

Dell EMC Unity: VMware Site Recovery Manager Best Practices

Abstract

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

January 2021

Revisions

Date	Description
January 2021	Initial release

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.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [1/19/2021] [Best Practices] [H18633]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	5
Audience	5
1 Introduction.....	6
1.1 Dell EMC Unity overview	6
1.2 Terminology	6
2 Setup prerequisites	9
2.1 Storage Replication Adapter.....	9
2.2 Dell EMC Unity and Unity XT	9
2.3 VMware vSphere and SRM	9
3 Site Recovery Manager architecture.....	10
3.1 Array-based replication: single protected site	10
3.2 Array-based replication: dual protected site	11
3.3 vSphere replication: single protected site.....	12
3.4 vSphere replication: dual protected site	13
4 Unisphere configuration	14
4.1 Unisphere availability.....	14
4.2 Unisphere logins	14
4.3 Creating dedicated SRA access accounts	15
4.4 Modifying SRM settings for larger environments.....	16
5 Replication configuration	18
5.1 Synchronous replication	19
5.2 Asynchronous replication and snapshots.....	21
5.3 VMware NFS datastore replication.....	25
5.4 Snapshots and application consistency	27
5.5 Custom recovery tasks	28
6 Site Recovery Manager configuration	29
6.1 SRA installation	29
6.2 Configuring the array managers.....	29
6.3 Creating array pairs	32
6.4 Array manager device discovery	32
6.5 Creating placeholder datastores.....	33

Table of contents

6.6	Protection group considerations	33
6.7	Recovery plan considerations	34
7	Recovery plan testing and running	35
8	Reprotect and failback	38
8.1	Reprotection	38
8.2	Failback	39
9	Conclusion	40
A	Additional resources	41
A.1	Technical support and resources	41
A.2	VMware support	41

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™ Unity™, 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 Dell EMC Unity 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 Unity. Besides 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 before beginning an SRM implementation.

1.1 Dell EMC Unity overview

In this constantly changing world of increasing complexity and scale, the need for an easy-to-use, intelligent storage system has only grown greater. Customers using new applications and solutions require dependable storage and are often tasked with the challenge of doing more with less. The Dell EMC Unity and Unity XT family addresses this challenge by packaging a powerful storage system into a cost and space-efficient profile.

Dell EMC Unity storage sets new standards for midrange storage. It offers a powerful combination of simplicity, modern design, affordable price point, and deployment flexibility—perfect for resource-constrained IT professionals in large or small companies. Dell EMC Unity systems are perfect for midsized deployments, remote office/branch office locations, and cost-sensitive mixed-workload environments. They are designed for all-flash, deliver the best value, and are available in several configurations. These configurations include purpose-built (all-flash or hybrid), converged deployment (through the Converged Infrastructure Portfolio), and software-defined virtual edition. With all-inclusive software, new differentiated features, internet-enabled management, and a modern design, Dell EMC Unity storage is where powerful meets simplicity.

1.2 Terminology

The following terms are used with Dell EMC Unity.

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 recovery point objective (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.

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

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

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

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

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.

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 Unity 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, that is 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.

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

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

Unisphere 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 schedule. 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 VMware vSphere 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.

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

2 Setup prerequisites

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

2.1 Storage Replication Adapter

You must install the Dell EMC Unity Storage Replication Adapter (SRA) on each SRM server. Dell EMC Unity offers SRAs for both the Photon operating-system-based SRM appliance and the Windows-based SRM installation. Also, Dell EMC Unity offers an SRA for block storage and an SRA for file storage. You can download the SRAs from the VMware website. Dell Technologies recommends 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: See the SRA release notes for specific requirements or features for the SRA.

2.2 Dell EMC Unity and Unity XT

SRM- and array-based replication of block or file datastores requires two Dell EMC Unity or Unity XT systems replicating between each other in one or both directions. Asynchronous and Synchronous replication is supported for both block and file datastores. This support includes MetroSync (synchronous replication of NAS servers and associated storage resources such as NFS datastores and snapshots). You can use SRM with vSphere replication to protect virtual machines that are backed by vVols.

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, check the [VMware Product Interoperability Matrix](#). SRM is supported with vCenter Server for Essentials, vCenter Server Foundation, and vCenter Server Standard.

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.

