

Dell EMC PowerStore: Replication Technologies

Abstract

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

January 2021

Revisions

Date	Description
April 2020	Initial release: PowerStoreOS 1.0
May 2020	Minor updates
August 2020	Minor updates
January 2021	Metro node updates

Acknowledgments

Author: Robert Weilhammer

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.

Acknowledgments

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [1/18/2021] [Technical White Paper] [H18153]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	4
Executive summary.....	6
Audience	6
1 Introduction.....	7
1.1 PowerStore overview.....	7
1.2 Terminology	8
2 Native asynchronous replication for block.....	10
2.1 Licensing.....	10
2.2 Theory of operation	10
2.2.1 Port configuration for asynchronous replication	10
2.2.2 Replication connection	11
2.2.3 Remote systems	12
2.2.4 Protection policies with replication rules.....	12
2.2.5 Replication session.....	14
2.3 Replication operations	16
2.3.1 Assign protection policy with replication	17
2.3.2 Pause and resume	17
2.3.3 Synchronize now	18
2.3.4 Planned failover.....	18
2.3.5 Unplanned failover.....	19
2.3.6 Reprotect	19
2.3.7 Unassign protection policy with replication.....	20
2.4 Supported replication configurations	20
2.5 Asynchronous replication in PowerStore Manager	21
2.5.1 Configuring replication.....	22
2.5.2 Storage network IPs and replication ports.....	22
2.5.3 Remote systems	24
2.5.4 Creating a protection policy with replication rule	26
2.5.5 Assign protection policy.....	27
2.5.6 Viewing the replication sessions.....	30
2.6 System limits.....	33
2.7 Interoperability	34

- 2.7.1 Other Dell EMC data protection solutions34
 - 2.8 Upgrades34
- 3 Integration with PowerStore35
 - 3.1 RecoverPoint for Virtual Machines35
 - 3.2 AppSync36
 - 3.3 Metro node.....37
- 4 Conclusion.....38
- A Replication maximum values.....39
- B Replication support across platforms40
- C Technical support and resources41

Executive summary

Dell EMC™ PowerStore™ provides native and non-native solutions to protect data and help organizations meet business goals for both data availability and protection. PowerStore provides native asynchronous replication solutions that can replicate data to other systems, whether they are at the same site or a remote facility. Remote replicas of data protect against outages on the main system. They 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 block replication
- Dell EMC RecoverPoint™ for Virtual Machines
- Dell EMC VPLEX™

Asynchronous block replication 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 also provides a visual representation of the configuration.

Dell EMC 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 EMC 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 to many other storage systems.

Dell EMC VPLEX offers continuous application data availability and transparent data mobility for block storage. VPLEX is placed into the data path between host and creates a flexible storage architecture.

Audience

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

1 Introduction

Data is one of the most valuable assets to an organization. It is accessed constantly by various applications and users, and is sometimes accessed directly by customers. This occurrence makes data a crucial part of day-to-day operations. Outages can occur at any time and can be restricted to a single system or an entire data center or location. Whether they are planned outages such as regular maintenance, or unplanned 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 which synchronizes data to a remote system within the same site or a different location. Replicating data helps to provide data redundancy and safeguards against storage system failures at the main production site. Having a remote disaster recovery (DR) site protects against system and site-wide outages. 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 once 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, which is measured in units of time, that may be lost due to an outage. This delta of time also affects the amount of data which 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.

1.1 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.

1.2 Terminology

Asynchronous replication: Replication method which 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 which 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): Unified snapshots that are created by the system and are 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.

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: 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: Read-write copy of a thin block storage resource (volume, volume group, or VMware vSphere VMFS datastore) that shares blocks with the parent resource.

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.

PowerStore CLI: Tool which can be installed on an operating system to manage a PowerStore system.

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 which a user provisions. It represents a SCSI logical unit.

Volume group: Storage instance which 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.

2 Native asynchronous replication for block

This section discusses the PowerStore native asynchronous replication feature. This feature allows users to create replication sessions for block storage resources between PowerStore systems. Supported storage resources for native asynchronous block replication are volumes, volume groups, and thin clones. The replication itself uses iSCSI through Ethernet (LAN) connections. All configuration and management operations in this section are shown in PowerStore Manager, but the PowerStore CLI and REST API may also be used. The following subsections discuss these topics:

- Licensing requirements for the native asynchronous replication feature
- How native asynchronous replication feature works
- Configurations supported for asynchronous replication
- PowerStore Manager configuration and management

2.1 Licensing

Asynchronous replication is supported on all PowerStore systems and is included at no extra cost.

2.2 Theory of operation

Native asynchronous replication has several physical and software components. Each of these components is discussed in the following sections. To fully configure asynchronous replication between systems, perform the following steps:

1. Check network configuration for iSCSI asynchronous replication traffic.
2. Configure the remote system connection between PowerStore arrays.
3. Configure replication sessions between systems that are based on protection policies and replication rules.

This section also discusses the various block resources which support asynchronous replication, and the replication modes. The following subsections outline the different functions and requirements, and how these components interact with each other. Configuration and management of these components is completed using PowerStore Manager.

2.2.1 Port configuration for asynchronous 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. The system uses the same storage network for host access to storage resources. Tagged ports for asynchronous remote replication can be modified in PowerStore Manager and are same for both nodes in PowerStore appliance. Tagging replication ports in PowerStore Manager is not related to VLAN tagging on network infrastructure. For asynchronous replication, replication is performed over Ethernet ports available on the system.

On PowerStore, the ports that are listed in Table 1 can be used for replication. When a 25GbE optical 4-port card or IO module is used, both 10Gb or 25Gb SFPs can be leveraged for replication.

Table 1 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 or 10/25 GbE optical
PowerStore T model	✓	✓
PowerStore X model	✓	✗

The figures in section 2.2.2 show examples of minimal cabling for replication between PowerStore T models (Figure 1), and between PowerStore T models and PowerStore X models (Figure 2). The link aggregated ports provide high availability, maximum throughput, and load balancing of replication traffic across physical ports in the aggregation. It is suggested to only tag replication interfaces on ports of the same type and speed. A successful asynchronous replication connection requires all asynchronous replication ports on a source system to see or route to all asynchronous replication ports on the destination system, and conversely.

When planning for replication using the default configuration, consider that the ports are also used for other traffic like I/O for iSCSI host access, migration, and file (PowerStore T 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 which was given in ICW or added in the networking section of PowerStore Manager.

2.2.2 Replication connection

Figure 1 shows an example configuration of an asynchronous replication connection between two physical systems. In the 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 1 shows cabling for a pair of PowerStore T model appliances using system bond0.




Figure 1 Native asynchronous replication using two PowerStore T model systems

Figure 2 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 2 Native asynchronous replication using PowerStore T model and PowerStore X model

2.2.3 Remote systems

Once the replication interfaces are created and cabled to the network on both systems, the remote system connection between the arrays can be made. Once a remote system is configured on one of the systems participating in replication, it is automatically created on the peer system. The verify and update operation is used to update the replication connection information about the system that it is issued on. This operation is performed on the replication connection itself, as opposed to an individual replication session. Verify and update can be used to test a replication connection to a remote system or update the replication information if changes to the system have been made. Verify and update should be issued to reestablish the replication connection to a remote system after an outage. A use case for verify and update is if a storage network IP address pool has been changed on the system.

2.2.4 Protection policies with replication rules

Remote replication between PowerStore systems uses policy-based protection. Asynchronous replication configuration is defined in replication rules. Protection policies allow the user to configure remote and local protection using replication, snapshot rules, or both. The policies combine one or more rules to fulfill the protection requirements for a storage resource on PowerStore. For a valid configuration, a protection policy must contain at least one protection rule, regardless if 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.

For minimal RPO compliance issues, replication cycles are scheduled at 50% of the RPO value and based on the hour. For example, a one-hour RPO leads to a replication event every 30 minutes to ensure 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. The events for the RPO are based on the configured time and not on the amount of data which 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 continues 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 3.

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




Figure 3 Replication scheduler events at 30-minute intervals for a one-hour RPO

2.2.5 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 volumes 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 storage resources in a replication session are also replicated in chronological order to the destination during initial and continuous synchronizations.

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

- Writes to a storage resource are saved to the source storage resource and acknowledged to the host before being replicated to the destination storage resource. Changes are retrieved using a snapshot-diff 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 that the user is willing to lose in a disaster, which is measured in time. 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 occurs. If replication is configured when a new resource is created, the synchronization is quick because no data must be copied to the destination storage resource. If a protection policy with replication is added to an existing storage resource, a full synchronization occurs between the source and destination storage resource. Writes occurring during the initial-synchronization period are not copied to the destination storage resource, but remain in the snapshot diff for a later synchronization.

