

Dell EMC PowerStore: VMware Site Recovery Manager Best Practices

Abstract

This document offers best practices for automated disaster recovery of virtualized workloads using Dell EMC™ PowerStore™ arrays, replication, and VMware® Site Recovery Manager™.

December 2020

Revisions

Date	Description
July 2020	Initial release
October 2020	SRA remote system requirement update
December 2020	Added vmware-dr.xml file location for Photon OS based appliance

Acknowledgments

Author: Jason Boche

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. [12/14/2020] [Best Practices] [H18425]

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	7
2 Setup prerequisites	10
2.1 Storage Replication Adapter.....	10
2.2 PowerStore	10
2.3 VMware vSphere and SRM	10
3 Site Recovery Manager architecture	11
3.1 Array-based replication: single protected site	11
3.2 Array-based replication: dual protected site	12
3.3 vSphere replication: single protected site.....	13
3.4 vSphere replication: dual protected site	14
4 PowerStore Manager configuration.....	15
4.1 PowerStore Manager availability	15
4.2 PowerStore Manager logins	15
4.3 Creating dedicated SRA access accounts	16
4.4 Modifying SRM settings for larger environments.....	16
5 Replication configuration	19
5.1 Asynchronous replication	19
5.2 Remote system network latency.....	21
5.3 Snapshots and application consistency	21
5.4 Custom recovery tasks	23
6 Site Recovery Manager configuration	25
6.1 SRA installation	25
6.2 Configuring the array managers	25
6.3 Creating array pairs	28
6.4 Array manager device discovery	28
6.5 Creating placeholder datastores.....	29
6.6 Protection group considerations	29

6.7	Recovery plan considerations	30
7	Recovery plan testing and running.....	31
8	Reprotect and failback.....	34
8.1	Reprotection	34
8.2	Failback	35
9	Conclusion.....	36
A	Additional resources.....	37
A.1	Technical support and resources	37
A.2	VMware support.....	37

Executive summary

Data-center consolidation by way of x86 virtualization is a trend that has gained tremendous momentum and offers many benefits. Although the physical nature of a server is transformed once it is virtualized, the necessity for data protection remains. Virtualization opens the door to new and flexible opportunities in data protection, data recovery, replication, and business continuity. This document offers best practices for automated disaster recovery of virtualized workloads using Dell EMC™ PowerStore™, replication, and VMware® Site Recovery Manager™ (SRM).

Audience

This document is intended for IT administrators, storage architects, partners, and Dell Technologies™ employees. This audience also includes individuals who may evaluate, acquire, manage, operate, or design a Dell EMC networked storage environment using PowerStore systems.

1 Introduction

This paper provides configuration examples, tips, recommended settings, and other storage guidelines to follow while integrating VMware Site Recovery Manager (SRM) with Dell EMC PowerStore. In addition to basic configuration, this document also answers frequently asked questions about VMware interactions with Site Recovery Manager.

We recommend reading the Site Recovery Manager documentation provided on [vmware.com](https://www.vmware.com) before beginning an SRM implementation.

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.

The PowerStore platform is available in two product models: PowerStore T models and the PowerStore X models. PowerStore T models are bare-metal, unified storage arrays which can service block, file, and VMware vSphere® Virtual Volumes™ (vVols) resources along with numerous data services and efficiencies. PowerStore X model appliances enable running applications directly on the appliance through the AppsON capability. A native VMware ESXi™ layer runs embedded applications alongside the PowerStore operating system, all in the form of virtual machines. This feature adds to the traditional storage functionality of PowerStore X model appliances, and supports serving external block and vVol storage to servers with FC and iSCSI.

1.2 Terminology

The following terms are used with PowerStore.

Appliance: Solution containing a base enclosure and attached expansion enclosures. The size of an appliance could be only the base enclosure or the base enclosure plus expansion enclosures.

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.

Fibre Channel (FC) protocol: Protocol used to perform IP and SCSI commands over a Fibre Channel network.

File system: Storage resource that can be accessed through file-sharing protocols such as SMB or NFS.

Internal snapshot (replication snapshot): Unified snapshots created by the system that 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 the destination storage system. The last successful internal read-only (RO) snapshots for source and destination storage resources and are used as a common base.

iSCSI: Provides a mechanism for accessing block-level data storage over network connections.

Network-attached storage (NAS) server: File-level storage server used to host file systems. A NAS server is required to create file systems that use SMB or NFS shares.

Network File System (NFS): An access protocol that allows data access from Linux® or UNIX® hosts on a network.

PowerStore base enclosure: Enclosure containing both nodes (node A and node B) and 25 NVMe drive slots

PowerStore cluster: Multiple appliances in a single grouping. Clusters can consist of one appliance or more. Up to four PowerStore T appliances can be clustered by adding appliances as required.

PowerStore Command Line Interface (PSTCLI): Tool which can be installed on an operating system to manage a PowerStore system. It allows a user to perform tasks on the storage system by typing commands instead of using the user interface.

PowerStore expansion enclosure: Enclosures that can be attached to a base enclosure to provide additional storage.

PowerStore Manager: An HTML5 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 node: Storage controller that provides the processing resources for performing storage operations and servicing I/O between storage and hosts. Each PowerStore appliance contains two nodes.

PowerStore REpresentational State Transfer (REST) API: Set of resources (objects), operations, and attributes that provide interactive, scripted, and programmatic management control of the PowerStore cluster.

PowerStore T model: Container-based storage system that is running on purpose-built hardware. This storage system supports unified (block and file) workloads, or block-optimized workloads.

PowerStore X model: Container-based storage system that runs inside a virtual machine that is deployed on a VMware hypervisor. Besides offering block-optimized workloads, PowerStore also allows users to deploy applications directly on the array.

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 recovery plan is run. For example, an RTO of one hour requires restoring data access within one hour after a disaster is declared and the disaster recovery plan performed.

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.

Server Message Block (SMB): Network file-sharing protocol, also known as CIFS, used by Microsoft® Windows® environments. SMB is used to provide access to files and folders to Windows hosts on a network.

Snapshot: Also called a unified snapshot, a snapshot is a point-in-time view of a storage resource or data stored on a storage resource. A user can recover files from a snapshot, restore a storage resource from a snapshot, or provide snapshot data access to a host. 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.

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 the system takes automatically using asynchronous replication.

Virtual Volumes (vVols): VMware storage framework which allows VM data to be stored on individual Virtual Volumes. This ability allows data services to be applied at a VM-granularity level while using Storage Policy Based Management (SPBM).

Volume: A block-level storage device that can be shared out using a protocol such as iSCSI or Fibre Channel. 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.

vStorage API for Array Integration (VAAI): VMware API that allows storage-related tasks to be offloaded to the storage system.

vSphere API for Storage Awareness (VASA): VMware API that provides additional insight about the storage capabilities in vSphere.

2 Setup prerequisites

Verify the solution requirements listed in this section before deploying or upgrading your environment.

2.1 Storage Replication Adapter

The PowerStore Storage Replication Adapter (SRA) must be installed on each SRM server. PowerStore offers SRAs for both the Photon operating-system based SRM appliance and the Windows-based SRM installation. You can download the SRAs from the VMware website. We recommend using the most current version of the SRA to ensure optimal compatibility and available features. See the SRA release notes and product documentation to determine SRA compatibility with SRM versions.

Note: Refer to the SRA release notes for specific requirements or features noted for the SRA. For example, at the time of this publication the SRA cannot work if more than one remote system is configured on PowerStore. This means that a 1:1 appliance replication relationship must be maintained between the protected site and the recovery site.

2.2 PowerStore

SRM- and array-based replication of block volumes requires two PowerStore appliances replicating between each other in one or both directions. You can replicate virtual machines that are based on PowerStore NFS and vVols using vSphere replication, or consider using RecoverPoint for Virtual Machines as an alternative.

2.3 VMware vSphere and SRM

Compatible versions of VMware SRM, VMware vCenter™ Server, and vSphere hosts are required. To see a list of software versions required for SRM to function, check the [VMware Product Interoperability Matrix](#). SRM is supported with vCenter Server for Essentials, vCenter Server Foundation, and vCenter Server Standard.

Note: At the time of publication, PowerStore X models deploy vSphere 6.7 Update 2.

3 Site Recovery Manager architecture

This section describes array-based replication architecture for single- and dual-protected sites.

3.1 Array-based replication: single protected site

This configuration (shown in Figure 1) is generally used when the secondary site does not have any virtual machines that SRM must protect. The secondary site exists solely for disaster-recovery purposes. The infrastructure at the recovery site must be available and online to run the SRM recovery plan.