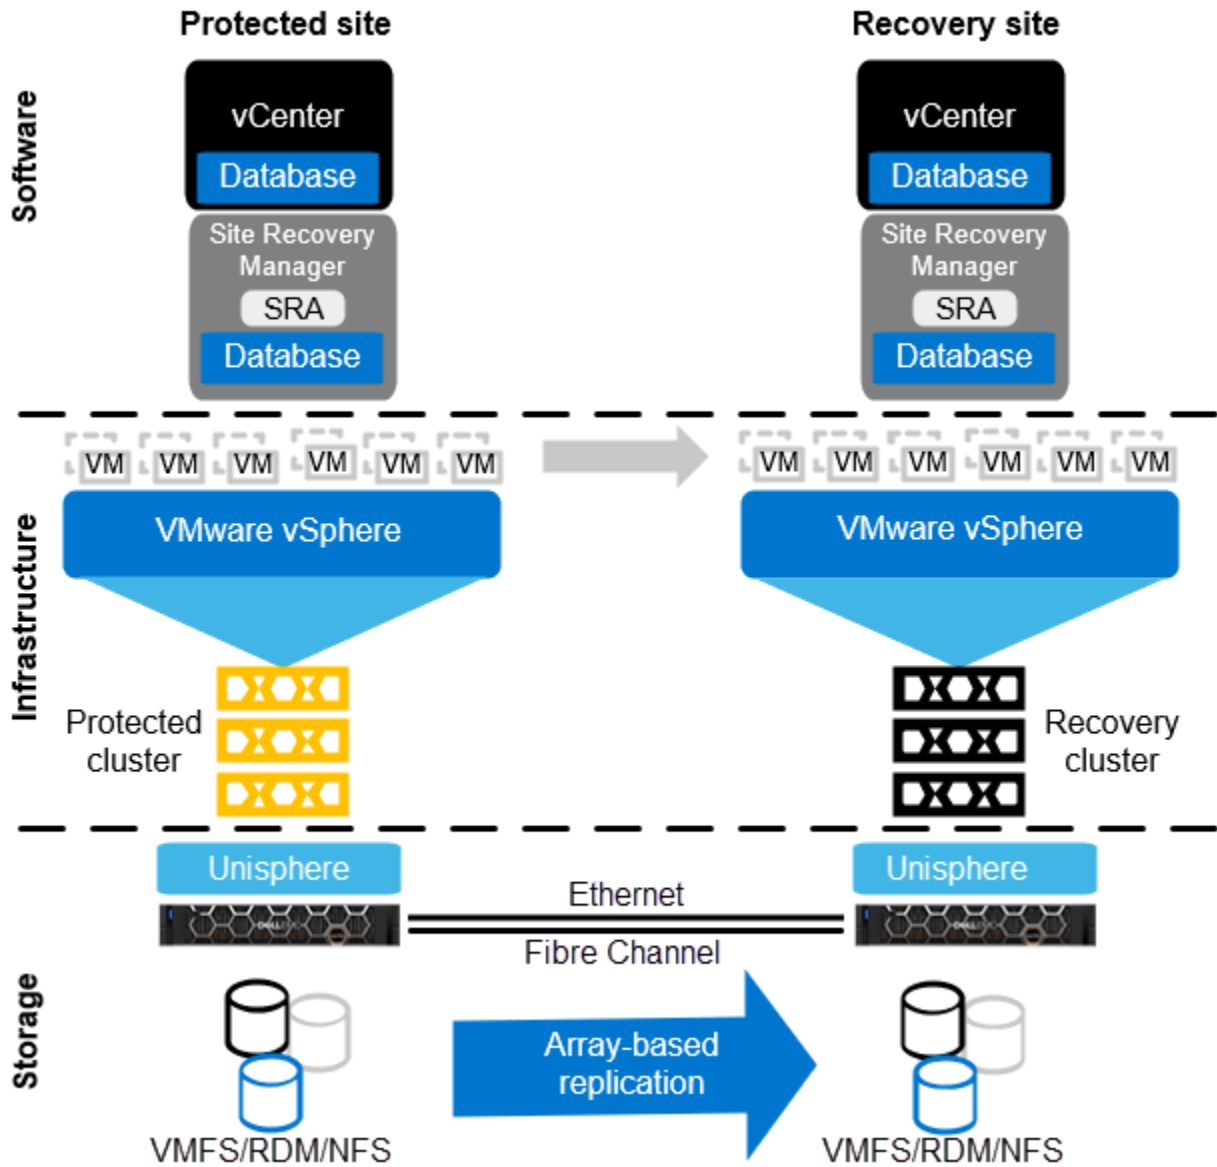


Figure 1 Architecture for a single protected site with array-based replication

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.

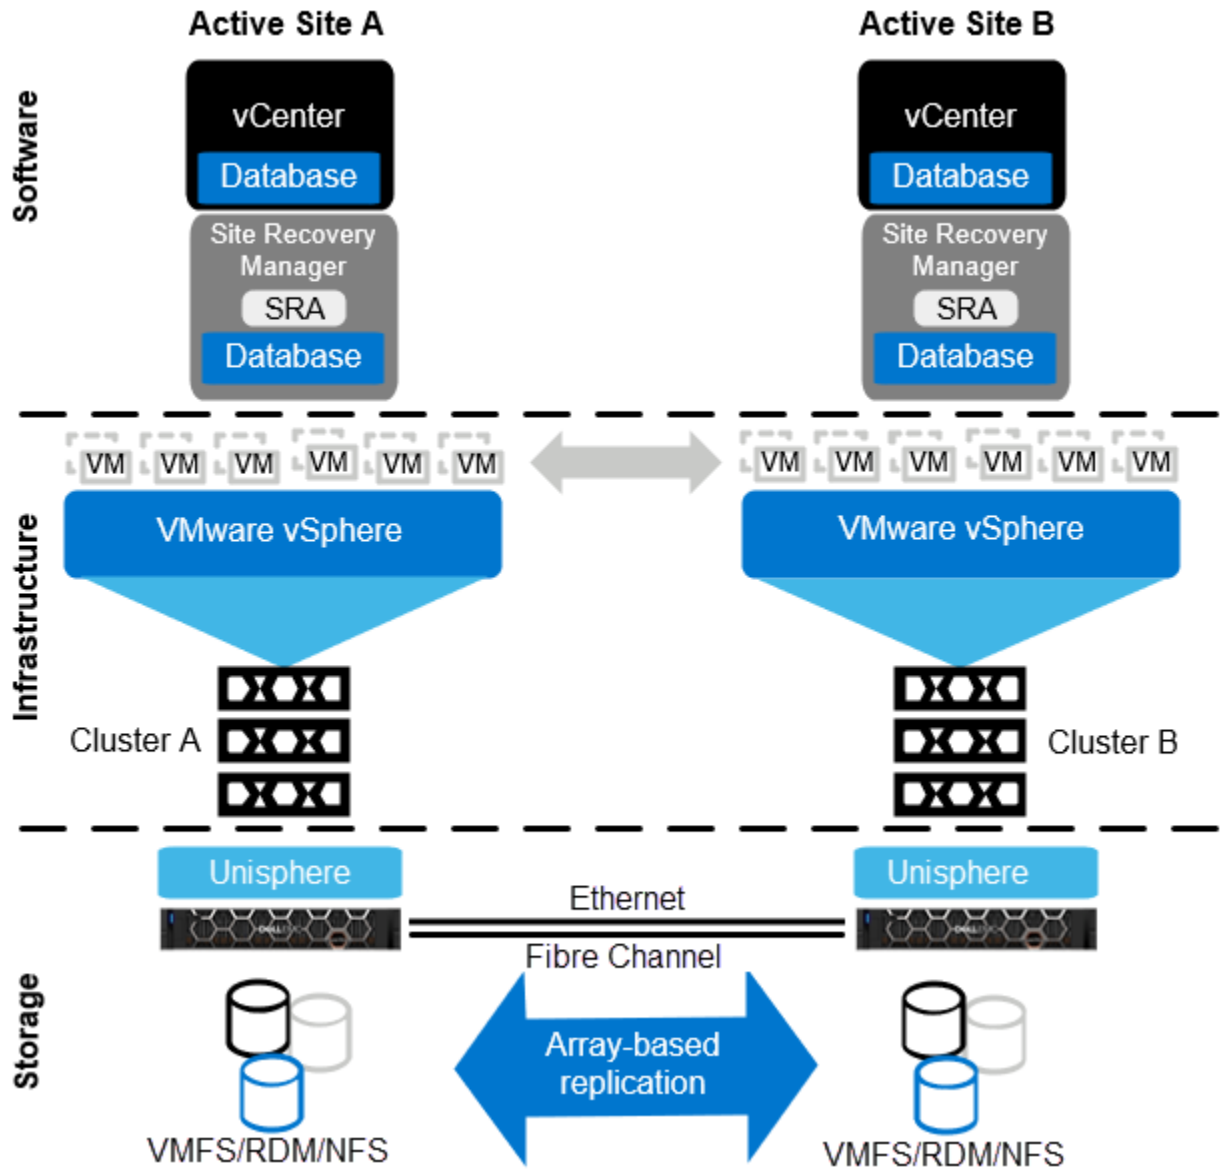


Figure 2 Architecture for a dual protected site with array-based replication

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.

