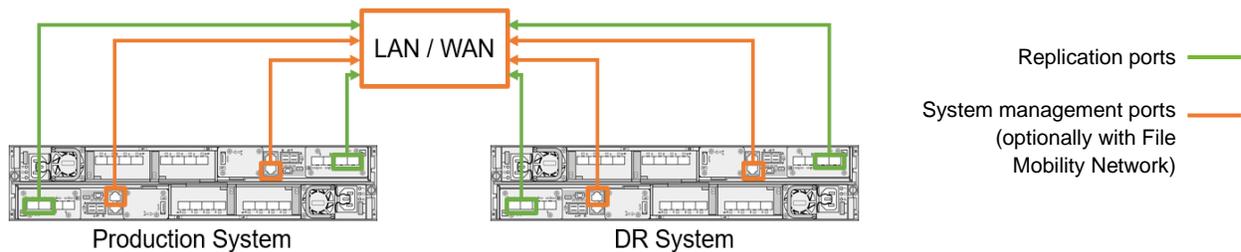


Figure 3. Native replication using two PowerStore T model systems

Remote systems

When the ports for replication traffic are configured and connected to the network, you can make a remote system connection between the arrays. After successful initialization, the remote system connection is automatically created on the peer system and can be used in both directions. The verify and update operation is used to update the replication connection information about the system on which it is issued. 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 check and reestablish (when required) the replication connection to a remote system after an outage. Running verify and update is a common use case when the storage network IP address pool has been changed by a network administrator.

All PowerStore native replication features rely on the same remote systems configuration.

- Asynchronous and synchronous block replication
- Metro Volume
- Asynchronous vVol replication
- Asynchronous and synchronous File replication

Remote systems configuration

Creating and managing replication in PowerStore Manager is easy and intuitive. All replication operations, including configuring of replication network ports, replication connections, and replication sessions, can be performed in the PowerStore Manager UI. With the help of wizards, replication can be configured by IT generalists or by advanced users. Replication can also be configured using the PowerStore Manager CLI or REST API. For more information about configuring and managing replication using the PowerStore Manager CLI, see the *Dell PowerStore Manager Command Line Interface Guide*. For more information about the REST API, use SwaggerUI (https://<PowerStore_Cluster_IP>/swaggerui) or see the *Dell PowerStore REST API Programmer's Guide*.

The following sections outline the remaining steps that are required to configure remote replication in PowerStore Manager. Each of the following operations is completed from a particular page in PowerStore Manager. Each page is discussed in detail in the following sections. For more information about using PowerStore Manager to configure and manage replication, see PowerStore Manager Online Help.

Storage network IPs and replication ports

When planning replication between two PowerStore arrays, consider the following:

- With PowerStore OS before 4.0, only one interface or system bond can be tagged for replication. If the system bond interfaces are connected to different switches, replication can continue even if one switch is down.
- PowerStoreOS 4.0 and later uses network groups for the remote system pair to specify the replication data networks. A remote system pair configuration can leverage multiple replication network groups to distribute the replication traffic.
- When file service is configured, configure a link aggregation interface for file services when using system bond for replication traffic.
- When host traffic is configured on an interface configured for replication, the available bandwidth is shared. This configuration can have an impact on host and replication performance.
- If host traffic and replication traffic are using the same network but different ports, configure host multipathing without using ports tagged for replication.
- In a configuration with multiple IP networks, the IP address ranges must not overlap with existing IP address ranges configured on the system.
- Only storage networks that are configured with an iSCSI purpose (OS <4.0) or replication purpose (OS 4.0 and later) can be used for replication.
- Depending on workload and data change rate on replicated volumes, a higher port speed might be required to steadily meet the RPO target for asynchronous replication. It is even more important when metro volumes or synchronous replication is set up for continuous replication.
- Starting with PowerStoreOS 3.0, it is also possible to use ports other than system bond for file I/O and to create additional link aggregation for file I/O.

Configure the replication network on each PowerStore cluster before setting them up as a remote system pair. The remote system pair configuration does not replicate the replication network configuration.

Replication data network configuration (PowerStoreOS 1.x-3.x)

This section shows the configuration for shared network ports as it has been supported since PowerStoreOS 1.x. A single storage network is used for host I/O or import, and replication-related data using the storage network. PowerStoreOS 2.x and later allows different storage networks for host access and replication data network.

Each port for a storage network configuration on a PowerStore requires its own IP address. When it is planned to extend an existing storage network, check the available storage IP addresses before creating interfaces. To verify the settings for storage network IPs, select **Settings > Networking > Network IPs**. Ensure that at least two storage network IPs for each appliance in the cluster configuration are unallocated for mapping new storage network ports which are distributed across the nodes. To tag new replication ports in PowerStore Manager, select **Hardware > Appliance-Name**, and select the **Ports** tab. All Ethernet ports and system-bond are eligible to be tagged as replication ports and are available in the ports list.

Figure 4 shows the PowerStore Manager Ports page. The figure also shows the default Link Aggregation ports (system-bond/bond0) that are set up on the system and are already tagged for replication using Default Storage Network, which was created beforehand. From this page, you can map ports to the storage network and change the tagging of replication data interfaces.

Node-Module-Name	Link State	Mapped for Storage	Tagged for Replication
BaseEnclosure-NodeA-EmbeddedModule-MgmtPort	🔗	⊘	⊘
BaseEnclosure-NodeA-bond0	🔗	✓ Default Storage Network	✓ Default Storage Network
BaseEnclosure-NodeA-4PortCard-FEPort0	🔗	⊘	⊘
BaseEnclosure-NodeA-4PortCard-FEPort1	🔗	⊘	⊘
BaseEnclosure-NodeB-EmbeddedModule-MgmtPort	🔗	⊘	⊘
BaseEnclosure-NodeB-bond0	🔗	✓ Default Storage Network	✓ Default Storage Network
BaseEnclosure-NodeB-4PortCard-FEPort0	🔗	⊘	⊘
BaseEnclosure-NodeB-4PortCard-FEPort1	🔗	⊘	⊘

Figure 4. Ports overview page

To change the replication data port to a port other than the system bond or vFE1 Port on port group TGT1, map a new set of ports to the storage network. It is only required to run these steps for a single node. PowerStore Manager configures the peer node in parallel.

The example in Figure 5 shows how to map a storage network:

1. Select the port.
2. Click **MAP STORAGE NETWORK**.
If only a single port is selected, PowerStore Manager automatically configures the corresponding port on the peer node.
3. Select the storage network to be mapped.
4. Confirm the selected mapping with **MAP STORAGE**.
5. To finish the configuration, confirm the following dialog.

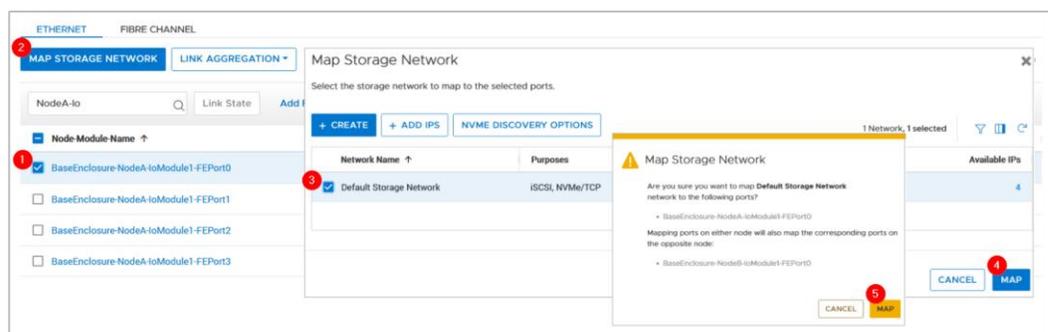


Figure 5. Map Storage Network

After ports are mapped to the storage network, select **MORE ACTIONS > Tag for Replication**, as shown in Figure 6. In the resulting window, click **TAG PORT** to finish the configuration. When it is set, the replication tag cannot be removed completely, but it is possible to reconfigure the replication tag for a different port or to a different storage

network. Similar to mapping, it is always a pair of ports that are tagged for replication— one port on Node A and a corresponding port on Node B.

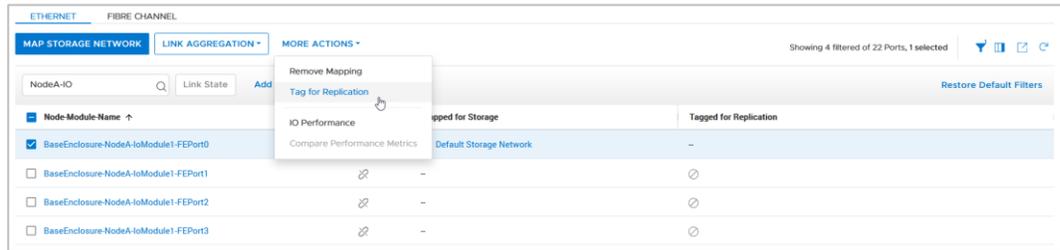


Figure 6. Tag port for replication

Individual networks for host and replication traffic

Starting with PowerStoreOS 2.0, multiple storage networks are supported. This feature allows users to separate host data from replication data either using same or different ports.

The following examples are using Default Storage Network and Replication Network (Figure 7) as already configured networks in PowerStore Manager (configured by selecting **Settings > Networking > Network IPs > Storage**).

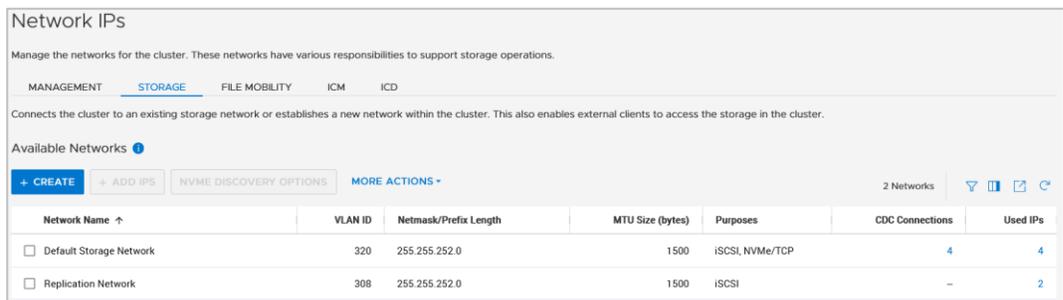


Figure 7. Multiple storage networks

Example 1: Two storage networks over a single port

When physical links for a storage network are not fully used by host data, it might be useful to set up a shared port for host data and replication data. To separate the traffic, it is required to set up VLANs on the switch ports. This example is using VLAN 320 for host access and VLAN 308 for replication traffic. The configured VLANs in PowerStore Manager must match the switch port configuration (VLAN tagging). As in previous sections, the port configuration in PowerStore Manager is available in the **Hardware > Appliance-Name > Ports** view. Figure 8 shows the current configuration where system bond is tagged for host I/O and replication using the mapped storage network **Default Storage Network**.

Node-Module-Name	Link State	Mapped for Storage	Tagged for Replication
BaseEnclosure-NodeA-1oModule1-FEPort0		✓ Default Storage Network	-
BaseEnclosure-NodeA-bond0		✓ Default Storage Network	✓ Default Storage Network
BaseEnclosure-NodeA-4PortCard-FEPort0		⊗	⊗
BaseEnclosure-NodeA-4PortCard-FEPort1		⊗	⊗
BaseEnclosure-NodeB-1oModule1-FEPort0		✓ Default Storage Network	-
BaseEnclosure-NodeB-bond0		✓ Default Storage Network	✓ Default Storage Network
BaseEnclosure-NodeB-4PortCard-FEPort0		⊗	⊗
BaseEnclosure-NodeB-4PortCard-FEPort1		⊗	⊗

Figure 8. Single network configuration

For replication tagging, configuring the additional storage network is required. **Replication Network** as the second Storage Network for the port pair was created in network settings in advance. Because the port configuration is the same on partner nodes, it is only required to select one single port for configuration and use the **MAP STORAGE NETWORK** button. In the selection window that is displayed, choose the **Replication Network** and continue with **MAP NETWORK** (Figure 9).

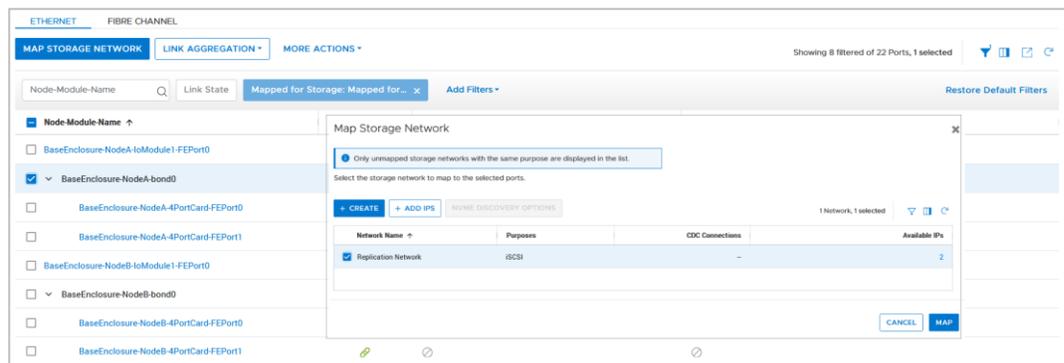


Figure 9. Map Storage Network

When the Map Storage Network dialog is confirmed, the port overview column **Mapped for Storage** changes to **2**, which indicates that two storage networks are mapped and using this port. Hovering over the number shows the IP address information for the ports (Figure 10).

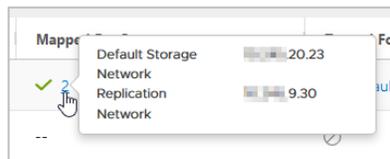


Figure 10. Detailed port information

The mapped network can be tagged as a replication network as in a single network configuration. Now you can choose the network used to tag the selected port, as shown in Figure 11.

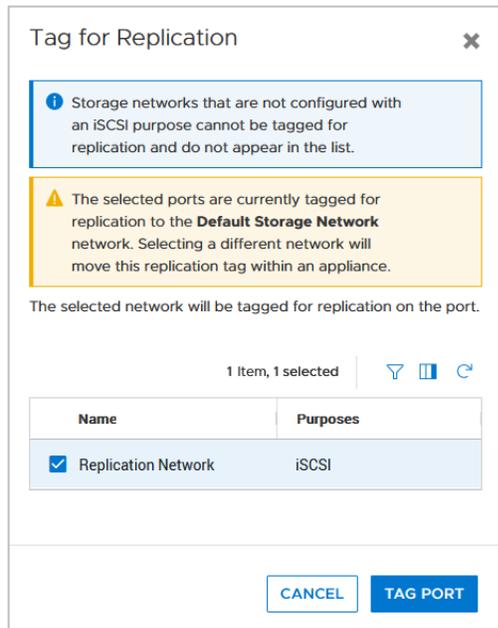


Figure 11. Tag for Replication – Network selection

After the dialog to perform configuration on both nodes is confirmed, the tagged network for replication changes to the new **Replication Network** (Figure 12).