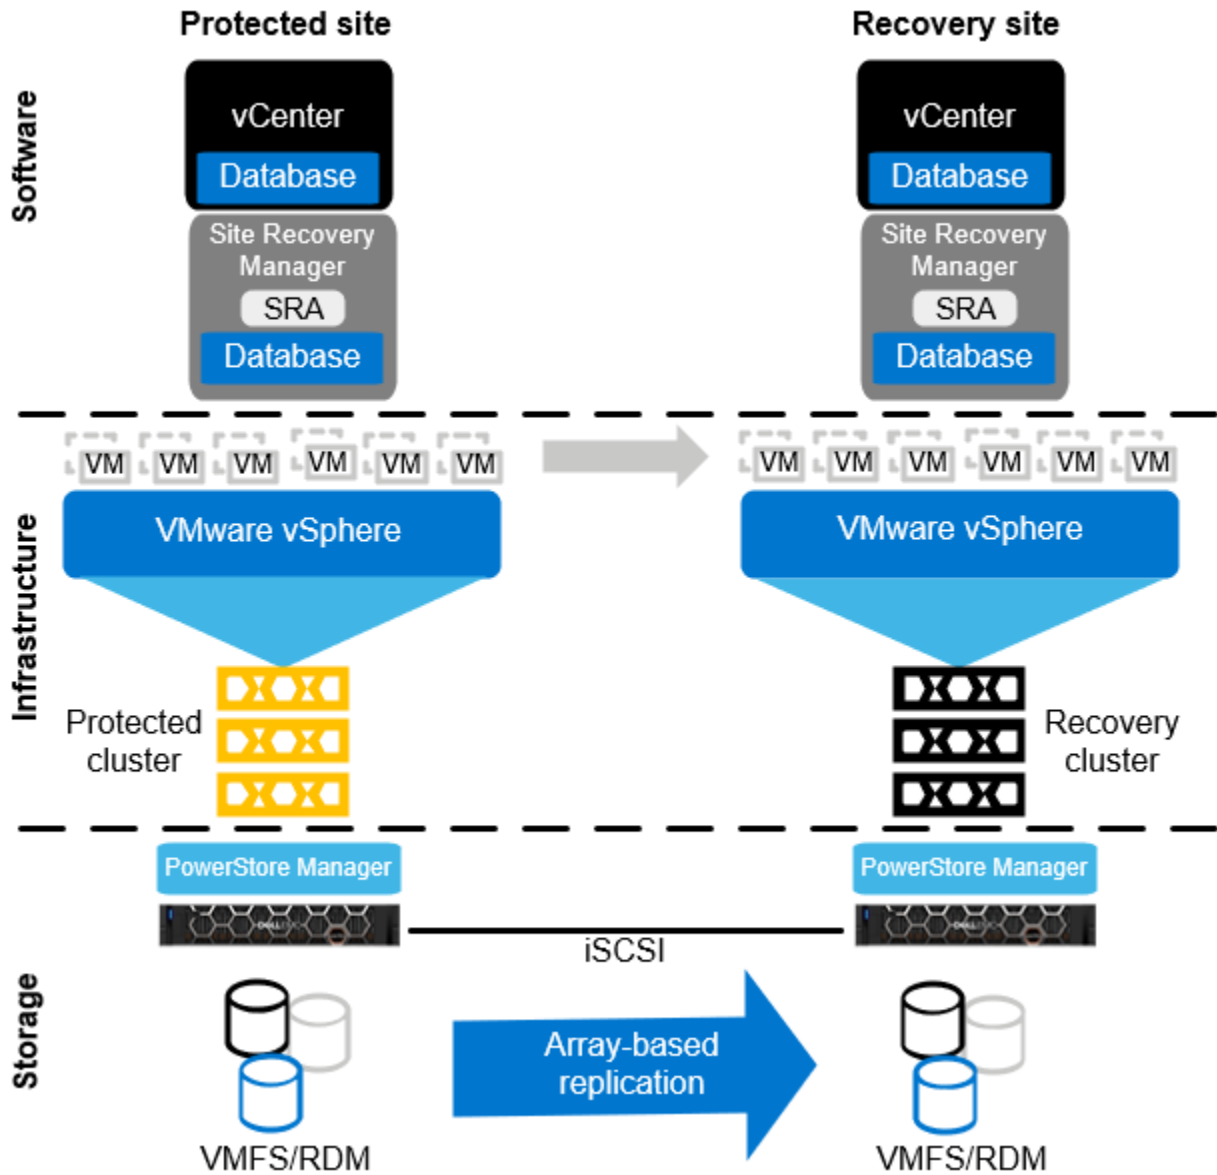


Figure 1 Architecture for a single protected site with array-based replication

3.2 Array-based replication: dual protected site

This configuration (shown in Figure 2) is generally used when both sites have virtual machines that need to be protected by SRM. Each site replicates its virtual machines to the opposing site where they can be recovered.

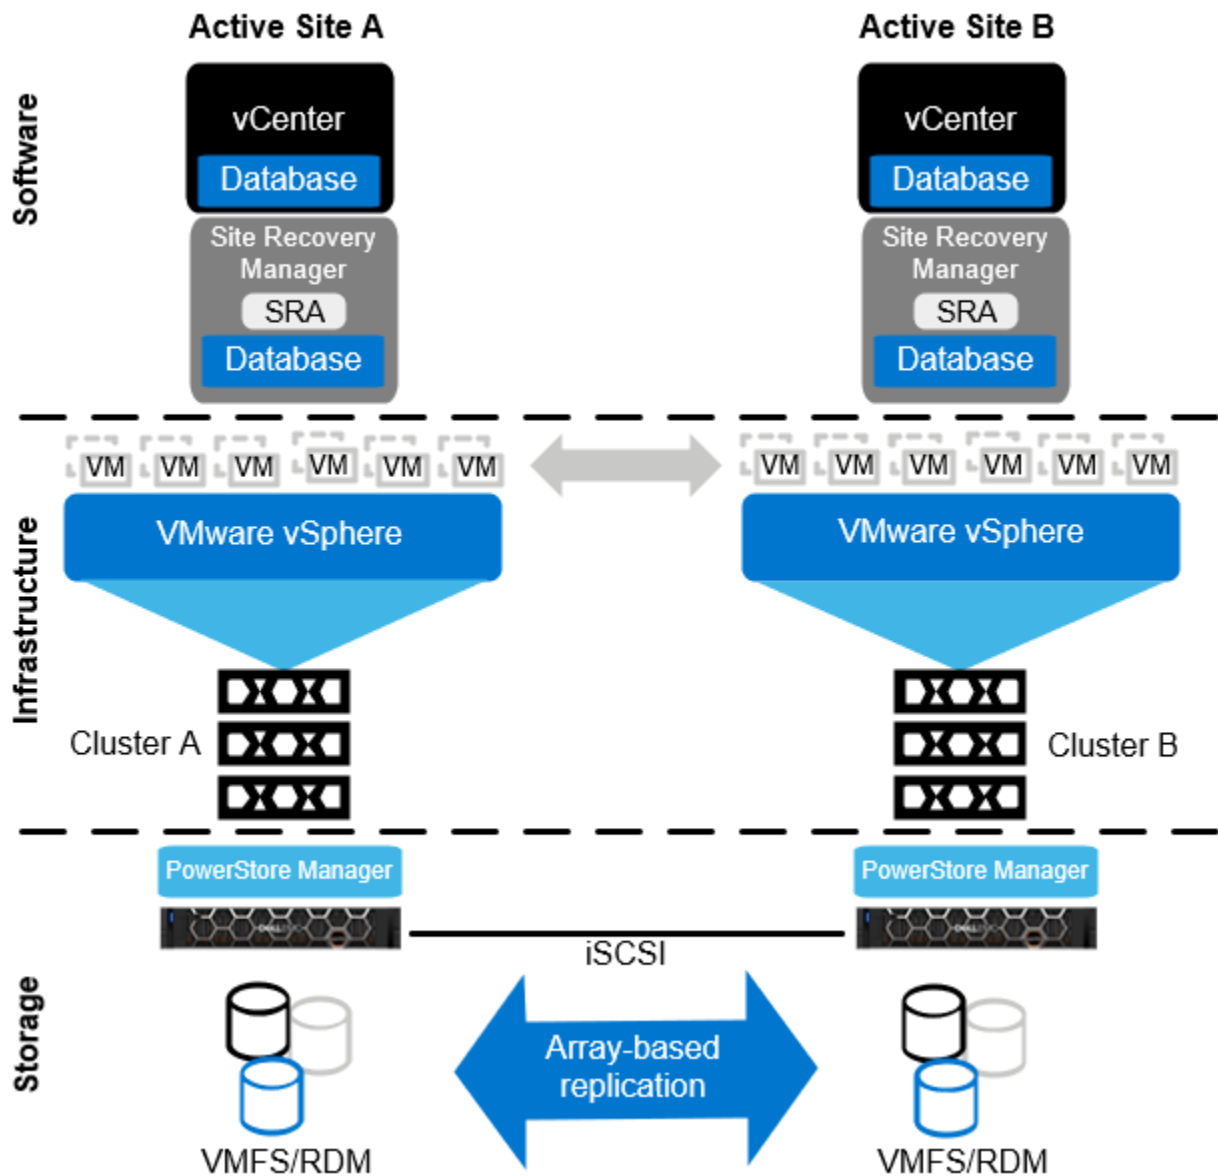


Figure 2 Architecture for a dual protected site with array-based replication

3.3 vSphere replication: single protected site

vSphere replication can be used in addition to or in place of array-based replication (see Figure 3). Here are two of the main advantages of vSphere replication over array-based replication:

- It enables a granular selection of individual powered-on VMs to be replicated instead of entire datastores of VMs.
- vSphere datastore objects abstract the underlying storage vendor, model, protocol, and type. This behavior means that replication can be carried out between different array models and protocols, including local storage.

vSphere replication, along with other feature support for vSphere replication added in SRM 5.1, makes SRM appealing and adaptable as a DR solution for organizations with storage or budget constraints.