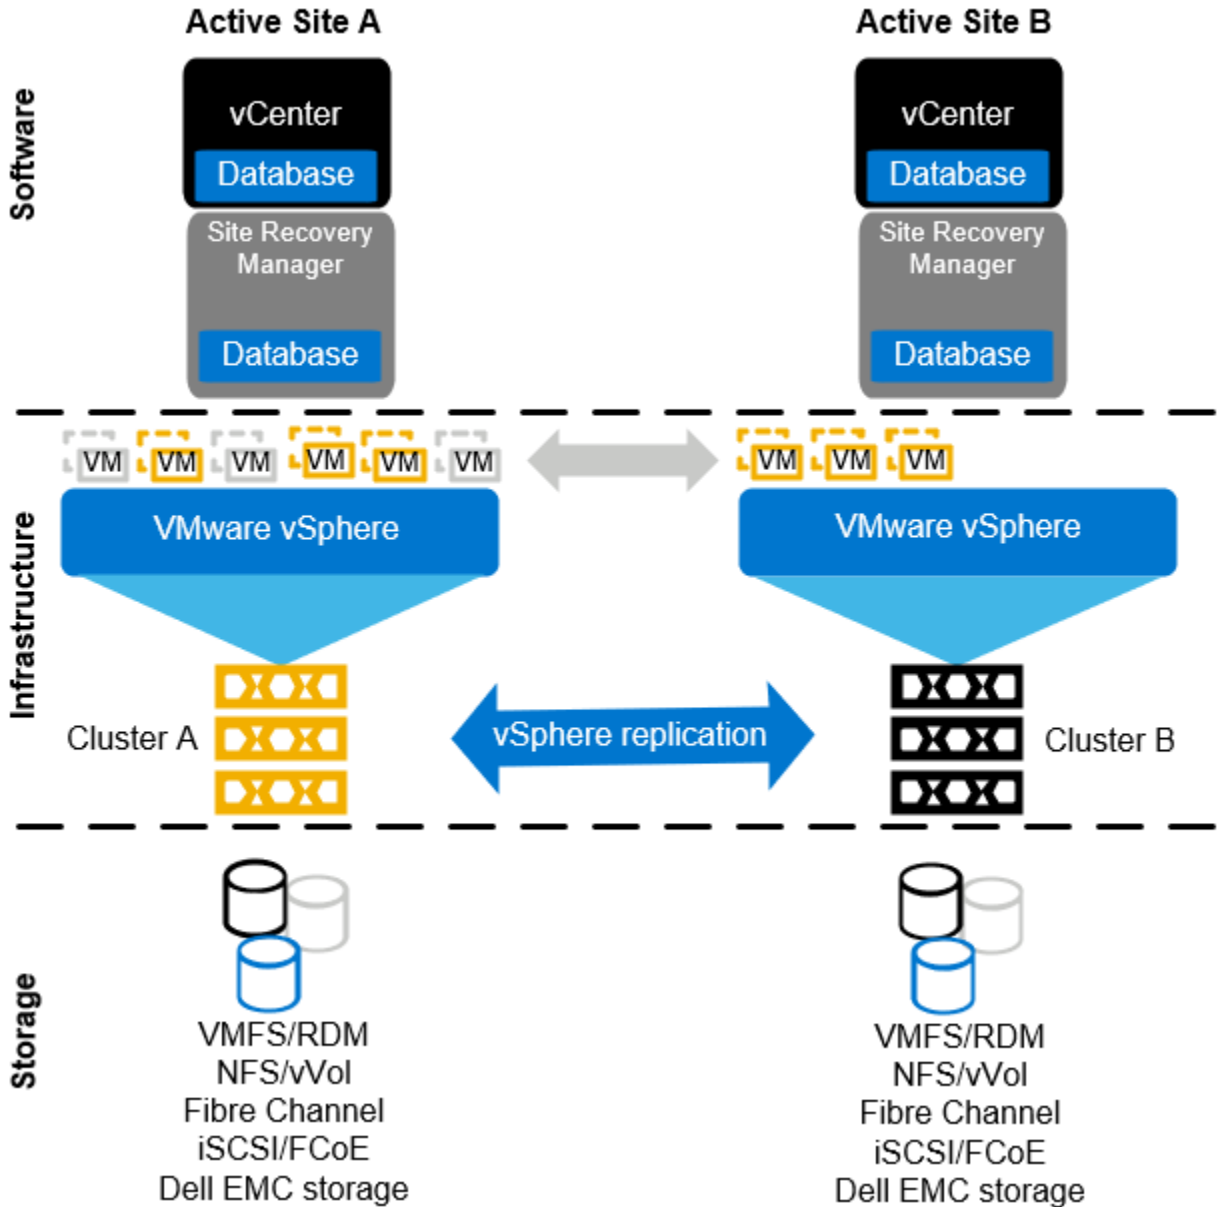


Figure 4 Configuration for a dual-protected site with vSphere replication

Note: You can use vSphere replication to replicate virtual machines that are backed by vVols.

4 Unisphere configuration

This section provides best practices for configuring Unisphere.

4.1 Unisphere availability

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

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

4.2 Unisphere logins

For SRM to function, the SRA must use login credentials that have rights to the respective Dell EMC Unity systems that are replicating the virtual-machine volumes.

Keep in mind that each Dell EMC Unity system, whether it is at the protected or remote site, maintains its own user-access database. Credentials are required for Dell EMC Unity systems at both sites. For example, if Dell EMC Unity Array 1 (U1) is replicating virtual-machine volumes to Dell EMC Unity Array 2 (U2), the credentials that the SRAs use must have administrator privileges to both systems U1 and U2. Figure 5 shows the default administrator credential that is used to manage the Dell EMC Unity system.

