

Dell PowerStore: Replication Technologies

October 2022

H18153.7

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.

Dell Technologies

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 © 2020-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA October 2022 H18153.7.

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 | 5 |
| Remote system configuration..... | 10 |
| Native asynchronous replication..... | 21 |
| Metro Volume | 45 |
| Asynchronous replication for vVol based VMs..... | 45 |
| System limits..... | 48 |
| Integration with PowerStore..... | 48 |
| Conclusion..... | 50 |
| Appendix A: Replication support across platforms | 52 |
| Appendix B: Technical support and resources..... | 53 |

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.

This white paper describes the following replication technologies for PowerStore:

- Native asynchronous replication for Block and File
- Native asynchronous replication for vVol based VMs
- Dell RecoverPoint for Virtual Machines

Native synchronous replication is available with Metro Volumes. The white paper *Dell PowerStore: Metro Volume* describes this feature in detail. Asynchronous replication and Metro Volume 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.

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.

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.

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 | Description |
|---------------|--|
| April 2020 | Initial release: PowerStoreOS 1.0 |
| May 2020 | Minor updates |
| August 2020 | Minor updates |
| January 2021 | Metro node updates |
| April 2021 | PowerStoreOS 2.0 updates including failover test |
| November 2021 | Template update |

| Date | Description |
|--------------|------------------|
| July 2022 | PowerStoreOS 3.0 |
| October 2022 | PowerStoreOS 3.2 |

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

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.

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 | Stands for 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 | 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 | Amount of data, represented in MB/s, which can be transferred in a given period. |
| Common base | 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 | 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 | 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. |

| Term | Definition |
|------------------------------------|--|
| PowerStore CLI | 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. |
| 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 | 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. |
| 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 | 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. |
| User snapshot | 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 | Block-based storage resource that a user provisions. It represents a SCSI logical unit. |

| Term | Definition |
|--------------|--|
| Volume group | 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 a number of replication approaches, but two methods stand out as highly recognized in the storage industry: asynchronous and synchronous. PowerStore supports asynchronous replication of block and file volumes and synchronous replication with Metro Volume (block).

Synchronous replication

Synchronous replication guarantees data consistency (zero data loss) between the replication source and destination volumes during normal operation. This is achieved by ensuring write I/O commitments at the replication source and destination before a successful write acknowledgement is sent back to the host and the requesting application. Synchronous replication provides a blend of data consistency and high availability with uniform storage presentation.

Note: PowerStore uses a Symmetric Active/Active Metro Volume architecture. This means that either volume may be a synchronous replication source and either volume may be a synchronous replication destination. Synchronous replication happens bidirectionally between clusters that support Metro Volumes.

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 also applies to Metro Volumes built on top of synchronous replications. 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.

The figure below 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 acknowledgement is sent to the application or server.

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

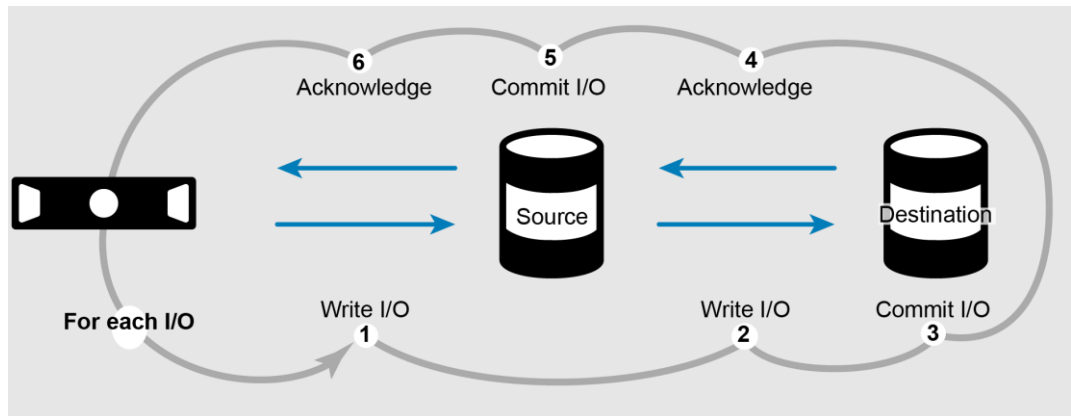


Figure 1. Block storage synchronous replication write I/O sequence

Note: PowerStore supports synchronous replication as part of the Metro Volume feature introduced in PowerStore OS 3.0. For more information, please refer to 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 acknowledgement 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 volume snapshot rules and 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. 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.

The figure below 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 acknowledgement 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.

- A batch acknowledgement is sent to the source.

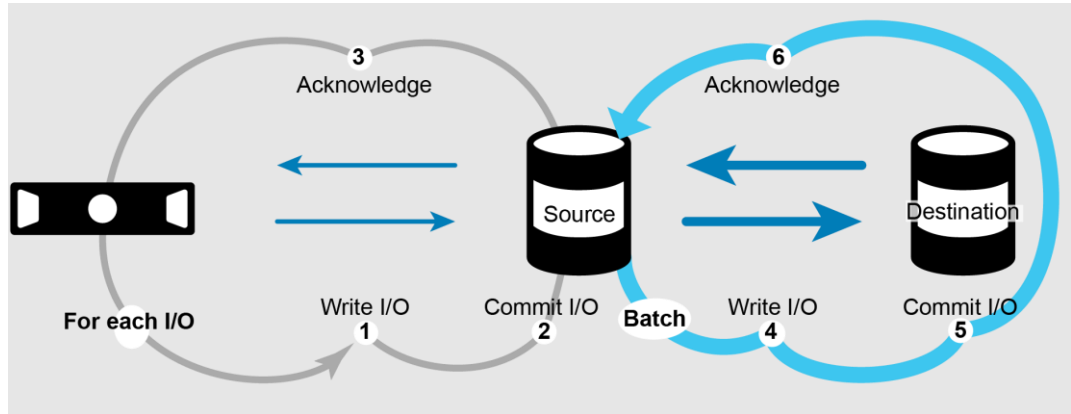


Figure 2. Block storage asynchronous replication write I/O sequence

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 and contains information about related network information for management and data connection. After a remote system is set up, the remote system configuration can be used on both participating PowerStore clusters for replication in any direction for available capabilities. Capabilities represent the available replication in PowerStoreOS and can be identified as “Block”, “vVol”, “Metro”, and “File”. PowerStoreOS supports up to 16 different remote system pairs.

Overview and prerequisites

PowerStore’s 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:

- Verify that a storage network tagged 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 management traffic it is required to have network connectivity between the participating PowerStore clusters on the management interface. 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.
- 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 mask. 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 subsections outline the different functions and requirement details, and how these components interact with each other. Use PowerStore Manager to configure and manage these components.

Port configuration for replication

Ports 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 on a PowerStore T model, and port vFEPort1 in group PG_Storage_TGT1 on a PowerStore X model appliance. 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 in the 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.

On PowerStore, the ports that are listed in Table 2 can be used for replication. When a 25GbE optical 4-port card or IO module is used, both 10Gb and 25Gb SFPs can be leveraged for replication.

Table 2. Supported ports for replication

| Model | 4-port card: 10 GbE BaseT or 10/25 GbE optical | I/O modules 0 and 1: 10 GbE BaseT, 10/25 GbE optical, or 100GbE |
|--------------------|--|---|
| PowerStore T model | Yes | Yes |
| PowerStore X model | Yes | No |

The figures in the next section [Error! Reference source not found.](#) show examples of minimal cabling for replication between PowerStore T models (Figure 3), and between PowerStore T models and PowerStore X models (Figure 4). 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. It is recommended to only tag replication interfaces on ports of the same type and speed. 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 the default configuration as outlined below, consider that the ports might also be used for other traffic such as I/O for block storage host access, migration, file (PowerStore T models only), or vMotion/VM traffic (PowerStore X models only). If these features are used, it is recommended to plan replication using dedicated interfaces. Extra ports use dynamic storage IP configuration from the range that was given in ICW (PowerStore X Model only) or added afterwards in the networking section of PowerStore Manager.

Replication connection

Figure 3 shows an example configuration of a replication connection between two physical systems. In both of the following figures, the source of the replication session is

the **Production System**, and the destination is the **DR System**. For each of these example systems, the default port configuration is used as replication ports. Figure 3 shows cabling for a pair of PowerStore T model appliances using system bond0.




Figure 3. Native replication using two PowerStore T model systems

Figure 4 shows cabling between PowerStore X model and PowerStore T model appliances. With PowerStore X model arrays, ports are used for replication management and replication data traffic.




Figure 4. Native replication using a PowerStore T and a PowerStore X model system

Remote systems

When the ports for replication traffic are tagged 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 reestablish 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 block replication
- Metro Volume
- Asynchronous vVol replication
- Asynchronous File

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:

- 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.
- 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 tagged for replication, the available bandwidth is shared. This configuration can have an impact on 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 can be tagged 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 and for continuous replication of Metro Volumes.