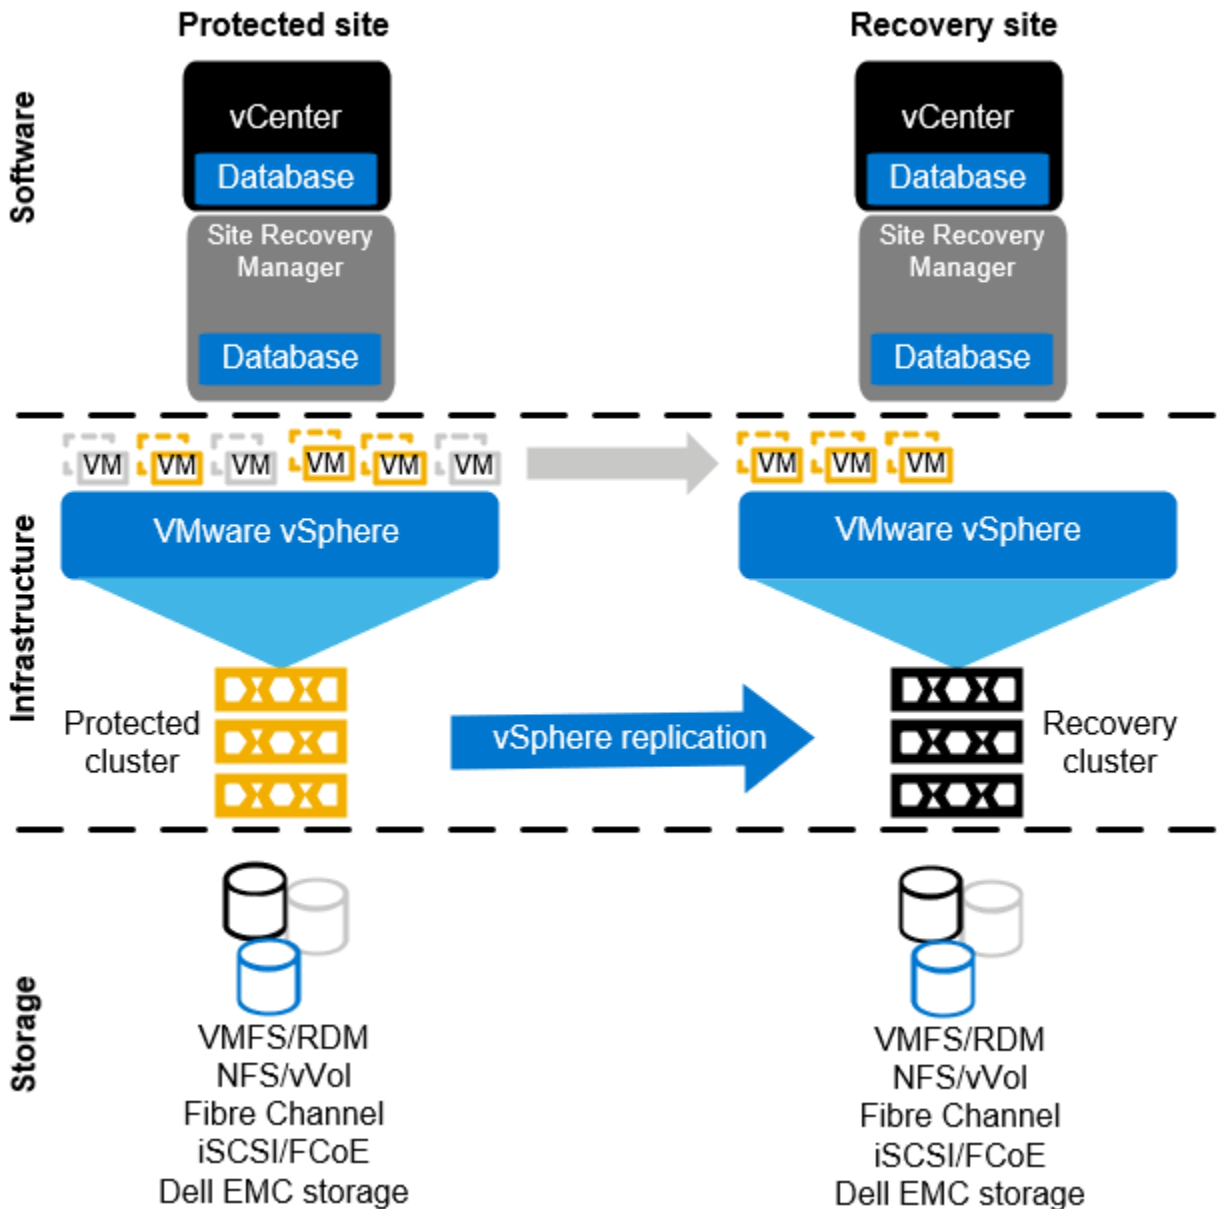


Figure 3 Configuration for a single protected site with vSphere replication

3.4 vSphere replication: dual protected site

vSphere replication also supports the active/active site model (see Figure 4). In each vSphere replication architecture diagram, replication is handled by the vSphere hosts that use the vSphere network stack. An array-based SRA is not present in vSphere replication architecture. These figures do not represent all the components of vSphere replication. A deployment of vSphere replication consists of multiple virtual appliances at each site and on each vSphere host that handles the movement of data between sites. Go to [VMware Documentation](#) for a detailed look at vSphere replication.

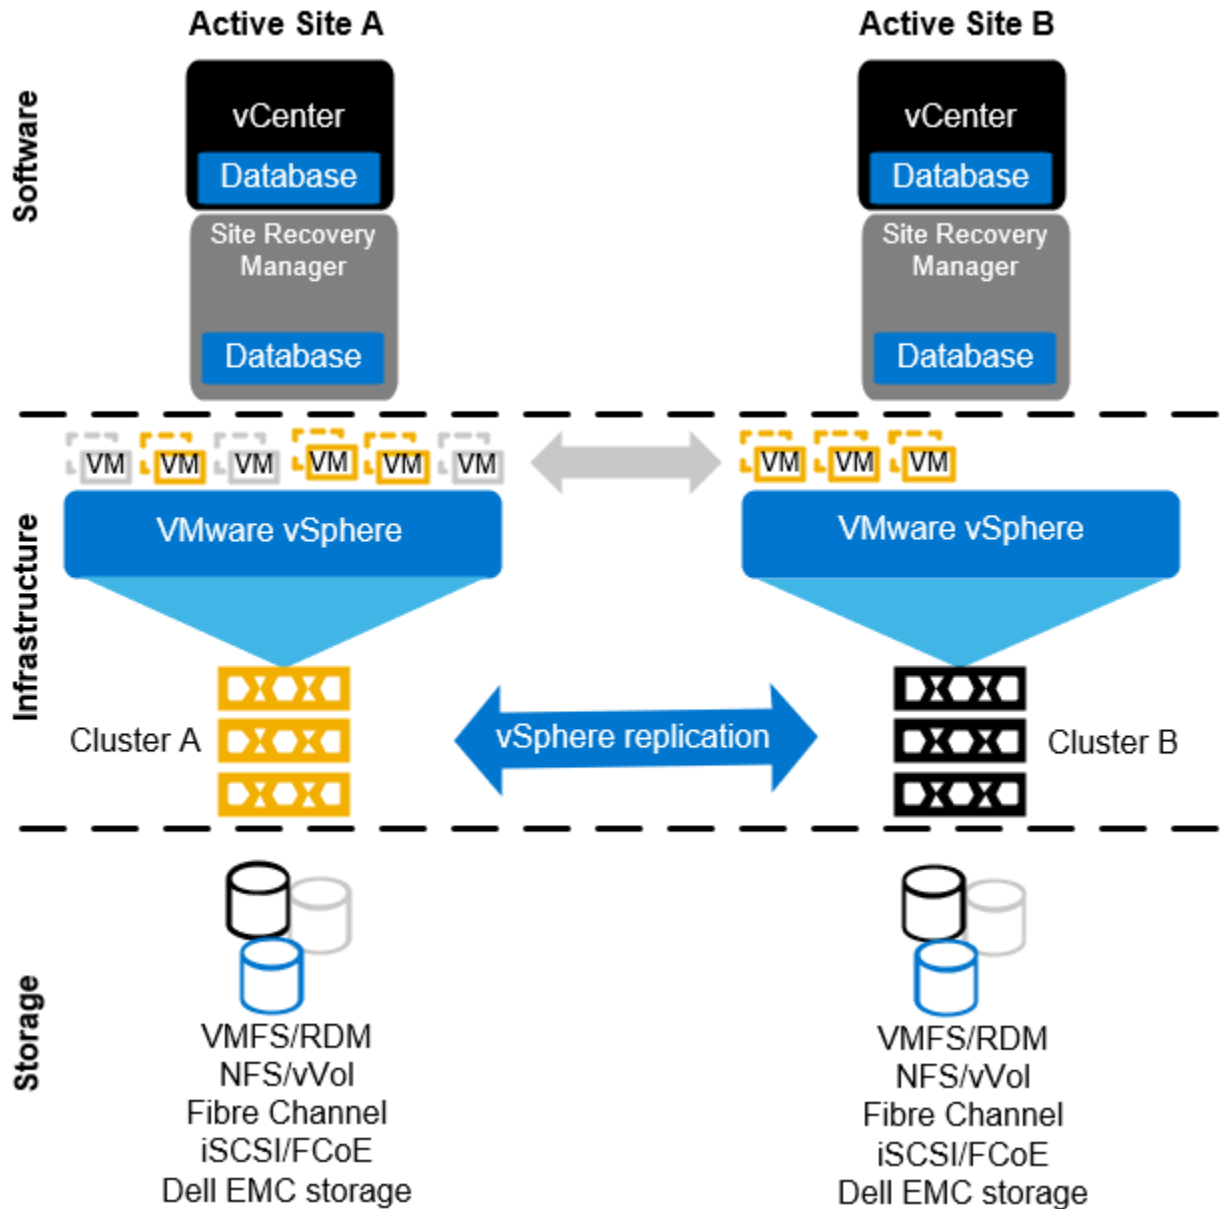


Figure 4 Configuration for a dual-protected site with vSphere replication

Note: You can use vSphere replication to replicate virtual machines that are based on PowerStore NFS and vVols. Alternately, consider using RecoverPoint for Virtual Machines.