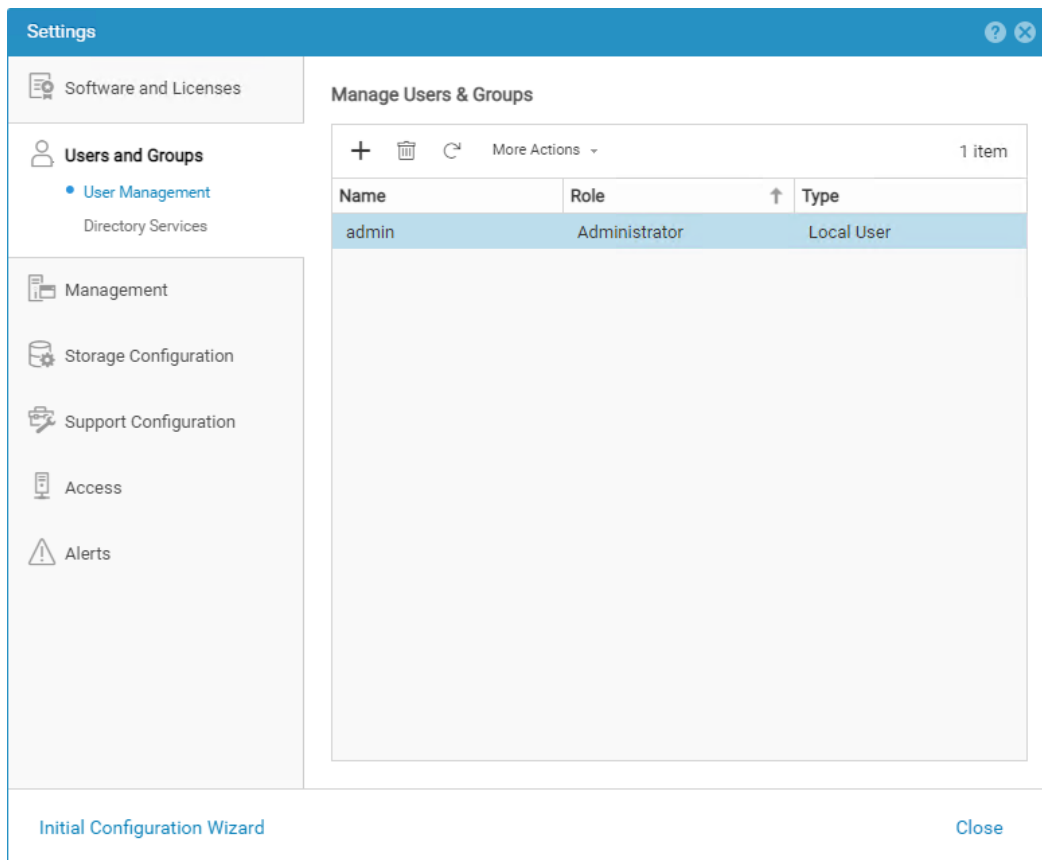


Figure 5 Manage Users & Groups menu in Unisphere

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 Storage Administrator privileges at a minimum, and the password assigned must meet Dell EMC Unity 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 an account in Unisphere named **srmadmin**.

The screenshot shows the 'Create User or Group' dialog box with the 'Specify User Information' step selected. The left sidebar shows a progress indicator for 'User or Group Type' (checked) and 'User Information' (selected). The main area contains three input fields: 'Username' with the value 'srmadmin', 'Password' with masked characters, and 'Confirm Password' with masked characters. At the bottom right are 'Cancel', 'Back', and 'Next' buttons.

The screenshot shows the 'Create User or Group' dialog box with the 'Select a Role' step selected. The left sidebar shows 'User or Group Type' and 'User Information' (checked), and 'Role' (selected). The main area lists several roles with radio buttons: 'Administrator', 'Storage Administrator' (selected), 'Operator', 'VM Administrator', and 'Security Administrator'. Each role has a brief description of its permissions. At the bottom right are 'Cancel', 'Back', and 'Next' buttons.

You can now use the **srmadmin** account within the SRM Array Manager configuration.