Once the initial synchronization is complete, 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. 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 4 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 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.

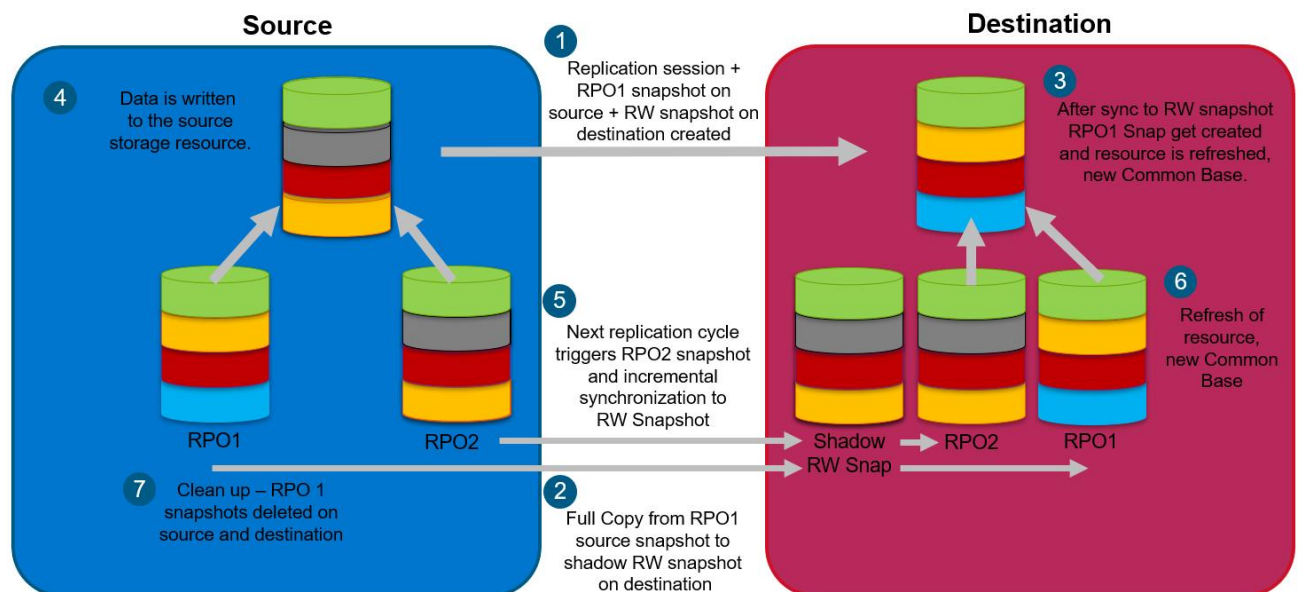


Figure 4 Asynchronous replication theory

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

Asynchronous replication operates in the same way for volumes and thin clones 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 which 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.

2.3 Replication operations

Several operations are available to manipulate replication sessions as needed. Not all operations are always available, because some depend 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.

2.3.1 Assign protection policy with replication

It is easy to configure asynchronous replication on Dell EMC PowerSystems by assigning an existing protection policy to a storage resource. In the **Storage Resource Overview**, click **MORE ACTIONS > Assign Protection Policy** (see Figure 5).

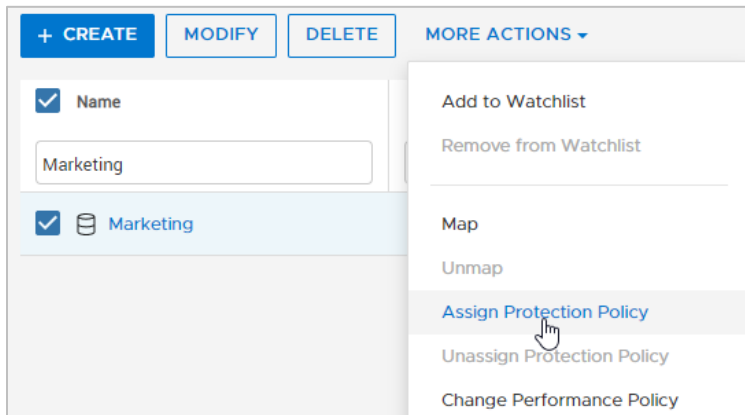


Figure 5 Assign Protection Policy

2.3.2 Pause and resume

The **PAUSE** and **RESUME** functions can stop and start replication between the resources for a particular replication session (see Figure 6). In PowerStore Manager, the pause operation is issued from the source or destination system. If the session is paused while a sync 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 synchronizations to the destination storage resource continue 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.

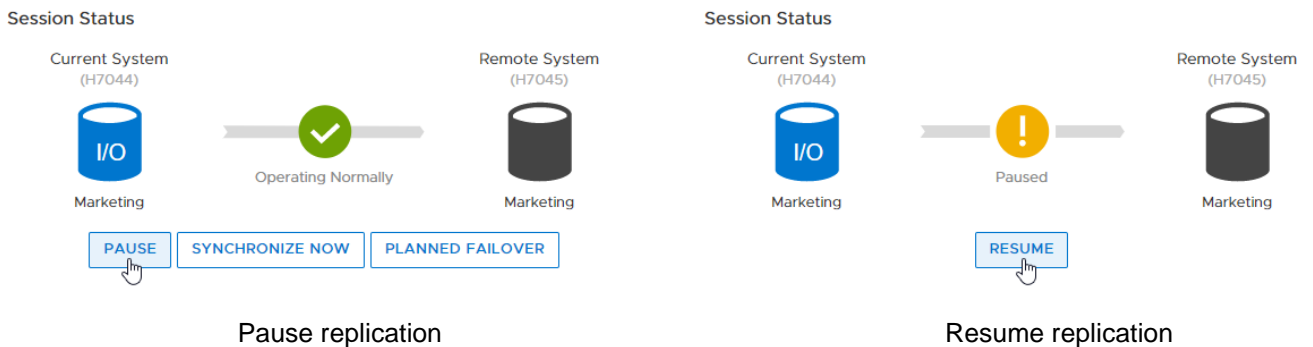


Figure 6 Pause and resume replication

2.3.3 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 active and an update is not occurring, a **SYNCHRONIZE NOW** operation can be issued to synchronize the latest changes to the destination resource (see Figure 7). After the sync operation is selected, all data that has changed since the last update is copied to the destination storage resource. Issuing a manual sync operation also updates the destination image size when the source image size has been changed.

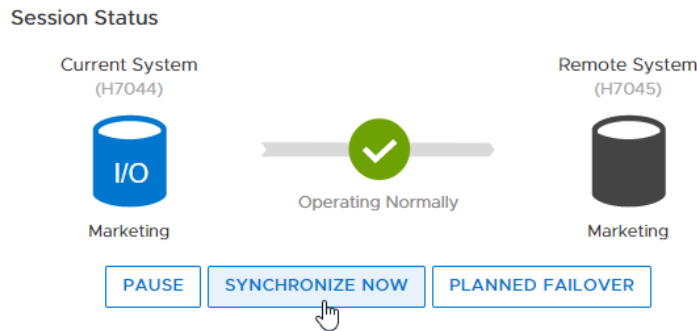


Figure 7 Synchronize Now

2.3.4 Planned failover

With asynchronous replication, the **PLANNED FAILOVER** option allows a failover without losing data (see Figure 8). The replication session fails over after completing a synchronization between the images. 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 makes a short period of data unavailable 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 with sync. When issuing a planned failover, it is suggested to quiesce I/O to the source image first. After the synchronization completes, the destination storage resource is available for production I/O and the original source no longer allows read/write I/O. If host access is configured on the destination resource, hosts can access the data currently. If **Reprotect after failover** is not selected when initiating the failover, replication does not resume in either direction when a planned failover is used.

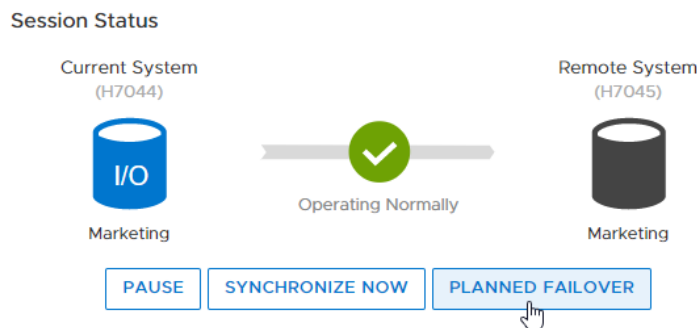


Figure 8 Planned Failover

2.3.5 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 available. When **FAILOVER** is selected on a destination resource of a replication session (Figure 9), 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 to the original source when the source becomes available.