4 PowerStore Manager configuration

This section provides best practices for configuring PowerStore Manager.

4.1 PowerStore Manager availability

As described in section 3, PowerStore Manager is a critical piece in the SRM infrastructure because it processes all calls from the SRA and performs the storage-related workflow tasks at the recovery site.

PowerStore Manager is natively integrated and deployed with each PowerStore appliance, so there are no architectural decisions required regarding where to deploy PowerStore Manager. If the recovery-site PowerStore appliance is healthy and available, the requirement for PowerStore Manager availability is met. Ensure that monitoring and alerting processes are in place for each PowerStore appliance.

4.2 PowerStore Manager logins

For SRM to function, the SRA must use login credentials that have rights to the respective PowerStore appliances that are replicating the virtual-machine volumes.

Keep in mind that each PowerStore appliance, whether it is at the protected or remote site, maintains its own user-access database. Credentials are required for PowerStore appliances at both sites. For example, if PS-10 is replicating virtual-machine volumes to PS-4, the credentials that the SRAs use must have administrator privileges to both appliances PS-10 and PS-4. Figure 5 shows the default admin credential that is used to manage the PowerStore appliance.

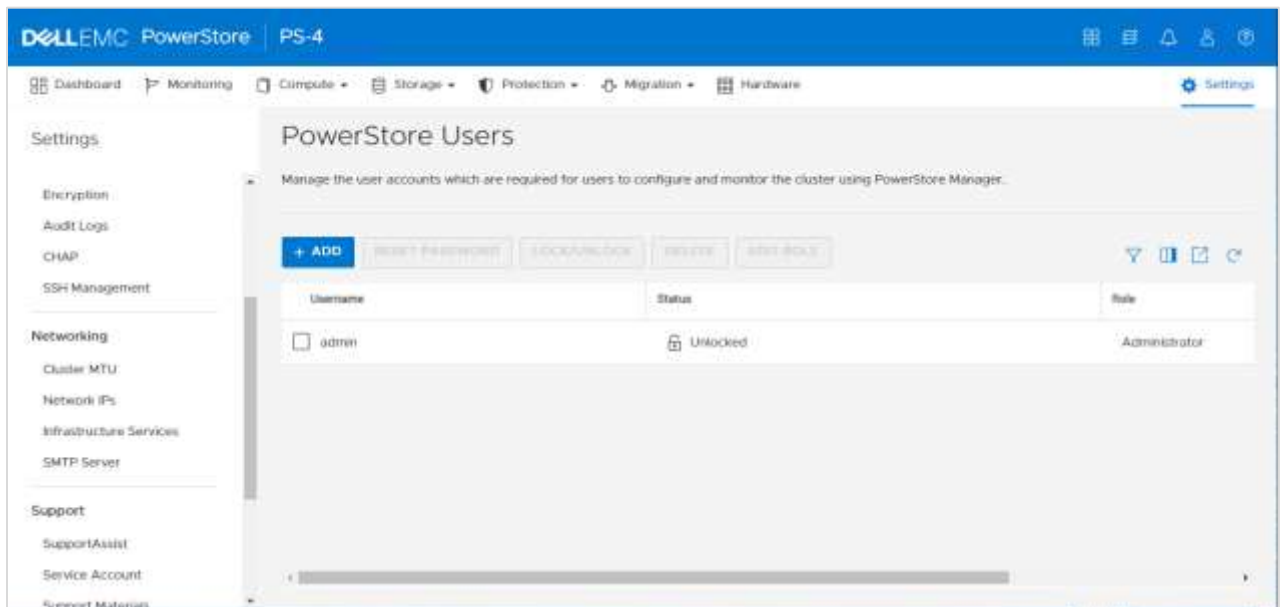


Figure 5 PowerStore Users menu in PowerStore Manager

4.3 Creating dedicated SRA access accounts

For the SRA to have uninterrupted access to both arrays, we recommend creating dedicated accounts for SRM. Using dedicated accounts on each array helps ensure that service is not unintentionally disrupted due to a password rotation, account lockout, account disablement, or account deletion.

Use these example steps to create dedicated accounts:

1. Create an account named **srmadmin** on both the protected-site array and the recovery-site array.

This account requires administrator privileges, and the password assigned must meet PowerStore password complexity requirements. For added security, create unique account names on each system with unique passwords. The account names and passwords are arbitrary.

2. Create a new account in PowerStore Manager named **srmadmin**.

The screenshot shows a 'Add User' dialog box with the following fields and values:

- Username:** srmadmin
- User Role:** Administrator
- New Password:** [Redacted] with a strength indicator of 'Very Strong' (green bar).
- Verify Password:** [Redacted]

Buttons at the bottom: CANCEL, APPLY.

The **srmadmin** account can now be used within the SRM Array Manager configuration.

Note: Each PowerStore appliance, whether it is at the protected or remote site, maintains its own user access database. Credentials are required for PowerStore appliances at both sites.

4.4 Modifying SRM settings for larger environments

VMware Site Recovery Manager ships with a default configuration that is tuned for a large cross-section of environments. However, each environment is unique in terms of architecture, infrastructure, size, and recovery time objective (RTO). Larger or more-complex SRM environments may require tuning adjustments in SRM (listed in the following bullet points) for SRM workflows to carry out their assigned tasks properly. For more information about making adjustments to accommodate such environments, see the SRM documentation section [Modify Settings to Run Large Site Recovery manager Environments](#).

- `storage.commandTimeout` – Min: 0 Default: 300

This option specifies the timeout allowed (in seconds) for running SRA commands in array-based-replication-related workflows. Increasing this value is typically required for larger environments. Recovery plans with many datastores to manage may fail if the storage-related commands take longer than five minutes to complete. For larger environments, increase this value (for example, to 3600 or higher) in the advanced SRM settings.

- `storage.maxConcurrentCommandCnt` – Min: 0 Default: 5

This option specifies the maximum number of concurrent SRA operations allowed.

- `storageProvider.hostRescanRepeatCnt` – Min: 0 Default: 1

This option specifies the number of additional host rescans during test, planned-migration, and recovery workflows. This feature was not available in SRM 5.0 and was reintroduced in SRM 5.0.1. Increase this value (for example, to 2 or higher) in the advanced SRM settings.

- `storageProvider.hostRescanTimeoutSec` – Min: 0 Default: 300

This option specifies the timeout allowed (in seconds) for host rescans during test, planned migration, and recovery workflows. Recovery plans with many datastores or hosts will fail if the host rescans take longer than five minutes to complete. Increase this value (for example, to 600 or higher) in the advanced SRM settings.

- `defaultMaxBootAndShutdownOpsPerCluster` – Default: off

This option specifies the maximum number of concurrent power-on operations performed by SRM at the cluster object level. To enable the option globally, specify a numerical value (such as 32) by modifying the `vmware-dr.xml` file. You can add this option anywhere in the `<config>` section, and restart the Site Recovery Manager Server service after making a change.

```
<config>
  <defaultMaxBootAndShutdownOpsPerCluster>32
</defaultMaxBootAndShutdownOpsPerCluster>
</config>
```

You can configure this value per cluster by editing the **`srmMaxBootShutdownOps`** in vSphere DRS Advanced Options. This value overrides a value specified in the `vmare-dr.xml` file.

- `defaultMaxBootAndShutdownOpsPerHost` – Default: off

This option specifies the maximum number of concurrent power-on operations performed by SRM at the host object level. To enable this option, specify a numerical value (such as 4) by modifying the `vmware-dr.xml` file. You can add this option anywhere in the `<config>` section, and restart the Site Recovery Manager Server service after making a change.

```
<config>
  <defaultMaxBootAndShutdownOpsPerHost>4
</defaultMaxBootAndShutdownOpsPerHost>
</config>
```

The **vmware-dr.xml** file is located in the **config** directory that resides in the SRM installation folder. The specific location varies depending on the operating system and SRM version. For example:

Windows:

C:\Program Files\VMware\VMware vCenter Site Recovery Manager\config\vmware-dr.xml

Virtual Appliance:

/opt/vmware/srm/conf/vmware-dr.xml

5 Replication configuration

PowerStore replication, in coordination with Site Recovery Manager (SRM), can provide a robust and scalable disaster-recovery solution. Since each snapshot and replication strategy affects recovery differently, choosing the correct protection policy to meet business requirements is important. 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 replication uses iSCSI through Ethernet (LAN) connections. When 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 screenshot shows a dialog box titled "Add Remote System" with a close button (X) in the top right corner. The form contains the following fields:

- Type:** PowerStore
- Management Cluster IP Address:** 100.88.130.40
- Description (Optional):** An empty text area.
- Network Latency:** A dropdown menu currently showing "Low" with an information icon (i) to its right.
- Username:** admin
- Password:** A masked password field represented by a series of dots.

At the bottom of the dialog, there are two buttons: "CANCEL" and "ADD".

Figure 6 Adding a remote system for replication

5.1 Asynchronous replication

With asynchronous replication, the I/O must be committed to and acknowledged by the source system so the data can be transferred to the destination in an independent timeframe. Supported storage resources for native asynchronous block replication are volumes, volume groups, and thin clones.

Note: Volume groups are treated as single entities when they are replicated. For virtual machines or tiered applications spanning multiple volumes, consider using volume groups to tie snapshot and replication schedules to the entire application group of volumes. This practice ensures point-in-time consistency across the volumes replicated to the recovery site.