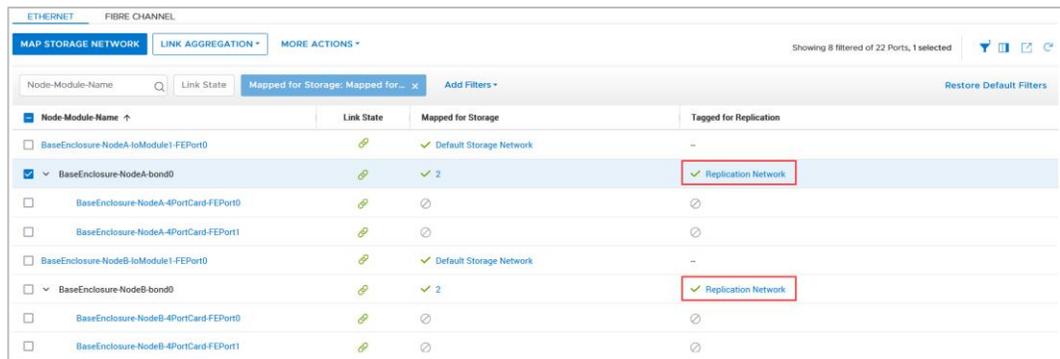


Figure 12. Single port configuration with dedicated network tagged for replication

Example 2: Separated host and replication networks

Note: This example uses the system bond as a replication port, which might not be the optimal configuration for all use cases.

In some use cases, it might be useful to separate the replication data network from production host traffic by using different physical ports and networks. The configuration is similar to Example 1. The difference is to map the **Replication Network** to another port than the **Default Storage Network**, which is used for host traffic. The example in Figure 13 shows a selected port with a dialog box to select the storage network for mapping.

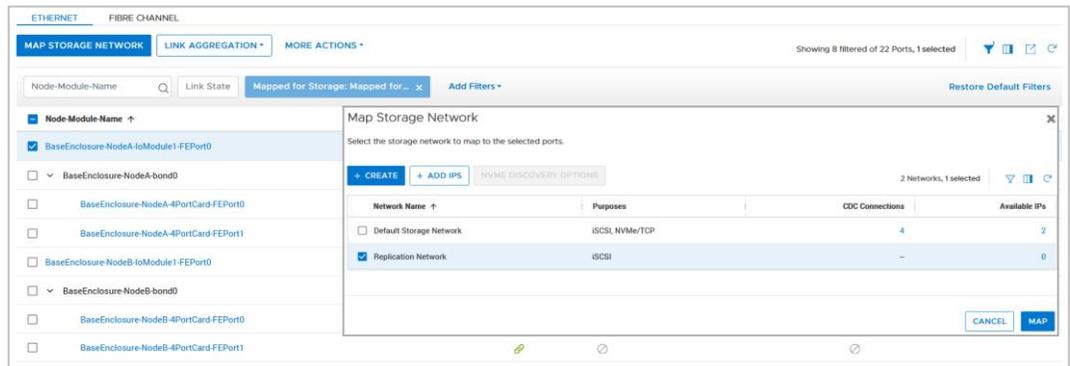


Figure 13. Map replication network to a new network port

After the configuration is finished, it is possible to tag the new storage network port for replication. Because only one storage network is configured in our example, there is no additional dialog to select the network. The port is tagged with **Replication Network** after the configuration is confirmed. [Figure 14](#) shows the configuration for that example.

Node-Module-Name	Link State	Mapped for Storage	Tagged for Replication
<input checked="" type="checkbox"/> BaseEnclosure-NodeA-IoModule1-FEPort0		<input checked="" type="checkbox"/> Replication Network	<input checked="" type="checkbox"/> Replication Network
<input type="checkbox"/> BaseEnclosure-NodeA-bond0		<input checked="" type="checkbox"/> Default Storage Network	-
<input type="checkbox"/> BaseEnclosure-NodeA-4PortCard-FEPort0		<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> BaseEnclosure-NodeA-4PortCard-FEPort1		<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> BaseEnclosure-NodeB-IoModule1-FEPort0		<input checked="" type="checkbox"/> Replication Network	<input checked="" type="checkbox"/> Replication Network
<input type="checkbox"/> BaseEnclosure-NodeB-bond0		<input checked="" type="checkbox"/> Default Storage Network	-
<input type="checkbox"/> BaseEnclosure-NodeB-4PortCard-FEPort0		<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> BaseEnclosure-NodeB-4PortCard-FEPort1		<input type="checkbox"/>	<input type="checkbox"/>

Figure 14. Port configuration with dedicated host and replication storage network

Replication data network configuration (PowerStoreOS 4.0 and later)

This section addresses the new replication data network confirmation in PowerStoreOS 4.0 and later. A replication data network relies on a TCP connection over one or multiple networks for the replication, consisting of a network port and an IP configuration. An IP storage network leveraged for replication requires setting its purpose to *Replication*. The purpose can be set when a new storage network is created or when modifying an existing storage network in **Settings > Networking > Network IPs > Storage Networks** ([Figure 15](#)).

Create Storage Network

Network Details	Network Details
Mapping for Appliance-PS1200T-4618	Enter information about this network.
Network Addressing	Name <input type="text" value="3000 Replication"/>
Summary	Purposes <input type="checkbox"/> Storage (iSCSI) <input type="checkbox"/> Storage (NVMe/TCP) ⓘ <input checked="" type="checkbox"/> Replication

Figure 15. Create Storage Network with replication purpose

Creating the Storage Network in the next step allows the user to map the network to a port. A previous set purpose is adopted to the port during the mapping process and is available in the column **Port Assigned Purposes** (Figure 16).

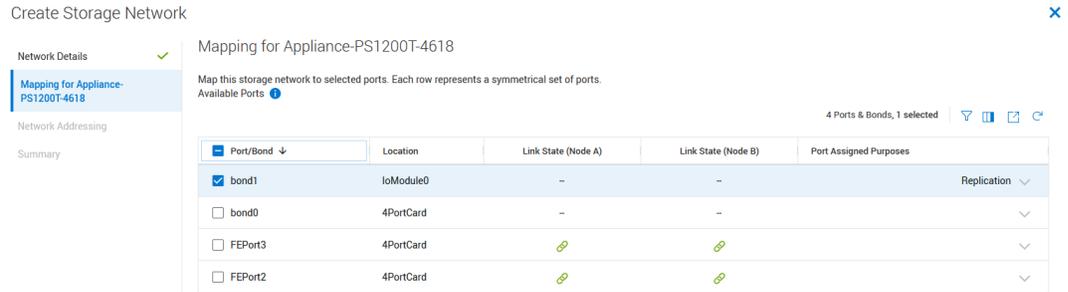


Figure 16. Replication network to port mapping

After the configuration wizard finishes, the storage network overview shows the network name, VLAN ID, and the configured purpose for the individual storage networks in the column **Purposes** (Figure 17).

Network Name	VLAN ID	Purposes
<input type="checkbox"/> 142 Hybrid	-	Storage (NVMe/TCP), Storage (iSCSI), Replication
<input type="checkbox"/> 3000 Replikation	3000	Replication
<input type="checkbox"/> 3001 Storage	3001	Storage (NVMe/TCP), Storage (iSCSI)

Figure 17. Storage Networks with different purposes

The current usages of a port are determined by the storage network(s) to which it is mapped. This information can be found in the PowerStore Manager ports view **Hardware > [Appliance Name] > Ports** (Figure 18).

Node-Module-Name	Link State	Networks Mapped	Current Usages
> BaseEnclosure-NodeA-bond1	-	2	iSCSI, Replication
> BaseEnclosure-NodeA-bond0	-	142 Hybrid	iSCSI, Replication, File Data, ICM, ICD

Figure 18. Ports view with usages

The **Current Usages** column in the **Ports** view shows the usages of the individual ports. For instance, port BaseEnclosure-Node-A-bond0 has 142 Hybrid network mapped. The port is used for iSCSI, Replication, File Data, ICM, and ICD. The number “2” in the **Networks Mapped** column indicates two mapped networks for port BaseEnclosure-Node-A-bond1 and the link behind the number leads to a slide out that shows all configured IP networks for that port (Figure 19).

Network Name	Network Purposes	IP Address Purposes
3000 Replikation	Replication	External Replication iSCSI, Replication
3001 Storage	Storage (NVMe/TCP), Storage (iSCSI)	iSCSI Target

Figure 19. Storage networks mapped to a port

All network port configurations for a PowerStore appliance are symmetric for both nodes. When changing the configuration on one node port of the appliance, it immediately affects the same port on the other node. Only ports with Replication usage can be used for a replication network group in a remote system pair configuration.

Remote systems

The next step in configuring remote replication is to create a remote systems pair with another system. Configuring a remote system pair is only required once for both systems. After configuration, the remote system pair is available on both participating PowerStore clusters. For replication management, a private replication connection using the management ports is set up. To set up a replication connection, select **Protection > Remote Systems** to start adding remote systems (see [Figure 20](#)).

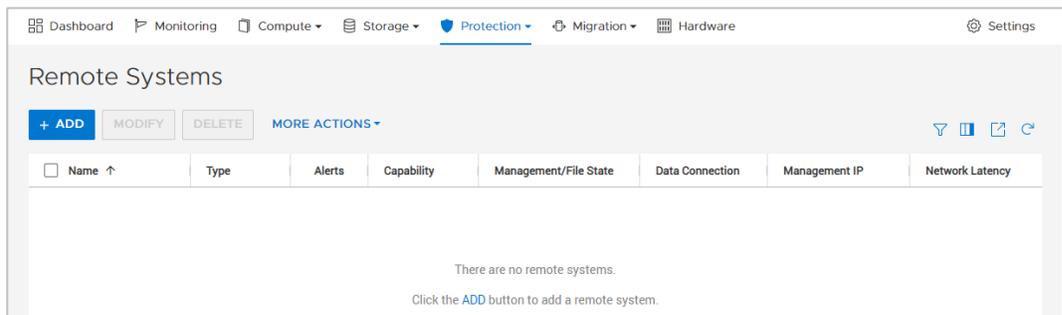


Figure 20. Remote systems replication setup

To define a new remote system, click **ADD** as shown in [Figure 20](#). The **Add Remote System** window that is displayed ([Figure 21](#)) requires the following information:

- Management Cluster IP Address
- Username and Password of the remote system
- Network Latency setting

Figure 21. Add Remote System

Replication traffic can be tuned for higher efficiency depending on the expected network latency. When network latency between the remote systems is unknown, use the ping utility to determine the latency. For PowerStoreOS releases 1.x and 2.x, use **Low** when the expected latency is less than 5 milliseconds, otherwise use **High**. PowerStoreOS 3.0 and later allows a more granular setting of network latency, as shown in [Figure 21](#). Metro and synchronous replication require a low latency network (<5ms).

The provided credentials for a configured user are not stored on the system and are only used for the relationship setup. After the relationship is set up, PowerStore uses SSL certificate-based authentication. When the required fields are entered, click **ADD**. Because the management connection for the remote systems pair uses SSL encryption, it is required to confirm the remote SSL certificate. After the configuration task is finished, the new remote system is listed on both sides. If using bi-directional replication, the same remote systems pair can be used for replication sessions from the opposite systems.

After a remote system is set up, select **MORE ACTIONS > Verify and Update**. This action verifies that the selected replication connection still exists with the remote system, and it updates the connection details if any changes were made. [Figure 22](#) shows the Remote Systems Overview. The **Capability** column indicates the supported types of replications for the remote system pair. The **Management/File State** and **Data Connection** columns indicate the link status.

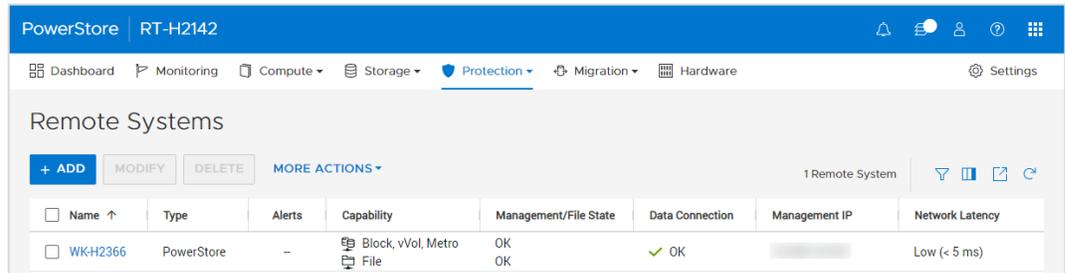


Figure 22. Remote Systems Overview

Replication data network

For PowerStoreOS releases 1.x and 2.x, the iSCSI protocol is used for replication data traffic. PowerStoreOS 3.0 and later leverages a proprietary TCP-based protocol for replication data traffic which is mandatory for a metro configuration. The TCP-based replication protocol improves replication performance between systems with network impairment, such as high latency or packet loss. Replication between earlier releases and PowerStoreOS 3.x is supported and relies on the iSCSI protocol. Starting with PowerStoreOS 4.0 TCP-based replication is the only supported protocol for replication. An NDU to 4.0 or later is blocked when a remote system pair is still configured with an iSCSI based replication protocol. PowerStore Manager will show an alert that the iSCSI based protocol is used, and provides a link for the update to the TCP based replication protocol. Each latency category uses a different network port number. For replication across network borders, adjusting network ACL or network firewall rules might be required to allow replication traffic. [Table 3](#) shows an overview of different network latency settings and the network port that is used on PowerStore.

Table 3. Remote systems network latency overview

PowerStoreOS versions	Network latency between remote systems	Port #
PowerStoreOS 1.x, 2.x	Low (default) < 5 milliseconds	3260
	High >= 5 milliseconds	3261
PowerStoreOS 3.0 and later	Low (default) < 5 milliseconds	13333
	Low Medium >= 5 and < 20 milliseconds	13334
	Medium >= 20 and < 60 milliseconds	13335
	Medium High >= 60 and < 120 milliseconds	13336
	High >= 120 milliseconds	13337

PowerStoreOS 4.0 and later is no longer limited to the single storage network which is tagged as the replication network. One or more replication networks can be configured in a network replication group, on a remote system pair granularity. The replication network pair is automatically established by the remote system pairing when both PowerStore systems use the same IP network port for replication. When more network ports do have replication usage, a manual configuration is required. After an NDU from an earlier release, a post upgrade service starts the migration of the replication purpose tagged storage network towards the new concept of network group for the replication. After

selecting a remote system pair (Figure 23), the replication network configuration is in the **Protection > Remote Systems** overview, in the **More Actions** pull down.

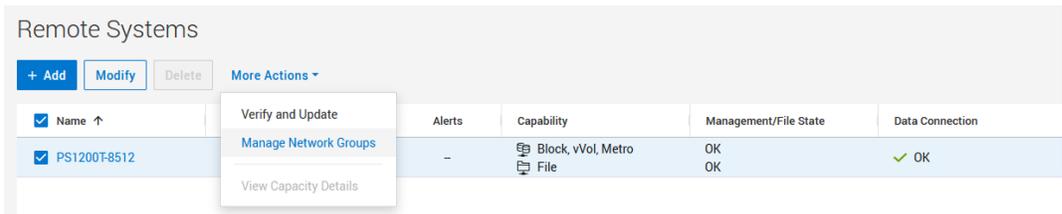


Figure 23. Remote Systems – Manage Network Groups

For a remote system pair there must be at least one network group configured. The slide out to manage network groups allows you to modify existing network groups, as well as to create and delete a network group (Figure 24).