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. Starting with PowerStoreOS 3.0, it is also possible to use different ports than system bond for file I/O and to create additional link aggregation for file I/O.

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, click 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, click Hardware > Appliance-Name > Ports tab. All Ethernet ports

and system-bond are eligible to be tagged as replication ports and are available in the ports list.

Figure 5 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, ports can be mapped to the storage network and you can 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 5. 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 6 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 6. Map Storage Network

After ports are mapped to the storage network, click **MORE ACTIONS** > **Tag for Replication** as shown in Figure 7. 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.

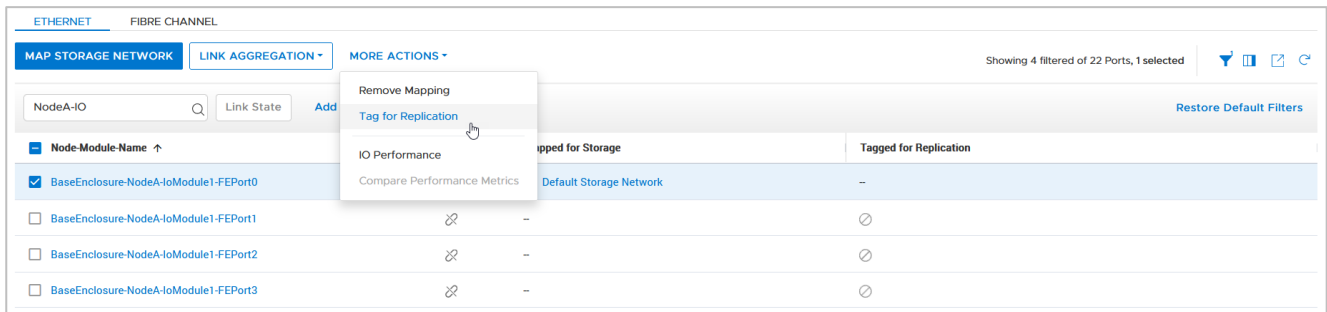


Figure 7. Tag port for replication

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** as already configured networks in PowerStore Manager (Figure 8 / Settings > Networking > Network IPs > Storage).

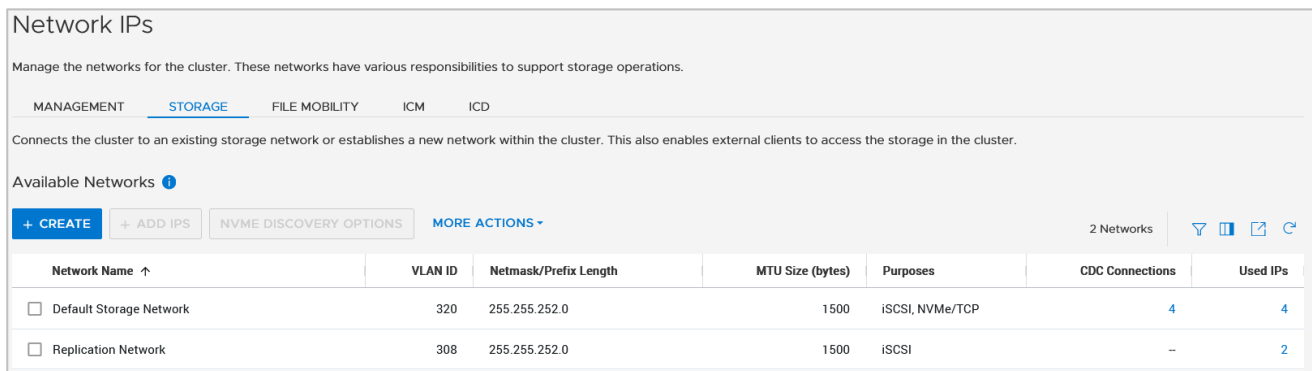


Figure 8. Multiple storage networks

Example 1: Two storage networks over a single port

When physical links for 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 9 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 9. Single network configuration

For replication tagging, it is required to configure the additional storage network. **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 appears, choose the **Replication Network** and continue with **MAP NETWORK** (Figure 10).

The 'Map Storage Network' dialog box is open, showing a list of available storage networks. The 'Replication Network' is selected. The dialog includes options for '+ CREATE', '+ ADD IPS', and 'NVM DISCOVERY OPTIONS'. A table within the dialog shows the selected network details:

| Network Name | Purposes | CDC Connections | Available IPs |
|---------------------|----------|-----------------|---------------|
| Replication Network | iSCSI | -- | 2 |

The 'MAP' button is highlighted in blue.

Figure 10. Map Storage Network

When the Map Storage Network dialog is confirmed, the port overview column **Mapped for Storage** turns into number **2**, which indicates that two storage networks are mapped and using this port. IP address information for the ports is available when hovering above the number (Figure 11).

The tooltip shows the following information:

| | |
|-------------------------|-------|
| Default Storage Network | 20.23 |
| Replication Network | 9.30 |

Figure 11. 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 12.

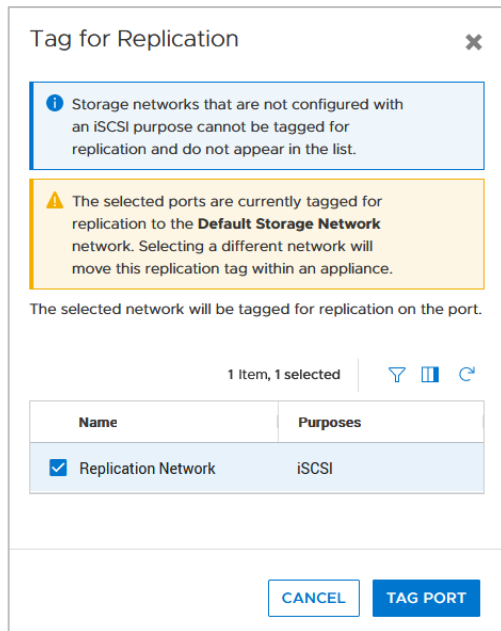


Figure 12. Tag for Replication – Network selection

After the dialog to perform configuration on both nodes is confirmed, the tagged network for replication has changed to the new **Replication Network** (Figure 13).

| Node-Module Name | Link State | Mapped for Storage | Tagged for Replication |
|--|------------|---------------------------|------------------------|
| <input type="checkbox"/> BaseEnclosure-NodeA-IOModule1-FEPort0 | 🟢 | ✓ Default Storage Network | - |
| <input checked="" type="checkbox"/> BaseEnclosure-NodeA-bond0 | 🟢 | ✓ 2 | ✓ Replication Network |
| <input type="checkbox"/> BaseEnclosure-NodeA-4PortCard-FEPort0 | 🟢 | ⊘ | ⊘ |
| <input type="checkbox"/> BaseEnclosure-NodeA-4PortCard-FEPort1 | 🟢 | ⊘ | ⊘ |
| <input type="checkbox"/> BaseEnclosure-NodeB-IOModule1-FEPort0 | 🟢 | ✓ Default Storage Network | - |
| <input type="checkbox"/> BaseEnclosure-NodeB-bond0 | 🟢 | ✓ 2 | ✓ Replication Network |
| <input type="checkbox"/> BaseEnclosure-NodeB-4PortCard-FEPort0 | 🟢 | ⊘ | ⊘ |
| <input type="checkbox"/> BaseEnclosure-NodeB-4PortCard-FEPort1 | 🟢 | ⊘ | ⊘ |

Figure 13. Single port configuration with dedicated network tagged for replication

Example 2: Separated host and replication networks

Note: This example is using 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, with the difference to map the **Replication Network** to another port than the **Default Storage Network**, which is used for host traffic. The example in Figure 14 shows a selected port with dialog to select the storage network for mapping.

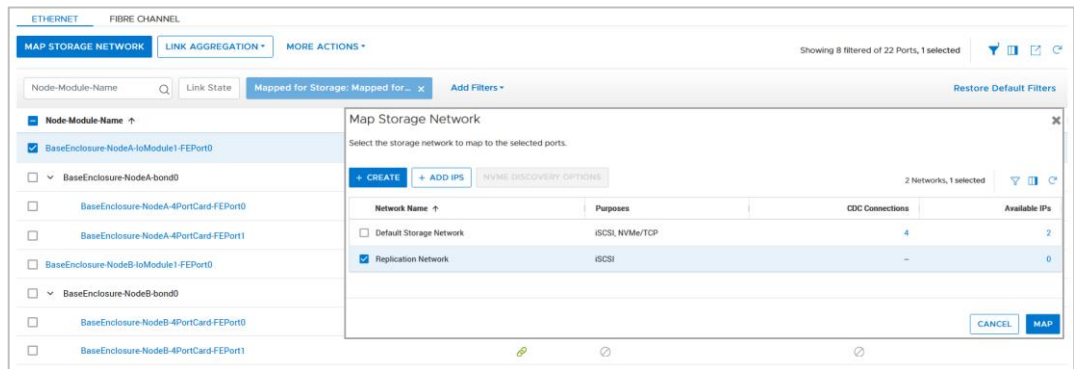


Figure 14. Map replication network to a new network port

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 dialog is confirmed. Figure 15 shows the configuration for that example.