Note: Each Dell EMC Unity system, whether it is at the protected or remote site, maintains its own user access database. Credentials are required for Unity systems 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>
```


To configure this value per cluster, in **vSphere DRS Advanced Options**, edit the **srmMaxBootShutdownOps** value. This value overrides a value specified in the vmware-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 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

Dell EMC Unity 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. Dell EMC Unity supports synchronous and asynchronous replication for both block and file volumes. Dell EMC Unity also supports manual replication, which does not automatically update a destination image with changes on the source. You can configure Dell EMC Unity replication using Unisphere, Unisphere CLI, or REST API.

When the replication interfaces are created and cabled to the network on both systems, you can make the remote system replication connection between the arrays (Figure 6). Fibre Channel or Ethernet interfaces are used depending on synchronous or asynchronous replication. After you configure a remote system on one of the systems participating in replication, it is automatically created on the peer system.

Figure 6 Creating a replication connection to a remote system

After the replication connection is established between two Dell EMC Unity systems, storage resources consumed by vSphere can be replicated. See the following sections to determine which storage resources can be replicated based on the replication method.

Note: See the document [Dell EMC Unity: Replication Technologies](#) for complete coverage of replication deployment, configuration, and considerations.

5.1 Synchronous replication

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

Native synchronous replication is supported on the following storage resources:

- LUNs
- Consistency groups
- VMware VMFS datastores
- Thin clones

Also, MetroSync replication is supported on the following storage resources:

- NAS servers
- File systems
- VMware NFS datastores

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

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

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

The figure consists of two screenshots of the 'Create a Session' configuration wizard in Unisphere.

Top Screenshot: Provide a Replication Mode and RPO

- Replication Settings** (selected in the left sidebar):
 - Destination
 - Summary
 - Results
- Replication Mode:** Synchronous (dropdown menu)
- Replicate To:** U2 (dropdown menu)
- Help Text:**
 - To create a remote replication connection, navigate to the Protection & Mobility > Replication > Connections section.
 - To create interfaces for use with remote replication, navigate to the Protection & Mobility > Interfaces section.
 - Ensure that Ethernet ports, link aggregations and FSN ports with same names are available on destination, otherwise use override to change port assignment to NAS interfaces.
- Buttons:** Cancel, Next

Bottom Screenshot: Provide destination storage resource info

- Replication Settings** (checked in the left sidebar)
- Destination** (selected in the left sidebar):
 - Summary
 - Results
- Name:** srmids1 (text input)
- Destination System:** U2 (text input)
- Pool:** Extreme Performance (Dynamic, Extreme Performance Pool, 14) (dropdown menu)
- Size(GB):** 750.0 (text input)
- Thin
- Data Reduction
- Buttons:** Cancel, Back, Next

Figure 7 Using Unisphere to replicate a VMFS datastore synchronously

5.2 Asynchronous replication and snapshots

Asynchronous replication is primarily used to replicate data over long distances, but it can also be used to replicate storage resources between pools within the same system. Asynchronous replication does not impact host I/O latency since host writes are acknowledged when they are saved to the local storage resource. Because write operations are not immediately replicated to a destination resource, all writes are tracked on the source. This data is replicated during the next synchronization. Asynchronous replication introduces the concept of a Recovery Point Objective (RPO). The RPO is the acceptable amount of data, measured in units of time, which may be lost due to a failure. This delta of time also affects the amount of data which needs to be replicated during the next synchronization and the amount of potential data loss if a disaster occurs. Asynchronous replication image synchronizations are triggered by a user-defined RPO (individually configurable per volume) or at any time manually by the user.