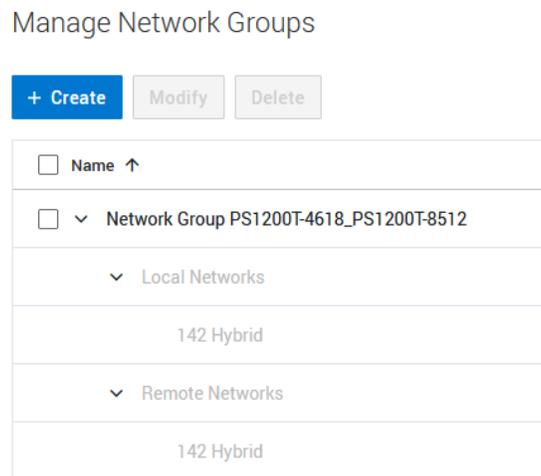


Figure 24. Manage Network Groups

When you create a new replication network group, the wizard shows the available local and remote networks you can select for replication (Figure 25). Remote networks are fetched during the remote system pairing and are synchronized when verifying and updating a remote system pair configuration (**More Actions – Verify and Update**).

Create Network Group



Name

3000 Replication Network Group

Local Networks

1 Local Networks, 1 selected



<input checked="" type="checkbox"/> Name ↑	VLAN ID
<input checked="" type="checkbox"/> 3000 Replication	3000

Remote Networks

1 Remote Networks, 1 selected



<input checked="" type="checkbox"/> Name ↑	VLAN ID	Netmask/Prefix Length	Gateway
<input checked="" type="checkbox"/> 3000 Replication	3000	24	–

Figure 25. Create replication network group wizard

PowerStore has no limitation on which IP ports and networks are used for the replication network groups, as long as Layer 2 (switched) or Layer 3 (routed) communication is possible between systems in a remote system pair configuration. The networks shown in the dialog only rely on the configured purpose and resulting replication port usage. That flexibility allows for multiple different replication networks for the remote system pair and enables additional advanced configuration to control the network used for replication traffic. Bear in mind that replication networks' performance for a metro or synchronous replication configuration have a direct influence on host IO performance.

Use case examples:

- A remote system pair can leverage multiple groups of different IP networks on a potentially different physical network infrastructure for a dual fabric architecture used for replication.
- Even though you may plan to use the same IP networks to host front end access and replication, it is possible to set up distinct physical ports for the individual workloads.
- Different remote system pairs can use individual non-overlapping ports and networks for replication.

File mobility network

File replication requires an additional file mobility network configuration for control file related traffic between the clusters. The file mobility network resides in the same subnet as the PowerStore cluster management network and requires three additional IP addresses in that range for each PowerStore cluster. While PowerStoreOS 3.0 does not support any changes when file replication is configured, PowerStoreOS 3.2 supports deleting and changing the file mobility network in a paused state without needing to delete the existing replication sessions. Even though deleting the file mobility network is supported, it is required when replication sessions are activated again.

The configuration for the file mobility network (Figure 26) can be found on the **File Mobility** tab at **Settings > Networking > Network IP**. When the initial network configuration is finished, map the file mobility network to the PowerStore management ports of the appliance. For the **Reconfigure** or **Delete** tasks (PowerStore 3.2 and later), a dialog is displayed to confirm that no active file migrations or replication sessions are in place.

Create File Mobility Network
✕

File mobility network is a pre-requisite infrastructure needed for exchange of control traffic between file clusters in replication and file import environments.

i The file mobility network can only communicate through the management network. All the file mobility network settings must be on the same subnet as the management network.

Use VLAN tagging

Netmask/Prefix Length

255.255.252.0

Gateway (Optional)

.1

Cluster IP

File Mobility Network IPs **i**

Enter IP address, range, comma separated ADD

Address

The list is empty.

0 IPs provided, 2 IPs required (minimum)

Network MTU Size **i**

1500

CANCEL
CREATE

Figure 26. Create file mobility network

Native replication for block and file

Introduction

This section describes the PowerStore native asynchronous (PowerStoreOS 3.0 and later) and synchronous replication feature (PowerStoreOS 4.0 and later) which allows users to create replication sessions for block and file storage resources between PowerStore systems. Synchronous replication for Metro is covered in the white paper [Dell PowerStore: Metro Volume](#). Supported storage resources for native replication are volumes, volume groups, thin clones, and NAS servers which includes the underlying file systems. The replication itself uses iSCSI or the optimized Dell proprietary TCP-based replication protocol (PowerStoreOS 3.0 and later) through Ethernet (LAN) connections. All configuration and management operations in this section are shown in PowerStore Manager, but you can also use the PowerStore CLI and REST API. The following subsections describe these topics:

- Licensing requirements for the native replication features
- How the native asynchronous replication feature works
- Enhancements to asynchronous replication to get synchronous replication
- Supported configurations
- PowerStore Manager configuration and management

Licensing

Supported replication technologies depend on the running version of PowerStoreOS. When supported, the native asynchronous and synchronous replication features are included in the base license at no additional cost.

Theory of operation

Protection policies with replication rules

Remote replication between PowerStore systems relies on a remote system configuration and uses policy-based protection. Protection policies allow the user to configure remote and local protection using replication rules, snapshot rules, or remote backup rules. The policies combine one or more rules to fulfill the protection requirements for a storage resource on PowerStore. A protection policy must contain at least one protection rule, regardless of whether it is a local or remote protection rule. Each protection policy can contain up to one replication rule, up to four snapshot rules and up to one remote backup rule. A remote backup rule can only be combined with snapshot rules and an asynchronous replication rule.

Table 4. Maximum rules per policy

Volume replication type	Snapshots	Asynchronous replication	Synchronous replication	Remote Backup
Asynchronous replication	up to 4	1	-	up to 1
Synchronous replication	up to 4	-	1	-
Metro volumes	up to 4	-	-	-

The replication rule defines the parameter for the replication on PowerStore and is set up on the source array. Even when the rule is synchronized to remote systems when it is added to a protection policy, it is not possible to edit a replication rule on the remote system. It is also not possible to view it in the replication rule overview in PowerStore Manager. The required information for creating a rule includes the partner remote system, replication type with RPO, and alert threshold for the planned replication session. A zero “0” RPO also indicates the native synchronous replication in PowerStore Manager. When a protection policy with a replication rule is assigned to a storage resource, the configured RPO in the rule is used to set up the replication. After initial synchronization, the replication session sets up a mirror from source to destination for synchronous replication. For asynchronous replication, the replication session sets up an internal event scheduler for the recurring replication of the storage resource, based on the RPO setting as defined in the replication rule.

Asynchronous replication

To minimize the chance of RPO compliance issues in an asynchronous replication configuration, replication cycles are scheduled at 50% of the RPO value. For example, a one-hour RPO leads to a replication event every 30 minutes to provide enough overlapping to meet the target of a one-hour RPO. The scheduled RPO events for the following example are at x:00 and x:30 every hour. PowerStore optimizes the replication schedules to serialize the individual synchronization events. The events for the RPO are based on the configured RPO time and not on the amount of data that is written on the source storage resource.

Each storage resource can only have one active replication synchronization at a time. For example, the event scheduler cannot initiate a replication at a given time because replication is paused, or a previous replication has not finished. In this case, the schedule is skipped, and replication proceeds with the next planned replication.

The alert threshold defines the time when an alert is triggered after a target RPO is missed during continuous replication. There is no event that is triggered when the initial replication needs more time to complete. When a compliance alert is raised for replications using an RPO of more than five minutes, it is cleared when the next replication cycle finishes successfully.

The following steps and [Figure 27](#) illustrate the workflow of a storage resource that is scheduled with a target RPO of one hour and an alert threshold of zero minutes:

1. The initial synchronization is completed successfully within the 30-minute RPO window, and the first scheduled RPO synchronization cycle begins.
2. An RPO snapshot for the second regular replication cycle is performed at 11:00, and the synchronization finishes successfully within 30 minutes. The replicated snapshot meets the RPO target by 12:00.
3. The next replication cycle at 11:30 is not finished replicating to the target after reaching the 12:00 limit for the 11:00 RPO schedule (step 2). The 12:00 scheduler event is skipped, and an RPO compliance alert is raised at 12:05.
4. When the 11:30 replication finishes at 12:10, the RPO target is achieved again until 13:10. Since the raised alert needs at least one successful replication within the timeframe to be cleared, the alert remains active until it is cleared. In our example, the alert is cleared after the 12:30 replication finishes at 13:00.

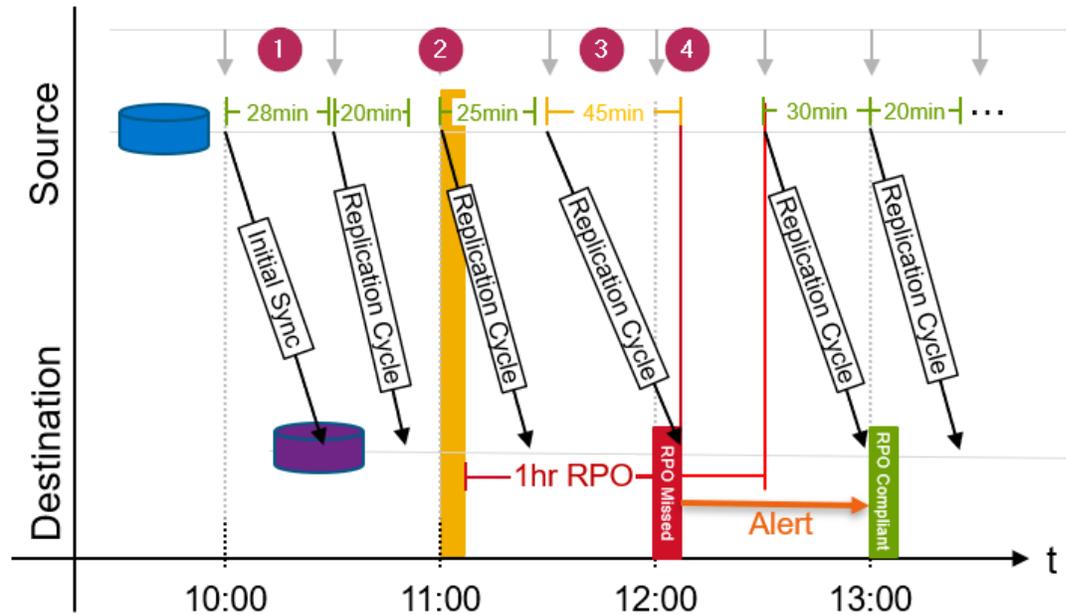


Figure 27. Replication scheduler events at 30-minute intervals for a 1-hour RPO

Replication session

Assigning a protection policy with replication rule to a storage resource creates the replication session. The replication session operates the scheduling and replication from the source resources to the target storage resources. When a replication session is created in PowerStore, a storage resource of the same size and type is created on the destination system. PowerStore creates an individual RPO schedule or a mirror for each storage resource in that replication session. When the replication session is set up, existing scheduled or manual user snapshots of block storage resources are also replicated in chronological order to the destination during initial synchronizations. Snapshots created while an asynchronous replication is set up are replicated in the next RPO cycle. Snapshots of file storage sources in a replication session are not replicated to the destination.

Asynchronous replication synchronizations are triggered by a user-defined RPO or at any time manually by the user. The following characteristics define asynchronous replication:

- All 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. Changes are retrieved using a snapshot-differential operation and are replicated later.
- A user-defined RPO defines the maximum time between scheduled synchronizations.
- Between synchronizations, new data is only saved on the source storage resource. The RPO is the maximum amount of data measured in time that the user is willing to lose in a disaster or failure scenario. The RPO determines how often synchronizations occur at a minimum.
- Manual replication between RPOs operates the same as scheduled asynchronous replication.

When a replication session is created, and before the mirror or incremental cycles begin, a full synchronization of the source and destination storage resource is automatically initiated. If replication is configured when a new storage resource is created, the synchronization is quick because no user data needs to be copied to the destination storage resource. If a protection policy with replication is added to an existing storage resource, a full synchronization is initiated from the source to the destination storage resource. Writes occurring during the initial-synchronization period are not copied to the destination storage resource but remain in the snapshot differential for the next synchronization cycle.

RPO based asynchronous replication

When the initial synchronization is completed, a common base is established between the source storage resource and the destination. Host-write operations that occur after the initial synchronization are acknowledged with the host, and no data is replicated to the destination until the next synchronization cycle. On any recurring cycle, a new snapshot is created and all changes between the current and previous snapshots are replicated to the destination. A new common base is then established. If another replication is still running, either manually triggered or by the RPO event scheduler, the replication is skipped.

Asynchronous replication in PowerStore uses snapshots to maintain the common base images explained previously. The following steps and [Figure 28](#) show how snapshots are used with asynchronous and manually triggered replication.

1. When a replication session is created on a storage resource, a read-only internal RPO1 snapshot on the source system is created. On the destination system, a storage resource with same characteristics is created with an associated shadow read/write snapshot.
2. Data is replicated from the source RPO1 snapshot to the newly created destination shadow read/write snapshot. This replication is the initial synchronization of the source and destination storage resources and is a full copy of all the data.
3. When the remote read/write shadow snapshot is synchronized with the local RPO1 snapshot, an RPO1 snapshot is triggered on the destination. The RPO1 snapshots that are on the source and destination storage resources contain the same information and represent the point when the synchronization started. Snapshot RPO1 on each system is now a common base for the replication session. The remote storage resource is refreshed from the RPO1 snapshot, and the initial synchronization is completed.
4. Over time, the host application writes new data to the source storage resource.
5. The next update is either manually started or by the RPO with asynchronous replication. During the update, a new RPO2 snapshot is triggered to reflect the current, point-in-time view of the source storage resource. All changes that were made since the last update of the destination are copied to the destination shadow read/write snapshot.
6. After the incremental copy is complete, an RPO2 snapshot on the destination is created. This snapshot defines the new common base, and the remote storage resource is refreshed from that base.

7. Because the old, common-base compound of RPO1 snapshots on the source and destination are not relevant for upcoming replication cycles, the RPO1 snapshots are deleted. Only the RPO2 snapshots and shadow read/write snapshots remain.

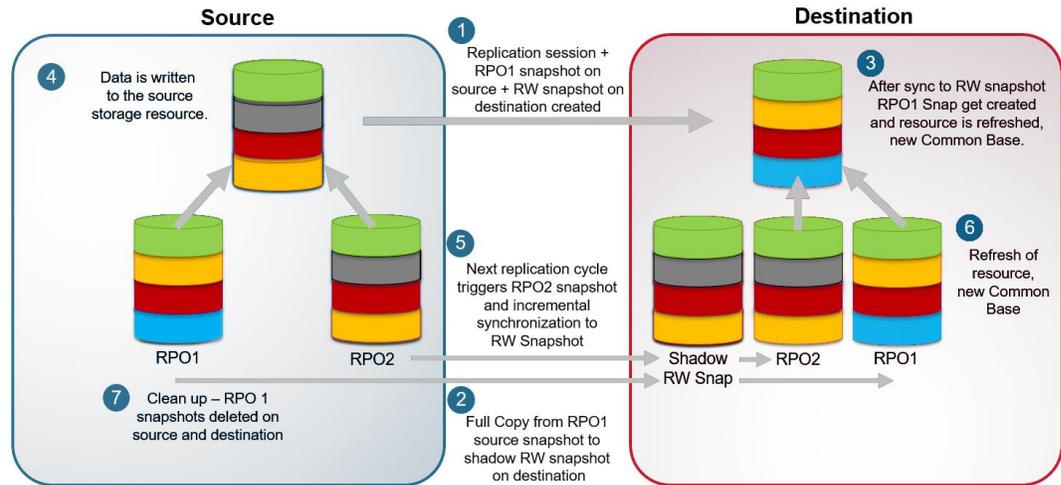


Figure 28. Asynchronous replication theory

Each time the replication interval (half of the RPO setting) is reached or a manual update is started, the common base image updates with the latest RPO snapshots.