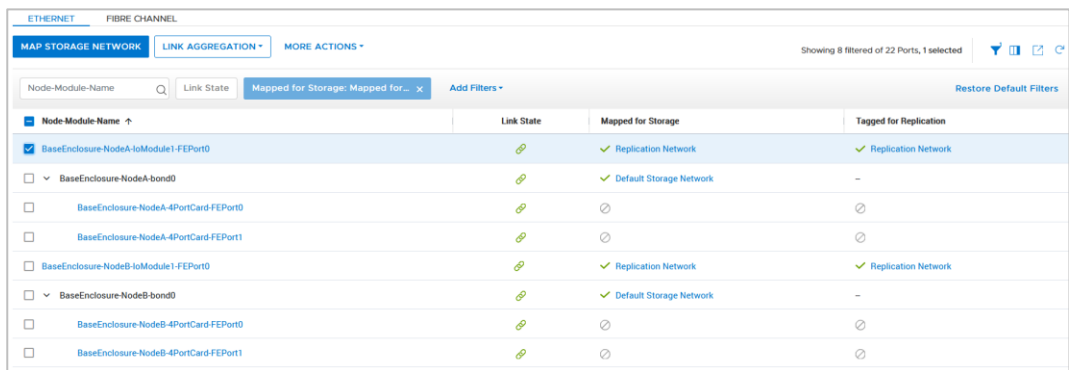


Figure 15. Port configuration with dedicated host and replication storage network

Remote systems

The next step in configuring remote replication is to create a remote systems pair with another system. This step configures a private replication connection using the management ports. To set up a replication connection, click **Protection > Remote Systems** to start creating remote systems (see Figure 16).

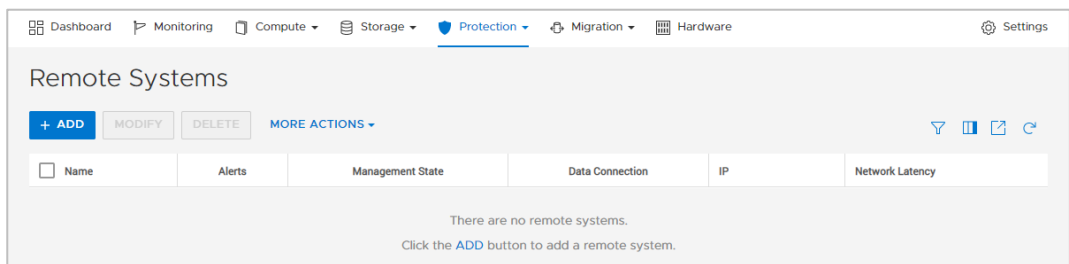


Figure 16. Remote systems replication setup

To define a new remote system, click **ADD** as shown in Figure 17. The new Add Remote System window appears (see Figure 17) and requires the following information:

- Management Cluster IP Address

- Username and Password of the remote system
- Network Latency setting

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 17.

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.

Figure 17. Add Remote System

After a remote system is set up, click **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 18 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 18. 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 TCP-based protocol for replication data traffic. 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 3.0 is supported and relies on the iSCSI protocol. Each latency category uses a different network port number. For replication across network borders it might be required to adjust network ACL or network firewall rules to allow replication traffic. Table 3 shows an overview of different network latency settings and the used network port on PowerStore.

| | 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 |

Table 3. Remote systems network latency overview

File mobility network

File replication requires an additional file mobility network configuration for control 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. While PowerStoreOS 3.0 does not support any changes when a 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 a deletion of the file mobility network is supported, it is required when replication sessions are activated again.

The configuration for file mobility network can be found in Settings > Networking > Network IP in the File Mobility tab (Figure 19). 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 appears 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

Gateway (Optional)

Cluster IP

File Mobility Network IPs **i**

| Address |
|--------------------|
| The list is empty. |

0 IPs provided, 2 IPs required (minimum)

Network MTU Size **i**

Figure 19. Create file mobility network

Native asynchronous replication

Introduction

This section describes the PowerStore native asynchronous replication feature which allows users to create replication sessions for block and file storage resources between PowerStore systems. Supported storage resources for native asynchronous replication are volumes, volume groups, thin clones, NAS servers, and 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 asynchronous replication feature
- How the native asynchronous replication feature works
- Configurations supported for asynchronous replication

- PowerStore Manager configuration and management

Licensing

Asynchronous replication is supported on all PowerStore systems and is included in base license at no extra 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. A replication rule defines the remote system and replication cycles for the asynchronous replication. Protection policies allow the user to configure remote and local protection using replication rules, snapshot rules, or both. 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 and up to four snapshot rules.

The replication rule defines the parameter for the asynchronous 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, RPO, and alert threshold for the planned replication session. Once a protection policy with a replication rule is assigned to a storage resource, the configured RPO in the rule will be used to set up the internal event scheduler for recurring replication of the storage resource.

To minimize the chance of RPO compliance issues, 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 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 example shows a storage resource that is scheduled with a target RPO of one hour and an alert threshold of zero minutes. See the following steps and Figure 20.

1. The initial synchronization is completed successfully within the 30-minute RPO window, and the first schedule RPO synchronization cycle begins.

- 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.
- 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.
- 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 20. 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 individual RPO schedules for each storage resource in that replication session. Scheduled or manual user snapshots of block storage resources in a replication session are also replicated in chronological order to the destination during initial and continuous synchronizations. 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 an asynchronous replication session is created, and before the 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.

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 21 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.

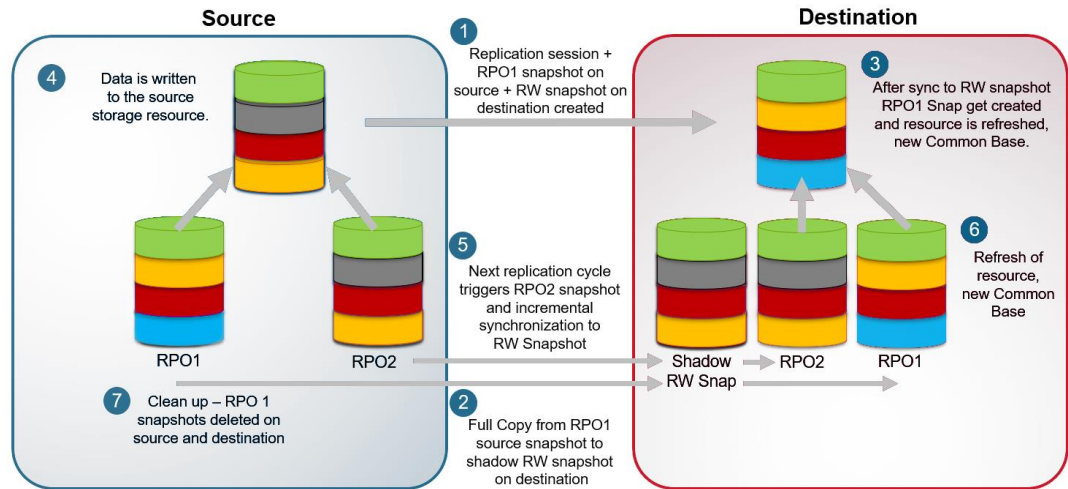


Figure 21. Asynchronous replication theory

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.

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

- Volumes
- Thin clones
- Volume groups
- NAS servers
- File systems

Asynchronous replication operates in the same way for volumes, volume groups, thin clones, and file resources on PowerStore. When asynchronous 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 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 sync.

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. This behavior means that a single replication session is created for the entire volume group no matter how many volumes it contains. 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. When replication is paused or resumed on a volume group, the replication operation affects the entire group. Check the **Write consistency order** option for the volume group to have a consistent replica at the volume-group level.

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 is 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 asynchronous remote replication to a system that is defined in remote systems. Click **Protection > Protection Policies** as shown in Figure 22.




Figure 22. Protection Policies

When the Protection Policies window appears (Figure 23), it is possible to create new and manage existing protection policies or rules. The following example creates a protection policy with replication to a previously configured remote system. In the Protection Policies window, click the **CREATE** button to begin the configuration.




Figure 23. 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 create a new rule in the **Replication Rules** area by clicking **CREATE**.

The following information is required. Each step corresponds with a number that is shown in Figure 24.

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

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

3. Replication Rule Name
4. Destination Remote System
5. RPO
6. Alert Threshold

Figure 24. 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 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. In that case, 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. For volumes, begin the process by clicking **Storage > Volumes**, and click the **CREATE** button (see Figure 25).