When an asynchronous replication session is created, a full synchronization of the source and destination storage resource occurs. If replication is configured when a new resource is being created, the synchronization is quick because no data needs to be copied to the destination storage resource. If 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 yet, but they are tracked for a later synchronization. After the initial synchronization is complete, a common base is established between the source storage resource and the destination. When creating a manual replication session, an initial synchronization does not automatically start. To synchronize the local and remote images, a manual sync must be initiated. Host write operations which occur after the initial synchronization are acknowledged with the host normally, and no data is replicated to the destination until the next sync. Manually, and later at the RPO interval for asynchronous replication, all changes made to the source storage resource since the last synchronization are replicated to the destination. A new common base is established. If a failure is encountered on the source, all data that is not copied to the destination is lost because the changes have not been copied to the destination.

Asynchronous replication in Dell EMC Unity storage uses Unified Snapshots to maintain the common base images. As detailed in the document [Dell EMC Unity: Replication Technologies](#), the following steps and Figure 8 show how Unified Snapshots are used with asynchronous and manual replication.

1. When a replication session is created on a storage resource, two internal Unified Snapshots are created on the source and destination storage resources. At time of creation, Snapshot 1 and Snapshot 2 on the source have the same contents as the source storage resource.
2. Data is replicated from Snapshot 1 to the newly created destination storage resource. This step is the initial synchronization of the source and destination storage resources and is a full copy of all the data.
3. After the initial synchronization is complete, Snapshot 1 is refreshed on the destination storage resource. Snapshot 1 on the source and destination storage resource contains the same information and represents the point in time in which the synchronization started. Snapshot 1 on each system is now a common base for the replication session.
4. Over time, the host application writes new data to the source storage resource.
5. During the next update, Snapshot 2 on the source storage resource is refreshed to reflect the current point-in-time view of the source storage resource. This step is either manually started or determined by the RPO with asynchronous replication. All changes that were made since the last update of the destination are copied to the destination storage resource.
6. After the incremental copy is complete, Snapshot 2 on the destination storage resource is refreshed to reflect the current information in the destination storage resource. Snapshot 2 on the source and destination contains the latest information, and the snapshots make up the latest common-base image for the replication session.

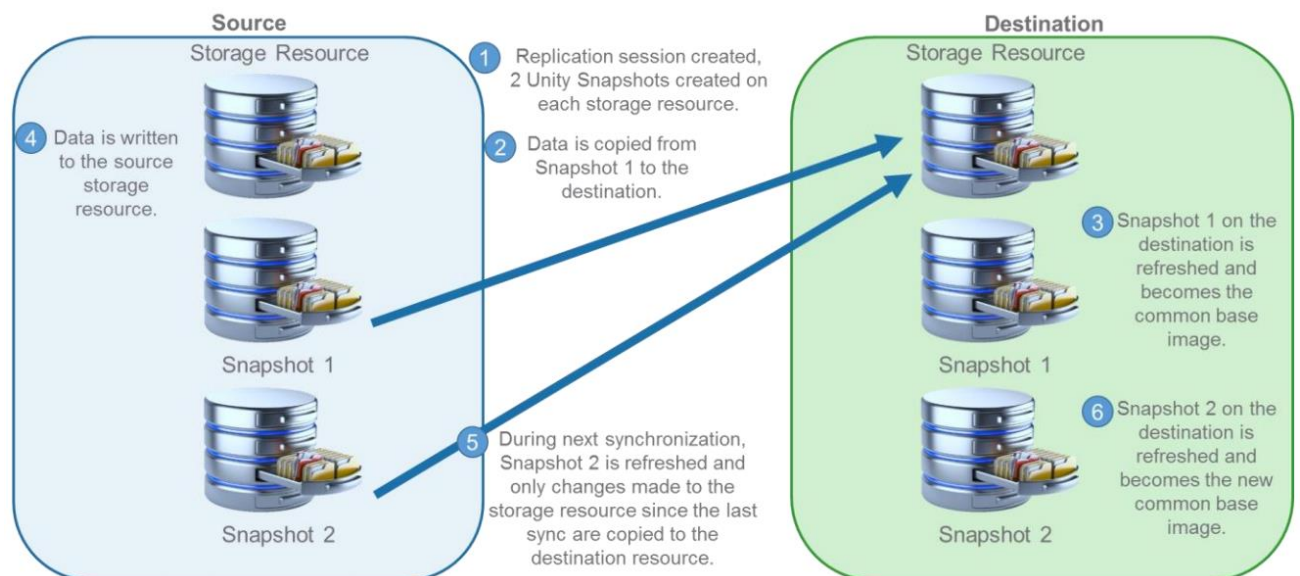


Figure 8 Asynchronous and manual replication

Each time the RPO is reached or a manual update is started, the common base image alternates between Snapshot 1 and Snapshot 2. Snapshots used for asynchronous replication behave the same as user Unified Snapshots and are based on redirect on write technology. Space required to preserve the point-in-time Snapshot is allocated from the same pool as the source storage resource. Although user snapshots and replication snapshots share the same technology, replication snapshots have restrictions on their usage. Replication snapshots can be viewed in Unisphere, but user operations such as a restore or mount operation are not allowed. Snapshots allocated for replication purposes do not count towards user-snapshot maximums.

Asynchronous replication is supported on the following storage resources:

- LUNs
- Thin clones
- Consistency groups
- VMware VMFS datastores
- File systems
- NAS servers
- VMware NFS datastores

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