Remote replication between PowerStore systems uses policy-based protection. Asynchronous replication configuration is defined in replication rules (see Figure 7). 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 at the protected site. For a valid configuration, a protection policy must contain at least one protection rule 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 PowerStore at the protected site. The required information for creating a rule includes the PowerStore system at the recovery site, the RPO, and the alert threshold for the planned replication session. When a protection policy with a replication rule is assigned to a storage resource, the configured RPO in the rule is used to set up the 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 are 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 this 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. Asynchronous replications usually have more flexible bandwidth requirements. This ability makes it the most common replication method for organizations that allow an RPO that is greater than zero (some amount of data will be lost when recovering from an asynchronous replication). Another benefit of PowerStore asynchronous replication is that the snapshots are transferred to the destination volume. By default, SRM recovers data from the most recent replicated snapshot. However, all snapshots replicated to the recovery site are available for recovery using PowerStore manager or APIs.

Create Protection Policy

Policy Properties

Name: Exchange Servers protection policy

Description (Optional):

Snapshot Rules

+ CREATE 1 Snapshot Rule, 1 selected

Name	Days	Frequency/Start Time	Retention	File Snapshot Access Ty...	Policies
hourly snap	Sunday, Monday, Tuesd...	1 hour	7 days	Protocol (Read-Only)	0

Replication Rules

+ CREATE 1 Replication Rule, 1 selected

Name	RPO	RPO Alert Threshold	Destination	Policies
hourly replication	1 hour	30 minutes	PS-4(100.88.130.40)	0

CANCEL CREATE

Figure 7 Creating a protection policy for an SRM protection group of volumes with one hour RPO

Note: Protection policies applied to a replicated volume carry over from the source to the destination after failover and re-protection workflows are run. Using the previous example, if the Exchange Server recovery plan is run and the volumes are failed over and re-protected, they maintain an hourly snapshot and hourly asynchronous replication to the peer storage array.

5.2 Remote system network latency

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 five 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.

5.3 Snapshots and application consistency

Asynchronous replication uses snapshots to provide point-in-time images as the source of RPO-based updates to the destination. These snapshots are used to maintain the common base images between the source and replicated resource across systems. Snapshots that replication creates and maintains are not visible in PowerStore Manager. When replication is configured, any snapshots that are created on the source resource are automatically replicated in chronological order to the destination system during the next RPO-based update (see Figure 8). There are several methods available for creating snapshots: PowerStore Manager, protection policies, PowerStore REST API, and PSTCLI. When replicated, SRM may use a thin

clone of the snapshot to present recovered data to the vSphere cluster. Snapshots created in PowerStore Manager or a protection policy are considered crash consistent. You can use other methods that result in application consistency within the snapshot. For example, where supported, you can use Dell EMC AppSync™ to create application-consistent snapshots. This practice ensures that all incoming I/O for a given application is quiesced and flushed before a snapshot is created. Another method is to use vSphere snapshots with quiescence captured inside a replicated PowerStore snapshot. Either of these examples results in application-consistent snapshots being replicated to the recovery site.

When using vSphere snapshots, there are two important facts to recognize:

- The VM is replicated to the destination site in a vSphere snapshot state. It should be addressed to prevent the VM from running continuously over a long time in a vSphere snapshot state.
- The application and data consistency are contained within the frozen-parent virtual disk, and crash-consistent data is contained in the delta virtual disk.

When the SRM recovery plan workflow is carried out, SRM registers the VM into inventory at the destination site. Then, it powers on the VM with no special attention given to the current snapshot state of the VM. This means that SRM powers on the VM using the delta, resulting in recovery from a crash-consistent state. To recover the VM from the frozen-parent disk with application and data consistency, revert the VM to the previous snapshot using the vSphere Snapshot Manager before powering on the VM. Once this process is done, you can delete (close) the snapshot and power on the VM. This process ensures the VM is powered on from its frozen-parent disk and the delta disk, and the crash-consistent data in it is destroyed.

If manually carrying out the previous process on a large scale, this can erode efforts made toward meeting the recovery plan RTO and is not the best use of SRM. In such instances, it is more efficient and consistent to script the snapshot management process using Microsoft PowerShell®. You can carry out this process as a pre-power-on step (or potentially post-power-on step) for the VM using a custom recovery task.

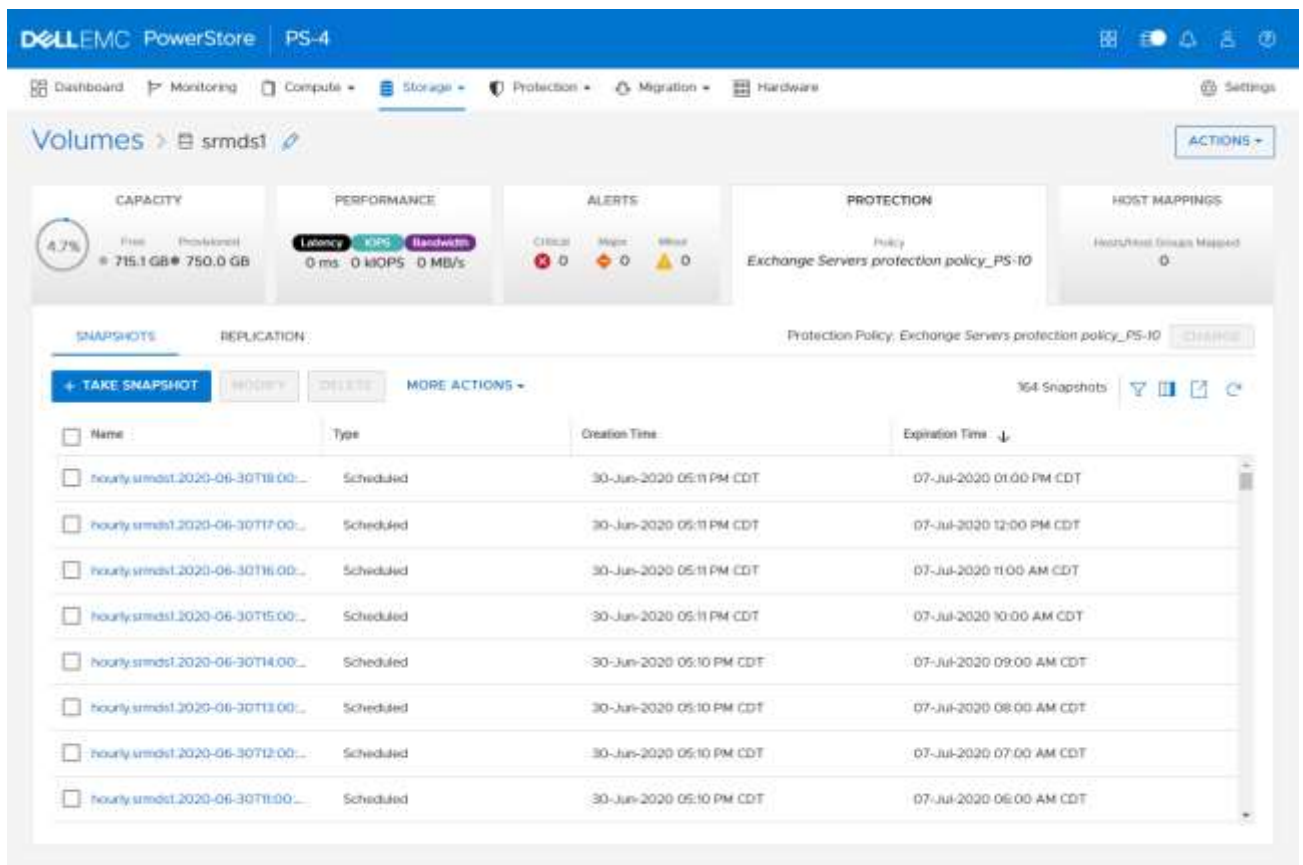


Figure 8 PowerStore Manager showing Exchange Server snapshots replicated to the recovery site, with the default snapshot retention policy of seven days

5.4 Custom recovery tasks

If the environment requires a custom recovery strategy, both Dell EMC storage and VMware have robust API sets to customize the recovery steps where needed. APIs include PowerShell cmdlets, PowerStore REST API, and PSTCLI. The APIs can be used for tasks such as managing snapshots, creating thin clones, mapping volumes, and managing replications. Within the same script, the VMware cmdlets can rescan HBAs, manipulate vDisks, add virtual machines to inventory, and perform most other tasks required for recovery (see Figure 9).

Add Step Before "Power on priority 1 VMs"

Type: Command on SRM Server
 Prompt (requires a user to acknowledge the prompt before the plan continues)

Name: Revert VMware Snapshots

Content: C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe C:\RevertSnapshots.ps1

Timeout: 5 minutes 0 seconds

CANCEL ADD

Figure 9 Custom recovery task in a Site Recovery Manager recovery plan

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

6 Site Recovery Manager configuration

This section provides guidance and best practices for configuring Site Recovery Manager.


6.1 SRA installation

The PowerStore Storage Replication Adapter (SRA) must be installed on each SRM server. PowerStore offers SRAs for both the Photon operating-system-based SRM appliance and the Windows-based SRM installation. You can [download the SRAs](#) from the VMware website. We recommend using the most current version of the SRA to ensure optimal compatibility and available features. See the release notes and product documentation to determine SRA compatibility with SRM versions. The video [Dell EMC PowerStore - Storage Replication Adapter Installation](#) demonstrates the installation process of the SRA on both operating systems.

Note: SRM supports installing multiple Storage Replication Adapters. This ability is beneficial when storage arrays of different types exist in the data center.

6.2 Configuring the array managers

To allow SRM to manage PowerStore storage, the SRA must be able to communicate with PowerStore. You can configure the array manager from the Array Managers module. You must add an array manager for each site in the unified interface (see Figure 10 and Figure 11).