Snapshots that are used for asynchronous replication operate the same as user snapshots and are based on redirect-on-write technology. Although user snapshots and replication snapshots share the same technology, replication snapshots have use restrictions. Although replication snapshots can be viewed in the PowerStore REST API and PowerStore CLI, user operations such as restore operations are not allowed. Snapshots that are allocated for replication purposes do not count toward user-snapshot maximums.

Synchronous replication (RPO=0)

For synchronous replication, the initial replication is similar up to the point where the common base is created. While asynchronous replication continues its replication based on the set RPO, additional steps are required to have the source and destination be in a continuous synchronized state.

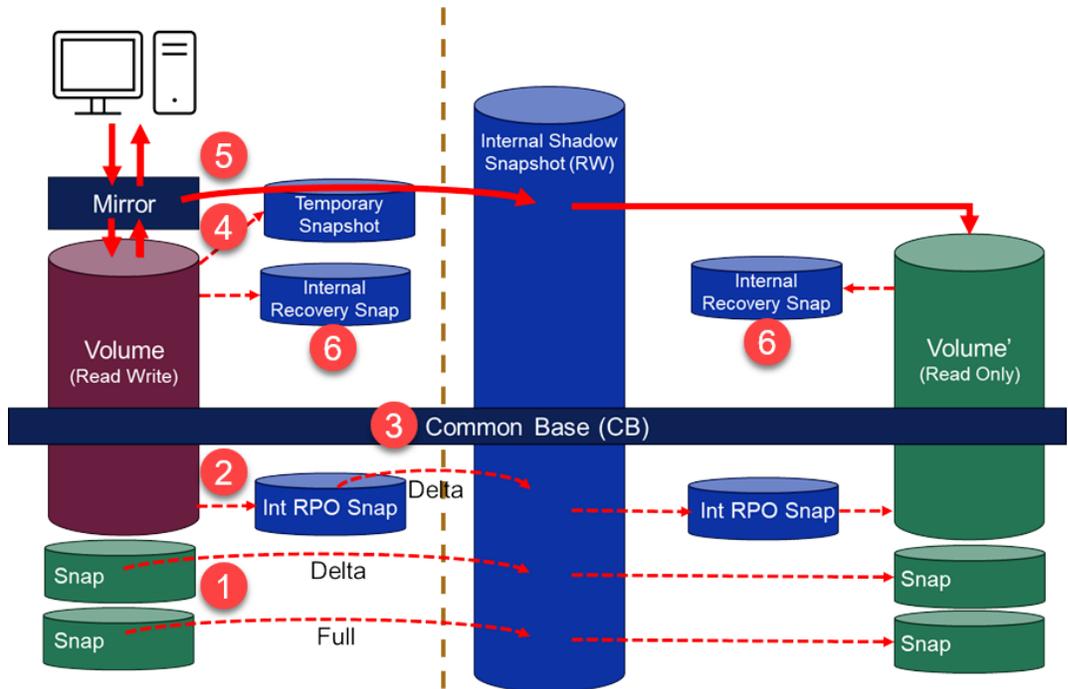


Figure 29. Synchronous replication in PowerStore 4.0 and later

After the empty destination volume and an internal writeable Shadow Snapshot is set up on the destination, the internal Shadow Snapshot gets all replicated data before being pushed to the destination volume. When destination and Shadow Snapshot are set up, the initial synchronization starts:

1. Only when a user snapshot exists on the source, all existing user snapshots are replicated in chronological order to an internal snapshot on the destination and applied to the destination volume. Only the first snapshot requires a full synchronization. For each possible subsequent snapshot, only the delta from the previously replicated snapshot is being replicated.
2. After all user snapshots are replicated, and when no snapshot of a previous replication pre-exists, an internal RPO snapshot is created on the source and replicated to the destination. Without a pre-existing snapshot, this replication of the internal RPO snapshot requires a full synchronization, otherwise only the delta from the previous snapshot is replicated.
3. When the RPO snapshot is fully replicated, a common base (CB) is created.
4. When the CB is created, a mirror starts mirroring all host writes for the volume to the internal shadow snapshot on the destination. All outstanding writes since the CB was created are kept in an internal temporary snapshot. The data from that temporary snapshot is synchronized along with the mirrored host IO to the internal shadow snapshot on the destination. When all outstanding data is synchronized, the temporary snapshot is deleted.
5. For the ongoing synchronous replication, the mirror continues to replicate all host writes immediately to the internal Shadow Snapshot on the destination, which updates the destination volume.

6. When the replication session is in an “Operating Normally” state, a recovery snapshot is created every 30 minutes on both sides.

Because asynchronous replication and synchronous replication for block are using a common base, it is possible to change between replication types without a new full synchronization. Changing replication type for a NAS server (File replication) is not supported and requires a full synchronization after the change.

In PowerStore, native synchronous replication is supported on the following storage resources:

- Volumes
- Thin clones
- Volume groups
- NAS servers with underlying file systems

The replication operates in the same way for volumes, volume groups, thin clones, and file resources on PowerStore. When replication is configured on a volume in PowerStore Manager, a single replication session is created, and the destination storage resource is created with the same size and type as the source storage resource. When configuring a replication session on a thin clone, the destination storage resource is a regular volume and not a thin clone. While asynchronous replication is configured, the volumes and thin-clone size can be extended, and the changes are reflected on the destination storage resource after the next synchronization cycle. For synchronous replicated storage resources, it is required to pause the replication session first. All changes are replicated after the session resumes.

Volume Groups

On PowerStore, a volume group is a storage resource that contains one or more volumes within a storage system. Volume groups help organize storage resources allocated for a particular host, hosts, or host groups. Volume groups are treated as a single entity when they are replicated. When replication is configured in PowerStore Manager for a volume group, the destination storage resource and its contents are created automatically. While a volume group is part of an asynchronous replication session, volumes within the volume group can be expanded. All changes to volumes within a volume group are reflected on the destination image after the next completed synchronization. Synchronous replicated volume group attributes and volume group members can be changed when the replication session is paused. Changes will be synchronized after the session is resumed. When replication is paused or resumed on a volume group, the replication operation affects the entire group.

There are two general operation models when using Volume Groups:

- Simplify management when using volume groups, when mapping hosts, and when working with protection rules
- Ensure write order consistency across multiple volumes even when replication is configured. For instance, when an application and database are using different volumes but need consistency.

Note: Select the **write-order consistency** option for the volume group to have a consistent replica at the volume-group level.

The decision to enable write-order consistency is based on some considerations:

- Only asynchronous replication supports volume groups without write-order consistency enabled
- Individual volumes in a volume group without write-order consistency can get different protection policies applied.
- Individual volumes in a volume group with write-order consistency cannot get a policy with synchronous replication enabled.
- When write-order-consistency is enabled, snapshots can only be created and restored for the whole volume group, not an individual volume.

File replication

File system and NAS server replication sessions are created by assigning a protection policy with a replication rule to a NAS server. Once applied to a NAS server, the NAS server and all underlying file systems are replicated to the destination system. An individual replication session is created for each file system associated with the NAS server being replicated and for the NAS server itself. File replication can only be applied, managed, and removed at the NAS server level. It is not possible to modify the replication state at the individual file system level. Any file systems created or deleted from the NAS server automatically have a replication session created or deleted as applicable. While user operations and management for file replication are handled at the NAS server level, each file system has its own replication session. This is a key distinction between how a NAS server and file systems replicate compared to a volume group and its member volumes.

Creating a protection policy with replication rules

To create a replication session in PowerStore Manager, first set up a protection policy with an underlying replication rule. A protection policy is a collection of different local or remote protection rules that are assigned to resources on the PowerStore cluster. Protection policies can contain between zero and four rules for scheduled snapshots. The policies also contain a single replication rule for remote replication to a system that is defined in remote systems. Select **Protection > Protection Policies**, as shown in [Figure 30](#).

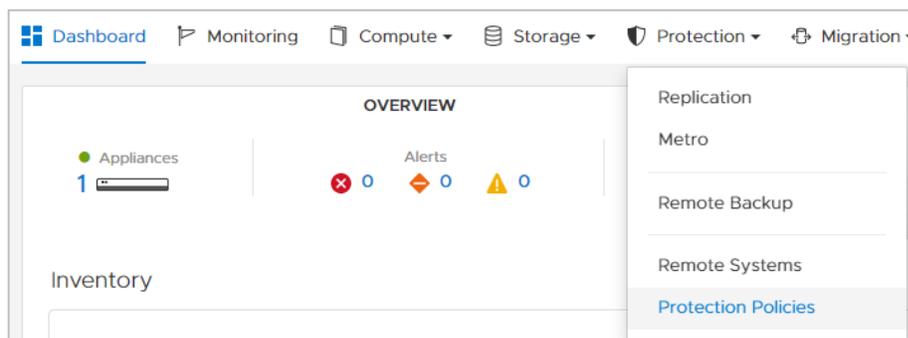


Figure 30. Protection Policies

In the **Protection Policies** window ([Figure 31](#)), you can create protection policies and rules and manage existing policies and rules. The following example shows how to create a protection policy with replication to a previously configured remote system. In the **Protection Policies** window, click **CREATE** to begin the configuration.

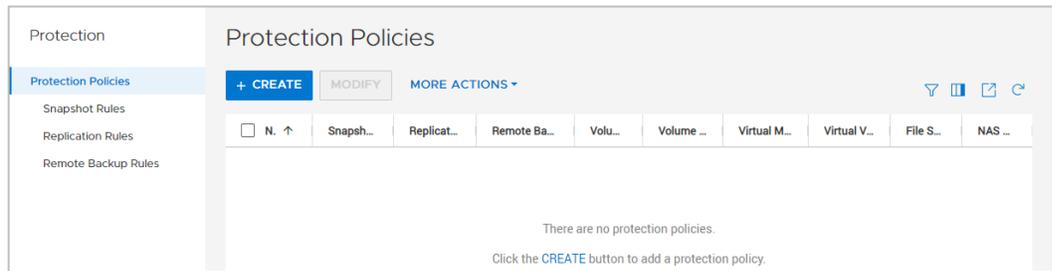


Figure 31. Protection Policies list

Because the Protection Policy is only the top-level object, which is assigned to a storage resource, only a policy **Name** is required. Use a meaningful name such as one that contains the remote system. Further down in the window, you can select an existing replication rule, or you can click **CREATE** to create a rule in the **Replication Rules** area.

The following information is required. Each step corresponds with a number that is shown in [Figure 32](#).

1. Enter a **Name** for the protection policy.
2. In the **Replication Rules** section, click **CREATE** to create a replication rule.

In the **Create Replication Rule** window, set the following:

3. Replication Rule Name
4. Destination Remote System
5. Replication Type – Asynchronous or Synchronous (PowerstoreOS 4.0 and later)
6. RPO – shows 0 (zero) for synchronous replication
7. Alert Threshold – Asynchronous replication only

Figure 32. Create replication rule

When all steps are finished, you can use the protection policy to protect storage resources with configured parameters.

Assign protection policy

The last step to establish a replication session is to assign the protection policy to a new or existing storage resource. This resource can include a volume, volume group, thin clone, or NAS server. A protection policy assigned directly to a file system will not implement any replication rules if they exist. To enact file system replication, the protection policy with the replication rule must be applied at the NAS server level. The following steps show assigning a protection policy on a volume. The required steps to assign a protection policy to a volume group, thin clone, or NAS server are the same. These steps can be applied either when creating or by modifying an existing storage resource.

Some limitations apply when creating the replication sessions. The replication session creates a storage resource with the same attributes on the destination as the source. Therefore, the name of the storage resource must not be used on destination. For example, it is not possible to create a replication session for a volume with the name **Volume** when a volume with the same name exists on the destination.

For volume groups configured **with write-order consistency**, all volumes inherit the protection policy as defined for the volume group. It is not possible to have individual policies set on different volumes. When a policy with replication is used for a volume group with write-order-consistency, only one replication session for the volume group is created. For volume groups **without write-order consistency**, members can have

different protection policies that result in individual replication sessions. Setting a protection policy for a whole group when write-order consistency (WOC) is configured, is only possible when no individual volume has a protection policy assigned. Because the replication configuration can differ between WOC and non-WOC volume groups, there are also restrictions on changing this volume group attribute if the group or any members are protected.

New storage resource with policy for replication

To create a replication session for a new storage resource, begin with the creation process. To begin the process for volumes, select **Storage > Volumes**, and click **CREATE** (see [Figure 33](#)).

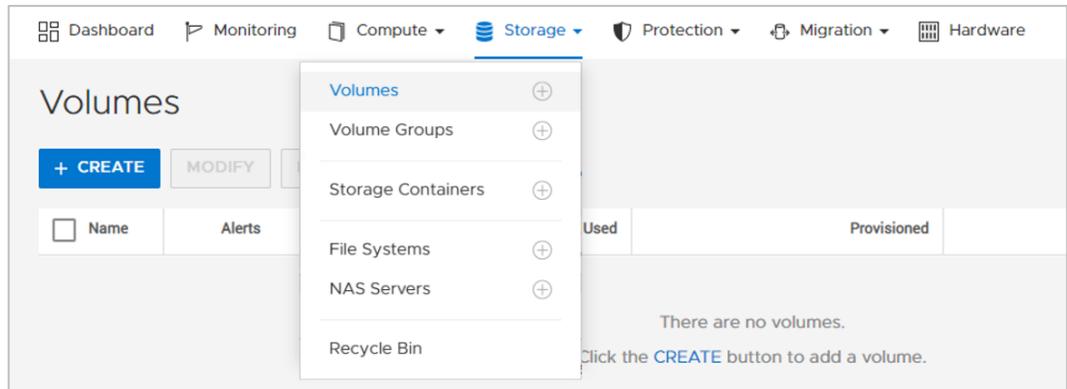


Figure 33. First step: Assign protection policy

In the **Create Volumes** window, enter the following information. Each number below corresponds with the number in [Figure 34](#).

1. **Name** for the new Volume
2. Application **Category**
3. **Application** name or type
4. **Quantity** of new volumes to be created
5. Volume **Size**
6. **Protection Policy**

To protect the new volumes with the protection policy, select the drop-down menu **Volume Protection Policy (Optional)**. This menu shows all local available protection policies.