PowerStore allows initiating an unplanned failover operation 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 section 2.3.4).

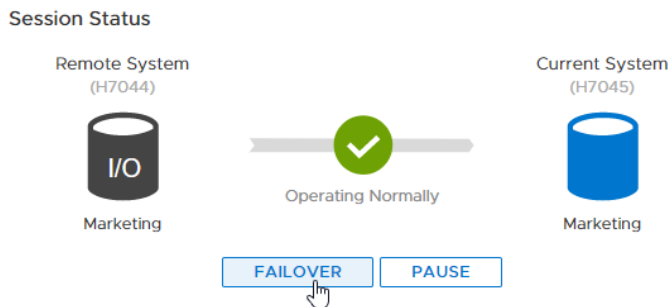


Figure 9 Unplanned Failover

2.3.6 Reprotect

After the Planned Failover or Failover option is used, the **REPROTECT** option (Figure 10) 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.

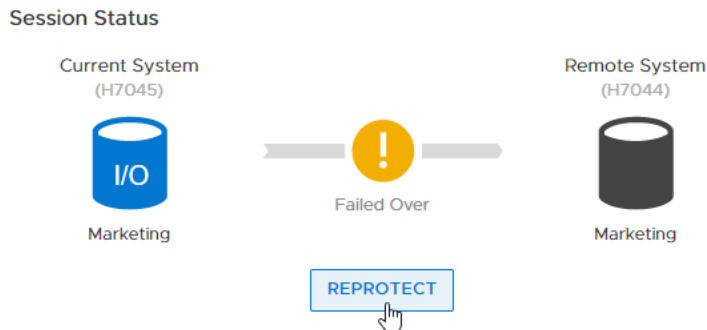


Figure 10 Reprotect

2.3.7 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. Figure 11 shows the option to **Unassign Protection Policy**. When there are no configuration issues and a detach 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.

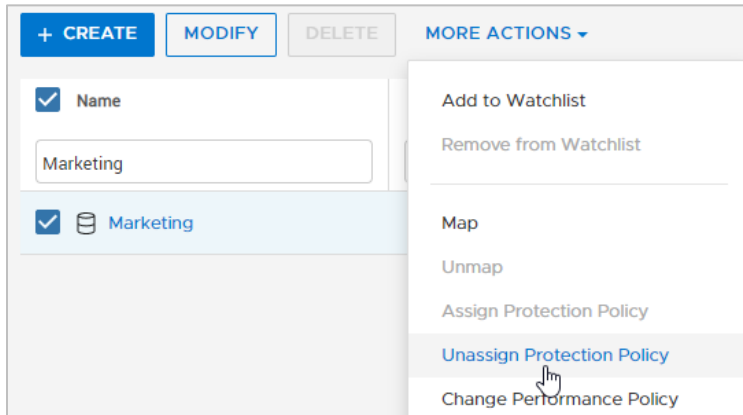


Figure 11 Unassign Protection Policy

2.4 Supported replication configurations

The PowerStore native asynchronous replication features allow supported storage resources to be replicated remotely between systems. This section outlines the supported configurations for asynchronous replication. For more information about which systems are supported for asynchronous replication, see appendix B.

The native asynchronous replication feature is supported in many different topologies, and 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 12 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.

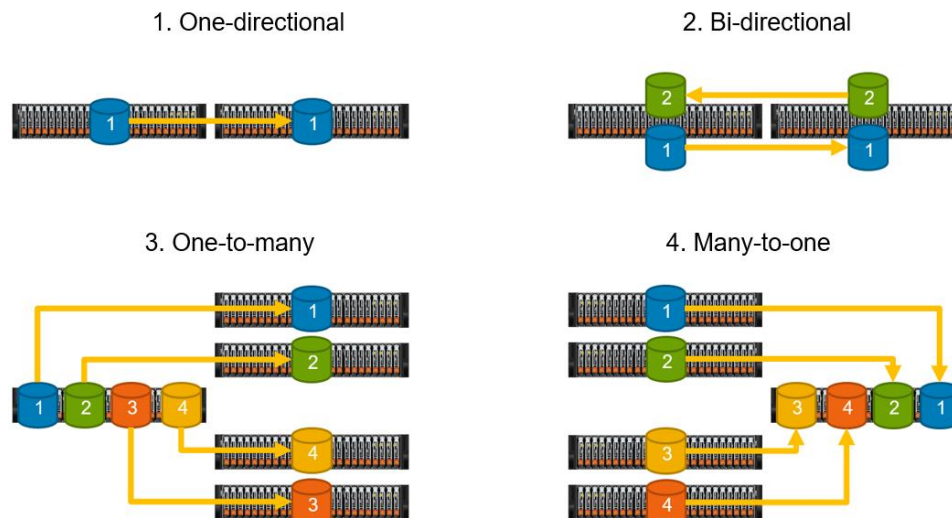


Figure 12 System-level asynchronous replication topologies

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 usually 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 12, 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.

2.5 Asynchronous replication in PowerStore Manager

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 quickly by IT generalists or advanced users alike. Replication can also be configured using the PowerStore Manager CLI or REST API. For more information about configuring and managing replication using PowerStore Manager CLI, see to the *Dell EMC PowerStore Manager Command Line Interface Guide* on the [PowerStore Info Hub](#).

For more information about REST API, use SwaggerUI (<https://<PowerStore>/swaggerui>) or see the *Dell EMC PowerStore REST API Programmer's Guide* found on the [PowerStore Info Hub](#).

2.5.1 Configuring replication

When remote asynchronous replication is configured, the following components must be available. See the corresponding sections listed for more details.

- Storage network IPs and replication ports (section 2.5.2)
- Remote systems connection (section 2.5.3)
- Protection policy with replication rule (section 2.5.4)
- Assign protection policy (section 2.5.5)

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.

2.5.2 Storage network IPs and replication ports

The replication ports may be shared network ports for host I/O or import, and replication-related data using the storage network ports. The systems come with tagged replication ports on Port0 and Port1 of the 4-port card (see section 2.2.1). No configuration change is required if planning to use these ports.

When new ports besides the default ports are used, it is recommended to check the available storage IP addresses before creating new interfaces. To verify the settings for storage network IPs, click **Settings > Networking > Network IPs**. Ensure that at least two storage network IPs are unallocated for mapping new storage network ports. To tag new replication ports in PowerStore Manager, click **Hardware > Appliance-Name > Ports** tab. All Ethernet ports are eligible to be tagged as replication ports and are available in the ports list.

Figure 13 shows the PowerStore Manager Ports page with the default tagging. The figure also shows the default Link Aggregation ports that are initially created on the system and are tagged for replication. From this page, ports can be mapped to the storage network and the tagging of replication data interfaces can be changed.

MAP STORAGE NETWORK

MORE ACTIONS

showing 8 of 26 Items

<input type="checkbox"/> Node-Module-Name ↑	Link St... (1) true	Link Aggregation - Name	Link Aggreg...	Link Aggreg...	Mapped for Stor...	Tagged for Replication
<input type="checkbox"/> BaseEnclosure-NodeA-4PortCard-FEPort1 BaseEnclosure-NodeA-4PortCard-FEPort0		BaseEnclosure-No...	LACP		✓	✓
<input type="checkbox"/> BaseEnclosure-NodeA-EmbeddedModule...		--	--	--	--	--
<input type="checkbox"/> BaseEnclosure-NodeA-IOModule0-FEPort0		--	--	--	⊖	--
<input type="checkbox"/> BaseEnclosure-NodeA-IOModule0-FEPort1		--	--	--	⊖	--
<input type="checkbox"/> BaseEnclosure-NodeB-4PortCard-FEPort0 BaseEnclosure-NodeB-4PortCard-FEPort1		BaseEnclosure-No...	LACP		✓	✓
<input type="checkbox"/> BaseEnclosure-NodeB-EmbeddedModule...		--	--	--	--	--
<input type="checkbox"/> BaseEnclosure-NodeB-IOModule0-FEPort0		--	--	--	⊖	--
<input type="checkbox"/> BaseEnclosure-NodeB-IOModule0-FEPort1		--	--	--	⊖	--

Figure 13 Interfaces page

To change the replication data interface to an interface other than the default system bond or vFE1 Port on port group TGT1, map a new set of interfaces to the storage network. To perform this required step on both nodes, select the interfaces and click **MAP NETWORK**. If only a single port is selected, PowerStore Manager automatically selects the corresponding port on the peer node. To finish this task, confirm the selected mapping as shown in Figure 14.