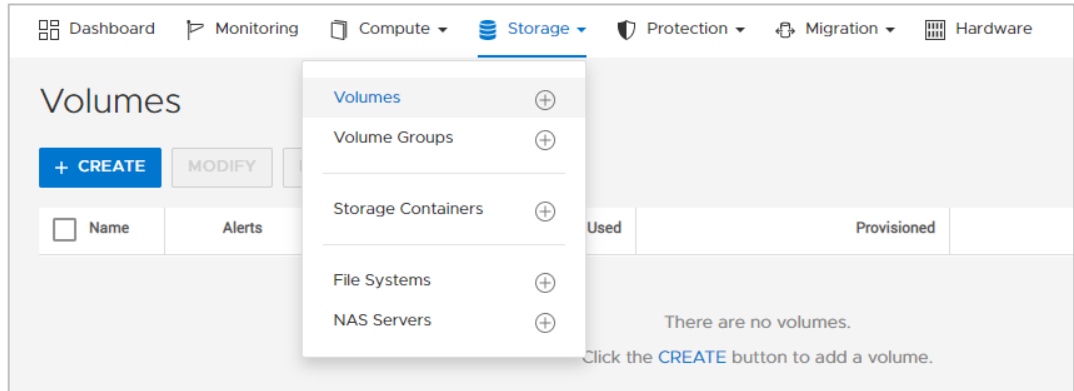


Figure 25. First step: Assign protection policy

In the **Create Volumes** window, enter the following information. Each number below corresponds with the number in Figure 26.

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, click the drop-down menu **Volume Protection Policy (Optional)**. This menu shows all local available protection policies.

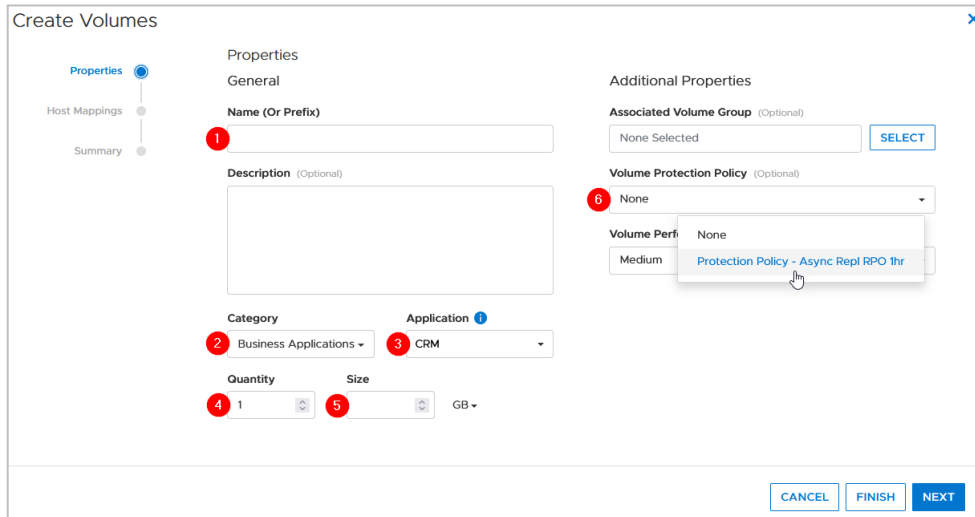


Figure 26. Second step: Assign protection policy

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 assigning a protection policy to an existing volume and are similar for existing volume groups, thin clones, or NAS servers. 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. Next, click the **PROTECT** drop-down menu and click **Assign Protection Policy** as shown in Figure 27.

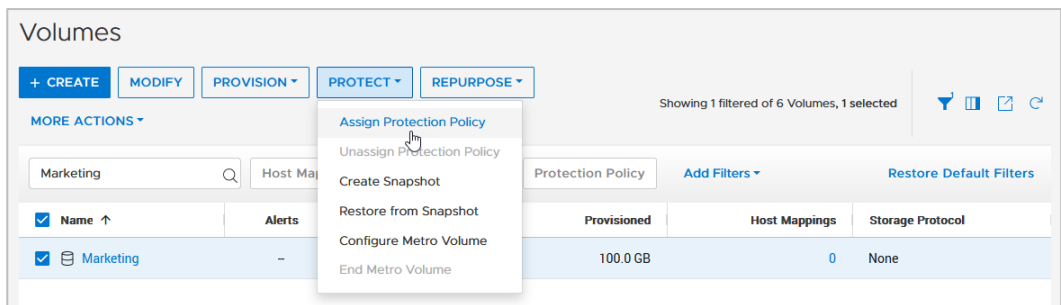


Figure 27. Assign protection policy to existing storage resource

In the following window, choose the appropriate protection policy and click **Apply** to apply it to the previously selected storage resources. After it is applied, the initial synchronization starts immediately.

Viewing the replication sessions

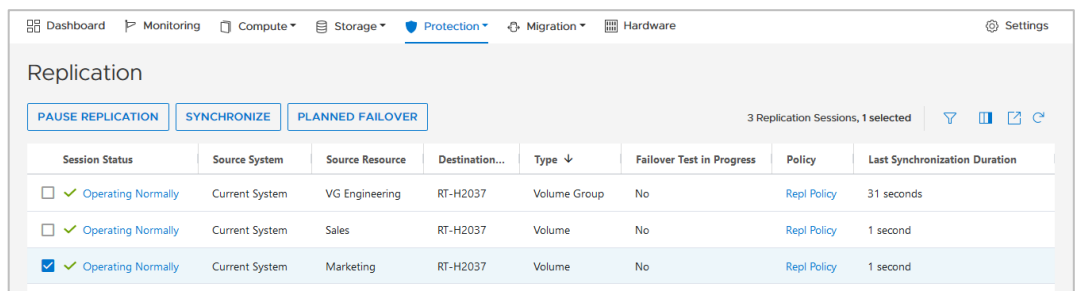
All replication sessions on the system can be viewed from the **Replication** page. To view this page in PowerStore Manager, under **Protection** click **Replication**. Figure 28 shows an example of the replication sessions overview with multiple replication sessions that are created on the system. This 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 will display “-” 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 28 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 cycle
- **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 28. Replication window for block source resource

For file replication, the NAS server replication session is shown at the top level. This session can be expanded by clicking the down arrow to the left of the session. Once expanded, all underlying file system replication sessions for that NAS server are shown (Figure 29). Note how 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.

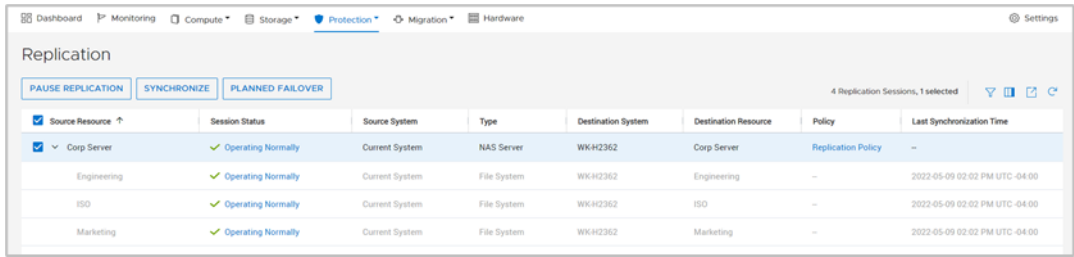


Figure 29. Replication window for NAS server and file systems

The Replication window for the destination resource shows different active buttons when a session is selected (Figure 30):

- **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.

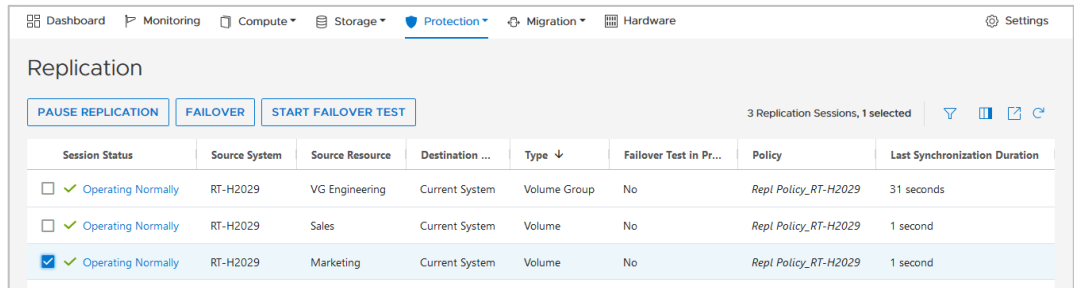


Figure 30. Replication window for destination resource

When the **Failover** button is clicked, a warning message appears saying 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 31. The local storage resource is always tagged with **Current System**.