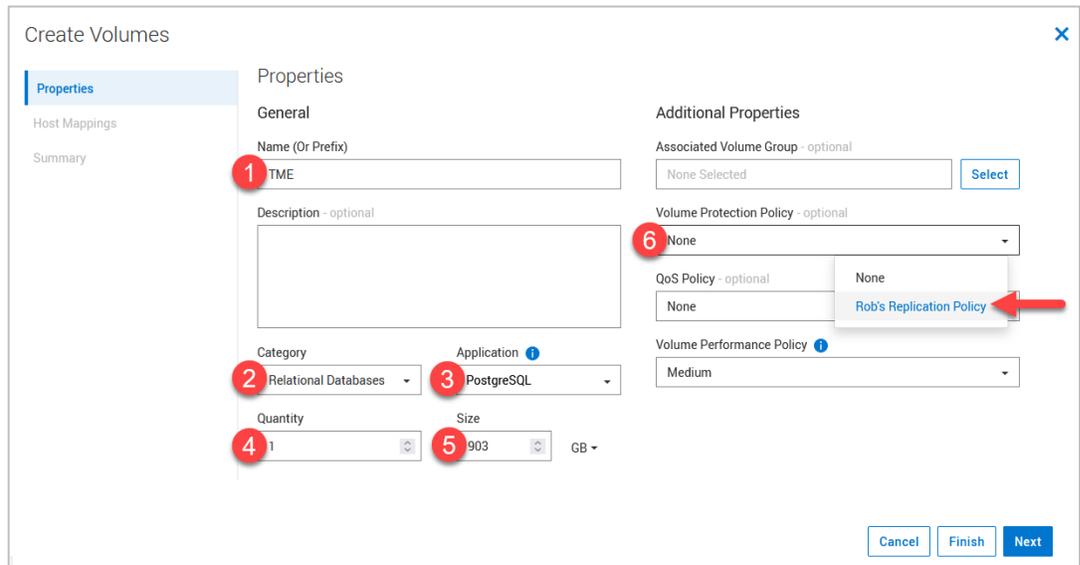


Figure 34. Second step: Assign protection policy

Complete the remaining steps to finish the configuration.

When a Volume Group with a protection policy and underlying replication rule is created, empty Volume Groups are created on the source and destination system before members are added. The members do not replicate until the next manual or RPO scheduled synchronization.

Protect existing storage resources

The following steps show how to assign a protection policy to an existing volume. The steps are similar for existing volume groups, thin clones, and NAS servers.

1. Open the **Storage Resource** page where the volumes or volume groups are listed, and select one or more resources to which to assign the protection policy.
2. Select **PROTECT > Assign Protection Policy**, as shown in [Figure 35](#).

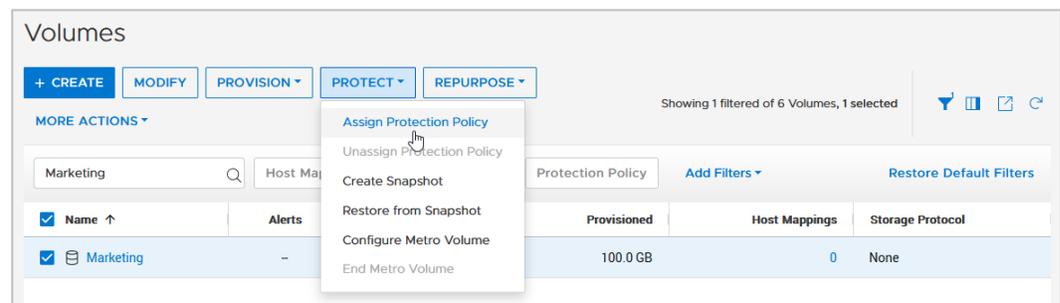


Figure 35. Assign protection policy to existing storage resource

3. Select the appropriate protection policy and click **Apply** to apply it to the previously selected storage resources.

After the policy is applied, the initial synchronization starts immediately.

Viewing the replication sessions

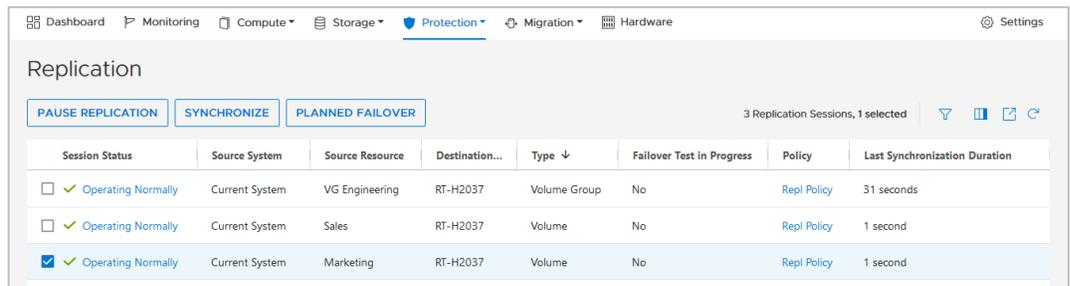
All replication sessions on the system can be viewed from the **Replication** page. To view this page in PowerStore Manager, select **Protection > Replication**. Figure 36 shows an example of the replication sessions overview with multiple replication sessions that are created on the system. This example shows the replication sessions for volumes and volume groups. A replicated thin clone is displayed in the same way as a volume. This page shows the information regarding each session and includes the following details:

- Replication Session Status.
- **Source System** including the source system and the source storage resource.
- Destination **System** including the destination system name and the destination storage resource.
- Resource **Type**.
- Protection **Policy**.
- **ETA** (estimated time) when the current synchronization will be finished. The ETA displays “-” if an active sync is not occurring.

Only one session can be selected at a time. The state of the selected session determines which buttons above the table are available. When no session is selected, the buttons are unavailable. Figure 36 shows the Replication page on the source system.

The Replication page for the source resource shows the following buttons:

- **PAUSE REPLICATION** to pause the replication
- **SYNCHRONIZE** to initiate a manual replication between regular RPO cycles
- **PLANNED FAILOVER** to manually initiate a failover during a planned maintenance window



Session Status	Source System	Source Resource	Destination...	Type ↓	Failover Test in Progress	Policy	Last Synchronization Duration
<input type="checkbox"/> ✓ Operating Normally	Current System	VG Engineering	RT-H2037	Volume Group	No	Repl Policy	31 seconds
<input type="checkbox"/> ✓ Operating Normally	Current System	Sales	RT-H2037	Volume	No	Repl Policy	1 second
<input checked="" type="checkbox"/> ✓ Operating Normally	Current System	Marketing	RT-H2037	Volume	No	Repl Policy	1 second

Figure 36. Replication window for block source resource

For file replication, the NAS server replication session is shown at the top level. You can expand the session by clicking the down arrow to the left of the session. Once the session is expanded, all underlying file system replication sessions for that NAS server are shown (Figure 37). The replication sessions for the file systems are disabled because file system replication sessions do not support individual management. All operations are performed at the NAS server level and are applied to every underlying file system replication session.

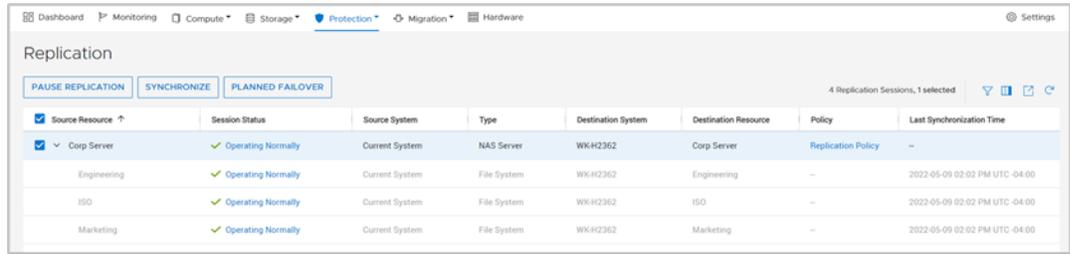


Figure 37. Replication window for NAS server and file systems

The Replication window for the destination resource shows different active buttons when a session is selected (Figure 38):

- **PAUSE** to pause the replication.
- **FAILOVER** to start an unplanned failover.
- **FAILOVER TEST** to initiate a failover test for block storage resources. The operation **FAILOVER TEST** is not supported for file resources and will not be shown when a NAS server is selected on the destination system.

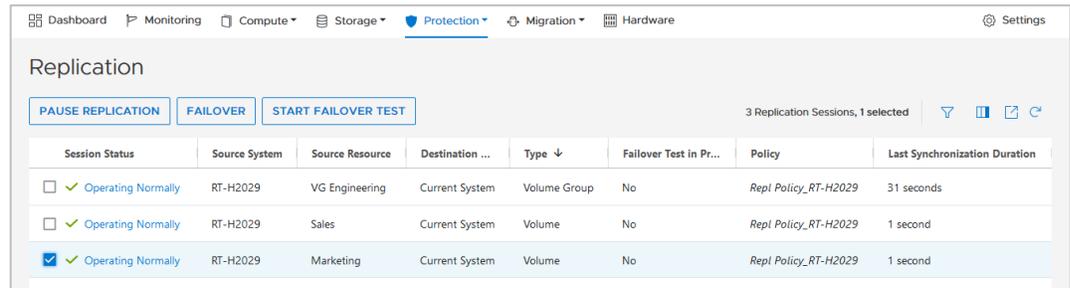


Figure 38. Replication window for destination resource

When you click **Failover**, a message is displayed to warn you that there is no final synchronization before the failover occurs. A planned failover must be run on the source if a final synchronization is needed.

For a more detailed view of the replication state, you can use the individual session states to view the selected replication session. This window displays a **Session Summary**, as shown in Figure 39. The local storage resource is always tagged with **Current System**.

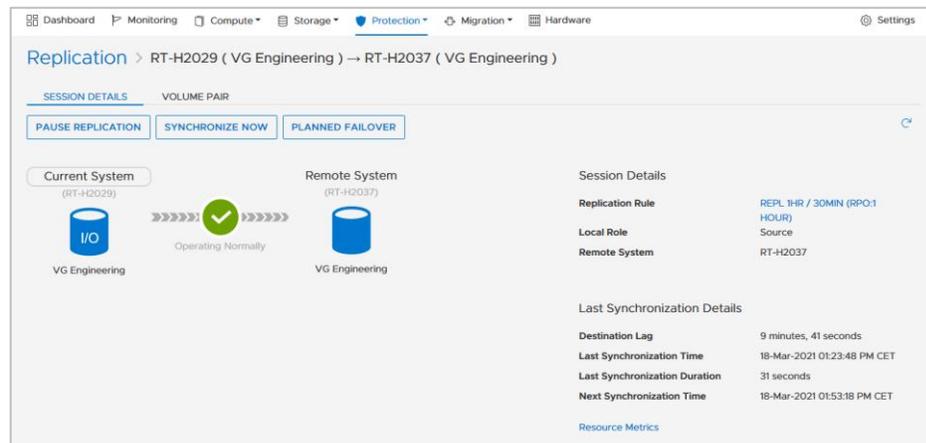


Figure 39. Replication Session Summary

When working in the volume or volume group properties pages, it is also possible to see and control the corresponding replication session. [Figure 40](#) shows the **Volume Group** replication page, which looks the same for a volume or a thin clone. To see the replication info, in the **Storage Resource** view, select the **Protection** tab, and then, as shown in [Figure 40](#), select the **Replication** tab.

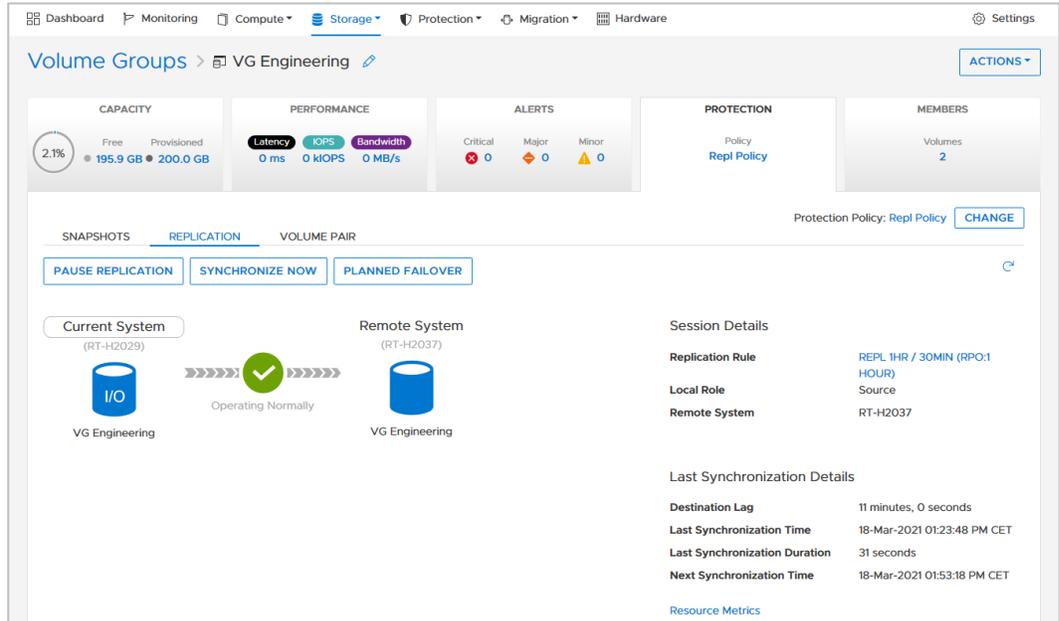


Figure 40. Volume Group details Storage Resource > Protection view

Besides the replication status, replication performance statistics are also available on the **Performance** tab ([Figure 41](#)) of the Storage Resource for volumes, volume groups, and thin clones. The following data is included:

- Replication Remaining Data
- Replication Bandwidth (Normalized)
- Replication Transfer Time

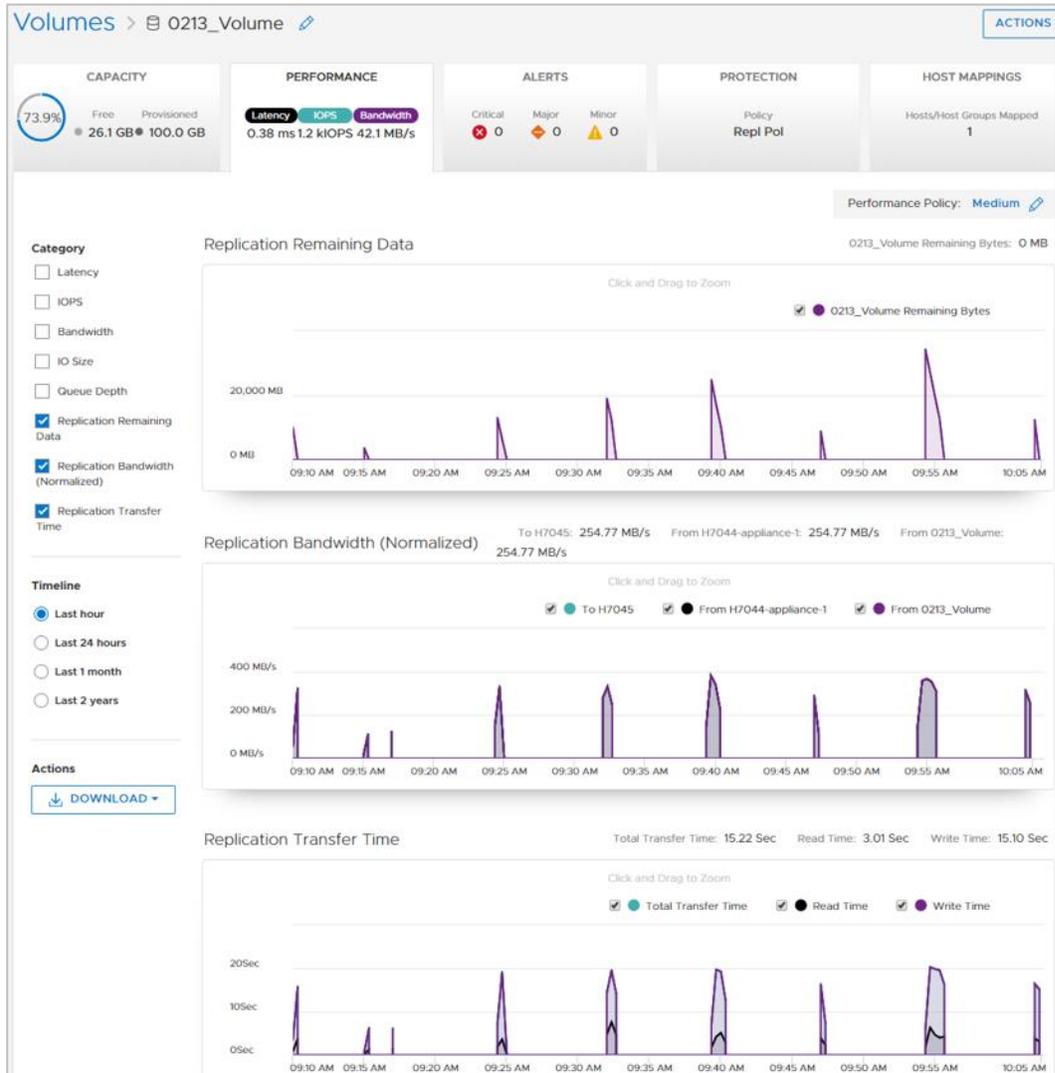


Figure 41. Replication performance view

Replication operations

Several operations are available to manipulate replication sessions as needed. Not all operations are always available, because some depend on the resource type and on the session being in a particular state. Also, certain operations perform differently depending on which system they are issued on—the source or destination. Only one replication operation can be issued and run at a particular time. Replication operations are available when browsing the storage resource details and then selecting the **PROTECTION > REPLICATION** tab or by browsing to the **Protection > Replication** section.