The screenshot shows the Site Recovery Manager interface with the following details:

- Page Title: Storage Replication Adapters
- Selected Site: vcsite1.techsol.local
- Section: RESCAN ADAPTERS
- Adapter Name: Dell EMC PowerStore SRA

Status	✓ OK
Version	1.0.0.52
Vendor	Dell inc.
Install Location	E:\Program Files\VMware\VMware vCenter ...
Vendor URL	https://www.dell.com/support
Supported Array Models	Dell inc., PowerStore
Supported Software	PowerStore 1.0
Stretched Storage	Not Supported

Figure 10 Examine the installed SRAs

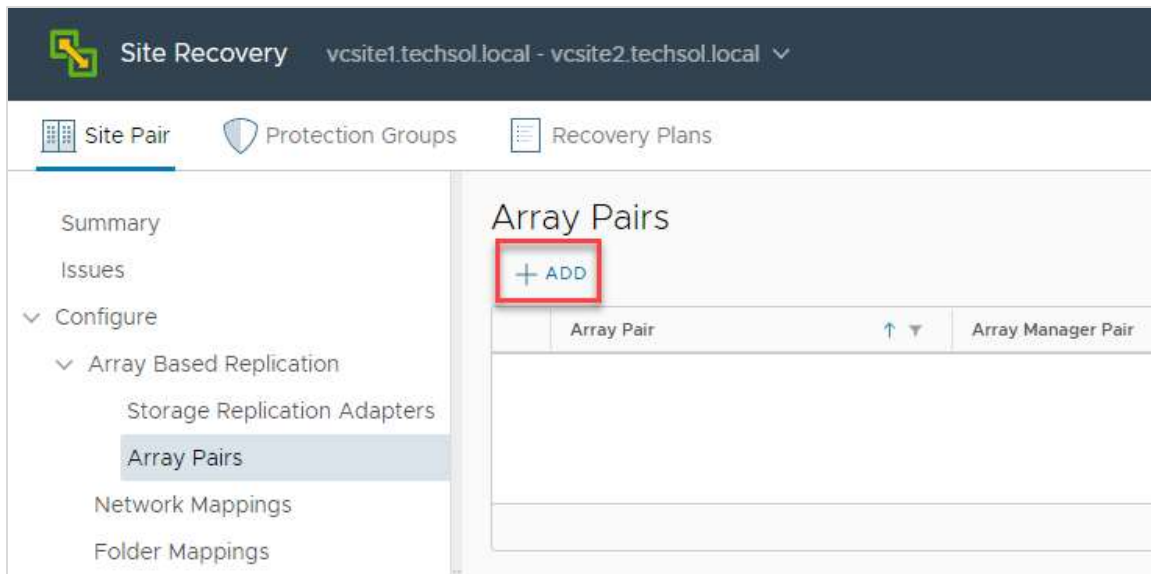
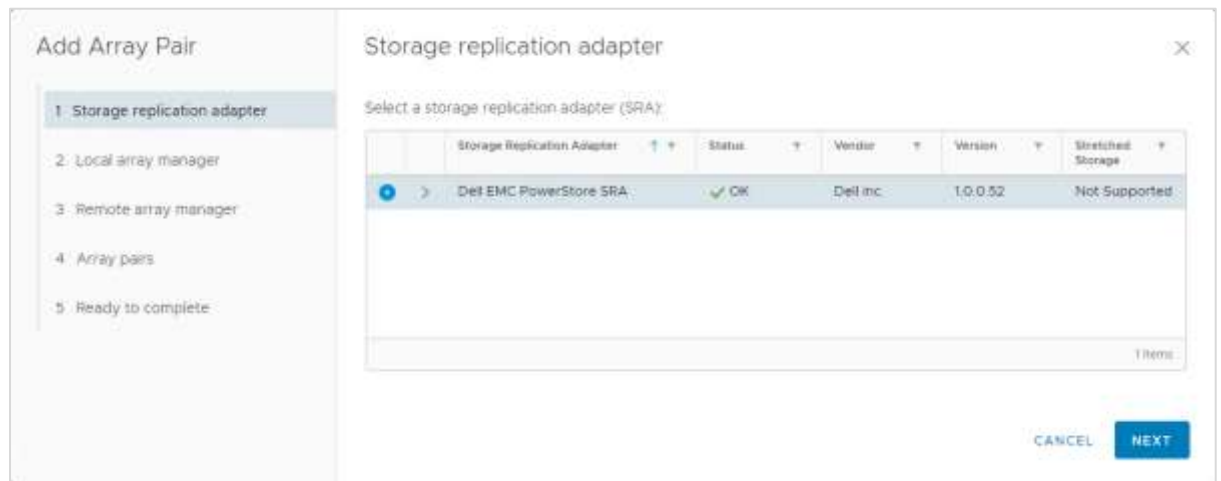


Figure 11 Adding an array manager

To perform the required process to configure the protected site array managers and the recovery site array managers for pairing, complete the following steps:

1. Choose the installed SRA.



2. Provide the **local** PowerStore connection parameters for the local array manager.

3. Provide the **remote** PowerStore connection parameters for the remote array manager.

6.3 Creating array pairs

When an array manager is added to each of the two sites in SRM, the arrays must be paired so that replicated volumes can be discovered by SRM as eligible devices (see Figure 12). In older versions of SRM, pairing was an action that was performed after the initial installation of SRM. However, as of SRM 5.8, you can perform pairing as part of the process of adding array managers to sites.

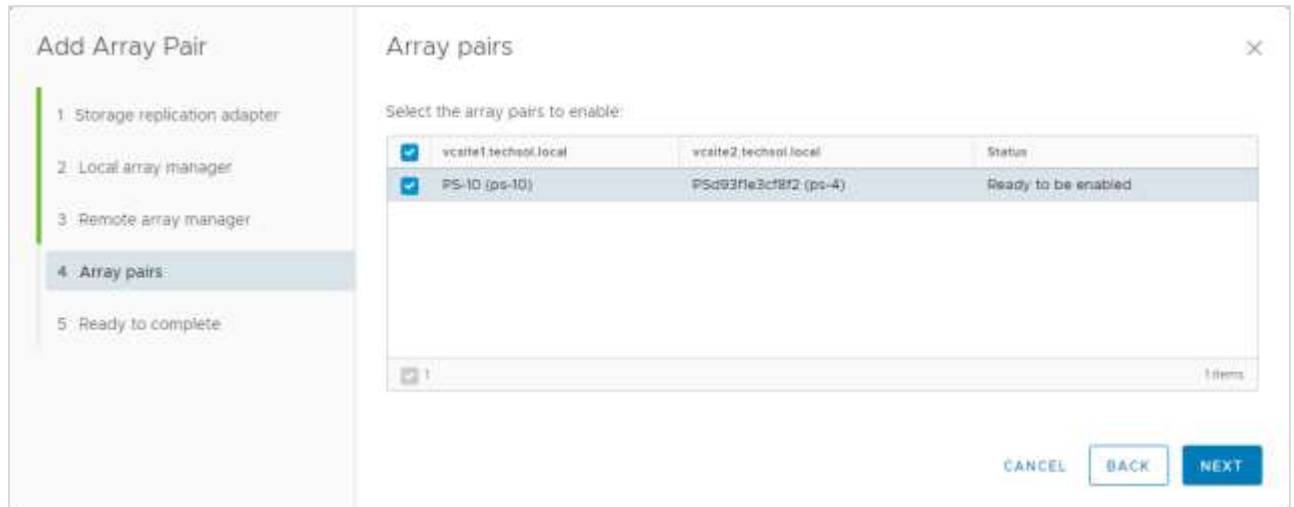


Figure 12 Creating array pairs

Arrays cannot be unpaired when there are downstream SRM dependencies such as protection groups.

6.4 Array manager device discovery

Whenever a new replicated datastore or RDM is added to the environment, the arrays should be rescanned within SRM for new devices. The array pair device discovery tool is located in the **Array Based Replication > Array Pairs** menu. Run the device discovery on both arrays to ensure a consistent list of devices. Non-replicated volumes are not discovered and displayed as eligible devices in SRM. Keep this in mind as a troubleshooting tip if datastores or RDMs are not listed as eligible devices in SRM. Conversely, SRM discovers all replicated volumes as devices, even vSphere does not use them. For example, replicated volumes could belong to other storage hosts such as physical Microsoft Exchange, SQL Server®, Oracle®, and file servers.

Select **Discover Devices** as shown in Figure 13 to invoke an SRA query of PowerStore to obtain the newest array-based replicated device information.

The screenshot shows the 'Array Pairs' section of the Site Recovery Manager interface. At the top, there are buttons for '+ ADD', 'ARRAY PAIR', and 'ARRAY MANAGER PAIR'. Below these is a table with columns: 'Array Pair', 'Array Manager Pair', and 'Last Array Manager Ping'. A single row is visible with a green checkmark, indicating a successful ping for the pair 'PS-10 ↔ PSd93fle3cf8f2' managed by 'ps-10 ↔ ps-4' on 6/29/2020 at 3:43:13 PM CDT. Below this table is a 'DISCOVER DEVICES' button, which is highlighted with a red box. Underneath the button is another table with columns: 'Device (vcsite1.techsol.local)', 'Datastore', 'Status', 'Device (vcsite2.techsol.local)', 'Protection Group', and 'Local Group'. Two rows are shown: 'srmds1' with 'Local: [srmds1]' and 'Forward' status, and 'srmds2' with 'Local: [srmds2]' and 'Forward' status.