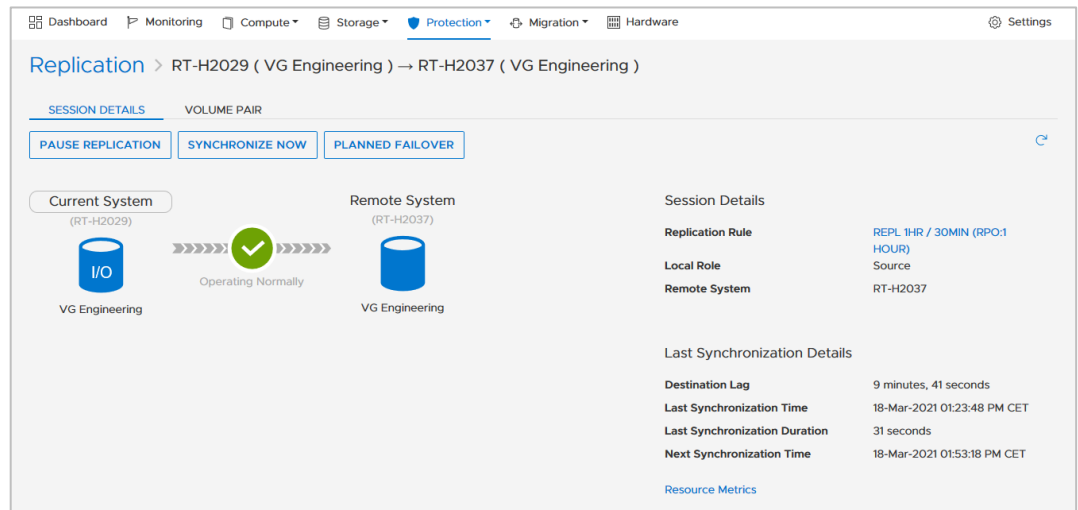


Figure 31. Replication Session Summary

When working in the volume or volume group properties pages, it is also possible to see and control the corresponding replication session. Figure 32 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, click the **Protection** tab > **Replication** tab in the following view in Figure 32.

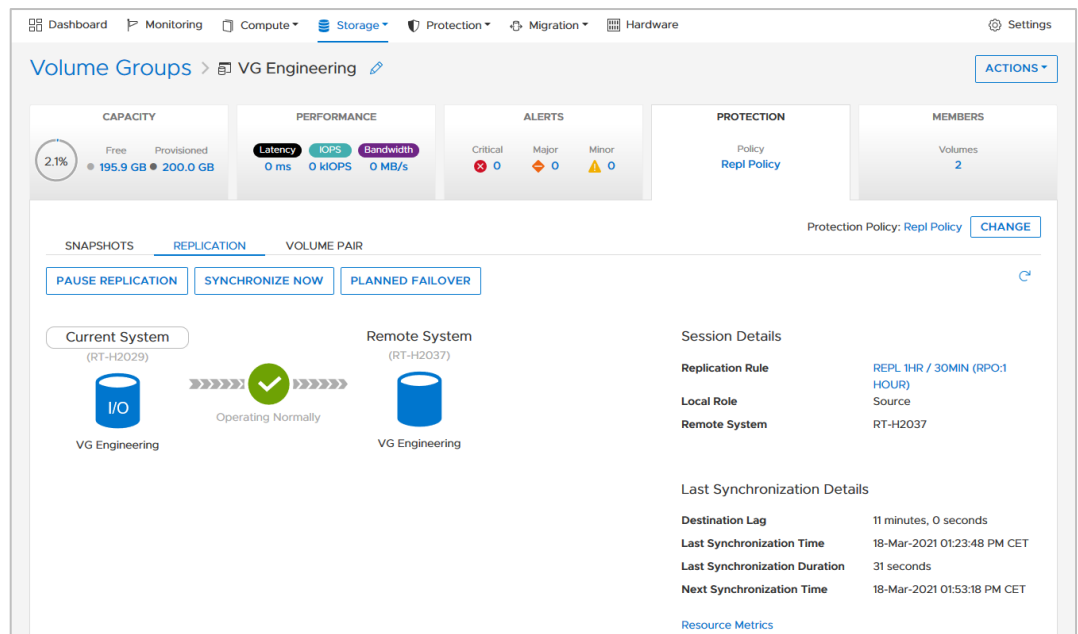


Figure 32. Volume Group details Storage Resource > Protection view

Besides the replication status, replication performance statistics are also available in the **Performance** tab (Figure 33) 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

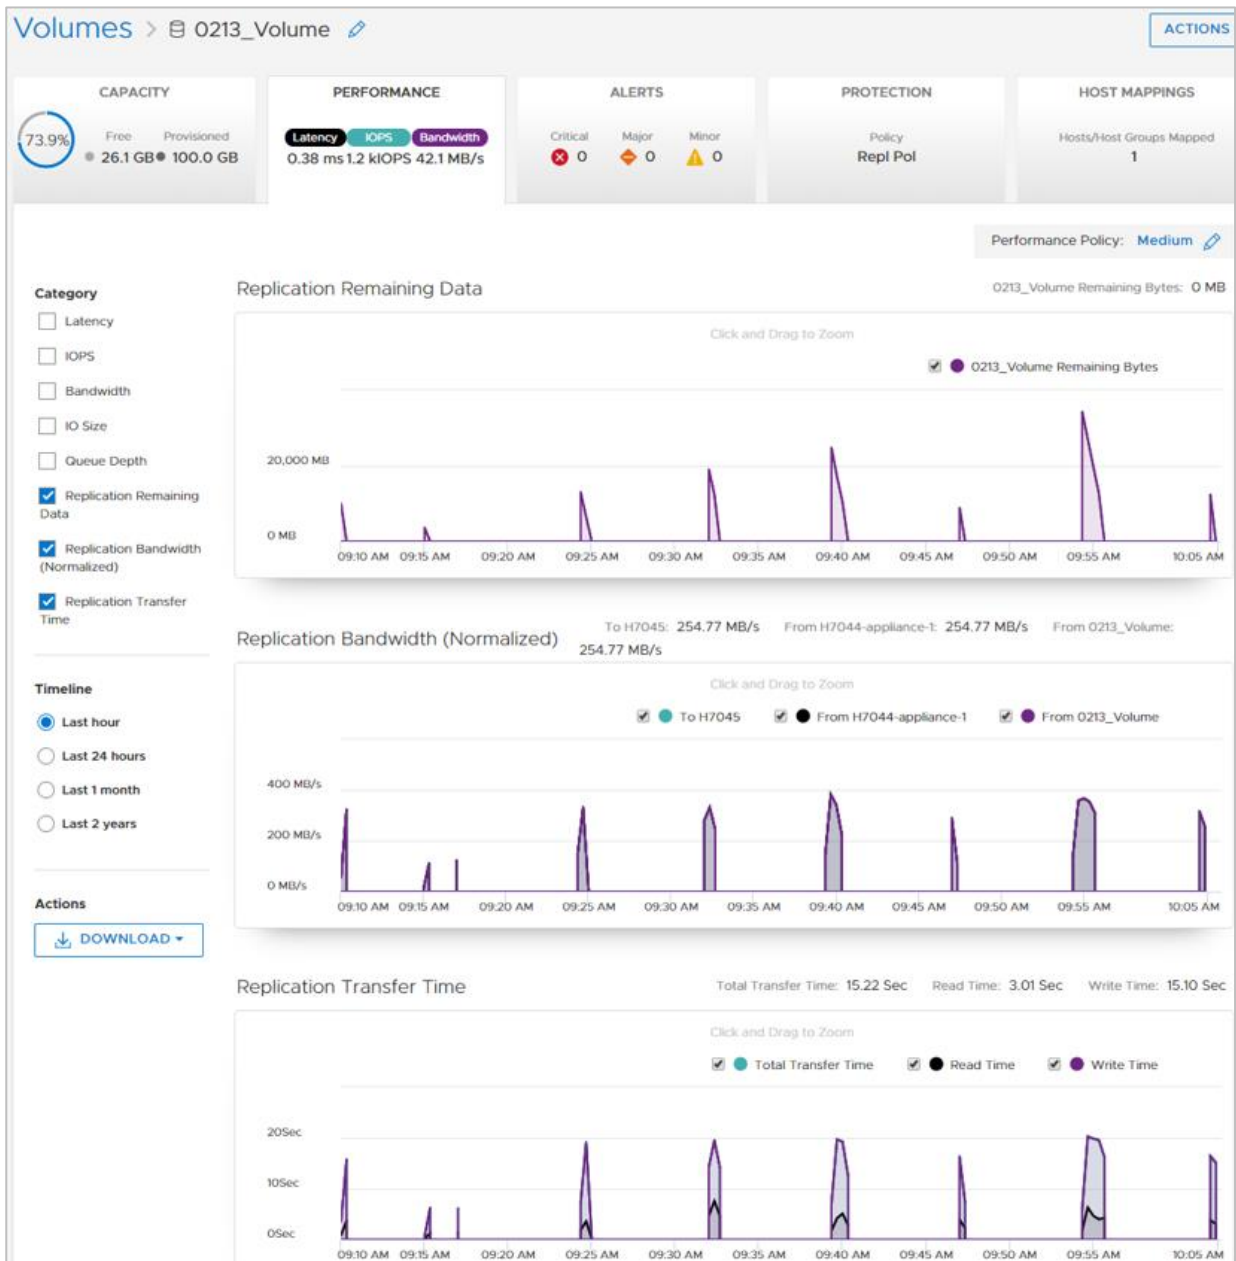


Figure 33. Replication performance view

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.