Table 5. Replication session operations

Operation	Asynchronous Replication	Synchronous Replication	Metro
Create session	✓ (assign policy)	✓ (assign policy)	✓ (configure Metro)
Pause and Resume	✓	✓	✓

Operation	Asynchronous Replication	Synchronous Replication	Metro
Synchronize now	✓		
Planned Failover (on source)	✓	✓	
Unplanned Failover (on destination)	✓	✓	
Failover Test	✓*	✓*	
Reprotect (after failover)	✓	✓	
Modify Metro Role			✓
End session	✓ (unassign policy)	✓ (unassign policy)	✓ (end Metro)

* PowerStore provides various capabilities to run failover tests for file replication, which are covered in the [Dell PowerStore: File Capabilities](#) white paper.

Create replication session

A replication session is created when a protection policy with an underlying replication rule is attached to a storage resource. Details are covered in the section [Assign protection policy](#).

Pause and resume

The **PAUSE** and **RESUME** functions can stop and start replication between the resources for a particular replication session (see [Figure 42](#) and [Figure 43](#)). In PowerStore Manager, the pause operation is issued from the source or destination system. If the session is paused while an initial sync or an incremental synchronization is in progress, all incremental changes on the destination are kept. All I/O is kept in a snapshot diff when the replication session is paused. When the session is resumed, replication resumes and the synchronizations to the destination storage resource continue from where they were paused. When a replication session is paused, it also pauses the scheduled RPO synchronizations. The resume operation can be issued on the source or destination system and does not change the replication direction.

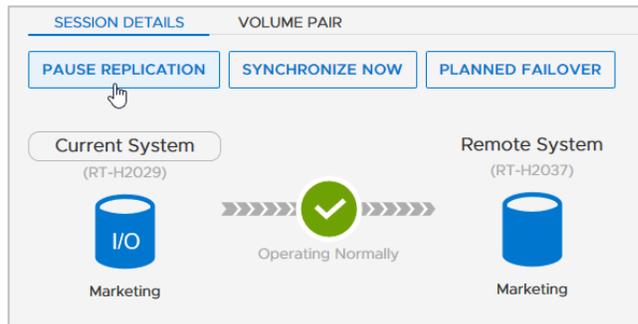


Figure 42. Pause replication

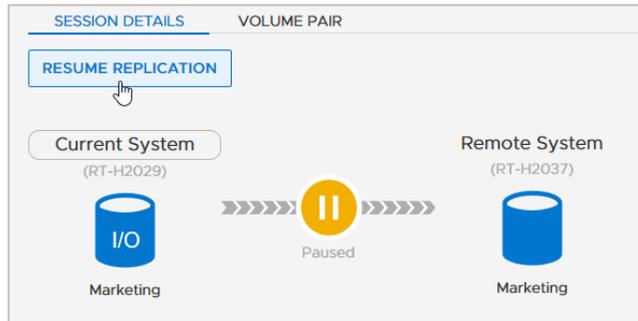


Figure 43. Resume replication

Synchronize now

With asynchronous replication, updates to a destination storage resource occur at a set interval that is based on the defined RPO. When replication is established and an update is not occurring, a **SYNCHRONIZE NOW** operation can be issued to synchronize the latest changes to the destination resource (see Figure 44). After the sync operation is selected, all data that has changed since the last update is copied to the destination storage resource.

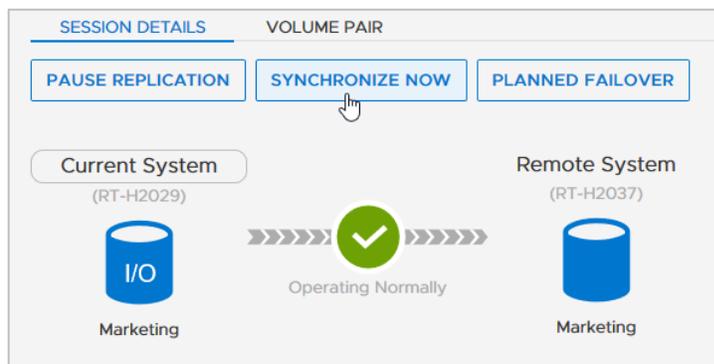


Figure 44. Synchronize now

Planned failover

A **PLANNED FAILOVER** operation allows for replicating the latest acknowledged host data on source volume while also performing a controlled failover (Figure 45). When initiating the operation, the following dialog also allows optionally selecting **Reprotect after failover**. When a planned failover starts, the replication session fails over after completing a synchronization between the volumes. The synchronization before failover

ensures that all data is replicated since the last RPO triggered or manual synchronization. The planned failover option is available on the source storage resource when the replication session is “Operating Normally” or a synchronization is in progress. It causes a short period of data unavailability during the failover operation. Before the Planned Failover operation is issued, it is suggested to issue a manual sync first. This action reduces the amount of data to copy during the planned failover. Quiesce I/O to the source volume before performing a planned failover. After the planned failover is completed, the destination storage resource is available for production I/O, and the original source no longer allows read/write I/O. If host access is configured on the destination resource, hosts can access the data. If reprotect after failover is not selected when initiating the failover, replication does not resume in either direction.

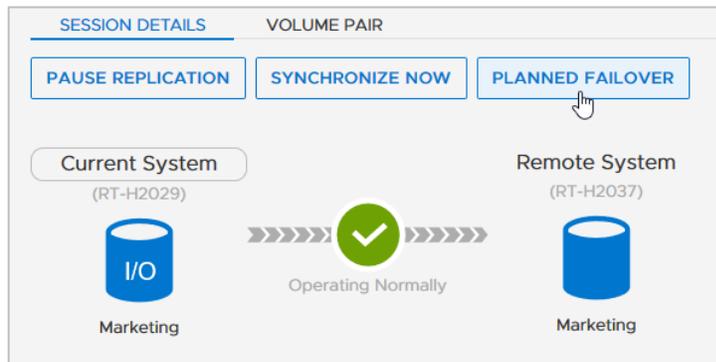


Figure 45. Planned failover

Unplanned failover

The unplanned failover option is only available on the destination of the replication session. This failover type fails over to the latest available common base image that exists at the target without any synchronization occurring beforehand. An unplanned failover assumes that a disaster has occurred on the production system, and the destination image is made read/write. When **FAILOVER** is selected on a destination resource of a replication session (Figure 46), read/write access is removed from the original source if the source is available to receive management commands. The replication session also pauses and does 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 back to the original source when the source becomes available. For file resources, **FAILOVER** is not supported on the destination resource if the source system and production NAS server are still online. If the source is still functioning, issue a **PLANNED FAILOVER** from the source.

PowerStore allows initiating an unplanned failover operation during a disaster scenario or even when the replication is in a **Paused**, **Failing Over**, or **Failed Over** state. Any changes made on the source system while the session is in these states might not be replicated to the destination. Because no final synchronization is performed, an unplanned failover can result in data inconsistency or data loss. It should only be initiated when the source system is not available anymore. Use a planned failover whenever possible (see [Planned failover](#)).

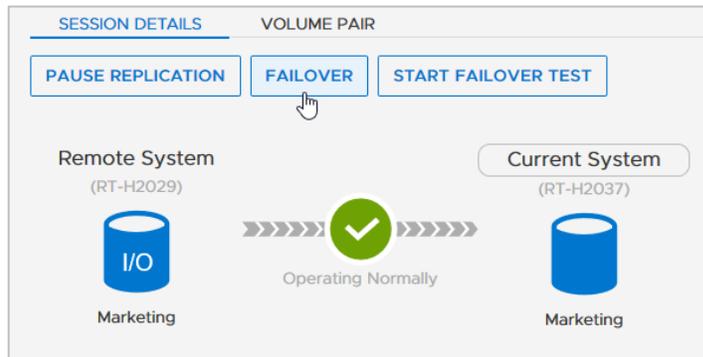


Figure 46. Unplanned failover

Reprotect

After the Planned Failover or Failover option is used, the **REPROTECT** option (Figure 47) becomes available on the new source system. It is also triggered after a planned failover with the reprotect operation is initiated. The reprotect operation starts the replication session and synchronization to the original source system. Because there might not be synchronized changes after an unplanned failover on the destination, it is recommended to take a snapshot on the remote system before initiating the reprotect operation.

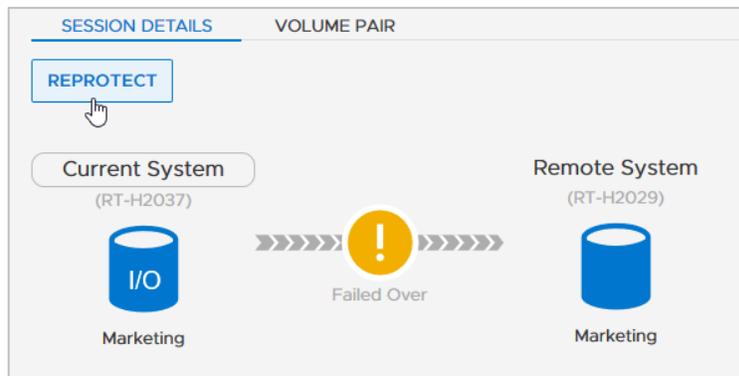


Figure 47. Reprotect

Unassign protection policy with replication

A replication session can be deleted on the source system by detaching the protection policy from the replicated storage resources or by removing the replication rule from a protection policy. Figure 48 shows the option to **Unassign Protection Policy**. When there are no configuration issues and an unassign operation is issued on the source system, the replication session is deleted from the source and destination systems. The destination storage resource is not automatically deleted when the replication session is deleted.

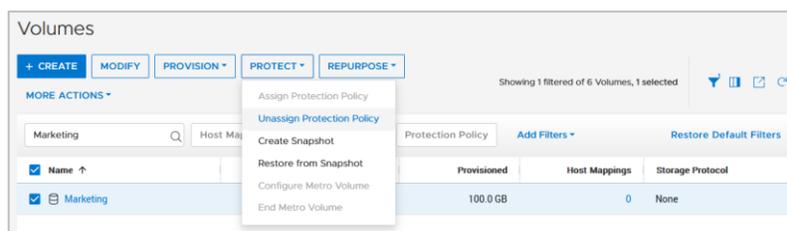


Figure 48. Unassign protection policy

Failover Test

This function allows testing the DR functionality and is only supported on volumes, volume groups, and thin clones. Dell PowerStore provides the Failover Test to enable R/W access to the DR site while production is still ongoing on the primary system (Figure 49).

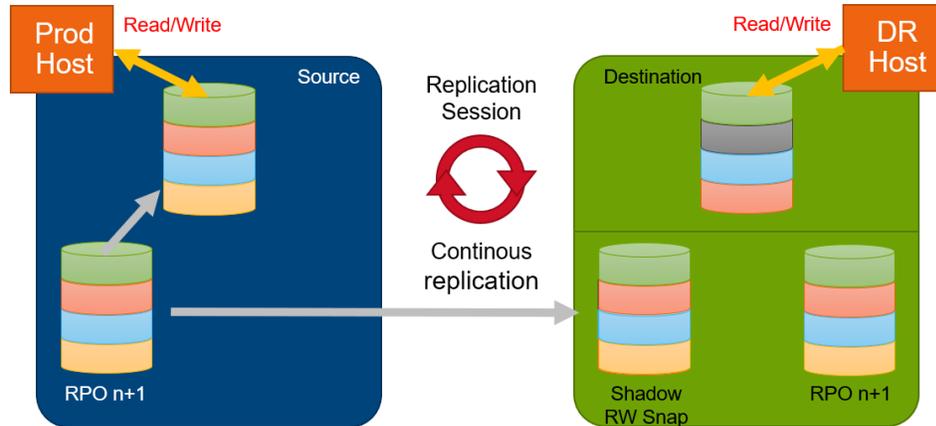


Figure 49. Active failover test

It is possible to start a failover test only on the replication destination (Figure 50) for each storage resource participating in a replication session.

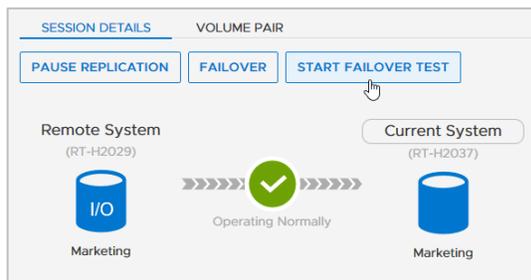


Figure 50. Start failover test

After **START FAILOVER TEST** is selected to initiate a failover test, you must select a snapshot, which will be used as the source of data for the DR test. You can select either the last successful synchronized RPO snapshot or any other existing manual or scheduled snapshot on the destination system for DR test (Figure 51).