Figure 13 Discover the devices of array pairs

6.5 Creating placeholder datastores

If not completed, create a small VMFS datastore at the disaster recovery site as a placeholder for VM configuration files. For each protected virtual machine, SRM creates a shadow VM at the recovery site. This VM serves as a placeholder for CPU, memory, and network resources that are required to perform a test, disaster recovery, or planned-migration plan.

Although this datastore must only be large enough to hold the configuration files for all the recoverable virtual machines, creating a standard-sized 500 GB datastore will suffice. PowerStore thinly provisions the volume, making this a space-efficient standard.

Note: The minimum PowerStore volume size is 1 MB. The minimum VMFS volume size is 2 GB. However, for practical reasons, the placeholder datastore should be at least 5 GB or larger. A smaller datastore will likely trigger ongoing vSphere datastore capacity alarms in the vSphere Client UI.

Typically, only one placeholder datastore per site is required. This is because the disaster recovery and migration processes unregister and reregister the recovered virtual machine with the .vmx file on the recovered volume. The placeholder volume does not need to be replicated or protected because VMware SRM places only transient data on this volume that can be easily regenerated within the UI.

6.6 Protection group considerations

With the placeholder datastore ready, you can create protection groups. Replicated datastore volumes are the foundation that protection groups are built upon. A protection group is effective immediately after being created. Once a VM is protected, it is essentially pinned to the datastore (or datastores) where the .vmx and .vmdk files reside. SRM does not support manually moving files that belong to a virtual machine off a datastore; the VM is not protected or replicated from its original datastore or datastores. Automated Storage DRS (SDRS) and VMware Storage vMotion® can be sparingly used with SRM-protected VMs if certain guidelines are followed. See the VMware *Site Recovery Manager Administration Guide*.

6.7 Recovery plan considerations

When creating recovery plans, a best practice to further automate DR failover or planned migration may be to add prompts or SRM server-side commands to the recovery plan. The SRM server-side commands could be application-specific or related to storage management and integrate a PowerStore REST API or PSTCLI script into the recovery plan. When the recovery plan runs, it pauses on prompts while SRM server commands are performed without a pause (see Figure 14).

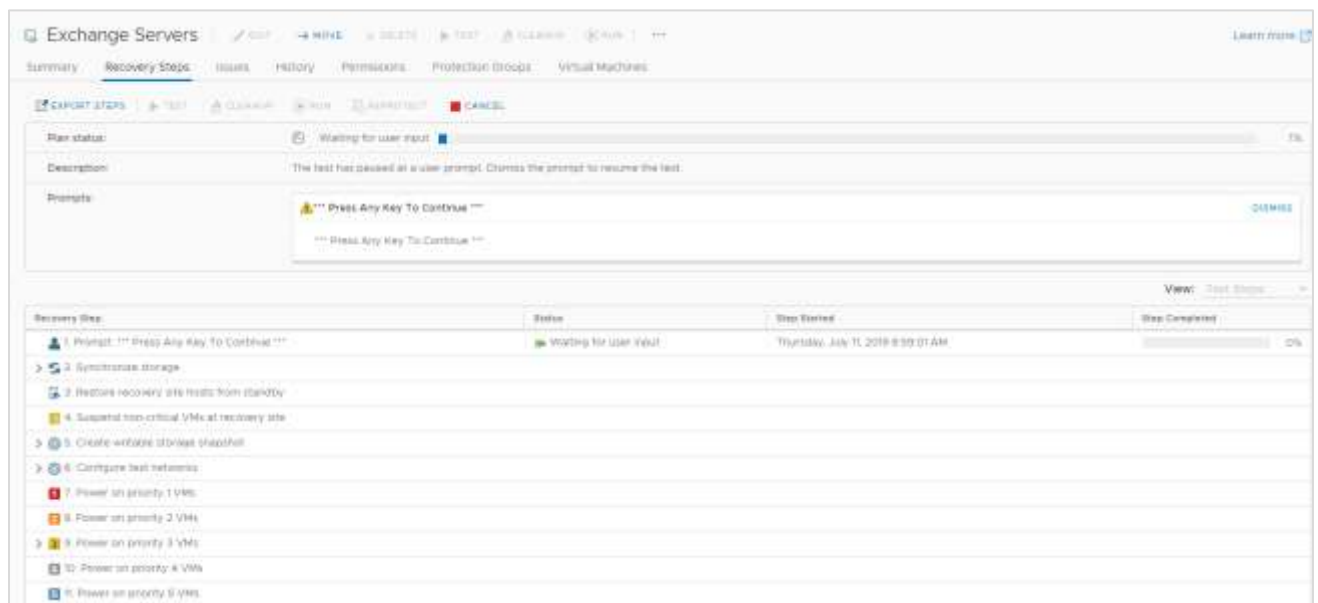


Figure 14 Recovery plan added step prompting to continue

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

7 Recovery plan testing and running

Testing the recovery plan is not disruptive to the storage replications, production volumes, and VMs because the test plan uses thin clone volumes from replicated snapshots at the recovery site. When testing a recovery plan, any tests, changes, or updates can be performed on the recovered VMs because they are discarded when the test recovery plan cleanup occurs. While the test plan is running, production VMs and replication continue to run without interruption.

To test a disaster recovery plan, right-click the recovery plan, and click **Test** (see Figure 15).

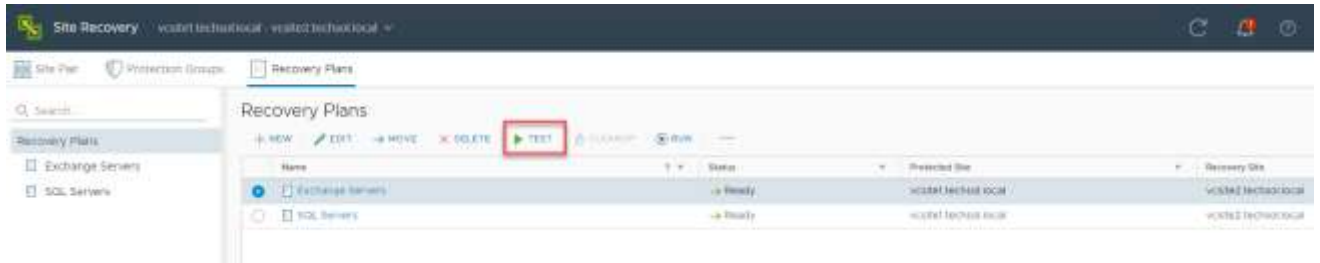


Figure 15 Testing a recovery plan

When testing or running recovery plans, SRM does not have integrated mechanisms to determine whether the replicated volumes are fully synced before the storage is prepared for recovery. In other words, data may be actively replicating to the secondary site which could influence the outcome of the recovery. As a best practice, check **Replicate recent changes to recovery site** when running a test plan. This action ensures that all data is successfully replicated to the secondary site (see Figure 16).

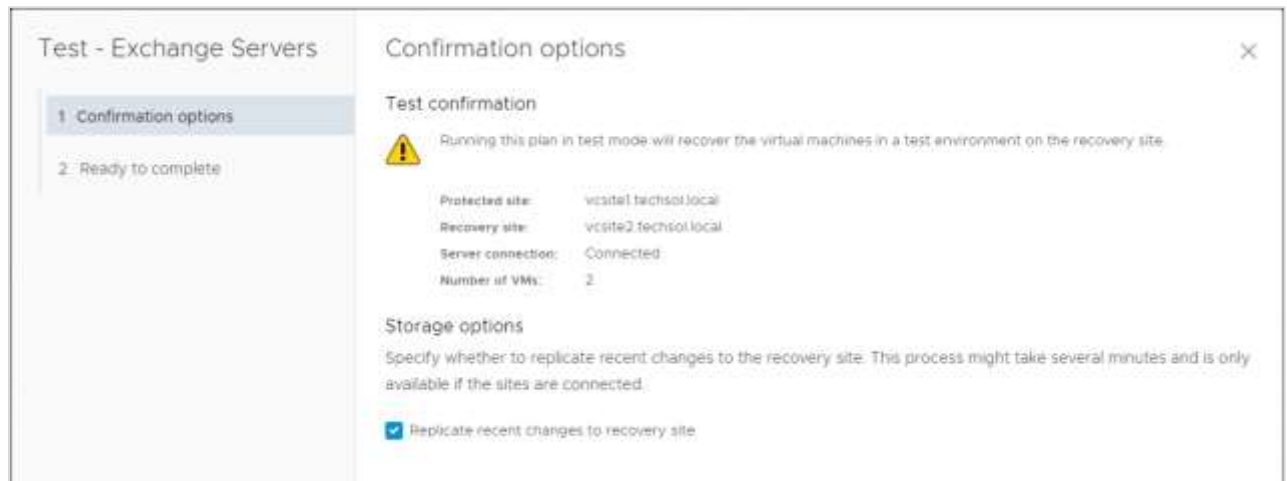
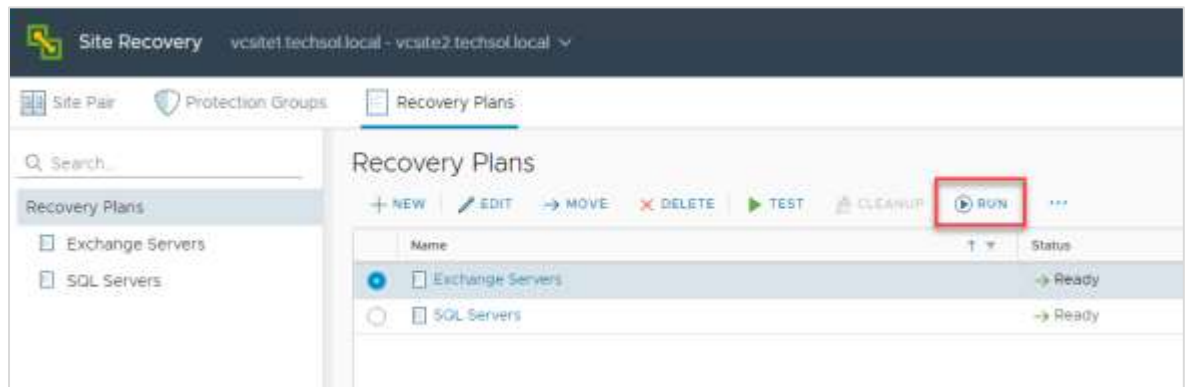


Figure 16 Replicate recent changes to recovery site during test plans

Note: The Replicate recent changes feature results in a longer running plan. The extra time is used to synchronize the volumes between sites. During a disaster-recovery cutover, the Replicate recent changes to recovery site option may or may not be available. For planned migrations using SRM, this step is required to proceed.