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 0 and Figure 35). 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.

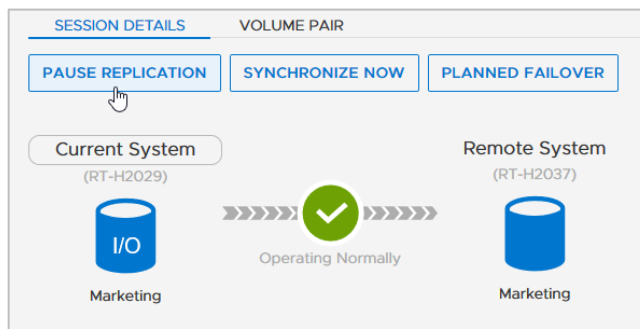


Figure 34. Pause replication

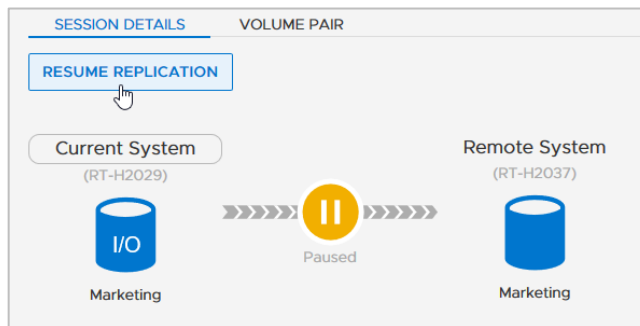


Figure 35. Resume replication

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 36). After the sync operation is selected, all data that has changed since the last update is copied to the destination storage resource.

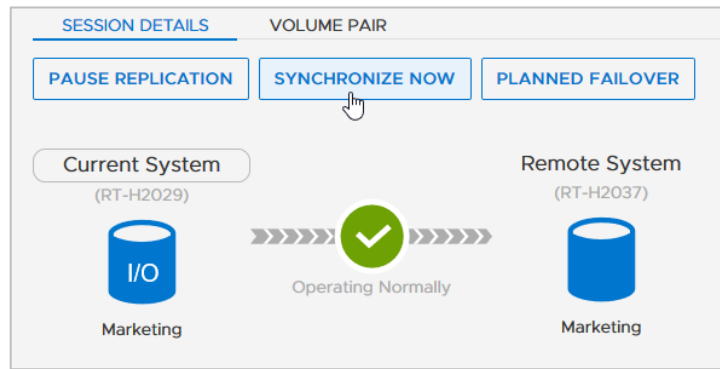


Figure 36. Synchronize now

Planned failover

A **PLANNED FAILOVER** operation allows for replicating the latest acknowledged host data on source volume while also performing a controlled failover (Figure 37). 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 completes, 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.

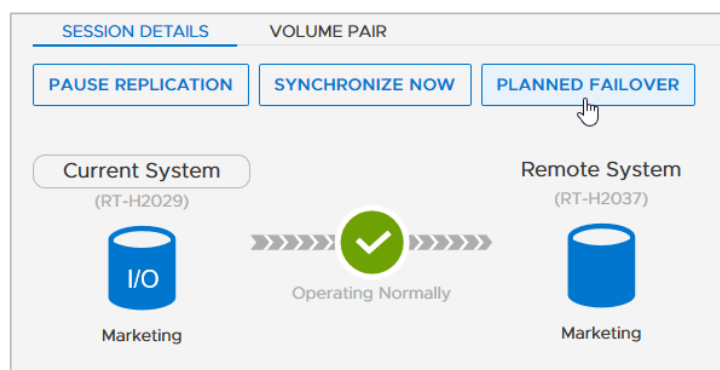


Figure 37. Planned failover

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 38), 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. Since no final synchronization is performed, an unplanned failover can result in data inconsistency or data loss. It should be only initiated when the source system is not available anymore. Use a planned failover whenever possible (see [Planned failover](#)).

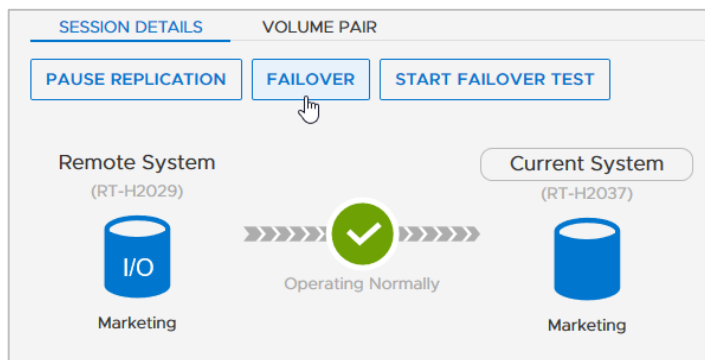


Figure 38. Unplanned failover

Reprotect

After the Planned Failover or Failover option is used, the **REPROTECT** option (Figure 39) 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. Since 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 the reprotect operation is initiated.

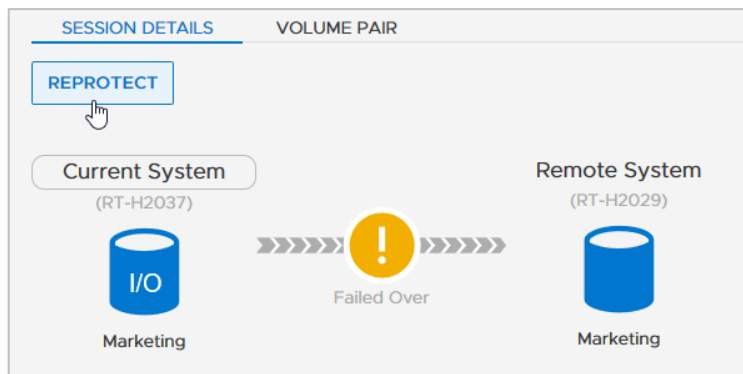


Figure 39. Reprotect

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 40 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 40. 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 41).

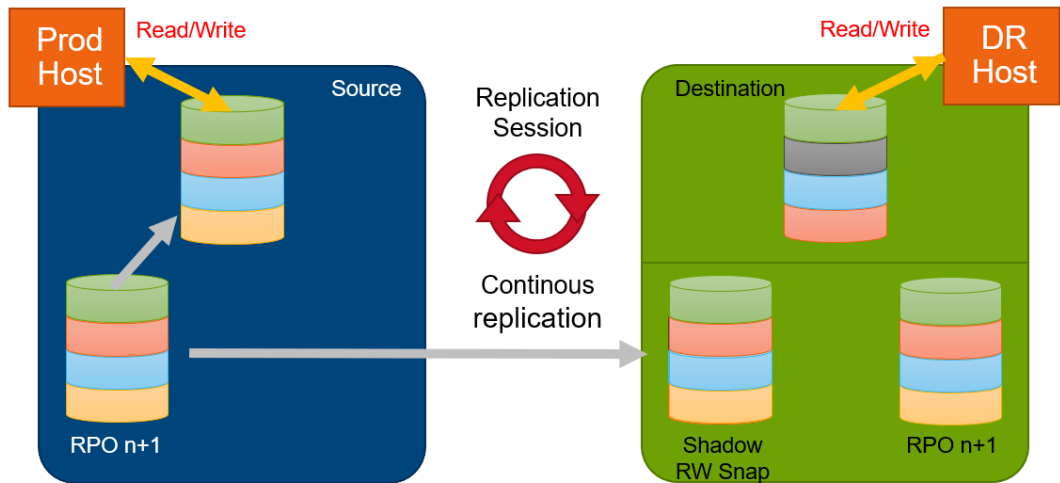


Figure 41. Active failover test

It is possible to start a failover test only on the replication destination (Figure 42) for each storage resource participating in a replication session.

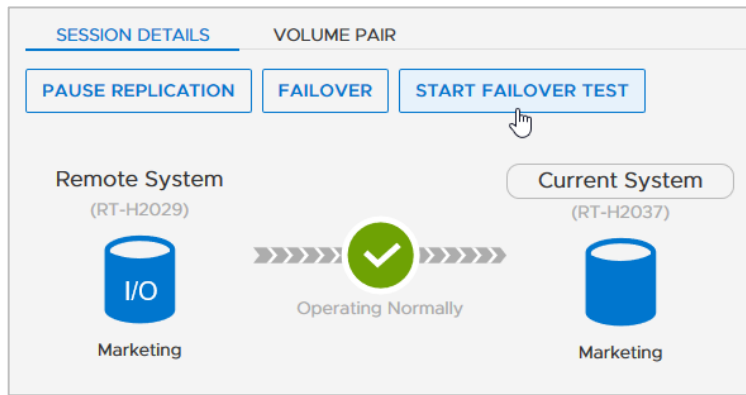


Figure 42. Start failover test

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 43).