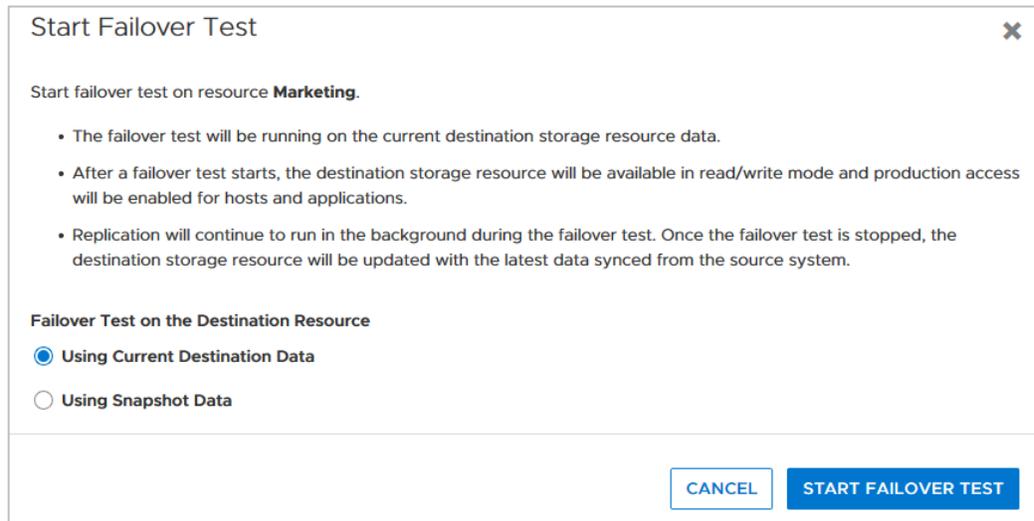


Figure 51. Select destination resource for failover test

When the failover test starts, the storage resource changes to Read/Write for the mapped host. While the failover test is activated, test data writes are stored in the mapped volume and replication continues in the background using a read/write snapshot. All updates from the replication source are baselined and kept in a replication snapshot. PowerStore has no limit on the duration of the DR Test.

The following section describes the options to stop a DR failover test:

- Stop the DR test, discard changes during the test and update the DR volume with the last replicated data
- Stop the DR test, take a snapshot of changed data, and update the DR volume with the last replicated data
- Fail over to the DR volume and continue production with the test data

When stopping a failover test, the access changes back to read-only for the DR Host. The PowerStore Manager provides an optional step to keep test data in a snapshot for later use before the DR host volume is updated with the last successful synchronized snapshot data (Figure 52). A snapshot of test data might be useful when test data should be used or analyzed later. Otherwise, the DR host volume is immediately updated with the last successful synchronized snapshot data.

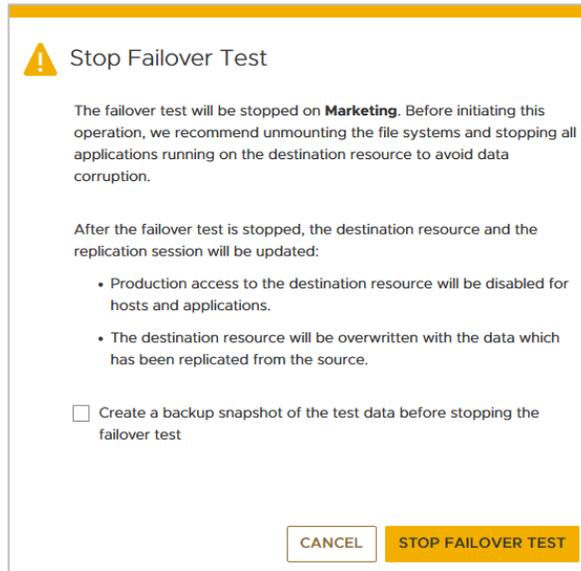


Figure 52. Stop Failover Test

If there is a real DR issue while the failover test is running, there is no further update of the destination volume from the source, and the test data is used for DR production. For this scenario, the operator must confirm the following dialog (Figure 53).

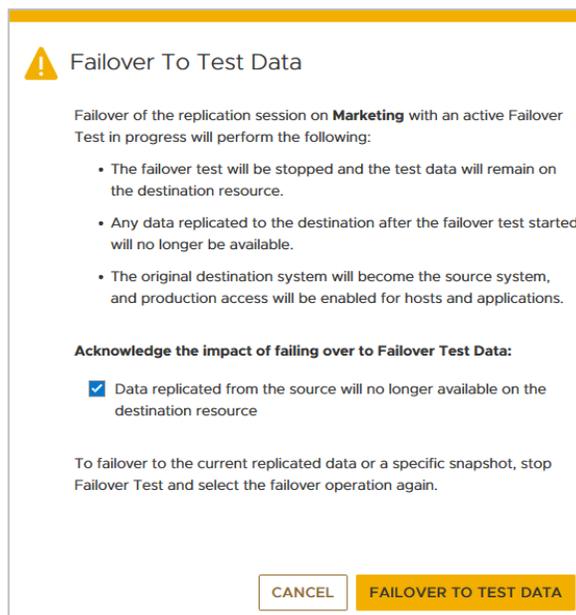


Figure 53. The Failover to Test Data dialog

Clone Destination NAS Server

PowerStore supports cloning the destination NAS server. This feature is designed to enable DR testing without any impact to the ongoing replication session or the production NAS server. It allows customers to confirm that an application can be brought online and write to a share hosted on the destination system.

On the destination system, the user selects the destination NAS server and selects **MORE ACTIONS > Clone**. A new name is provided, and then the user selects the file

systems that the user wants to create with the cloned NAS server. Any shares that exist on the selected file systems will also be cloned. When all information is provided, click **CREATE**.

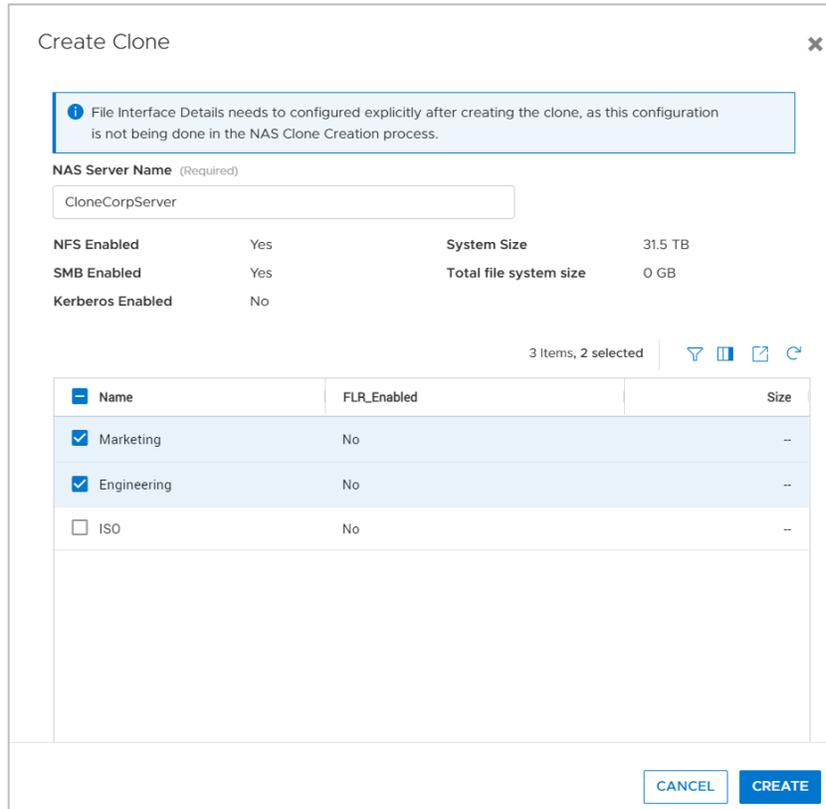


Figure 54. Modify destination NAS server IP address

The cloned NAS server is created without a file interface to ensure that there is no conflict with the production NAS server. In order to access the cloned file systems, a new file interface must be added to the cloned NAS server on the **NETWORK** page of the NAS server. The cloned NAS server is not domain joined automatically. If the cloned NAS server must be domain joined, a unique name needs to be specified before the join operation. After it is cloned, the new NAS server is a stand-alone resource and functions independently from the parent DR NAS server. The NAS server clone operation is not limited to DR testing, and source or even nonreplicated NAS servers support cloning. For more details, see the [Dell PowerStore: Snapshots and Thin Clones](#) white paper.

Modify destination

When replicating a NAS server, the destination NAS server may require different configuration settings than the source NAS server. PowerStore supports the ability to modify the destination NAS server and make these configuration changes before failing over. Therefore, if a failover needs to occur, the destination NAS server will be fully functional when it is promoted to a production instance. The following NAS server configuration options are available for modification on the destination:

- File interface
- DNS, NIS, and LDAP settings

- Virus check configuration
- Event publishing settings

To modify the destination NAS server, go to the **NAS Servers** page on the destination PowerStore system and click into the NAS server. Modify the settings directly on this NAS server. For example, to support a different IP address on the destination NAS server, select the interface on the **NETWORK** page and click **MODIFY**. Then select **Override** and enter the new destination IP address.

Figure 55. Modify destination NAS server IP address

Supported replication configurations

The PowerStore native replication features allow supported storage resources to be replicated remotely between systems. Table 6 shows the supported system configurations.

Table 6. Supported system configuration for replication

Replication Type	Source	Target	Block	File
Asynchronous	PowerStore T/Q model	PowerStore T/Q model	✓	✓
Synchronous	PowerStore T/Q model	PowerStore T/Q model	✓	✓
Synchronous - Metro*	PowerStore T/Q model	PowerStore T/Q model	✓	

* For details about metro, see the [Dell PowerStore Metro Volume](#) white paper.

This section also outlines the supported configurations for native asynchronous replication and synchronous replication. For more information about which systems are supported for asynchronous replication, see [Appendix A: Replication support across platforms](#).

The native replication feature is supported in many different topologies. Deployment models vary depending on the configuration requirements. At a system level, the following configurations are supported:

1. One-directional: A single-source system replicating to a single-destination system
2. Bi-directional: A two-system topology in which each system acts as a replication destination for the peer production resources
3. One-to-many: A system topology in which a single system replicates multiple resources, each to a different remote system (also known as “fan out”).
4. Many-to-one: A system topology in which multiple systems replicate their respective resources to a single system (also known as “fan in”).

Figure 56 shows these supported topologies. The figure uses volumes to represent the storage resources. Asynchronous replication allows for many different deployment models to meet the needs of an organization.

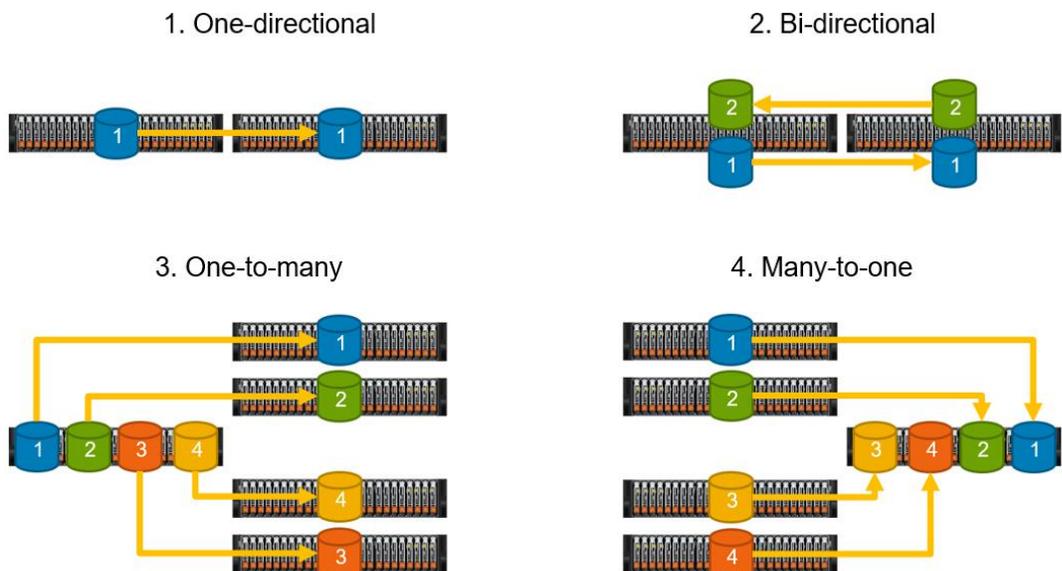


Figure 56. System-level asynchronous and synchronous replication topologies

The bi-directional replication topology is typically used when production I/O must be spread across multiple systems or locations. The systems may exist within a single data center or in different, remote locations. With this replication topology, production I/O from each system is replicated to the peer system. During 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 that both systems are always in use by production I/O.

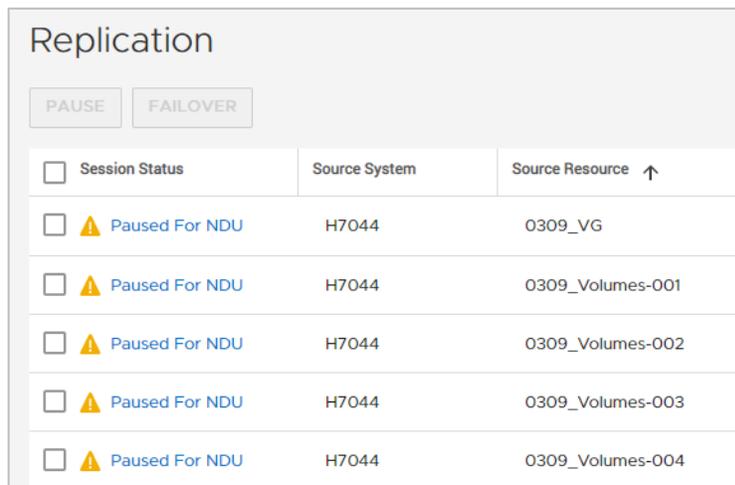
The one-to-many replication topology is deployed when production exists on a single system, but replication must occur to multiple remote systems. This replication topology can be used to replicate data from a production system to a remote location to provide local data access to a remote team. At the remote location, thin clones can be used to provide host access to the local organization or test team.

The many-to-one replication topology is deployed when multiple production systems exist and are replicating to a single system to consolidate the data. 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 a remote office branch office (ROBO) location.

For the one-to-many and many-to-one replication topology examples that are shown in [Figure 56](#), 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. This ability allows for more replication options than what is depicted in the figure.

Upgrades

Depending on the type and state of replication, a replication session can continue when upgrading the PowerStore system (NDU). During PowerStoreOS upgrades, asynchronous replication sessions are temporarily paused and display the status **Paused for NDU** (see [Figure 57](#)). The replication resumes after the upgrade has successfully finished.



The screenshot shows a 'Replication' interface with 'PAUSE' and 'FAILOVER' buttons. Below is a table with columns for Session Status, Source System, and Source Resource. All sessions are marked as 'Paused For NDU' with a yellow warning icon.

<input type="checkbox"/> Session Status	Source System	Source Resource ↑
<input type="checkbox"/> ⚠ Paused For NDU	H7044	0309_VG
<input type="checkbox"/> ⚠ Paused For NDU	H7044	0309_Volumes-001
<input type="checkbox"/> ⚠ Paused For NDU	H7044	0309_Volumes-002
<input type="checkbox"/> ⚠ Paused For NDU	H7044	0309_Volumes-003
<input type="checkbox"/> ⚠ Paused For NDU	H7044	0309_Volumes-004

Figure 57. Session status of paused for nondisruptive upgrade (NDU)

When a synchronous replication session was active when the NDU was initiated, it will keep running. Paused synchronous replication sessions remain paused and can resume after the NDU is finished.

Asynchronous replication for vVol based VMs

Introduction

PowerStoreOS 3.0 and later supports VASA 3.0 native storage-based asynchronous replication for vVol based VMs. This feature uses VMware Storage Policies and requires VMware Site Recovery Manager instances at both sites. The following section gives a brief overview how vVol replication is implemented in PowerStoreOS. See also the [Dell PowerStore: VMware Site Recovery Manager Best Practices](#) white paper for more information.