MAP STORAGE NETWORK

MORE ACTIONS

Items, 2 selected

⚠ Map Storage Network?

Are you sure you want to map a storage network to the following ports?

- BaseEnclosure-NodeA-IOModule0-FEPort0
- BaseEnclosure-NodeB-IOModule0-FEPort0

CANCEL

MAP NETWORK

<input type="checkbox"/> Node-Module-Name ↑	Link St...	Link Aggregation - Name	Link Aggreg...	Link Aggreg...	Mapped for Stor...	Tagged for Replic...
<input type="checkbox"/> BaseEnclosure-NodeA-4PortCard-FEPort1 BaseEnclosure-NodeA-4PortCard-FEPort0		BaseEnclosure-No...	LACP		✓	✓
<input type="checkbox"/> BaseEnclosure-NodeA-EmbeddedModule...		--	--	--	--	--
<input checked="" type="checkbox"/> BaseEnclosure-NodeA-IOModule0-FE...		--	--	--	⊖	--
<input type="checkbox"/> BaseEnclosure-NodeA-IOModule0-FE...		--	--	--	⊖	--
<input type="checkbox"/> BaseEnclosure-NodeB-4PortCard-FEPort0 BaseEnclosure-NodeB-4PortCard-FEPort0		BaseEnclosure-Nod...	LACP		✓	✓
<input type="checkbox"/> BaseEnclosure-NodeB-EmbeddedModule...		--	--	--	--	--
<input checked="" type="checkbox"/> BaseEnclosure-NodeB-IOModule0-FE...		--	--	--	⊖	--
<input type="checkbox"/> BaseEnclosure-NodeB-IOModule0-FE...		--	--	--	⊖	--

Figure 14 Map Storage Network

After ports are mapped to the storage network, click **MORE ACTIONS > Tag for Replication** as shown in Figure 15. In the resulting window, click **TAG PORT** to finish the configuration. The replication tag cannot be removed from all ports. A minimum of one pair of ports must be tagged for replication.

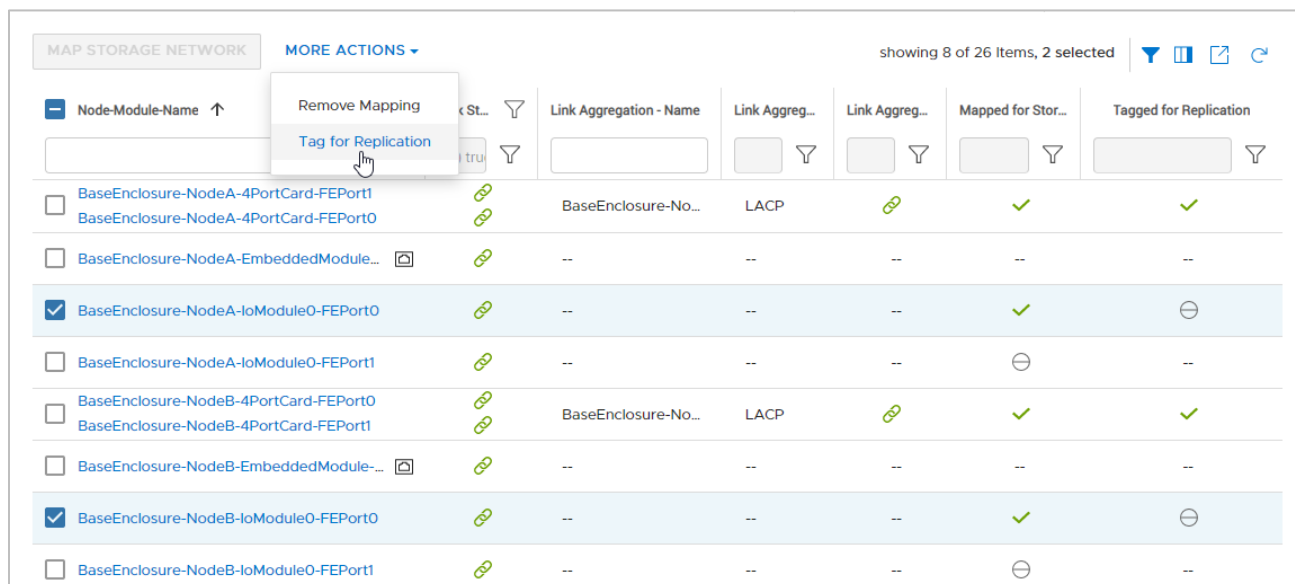


Figure 15 Tag port for replication

2.5.3 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).

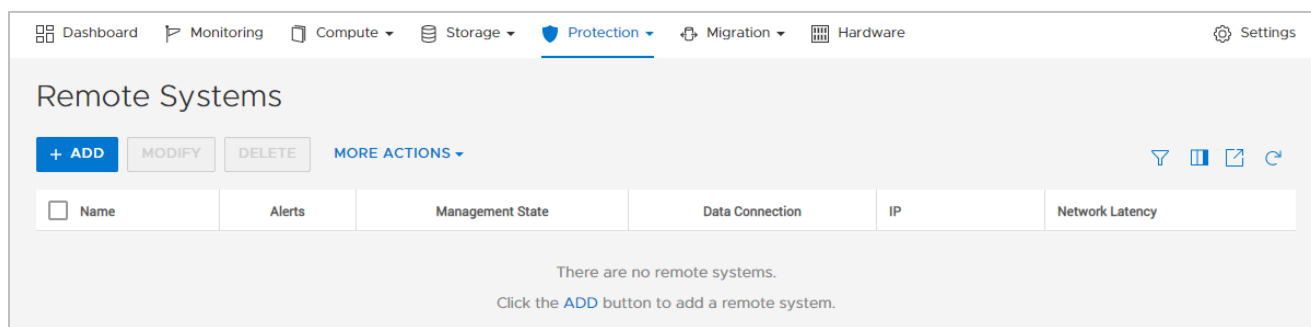


Figure 16 Remote systems replication setup

Once 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. To define new pair of remote systems, click **ADD** as shown in Figure 16.

The new Add Remote System window appears (see Figure 17) and requires following information:

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

Replication traffic can be tuned for higher efficiency depending on the expected network latency. If the network to the remote system has an expected latency below 5 milliseconds, keep the default value of **Low**. Otherwise, select **High**. Depending on the selection, different iSCSI portals with optimized buffer settings are used for the replication data traffic. The Low setting shares the host I/O iSCSI listening on port 3260, while the optimized portal for high latency uses the iSCSI portal listening on port 3261.

The provided credentials for a configured user are not stored on the system and are only used for the relationship setup. When the required fields are entered, click **ADD**. Since 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.

The screenshot shows the 'Add Remote System' window with the following fields and callouts:

- Type:** PowerStore
- Management Cluster IP Address:** A text input field with a red circle '1' next to it.
- Description (Optional):** A large text area.
- Network Latency:** A dropdown menu currently showing 'Low' with a red circle '4' next to it. A red arrow points from this field to a detailed view of the dropdown on the right.
- Username:** A text input field with a red circle '2' next to it.
- Password:** A text input field with a red circle '3' next to it.

At the bottom of the window are 'CANCEL' and 'ADD' buttons. The detailed view of the 'Network Latency' dropdown on the right shows the following options:

- Low (selected)
- High
- Low

Figure 17 Add Remote System

2.5.4 Creating a protection policy with replication rule

To create a replication session, first set up a protection policy with an underlying replication rule. A protection policy is a collection of different local or remote protection rules which is assigned to resources on the PowerStore cluster. Protection policies can contain several 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 18.

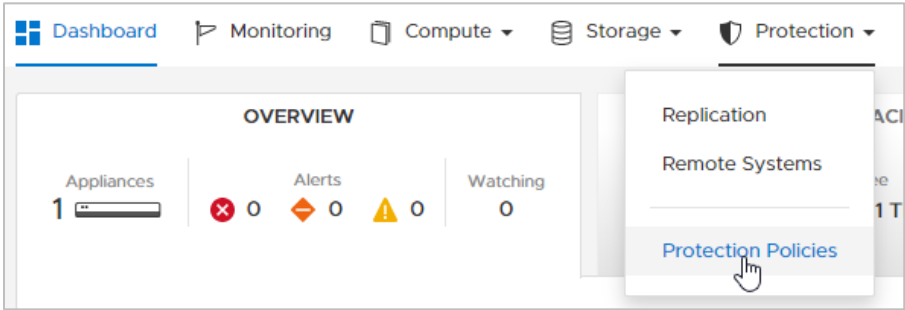


Figure 18 Protection Policies

When the Protection Policies window appears (Figure 19), 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.