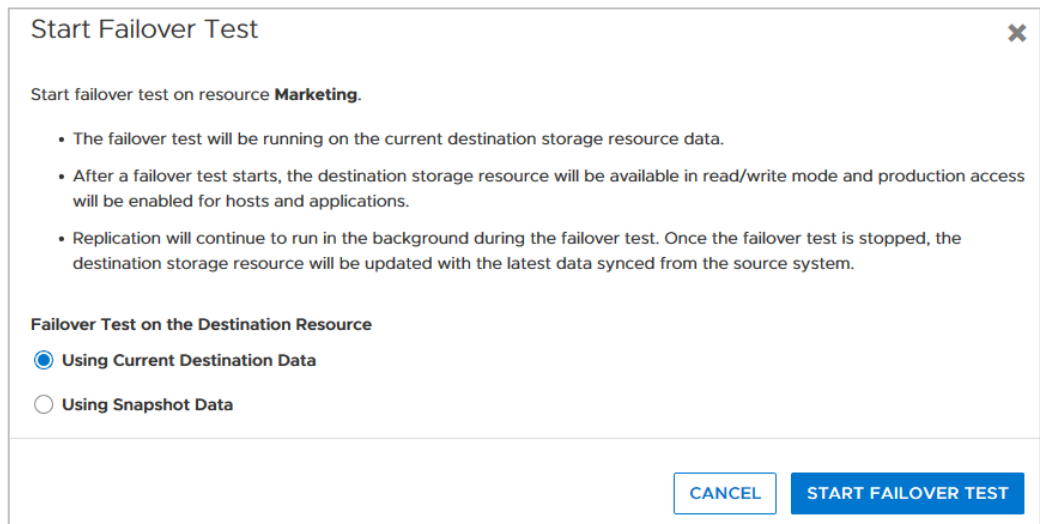


Figure 43. Select destination resource for failover test

As soon as the failover test starts, the storage resource changes to Read/Write for the mapped host. While failover test is activated, test data writes are stored in the mapped volume and replication continues in 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 44). 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.

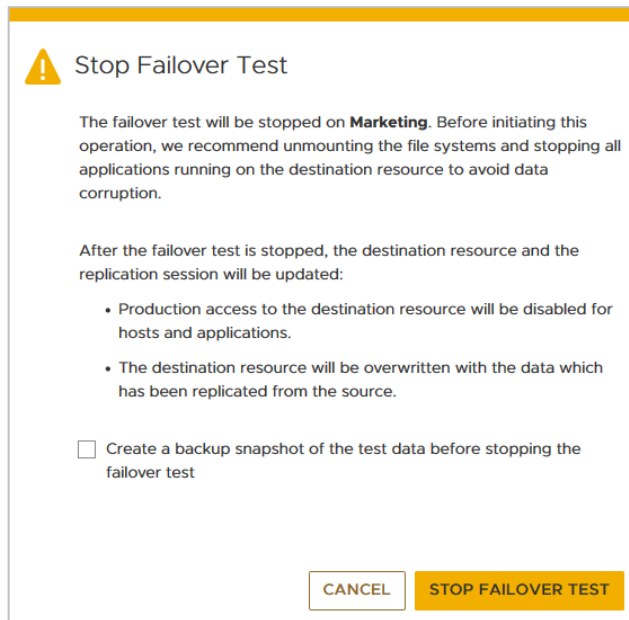


Figure 44. Stop Failover Test

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 has to confirm the following dialog (Figure 45).

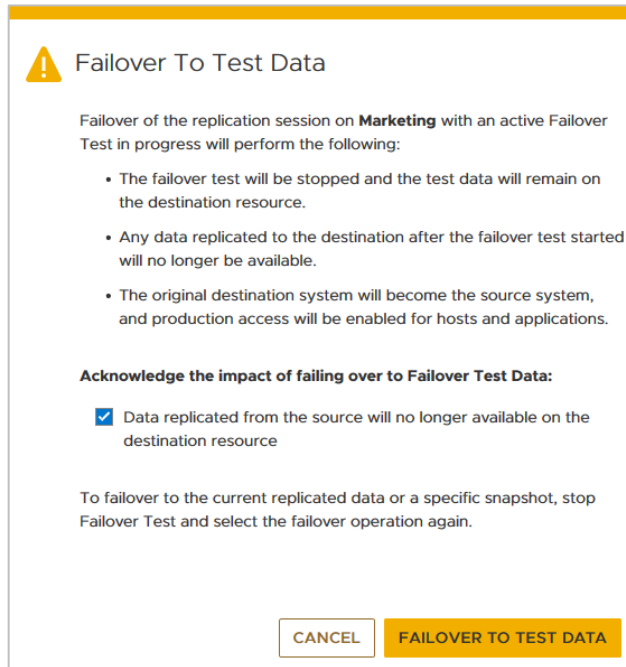


Figure 45. The Failover to Test Data dialog

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 clicks **MORE ACTIONS > Clone**. A new name is provided, and then the user selects the file systems that they would like created 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**.

Create Clone

File Interface Details needs to be configured explicitly after creating the clone, as this configuration is not being done in the NAS Clone Creation process.

NAS Server Name (Required)
CloneCorpServer

| | | | |
|------------------|-----|------------------------|---------|
| NFS Enabled | Yes | System Size | 31.5 TB |
| SMB Enabled | Yes | Total file system size | 0 GB |
| Kerberos Enabled | No | | |

3 Items, 2 selected

| Name | FLR_Enabled | Size |
|---|-------------|------|
| <input checked="" type="checkbox"/> Marketing | No | -- |
| <input checked="" type="checkbox"/> Engineering | No | -- |
| <input type="checkbox"/> ISO | No | -- |

CANCEL CREATE

Figure 46. 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 standalone 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 non-replicated 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. This ensures that 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 47. Modify destination NAS server IP address

Supported replication configurations

The PowerStore native asynchronous replication features allow supported storage resources to be replicated remotely between systems. Supported system configurations are:

Table 4. Supported system configuration for asynchronous replication

| Source | Target | Block | File |
|--------------------|--------------------|-------|------|
| PowerStore T model | PowerStore T model | ✓ | ✓ |
| PowerStore T model | PowerStore X model | ✓ | |
| PowerStore X model | PowerStore T model | ✓ | |
| PowerStore X model | PowerStore X model | ✓ | |

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

The native asynchronous 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
4. Many-to-one: A system topology in which multiple systems replicate their respective resources to a single system

Figure 48 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 48. System-level asynchronous 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 in always 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 48, 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

When upgrading the PowerStoreOS, the replication sessions are paused and show the status **Paused for NDU** (see Figure 49). The replication resumes after the upgrade has successfully finished.

The screenshot shows a 'Replication' management window 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 49. Session status of paused for nondisruptive upgrade (NDU)

Metro Volume

Introduction

This feature allows synchronous replicated active-active block volumes spanned across two PowerStore clusters running PowerStoreOS 3.0 or later. See the *Dell PowerStore: Metro Volume* white paper for more information.

Licensing

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

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 *PowerStore: VMware Site Recovery Manager Best Practices* white paper for VMware Site Recovery Manager 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 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. Combinations of all of the above

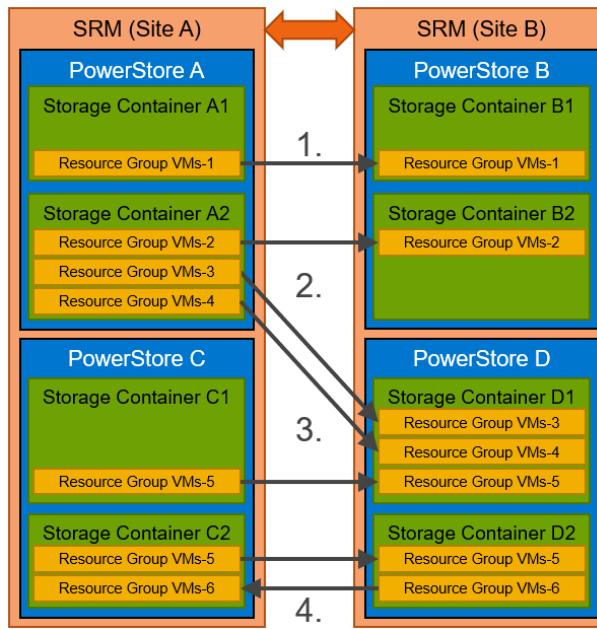


Figure 50. Supported replication flows

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, a manual synchronization is executed for all vVols in the replication group covered by the replication session.