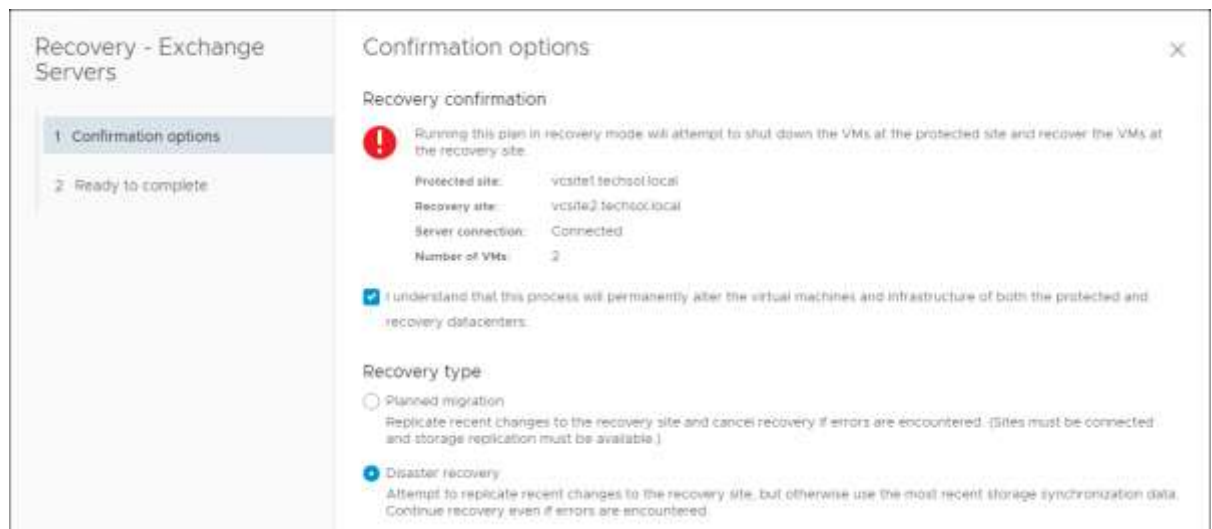
When choosing to run a planned migration or disaster recovery plan (as opposed to running a test), keep in mind this procedure is disruptive. It results in VMs being powered off at the primary site, replication mirrors being broken, and VMs being recovered at the secondary site.

1. In the event of a disaster or planned migration, right-click the recovery plan, and click **Run**.



Note: Before running a planned migration plan, run a test recovery of the plan.

2. Acknowledge the safety precaution message to run a live plan.



3. Review the success of the recovery plan after it completes.

The screenshot shows the 'Exchange Servers' interface with the 'Recovery Steps' tab selected. The plan status is 'Recovery complete'. The description states: 'The recovery has completed. Review the plan history to view any errors or warnings. You can now press Reprotect to configure protection in the reverse direction. Note that if you start to fallback the virtual machines to the original site, you must first run the plan in Reprotect mode. Once site protection is configured in reverse, you can run the plan in recovery mode to fallback the actual machines to the original site.'

Recovery Step	Status	Step Started	Step Completed
1. Pre-configure storage	Success	Monday, June 29, 2020 8:00:35 PM	Monday, June 29, 2020 8:01:08 PM
2. Shut down VMs at protected site	Success	Monday, June 29, 2020 8:01:09 PM	Monday, June 29, 2020 8:01:22 PM
3. Resume VMs suspended by previous recovery			
4. Restore recovery site hosts from standby	Success	Monday, June 29, 2020 8:01:22 PM	Monday, June 29, 2020 8:01:23 PM
5. Restore protected site hosts from standby	Success	Monday, June 29, 2020 8:01:23 PM	Monday, June 29, 2020 8:01:23 PM
6. Prepare protected site VMs for migration	Success	Monday, June 29, 2020 8:01:23 PM	Monday, June 29, 2020 8:01:26 PM
7. Synchronize storage	Success	Monday, June 29, 2020 8:01:26 PM	Monday, June 29, 2020 8:02:12 PM
8. Suspend non-critical VMs at recovery site			
9. Change recovery site storage to online	Success	Monday, June 29, 2020 8:02:12 PM	Monday, June 29, 2020 8:02:29 PM
10. Power on priority 1 VMs			
11. Power on priority 2 VMs			
12. Power on priority 3 VMs	Success	Monday, June 29, 2020 8:02:26 PM	Monday, June 29, 2020 8:02:30 PM
13. Power on priority 4 VMs			
14. Power on priority 5 VMs			

8 Reprotect and failback

After VMs are migrated to another site using the disaster-recovery or planned-migration features in SRM, they are in an active running state on the network at the alternate site. However, they are vulnerable to a site failure with no SRM protection. Previous versions of SRM required a manual reprotection of the VMs at the recovery site. Today, SRM automates the reprotection process and prepares the virtual machines for failback.

8.1 Reprotection

After protected VMs are migrated, or failed over to the secondary site as part of disaster recovery, the VMs are unprotected and are no longer replicated to a recovery site. Following the migration of protected virtual machines, SRM enables automating the reprotection of the VMs. The reprotection is carried out in a series of automated steps (see Figure 17).




Figure 17 Reprotect workflow of the Exchange Servers protection group being performed

During a reprotect, SRM commands the SRA to start storage replication for each of the datastores or volumes in the protection group. This action occurs in the opposite direction compared to the replication topology before the failover. The protection group that was originally set up at the primary site is migrated to the secondary site. Placeholder VMs that were originally set up at the secondary site are now created at the opposite site (the new recovery site) on its respective placeholder datastore (see Figure 18).

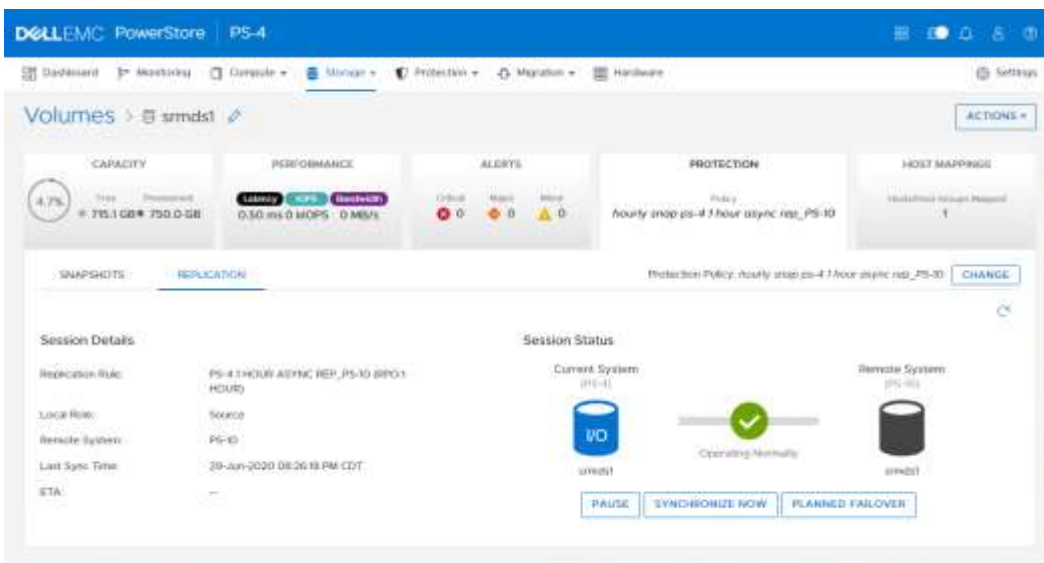


Figure 18 Replication direction and session status can be monitored in PowerStore Manager

8.2 Failback

Failback is an SRM term that describes the ability to perform a subsequent disaster recovery or planned migration after a successful recovery and reprotect. The benefit that failback introduced in SRM 5.x is the automated ability to move back and forth between sites with minimal effort. This capability facilitates several use cases including the ability to run production applications at the disaster recovery site, resource balancing, and improved disaster recovery infrastructure ROI.

Note: Before running the Failback recovery plan, run a test recovery of the plan.

9 Conclusion

VMware vSphere, Site Recovery Manager, and Dell EMC PowerStore combine to provide a highly available business platform for automated disaster recovery. This platform enables the best possible RTO and RPO, and supports planned migrations for your virtualized data center.

A Additional resources

A.1 Technical support and resources

[Dell.com/support](https://www.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.

PowerStore snapshot and replication-related resources:

- [Replication Technologies](#)
- [Snapshots and Thin Clones](#)

A.2 VMware support

For VMware support, see the following resources:

- [VMware.com](https://www.vmware.com)
- [VMware support](#)
- [Education and training](#)
- [Online documentation](#)
- [VMware communities](#)