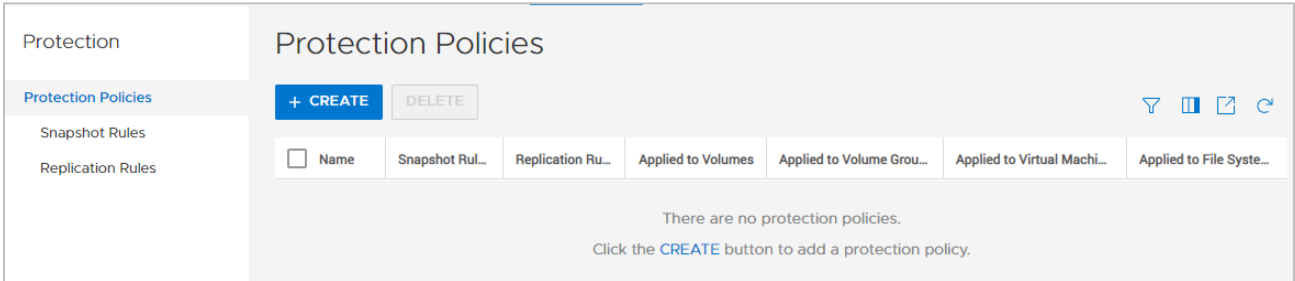


Figure 19 Protection Policies List

Since the Protection Policy is only the top-level object which is assigned to a storage resource, only a policy **Name** is required. It is recommended to use a meaningful name such as one that contains the remote system. Further down in the window, it is possible to 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 20.

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. Name
4. Destination Remote System
5. RPO
6. Alert Threshold

Figure 20 Create Replication Rule

Once all steps are finished, the protection policy can be used to protect storage resources with configured parameters.

2.5.5 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 could include a volume, volume group, or a thin clone. The following steps show assigning a protection policy on a volume. The required steps to assign a protection policy to a volume group or thin clone are the same. The 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. Since 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.

2.5.5.1 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 21).

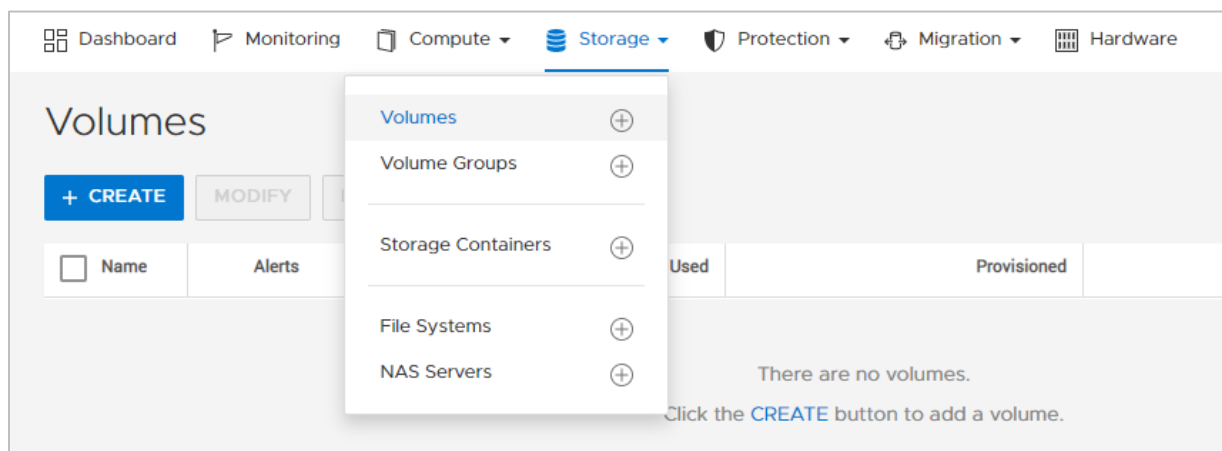


Figure 21 First step: Assign protection policy

In the **Create Volumes** window, enter the following information. Each number below corresponds with a number in Figure 22.

1. Name
2. Size
3. Quantity (if multiple volumes are created)

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. Click the **Protection Policy** item to view the assigned policy later.

Figure 22 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 source and destination system before members are added. The members do not replicate until the next manual- or RPO-scheduled synchronization.

2.5.5.2 Protect existing storage resources

Like creating a storage resource, the following steps show assigning a protection policy to an existing volume. The steps are similar for existing volume groups or thin clones also. Open the **Storage Resource** page where the volumes or volume groups are listed, and select one or more resources to assign the protection policy to. Next, click the **MORE ACTIONS** drop-down menu and click **Assign Protection Policy** as shown in Figure 23.

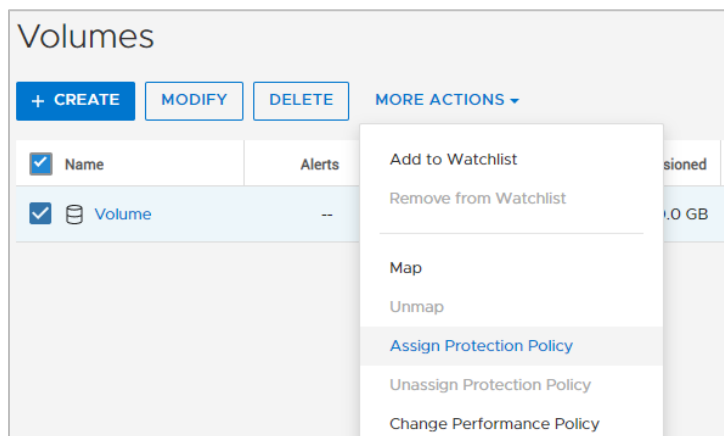


Figure 23 Assign protection policy to existing storage resource

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

2.5.6 Viewing the replication sessions

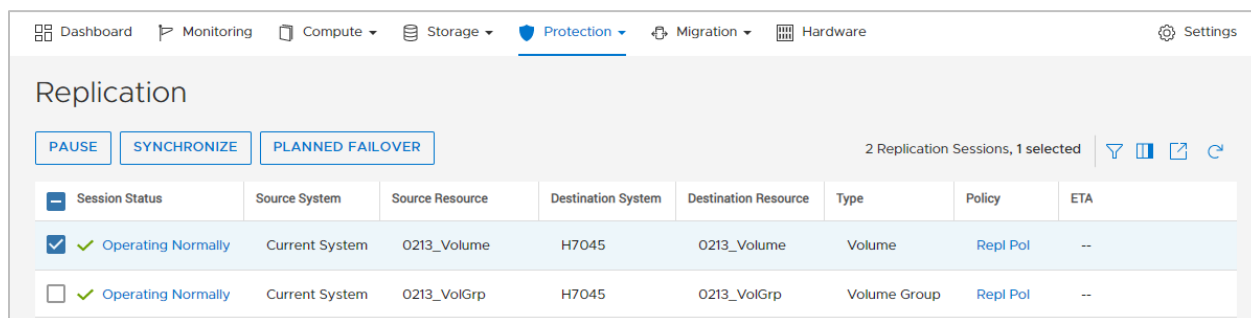
All replication sessions on the system can be viewed from the **Replication** page. To view this page in PowerStore Manager, under **Protection** click **Replication**. Figure 24 shows an example of the replication sessions overview with multiple replication sessions that are created on the system. This 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 (this shows -- when the session is not actively synchronizing)

Only one session can be selected at a time. The state of the selected session determines which buttons above the table are present. When no session is selected, the buttons are unavailable as shown in Figure 25. Figure 24 and Figure 25 show examples of the Replication page and show the differences between the source and destination.

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

- **PAUSE** to pause the replication
- **SYNCHRONIZE** to initiate a manual replication between regular RPO cycle
- **PLANNED FAILOVER** to start a planned failover



Session Status	Source System	Source Resource	Destination System	Destination Resource	Type	Policy	ETA
<input checked="" type="checkbox"/> Operating Normally	Current System	O213_Volume	H7045	O213_Volume	Volume	Repl Pol	--
<input type="checkbox"/> Operating Normally	Current System	O213_VolGrp	H7045	O213_VolGrp	Volume Group	Repl Pol	--

Figure 24 Replication window for source resource

There is an option to select if storage resource should be reprotected after the planned failover in a following window.