Licensing

Asynchronous replication of vVol based VMs is included at no extra cost for supported PowerStore clusters.

Theory of operation

The configuration of asynchronous replication for vVol-based VMs requires a remote system pair, as described in an earlier section configured for two PowerStore clusters running PowerStoreOS 3.0 or later. Each of the PowerStore clusters for vVol replication must have a registration in vCenter as a storage provider because the VASA 3.0 API is used to exchange information between the PowerStore cluster and the associated vCenter.

The VMware Storage Policy, which can be assigned to VMs in vCenter, leverages the same replication rules in PowerStore Manager as used for other PowerStore asynchronous replication sessions. Asynchronous replication for vVol based VMs also uses the same snapshot based asynchronous replication technology as native block replication, which is described in the section [native asynchronous replication](#).

When a VMware Storage Policy with PowerStore replication is assigned to a vVol-based VM, a replication session is created on PowerStore for the vVol resources in the same resource group. VMware resource groups can be selected when a VMware Storage Policy is configured for a VM. VMware SRM uses these VMware resource groups to manage the protected VMs in Replication Groups. An SRM Recovery Plan controls the PowerStore replication session for vVols in a replication group during test failover, failover, and reprotection. After a VM has a VMware Storage Policy assigned, and the Resource Group is in a Replication Group with a Protection Plan in SRM, a placeholder VM on destination vCenter and PowerStore is created. The storage container for the placeholder VM is part of the site pair configuration in SRM.

Supported replication flows

For replicating Resource Groups on PowerStore, different combinations of source and destination vVol Storage Containers are possible:

1. One or more Resource Groups from a single Storage Container to a single Storage Container on a different PowerStore cluster
2. One or more different Resource Groups on a single Storage Container to different Storage Containers on different PowerStore clusters
3. Resource Groups from Storage Containers on different PowerStore clusters to a single Storage Container
4. Multiple replications in different directions
5. Any combination of these replication flows

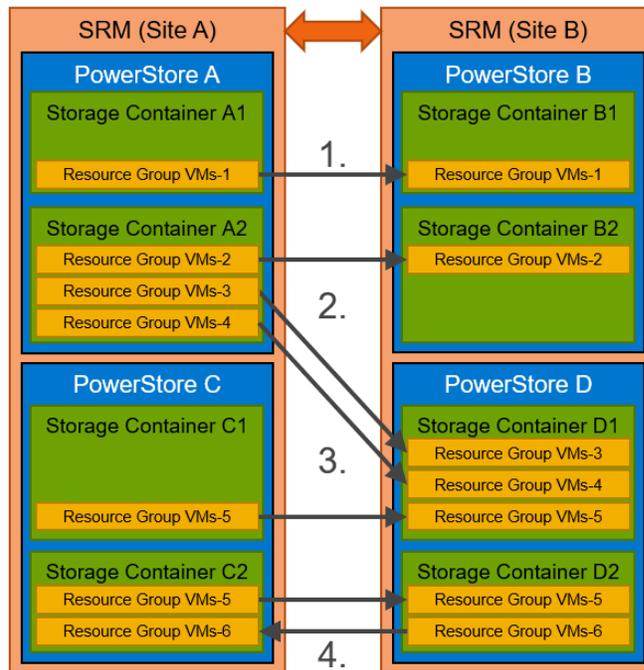


Figure 58. Supported replication flows

Replication operations

The main operations for protected VMs are available in VMware SRM only. This section gives an overview of available operations in PowerStore Manager and VMware Site Recovery Manager.

vVol replication operations in PowerStore Manager

Running an operation for a replication session in PowerStore Manager always affects all vVols in the same resource group. A resource group is configured during the protection of a VM when assigning the VMware Protection Policy in vCenter.

Synchronize

With Synchronize operations, all vVols in the replication group covered by the replication session are manually synchronized.

Pause

This operation pauses the replication session for all vVols in the replication group at the current state. After pausing a vVol replication session, the scheduled RPO replications are disabled.

Resume

The resume operation resumes a paused replication as it is and enables the schedules for RPO-based replications.

Additional resources

For more information about the vVol replication feature, see the [Dell PowerStore: VMware Site Recovery Manager Best Practices white paper](#).

Metro Volume

Introduction This feature allows synchronous replicated active/active block volumes to span two PowerStore clusters running PowerStoreOS 3.0 or later. Starting with PowerStoreOS 3.6, Metro Volumes can be configured with a Metro Witness to mitigate the risk of downtime during some failure scenarios. For more information, see the [Dell PowerStore: Metro Volume](#) white paper.

Licensing Metro Volume configuration is included at no additional cost for supported PowerStore clusters.

System limits

Up-to-date system limits For the most up-to-date system limits, see the Simple Support Matrix, available at <https://elabnavigator.dell.com>.

Integration with PowerStore

Interoperability PowerStore can integrate into other Dell data protection products, such as metro node, VPLEX, RecoverPoint for Virtual Machines, and Dell AppSync. All four products cover different layers for important applications. Metro node and VPLEX offer transparent, in-path data protection solutions for block storage, which can also be used for migration scenarios. RecoverPoint for Virtual Machines provides protection for virtual machines, and AppSync can help to build protection on the application layer. The following sections give a short introduction to the different data protection products.

RecoverPoint for Virtual Machines Along with the native replication options with physical PowerStore systems, RecoverPoint for Virtual Machines is also supported. It is used for disaster recovery and data-loss protection, protecting organizations from site outages due to unforeseen circumstances. It also protects against data loss due to corruption or human error. RecoverPoint for Virtual Machines helps with data-migration solutions and enables moving data between data centers and supported systems. It provides a DVR-like rollback function that allows data recovery to any point in time. It replicates VMs locally within the same PowerStore, or to remote systems. Replication solutions are designed to ensure the integrity of the replicated data at local and remote sites. Performance is not compromised when using RecoverPoint for Virtual Machines with PowerStore.

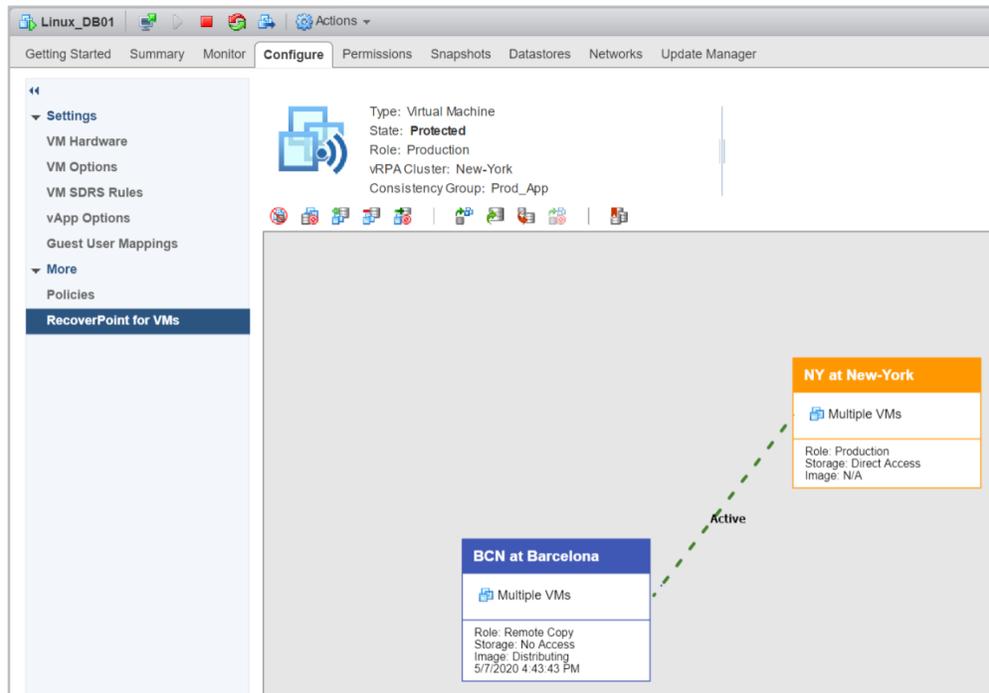


Figure 59. RecoverPoint for Virtual Machines

For more information about RecoverPoint, including RecoverPoint-specific concepts and management, see the *RecoverPoint Administrator's Guide* on [Dell Support](#).

AppSync

Dell AppSync simplifies, orchestrates, and automates the process of generating and consuming application-consistent copies of production data. The deep application integration of AppSync, coupled with the abstraction of underlying Dell storage and replication technologies, empowers application owners to satisfy copy demands for data repurposing, operational recovery, and disaster recovery, all from a single user interface. It can manage the protection, replication, and repurposing of databases and applications using integrated Copy Data Management (iCDM) and replication technologies across the Dell storage portfolio. AppSync supports Oracle, Microsoft SQL Server, Microsoft Exchange, VMware datastores, and other file systems. See the [Dell AppSync Simple Support Matrix](#) for information about supported features for specific environments.

With PowerStore, AppSync provides intuitive workflows to set up local and remote protection, and repurposing jobs. It provides end-to-end automation of all steps including application discovery and storage mapping, creating copies, and mounting or recovery of the copies to the target. AppSync supports both PowerStore T and X models and their snapshot and thin-clone technologies. If AppSync must create remote copies, it uses the native asynchronous replication feature of PowerStore. Currently, AppSync does not support PowerStore file storage, VMware vSphere Virtual Volumes (vVols), or integration with Dell metro node or VPLEX. See the [Dell AppSync Simple Support Matrix](#) for additional support information as it becomes available.

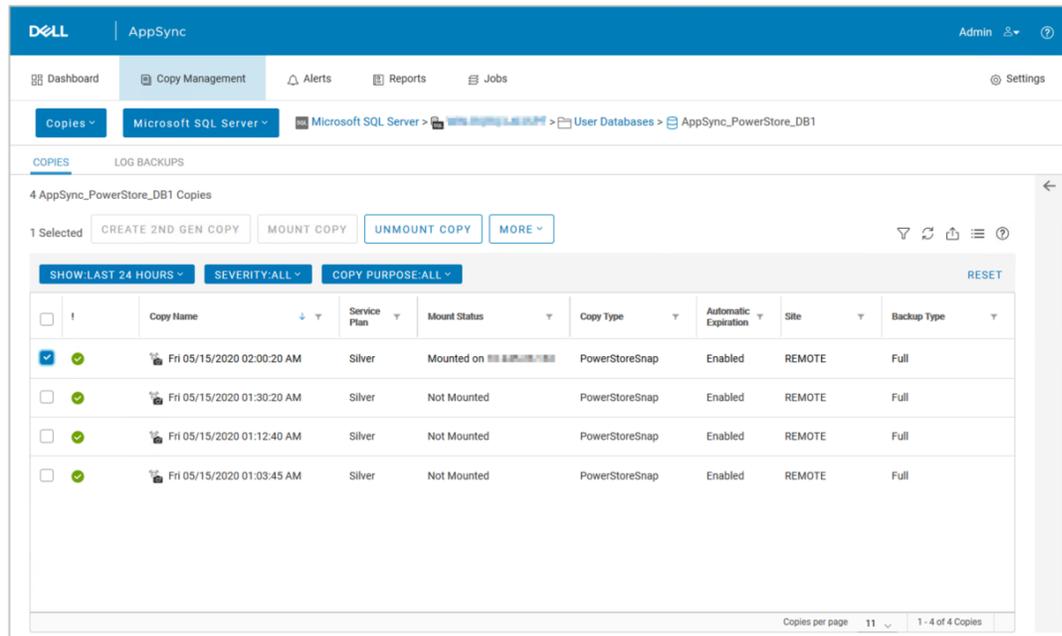


Figure 60. Dell AppSync

Metro node

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

Metro node can also be used in a three-site configuration. Customers can leverage metro node for synchronous replication between sites A and B. They can also add asynchronous replication between sites B and C.

PowerStore 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.

Array-based copy technologies are feature-rich, mature, robust, scalable, and purpose-built to provide enterprise-class replication services. Deploying VPLEX or metro node with storage frames does not diminish the function or value of array-based copy technologies. Underlying storage devices are left untouched by VPLEX or metro node, and the array replication technologies can continue to provide backup, business continuity, operational recovery, QA, or testing and development functionality.

For more information about metro node, see the following white papers:

- [Dell VPLEX: SAN Connectivity—Implementation planning and best practices](#)
- [Dell VPLEX: Leveraging Array Based and Native Copy Technologies.](#)

Conclusion

This paper describes the various native replication solutions that are provided with PowerStore. Configuring a data-protection solution helps guard against unforeseen situations, such as data loss or site-wide outages. PowerStore provides a remote data-protection solution to help minimize the costs that are associated with downtime and provides easy recovery in a disaster. With asynchronous replication solutions, data protection can be configured to meet the needs of the application and organization.

Native asynchronous replication is a data-protection solution that replicates storage resources remotely to other remote PowerStore systems. Asynchronous replication uses the PowerStore snapshot technology to provide consistent point-in-time replicas that can be used in a disaster. With asynchronous replication, no impact to host I/O is seen because data is not immediately replicated as it enters the system. Asynchronous replication uses a customizable RPO, which automatically replicates changes in data at consistent intervals. When data must be replicated over long distances, asynchronous replication can meet the needs of an organization.

PowerStore provides synchronous replication with native Metro Volumes or with the metro node solution. Native Metro Volumes are spanned across two PowerStore clusters and supported for a vSphere Metro Storage cluster configuration. For mapped hosts, a Metro Volume provides fully active/active workloads for high availability and load-balancing of data center resources.

RecoverPoint for Virtual Machines support allows PowerStore to use its enhanced replication features. Virtual machines running on PowerStore can be replicated locally or remotely to another supported system. With RecoverPoint functionality, such as point-in-time data recovery, PowerStore can be protected from disaster scenarios.

Appendix A: Replication support across platforms

Table 7 outlines asynchronous replication support across Dell storage platforms.

Table 7. Asynchronous replication support

Source	Destination	Asynchronous block	Asynchronous file ¹	Asynchronous vVol ¹	RecoverPoint for VMs
PowerStore T or Q models	PowerStore T or Q models	✓	✓	✓	✓
PowerStore T or Q models	Dell Unity	✗	✗	✗	✓
Dell Unity	PowerStore T or Q models	✗	✗	✗	✓

(1) PowerStore T model requires PowerStoreOS 3.0 or later.

Table 8 outlines synchronous replication support for block storage resources across Dell storage platforms.

Table 8. Synchronous replication support for block storage resources

Source	Destination	Metro Volume ¹ VMFS/vSphere	Metro Volume ² Linux/Windows	Dell Metro node
PowerStore T or Q models	PowerStore T or Q models	✓	✓	✓
PowerStore T or Q models	Dell Unity	✗	✗	✓
Dell Unity	PowerStore T or Q models	✗	✗	✓

(1) PowerStoreOS 3.0 and later

(2) PowerStoreOS 4.0 and later

References

The [Dell Technologies Storage Info Hub](#) provides expertise that helps to ensure customer success with Dell storage platforms.

White papers related to PowerStore data protection include:

- [Dell PowerStore: Metro Volume](#)
- [Dell PowerStore: VMware Site Recovery Manager Best Practices](#)
- [Dell PowerStore: Snapshots and Thin Clones](#)

[Dell.com/powerstoredocs](#) provides detailed documentation about how to install, configure, and manage Dell PowerStore systems.