Note: Consistency groups are treated as single entities when they are replicated. For virtual machines or tiered applications that span multiple volumes, consider using consistency 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.

Asynchronous replication for LUNs, thin clones, and VMware VMFS datastores function the same. When configuring asynchronous replication on a LUN in Unisphere, a single replication session is created. Then, the destination storage resource is created the same size and type as the source storage resource. When configuring replication session on a thin clone, the destination storage resource will be a regular storage resource and not a clone. While replication is configured, you can extend the size of LUNs, thin clones, and VMware VMFS datastores. The changes are reflected on the destination storage resource after the next sync. You can configure options such as the name or tiering policies of LUNs, thin clones, or VMware VMFS datastores differently between systems.

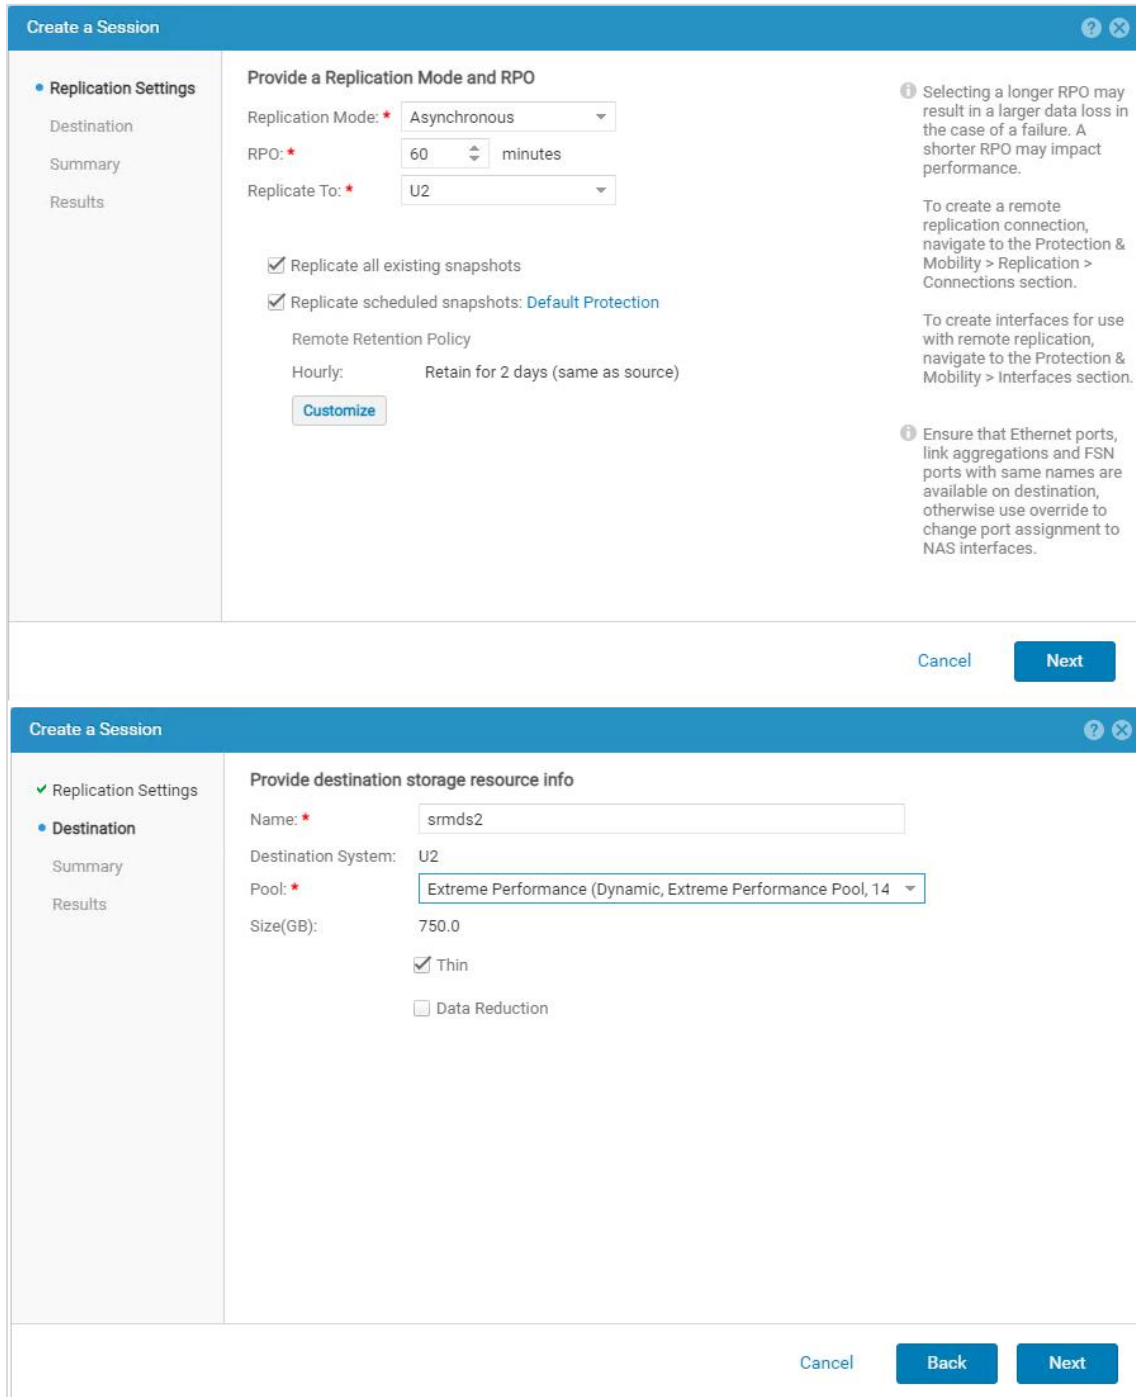


Figure 9 Using Unisphere to replicate a VMFS datastore asynchronously

5.3 VMware NFS datastore replication

Configuring VMware NFS datastore replication in Unisphere is similar to VMFS datastore replication with one addition. When replicating existing VMware NFS datastores, you must first configure replication for the NAS server that it is mounted on. If replication is configured in Unisphere for the NAS Server, all file systems and NFS datastores on the NAS server are also replicated to the destination. Later, you can delete replication sessions for resources which do not require replication. All replication sessions that are automatically configured when the NAS server is replicated have the same RPO. You can change the RPO for the individual replication sessions later. While replicated, all size changes to the file systems and NFS datastores are reflected on the destination after the next synchronization.

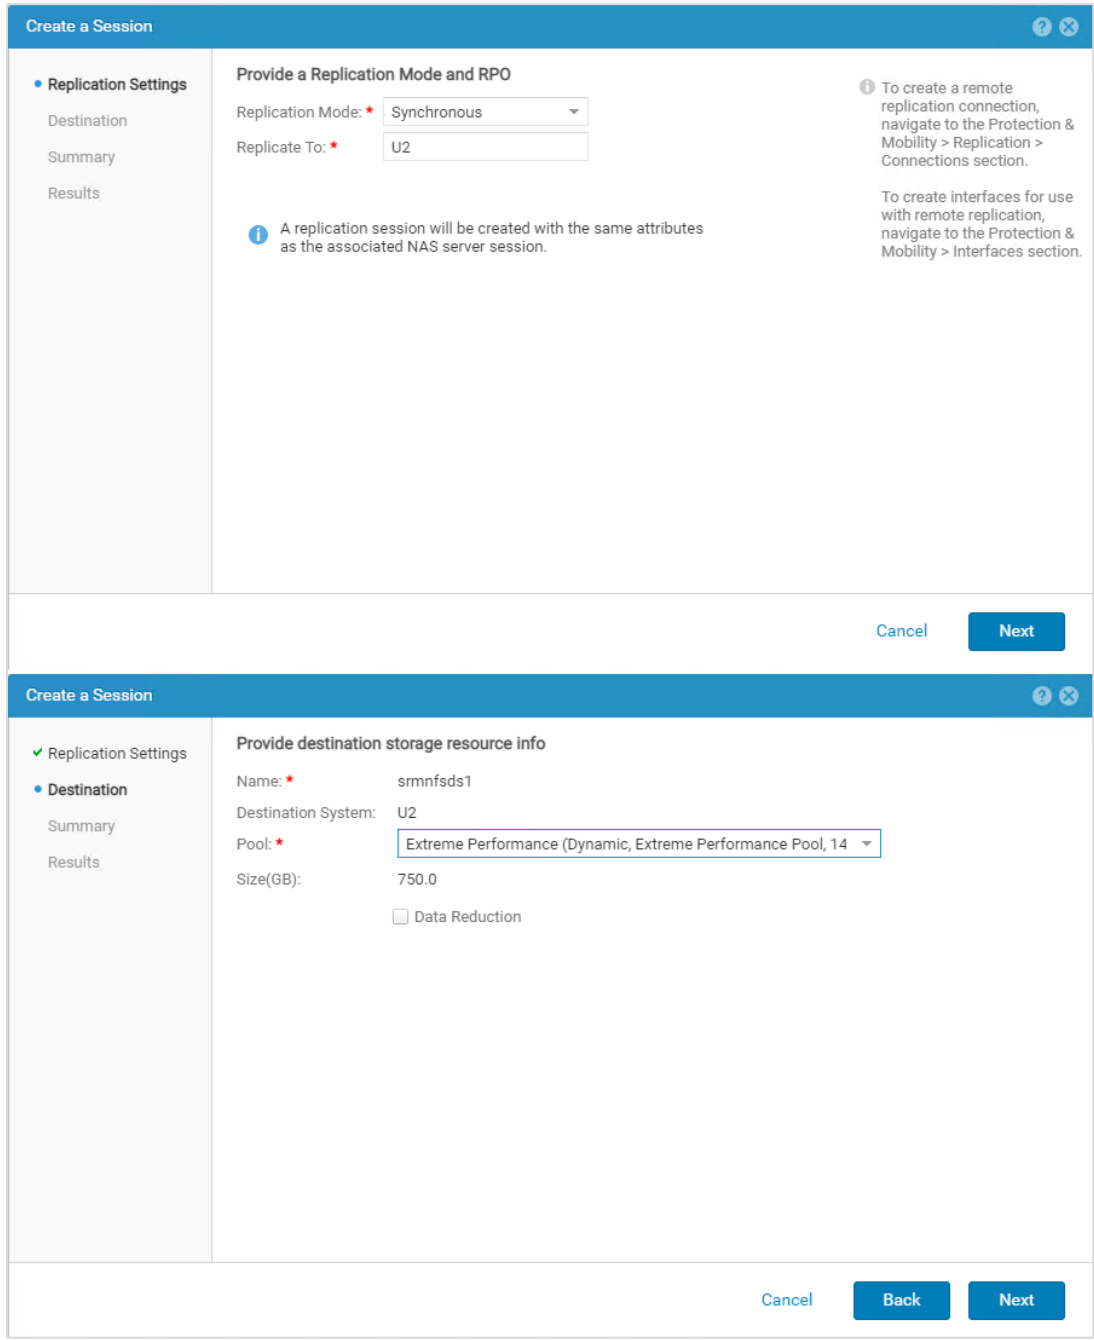


Figure 10 Using Unisphere to replicate an NFS datastore synchronously

? X
Create a Session

- Replication Settings
- Destination
- Summary
- Results

Provide a Replication Mode and RPO

Replication Mode: * Asynchronous

RPO: * 60 minutes

Replicate To: * U2

i A replication session will be created with the same attributes as the associated NAS server session.

Support Asynchronous Snap Replication

Replicate all existing snapshots

Replicate scheduled snapshots: [Default Protection](#)

Remote Retention Policy

Hourly: Retain for 2 days (same as source)

Customize

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

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

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

Cancel
Next

? X
Create a Session

- ✓ Replication Settings
- Destination
- Summary
- Results

Provide destination storage resource info

Name: * srmnfsds2

Destination System: U2

Pool: * Extreme Performance (Dynamic, Extreme Performance Pool, 14

Size(GB): 750.0

Data Reduction

Cancel
Back
Next

Figure 11 Using Unisphere to replicate an NFS datastore asynchronously

5.4 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 resources across systems. When replication is configured, any snapshots that are already created on the source resource may be replicated to the destination system (see Figure 12). There are multiple methods for creating snapshots: Unisphere, Unisphere CLI, or REST API. When replicated, SRM may use a thin clone of the snapshot to present recovered data to the vSphere cluster. Dell EMC Unity volume-based snapshots are considered to be 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 that is captured inside a replicated Dell Unity snapshot. Either of these examples results in application-consistent snapshots being replicated to the recovery site.

srmids2 Properties							
General	Host Access	Snapshots	Replication	Host I/O Limit			
Snapshots		+ - ↻ ✎ More Actions			4 items		
Snapshot Schedule	<input type="checkbox"/>	Name	State	↑ Taken	Auto-Delete	Taken By	Attached
Snapshot Access	<input type="checkbox"/>	UTC_2020-1...	Ready	12/10/2020, 2:00:01 AM	No	Default Prote...	No
	<input type="checkbox"/>	UTC_2020-1...	Ready	12/11/2020, 2:00:01 AM	No	Default Prote...	No
	<input type="checkbox"/>	8160437863...	Ready	12/10/2020, 3:12:27 PM	No	Replication	No
	<input type="checkbox"/>	8160437863...	Ready	12/10/2020, 3:12:28 PM	No	Replication	No