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

- **PAUSE** to pause the replication
- **FAILOVER** to start an unplanned failover

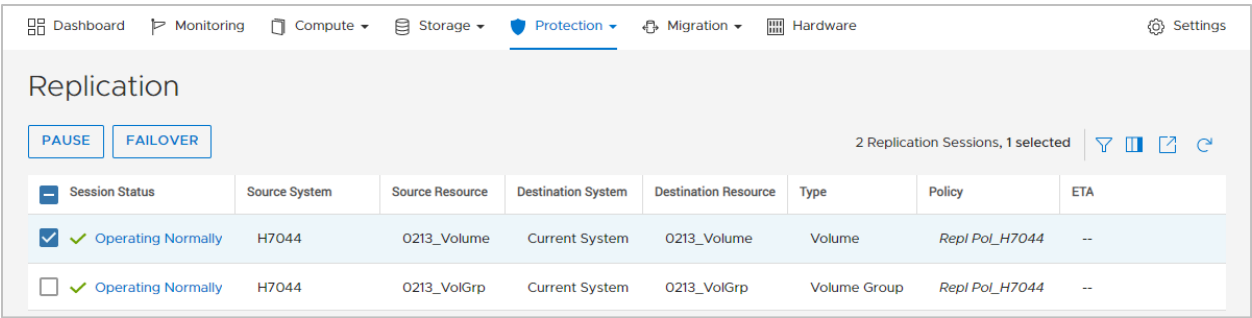


Figure 25 Replication window for destination resource

When the 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 source if a final synchronization is needed.

For a more detailed view of the replication state, it is possible to use the individual session states to view the selected replication session. This window displays a **Session Summary** as shown in Figure 26. The blue volume always references the local storage resource.

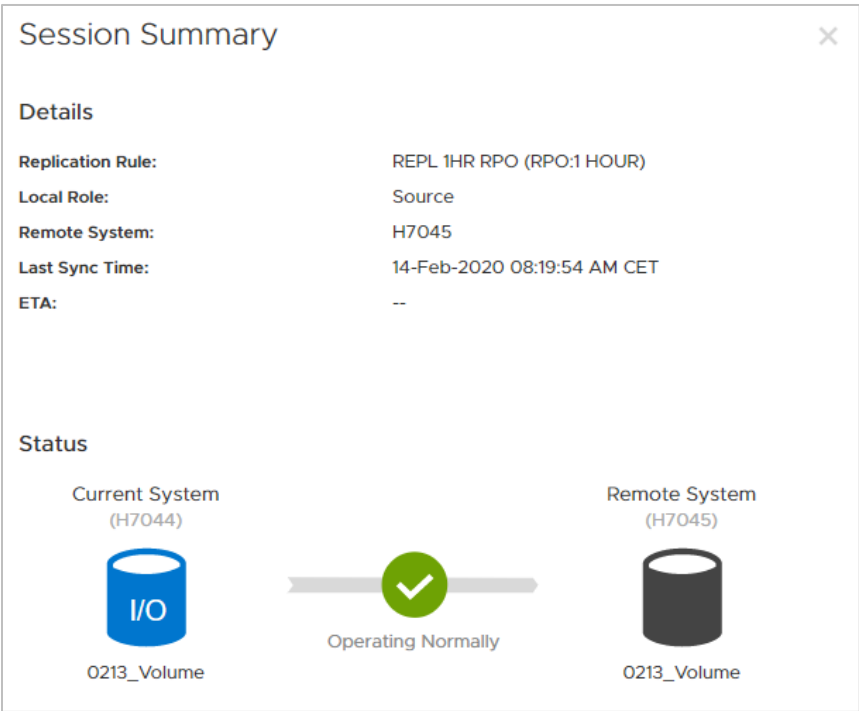


Figure 26 Replication Session Summary

When working in the Volumes or Volume Group properties pages, it is also possible to see and control the corresponding replication. Figure 27 shows the **Volumes** replication page, which looks the same for a volume group or a thin clone. To see the replication info, in the **Storage Resource** view, click the **Protection** tab > **Replication** tab in following view as shown in Figure 27.

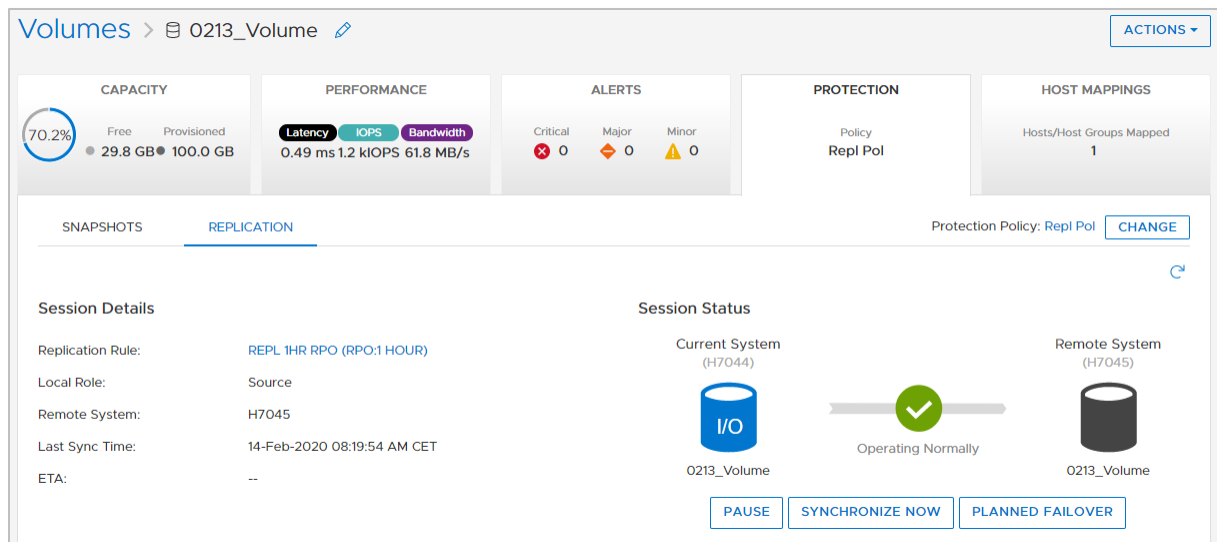


Figure 27 Storage Resource > Protection view

Besides the replication status, replication performance statistics are also available in **Performance** tab (Figure 28) of the Storage Resource and include the following:

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

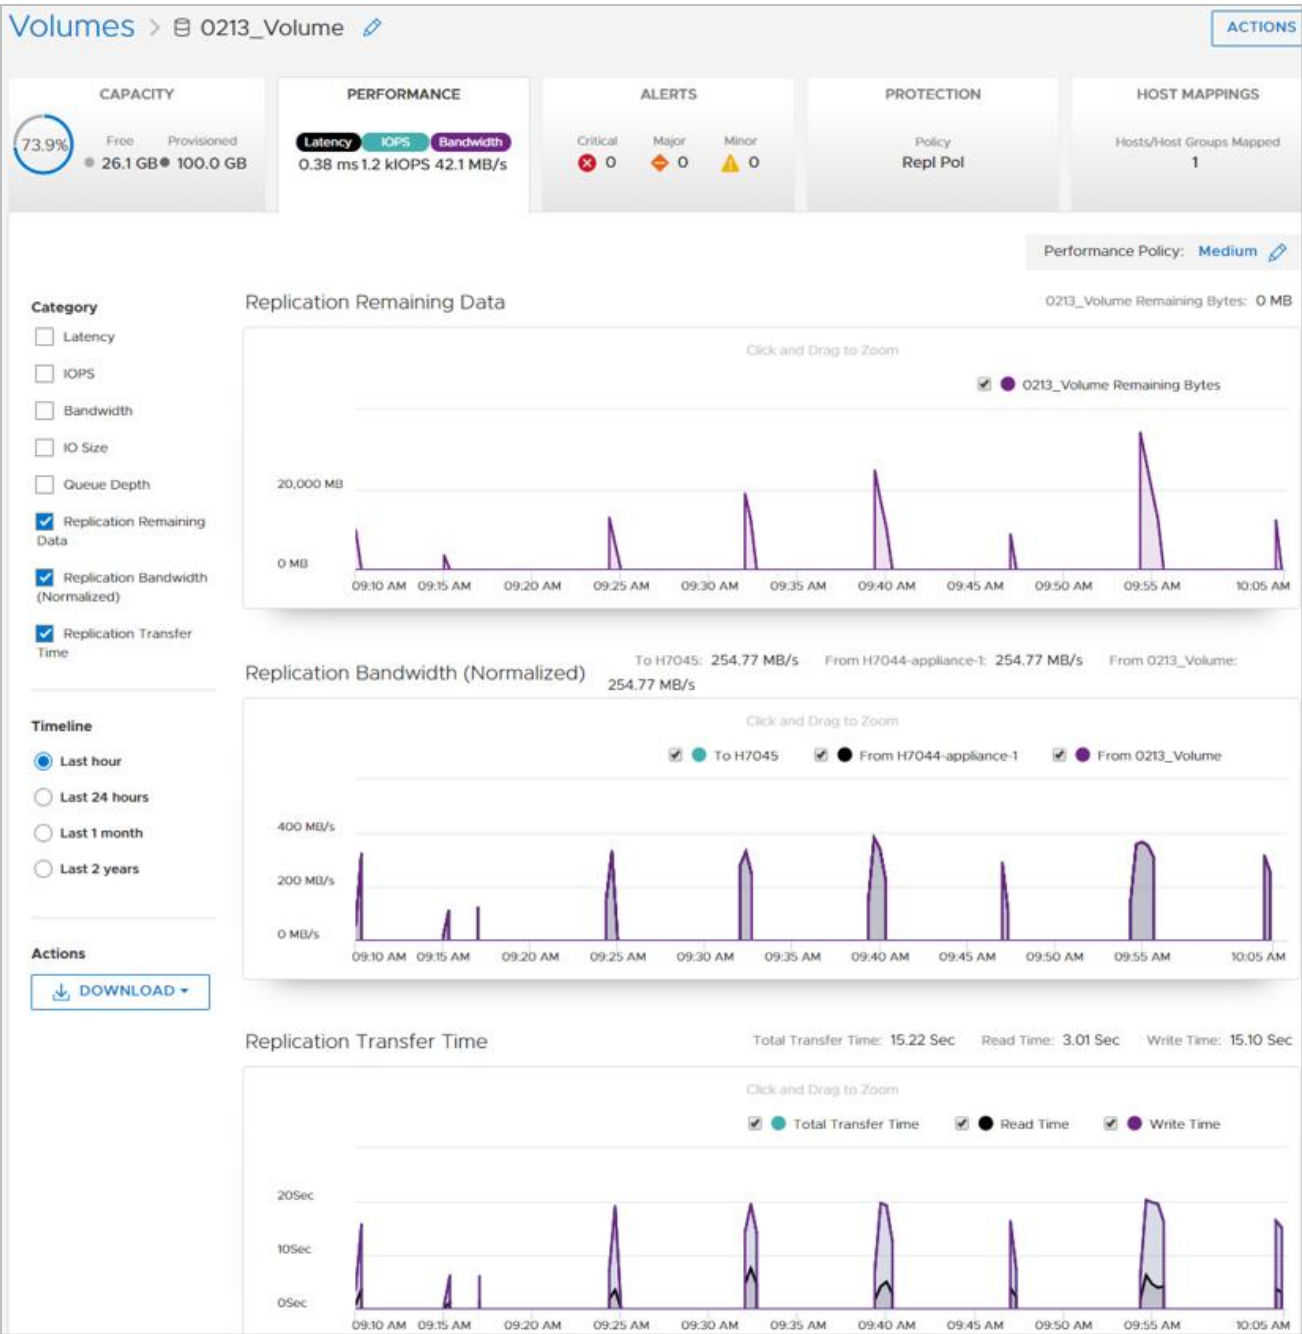


Figure 28 Replication performance view

2.6 System limits

For information about PowerStore replication maximums, see appendix A.

2.7 Interoperability

Native asynchronous replication is supported in either direction between PowerStore systems with PowerStoreOS version 1.0, regardless of the PowerStore model used. Native asynchronous replication can occur between the following systems:

- PowerStore T model and PowerStore T model
- PowerStore T model and PowerStore X model
- PowerStore X model and PowerStore X model

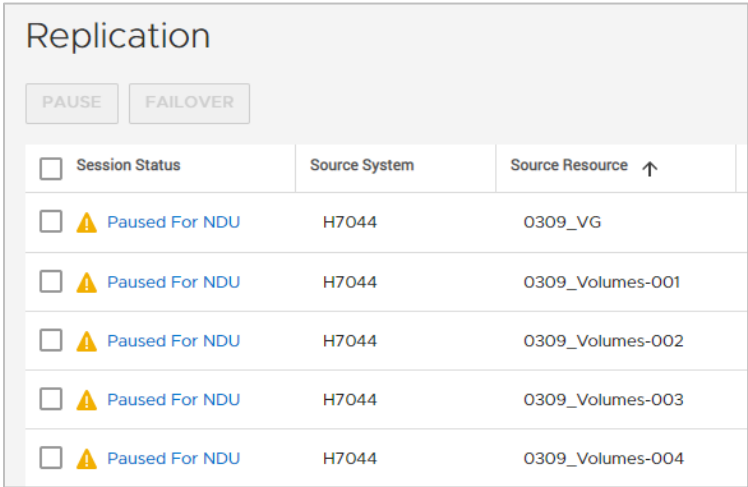
As media for replication data traffic, all Ethernet ports on the 4-port card are supported. PowerStore T models also support replication data traffic on Ethernet I/O modules.

2.7.1 Other Dell EMC data protection solutions

PowerStore can integrate into other Dell EMC data protection products like VPLEX, RecoverPoint for Virtual Machines, and Dell EMC AppSync™. All three products cover different layers for important applications. VPLEX offers 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.

2.8 Upgrades

When upgrading the PowerStoreOS, the replication sessions are paused and show the status **Paused for NDU** (see Figure 29). The replication resumes after the upgrade has successfully finished.



The screenshot shows a web interface titled "Replication". At the top, there are two buttons: "PAUSE" and "FAILOVER". Below these is a table with three columns: "Session Status", "Source System", and "Source Resource" with an upward arrow icon. The table contains five rows, each representing a paused replication session. Each row starts with a checkbox, followed by a yellow warning triangle icon and the text "Paused For NDU". The "Source System" column for all rows is "H7044". The "Source Resource" column lists different volumes: "0309_VG", "0309_Volumes-001", "0309_Volumes-002", "0309_Volumes-003", and "0309_Volumes-004".

<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 29 Session status of paused for non-disruptive upgrade (NDU)

3 Integration with PowerStore

3.1 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 which 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.