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 white paper *Dell PowerStore: VMware Site Recovery Manager Best Practices*.

System limits

System limits For the most up-to-date system limits, see the Simple Support Matrix, available on dell.com/powerstoredocs.

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.

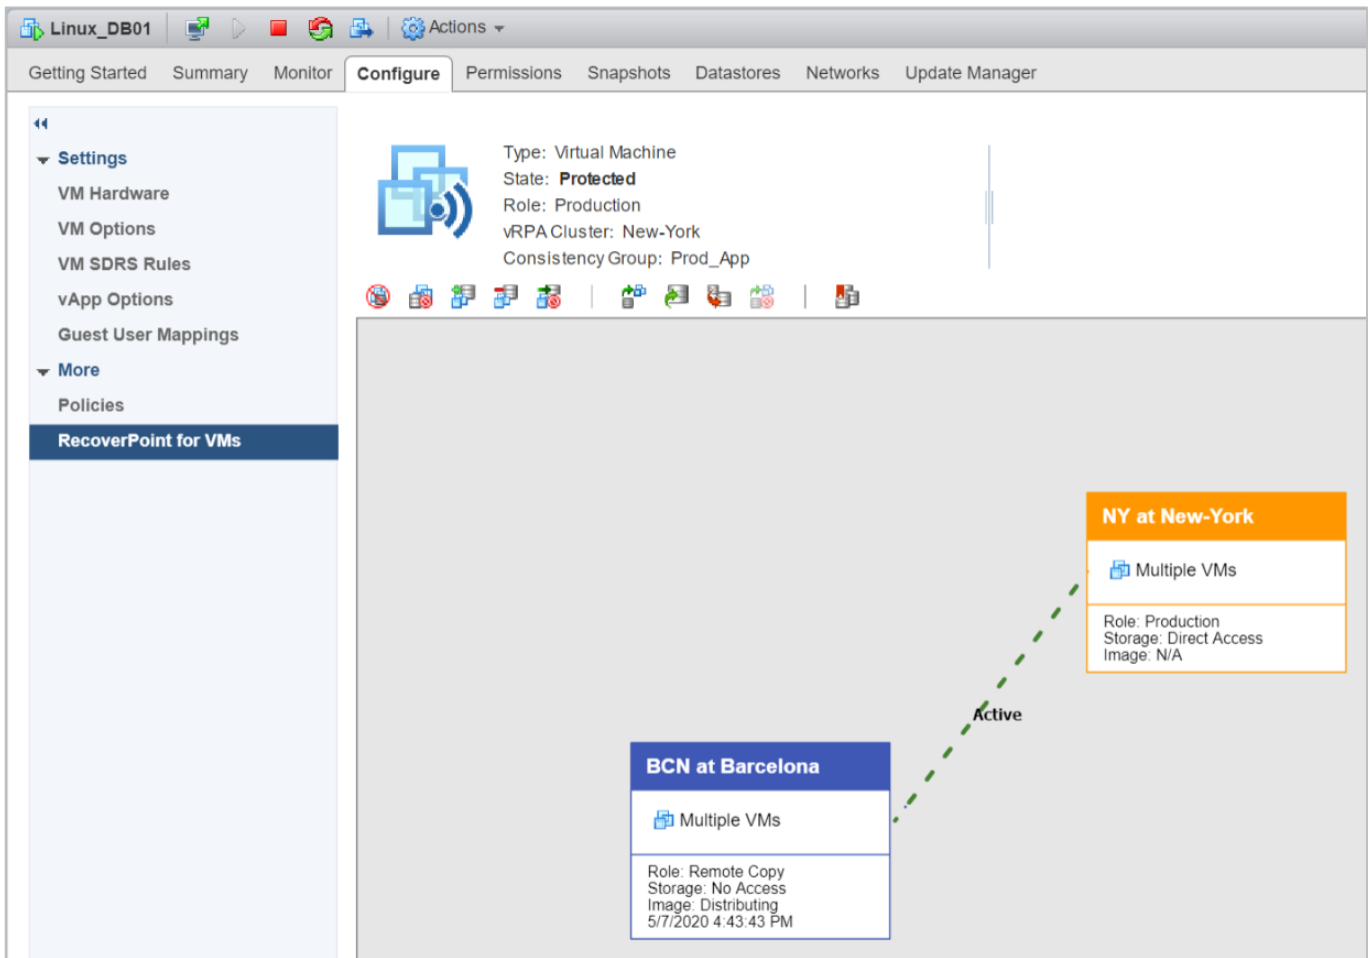


Figure 51. RecoverPoint for Virtual Machines

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.

In combination 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

Conclusion

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.

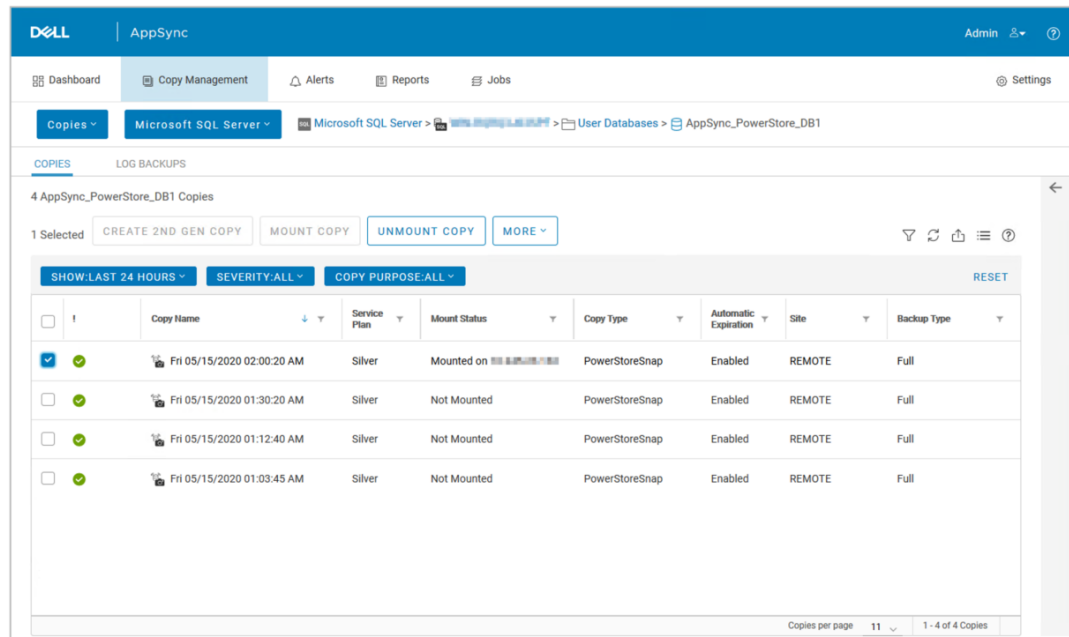


Figure 52. Dell AppSync

Metro node

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

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. For more information about metro-node, go to [Dell Support](#), and see the *Dell metro node best practices* white paper.

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 in conjunction 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 workload 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 5 outlines asynchronous replication support across Dell storage platforms.

Table 5. Asynchronous replication support

| Source | Destination | Asynchronous block | Asynchronous file ¹ | Asynchronous vVol ^{1,2} | RecoverPoint for VMs |
|--------------------------|--------------------------|--------------------|--------------------------------|----------------------------------|----------------------|
| PowerStore T models | PowerStore T models | ✓ | ✓ | ✓ | ✓ |
| PowerStore X models | PowerStore X models | ✓ | ✗ | ✓ | ✓ |
| PowerStore T or X models | PowerStore T or X models | ✓ | ✗ | ✓ | ✓ |
| PowerStore T or X models | Dell Unity | ✗ | ✗ | ✗ | ✓ |
| Dell Unity | PowerStore T or X models | ✗ | ✗ | ✗ | ✓ |

(1) PowerStore T model requires PowerStoreOS 3.0 or later

(2) PowerStore X model requires PowerStoreOS 3.2 or later

Table 6 outlines synchronous replication support for block storage resources across Dell storage platforms.

Table 6. Synchronous replication support for block storage resources

| Source | Destination | Metro Volume ¹ | Dell Metro node |
|--------------------------|--------------------------|---------------------------|-----------------|
| PowerStore T models | PowerStore T models | ✓ | ✓ |
| PowerStore X models | PowerStore X models | ✗ | ✓ |
| PowerStore T or X models | PowerStore T or X models | ✗ | ✓ |
| PowerStore T or X models | Dell Unity | ✗ | ✓ |
| Dell Unity | PowerStore T or X models | ✗ | ✓ |

(1) In a vSphere Metro Storage Cluster configuration, PowerStoreOS 3.0 or later

Appendix B: Technical support and resources

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

White papers related to PowerStore data protection:

- [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.