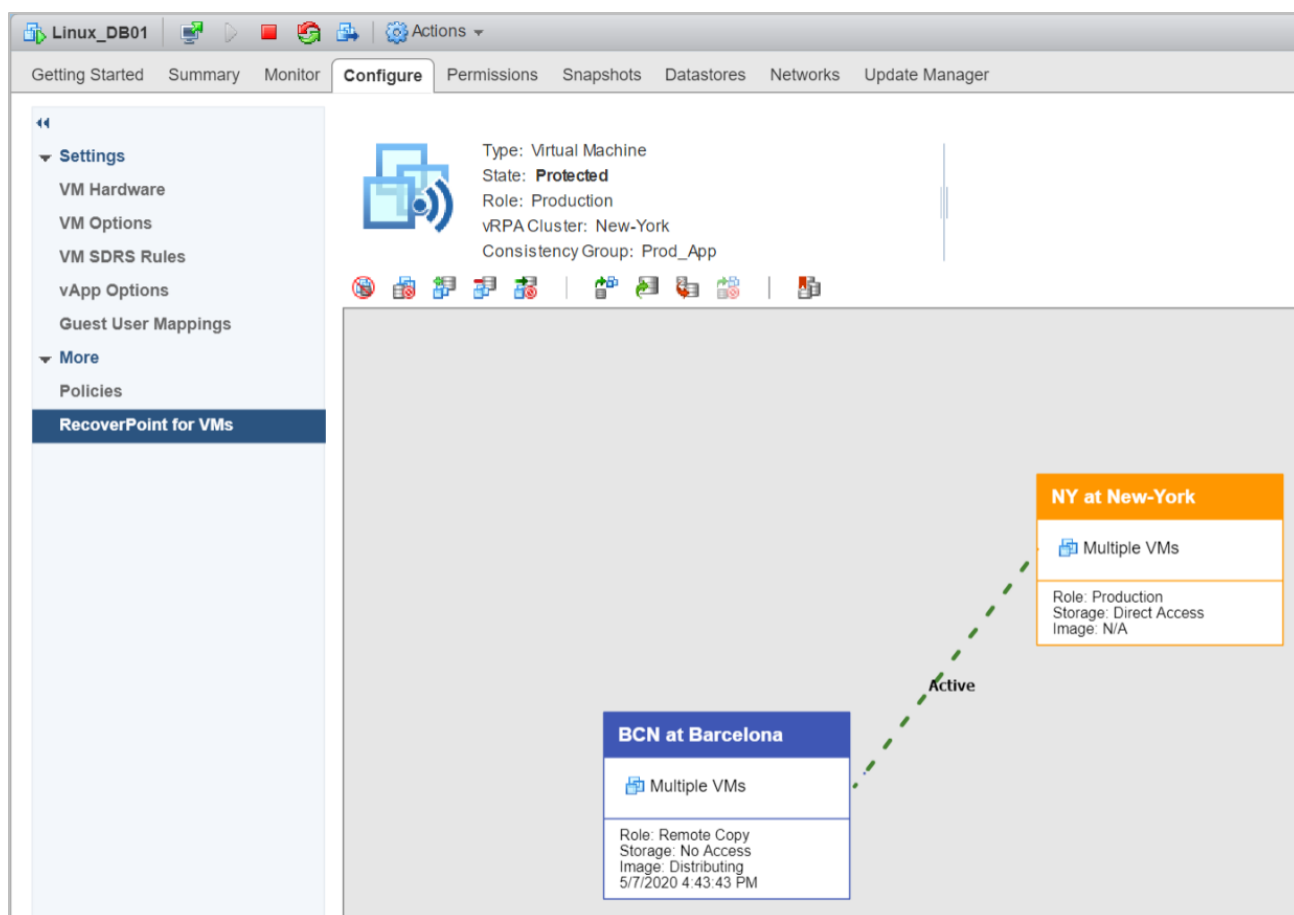


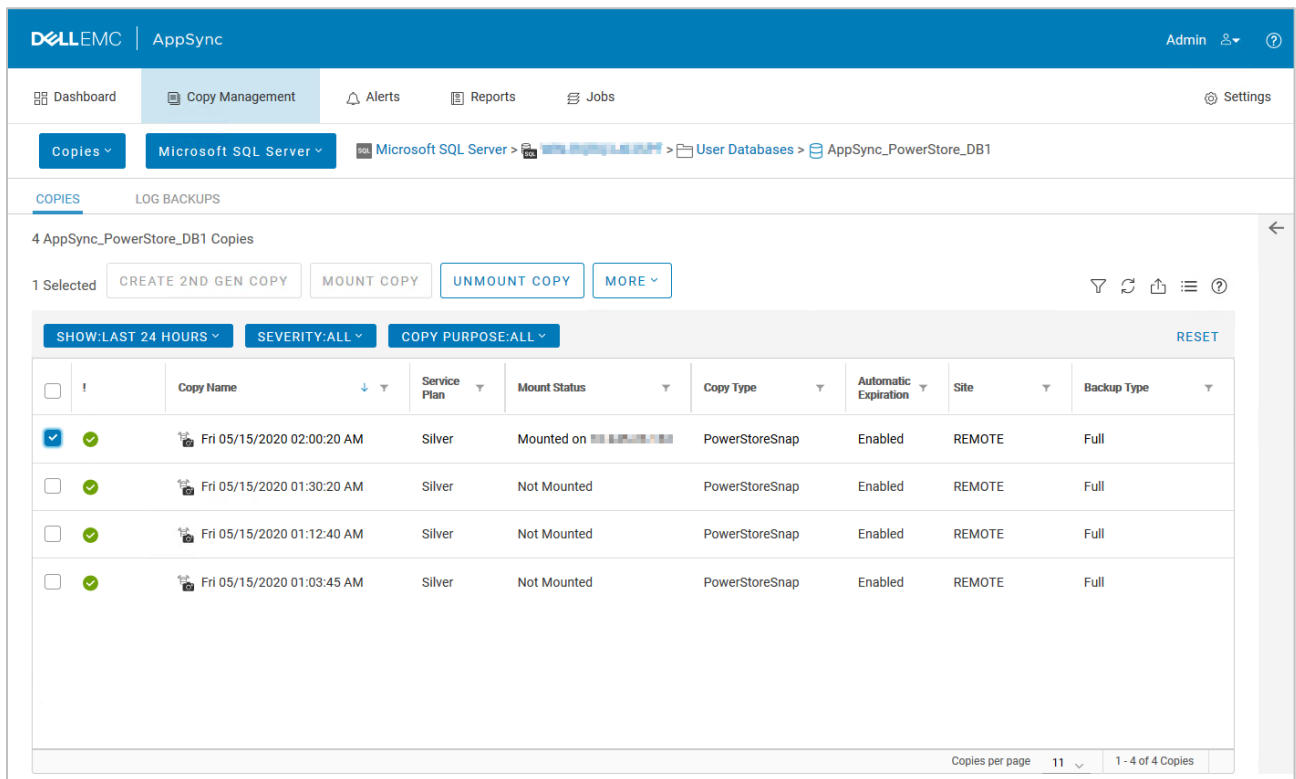
Figure 30 RecoverPoint for Virtual Machines

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

3.2 AppSync

Dell EMC 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 EMC 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 EMC storage portfolio. AppSync supports Oracle, Microsoft® SQL Server®, Microsoft Exchange®, VMware datastores, and other file systems. See the [Dell EMC 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 not support PowerStore file storage, VMware vSphere® Virtual Volumes™ (vVols), or integration with Dell EMC VPLEX®. See the *Dell EMC AppSync Simple Support Matrix* for additional support information as it becomes available.



The screenshot shows the Dell EMC AppSync web interface. The top navigation bar includes 'Dell EMC | AppSync', 'Admin', and a help icon. The main navigation menu has 'Dashboard', 'Copy Management', 'Alerts', 'Reports', 'Jobs', and 'Settings'. The 'Copy Management' section is active, showing a breadcrumb trail: 'Copies > Microsoft SQL Server > Microsoft SQL Server > User Databases > AppSync_PowerStore_DB1'. Below the breadcrumb, there are tabs for 'COPIES' and 'LOG BACKUPS'. The 'COPIES' tab is selected, showing '4 AppSync_PowerStore_DB1 Copies'. There are buttons for 'CREATE 2ND GEN COPY', 'MOUNT COPY', 'UNMOUNT COPY', and 'MORE'. A table of copies is displayed with columns: Copy Name, Service Plan, Mount Status, Copy Type, Automatic Expiration, Site, and Backup Type. The first copy is selected and mounted.

Copy Name	Service Plan	Mount Status	Copy Type	Automatic Expiration	Site	Backup Type
Fri 05/15/2020 02:00:20 AM	Silver	Mounted on	PowerStoreSnap	Enabled	REMOTE	Full
Fri 05/15/2020 01:30:20 AM	Silver	Not Mounted	PowerStoreSnap	Enabled	REMOTE	Full
Fri 05/15/2020 01:12:40 AM	Silver	Not Mounted	PowerStoreSnap	Enabled	REMOTE	Full
Fri 05/15/2020 01:03:45 AM	Silver	Not Mounted	PowerStoreSnap	Enabled	REMOTE	Full

Figure 31 Dell EMC AppSync

3.3 Metro node

Metro node is an external hardware/software add-on feature based on the Dell PowerEdge R640 platform 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 additional information about Metro node, go to [Dell Support](#) and reference the Metro node best practices white paper.

4 Conclusion

This paper discussed 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 block replication is a data-protection solution which replicates storage resources remotely to other remote PowerStore systems. Asynchronous replication uses the PowerStore snapshot technology to provide consistent point-in-time replicas which can be used in a disaster. With asynchronous replication, no impact to host I/O is seen since 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.

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.

A Replication maximum values

Table 2 outlines the PowerStore replication system maximum values. The table shows several system maximum values when using asynchronous replication. The maximum replication sessions include all replication sessions on the system. The replication destination storage resources count towards the system maximums, even though they are not host-accessible when they are a destination image.

Table 2 Replication maximum values

	PowerStore T model	PowerStore X model
Protection policies and snapshot or replication rules	32	32
max. Replication rules in a protection policy	1	1
Async replication sessions per appliance	125	125
Replicated volumes per appliance	500	500
Number of remote clusters	8	8
Snapshots per volume or thin clone/file system	256	256
Block snapshots per appliance	100,000	100,000
Snapshots per appliance	130,136	105,096
Recommended limit: max. replicated volumes with 5 min RPO	75	75

B Replication support across platforms

Table 3 outlines replication support across Dell EMC storage platforms.

Table 3 Replication support

Source	Destination	Async block replication	RecoverPoint for Virtual Machines	VPLEX
PowerStore T or X models	PowerStore T or X models	✓	✓	✓
PowerStore T or X models	Dell EMC Unity	X	✓	✓
Dell EMC Unity	PowerStore T or X models	X	✓	✓

C Technical support and resources

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

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

The [PowerStore Info Hub](#) provides detailed documentation on how to install, configure, and manage PowerStore systems.