Figure 12 Unisphere showing VMFS datastore snapshots asynchronously replicated to the recovery site, with the Default Protection snapshot schedule of seven days

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. After 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 you are manually carrying out the previous process on a large scale, this can erode efforts made toward meeting the recovery plan RTO. This scenario 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

complete this process as a pre-power-on step (or potentially as a post-power-on step) for the VM using a custom recovery task.

5.5 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, Unisphere CLI, and REST API. 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 13).

The screenshot shows a dialog box titled "Add Step Before 'Power on priority 1 VMs'". It contains the following fields and options:

- Type:** Radio buttons for "Command on SRM Server" (selected) and "Prompt (requires a user to acknowledge the prompt before the plan continues)".
- Name:** Text input field containing "Revert VMware Snapshots".
- Content:** Text area containing the command: `C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe C:\RevertSnapshots.ps1`.
- Timeout:** Spinners for "5" minutes and "0" seconds.
- Buttons:** "CANCEL" and "ADD" buttons at the bottom right.

Figure 13 Custom recovery task in a Site Recovery Manager recovery plan

Note: For more information about REST API, see the *Dell EMC Unity REST API Programmer's Guide* and the *Dell EMC Unity REST API Reference Guide* on the [Dell EMC Unity Info Hub](#).

6 Site Recovery Manager configuration

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

6.1 SRA installation

You must install the Dell EMC Unity Storage Replication Adapter (SRA) on each SRM server. Dell EMC Unity offers SRAs for both the Photon operating-system-based SRM appliance and the Windows-based SRM installation. Also, Dell EMC Unity offers an SRA for block storage and an SRA for file storage. 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.

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

6.2 Configuring the array managers

To allow SRM to manage Dell EMC Unity storage, the SRA must be able to communicate with the Dell EMC Unity system. 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 14 and Figure 15).

The screenshot shows the VMware Site Recovery Manager interface. The left sidebar contains navigation options: Summary, Issues, Configure (with sub-options for Array Based Replication, Storage Replication Adapters, Array Pairs, Network Mappings, Folder Mappings, Resource Mappings, Storage Policy Mappings, Placeholder Datastores, Advanced Settings, Permissions, and Recovery Plans History), and Recovery Plans History. The main content area is titled 'Storage Replication Adapters' and shows two sites: vcsite1.techsol.local and vcsite2.techsol.local. Below this, there are two sections: 'EMC Unity Block SRA' and 'EMC Unity File SRA'. Each section contains a table with the following data:

Property	Value
Status	✓ OK
Version	5.0.4.146
Vendor	EMC
Install Location	C:\Program Files\VMware\VMware vCenter Site Recove...
Vendor URL	http://www.emc.com
Supported Array Models	EMC, Unity 600 EMC, Unity 500 MORE
Supported Software	Remote Replication 1.0
Stretched Storage	Not Supported

Figure 14 Examine the installed SRAs

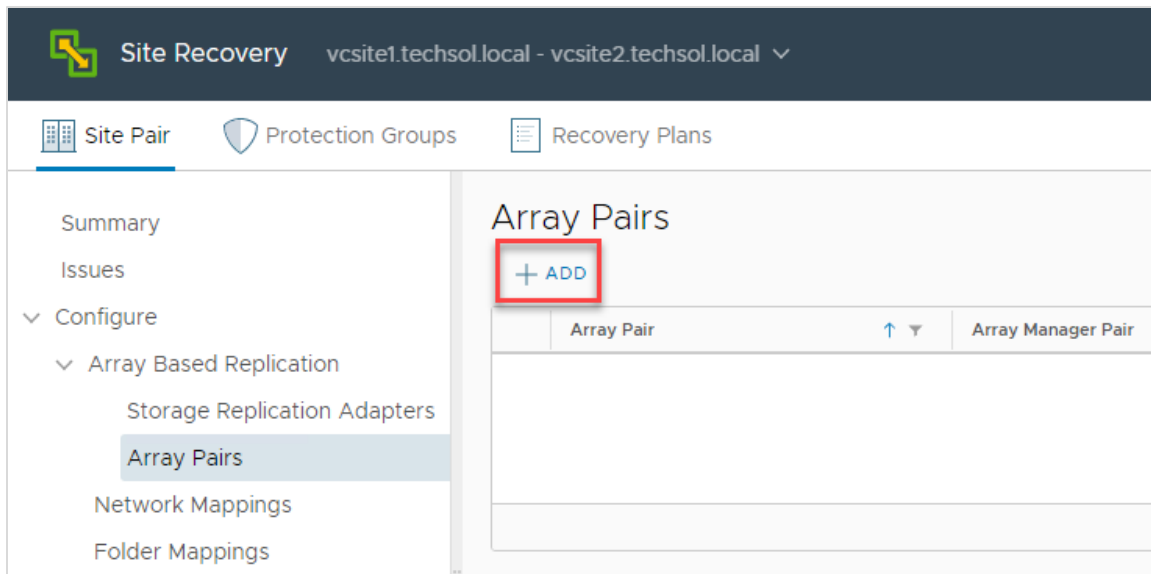
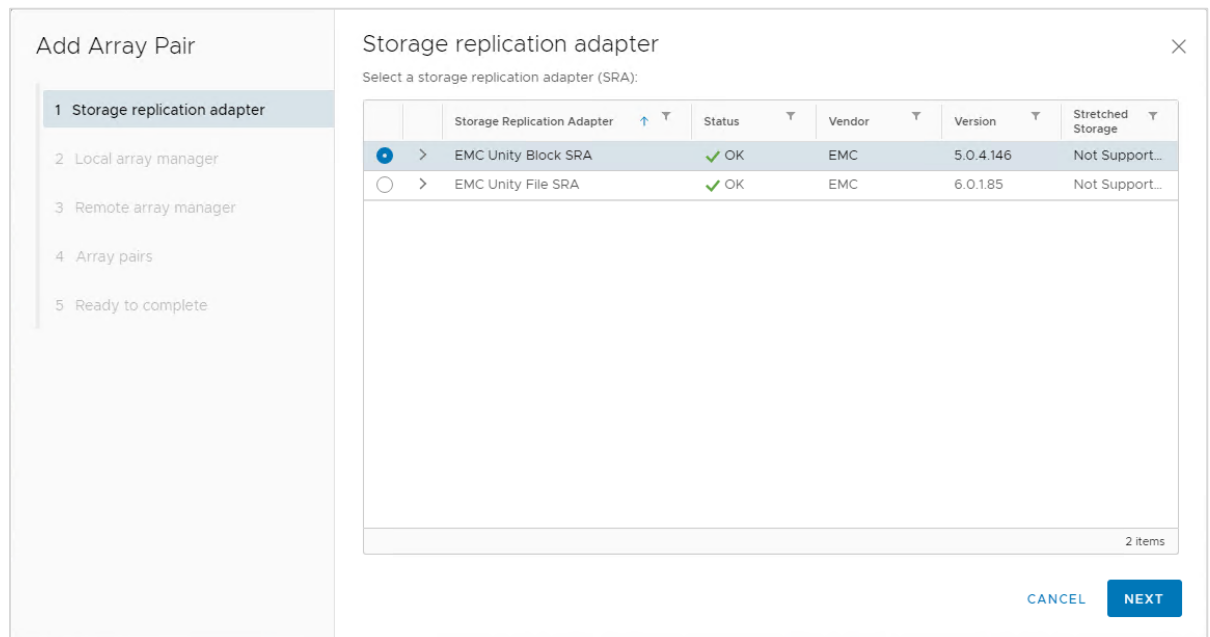


Figure 15 Adding an array manager

Complete the following steps to perform the required process to configure the protected-site array managers and the recovery-site array managers for pairing:

1. Choose the installed SRA.



2. Provide the **local** Dell EMC Unity connection parameters for the local array manager.

3. Provide the **remote** Dell EMC Unity connection parameters for the remote array manager.

Note: Each array pair uses a specific array manager. If both block and file SRAs are used with SRM, two array pairs are created (one for block, one for file). It is helpful if the array pair name identifies the protocol or array manager it is using as shown in the example above.

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 SRM can discover replicated volumes as eligible devices. See Figure 16. In older versions of SRM, pairing was an action that was performed after configuring the array managers. However, as of SRM 5.8, you can perform pairing as part of the process of adding array managers to sites.

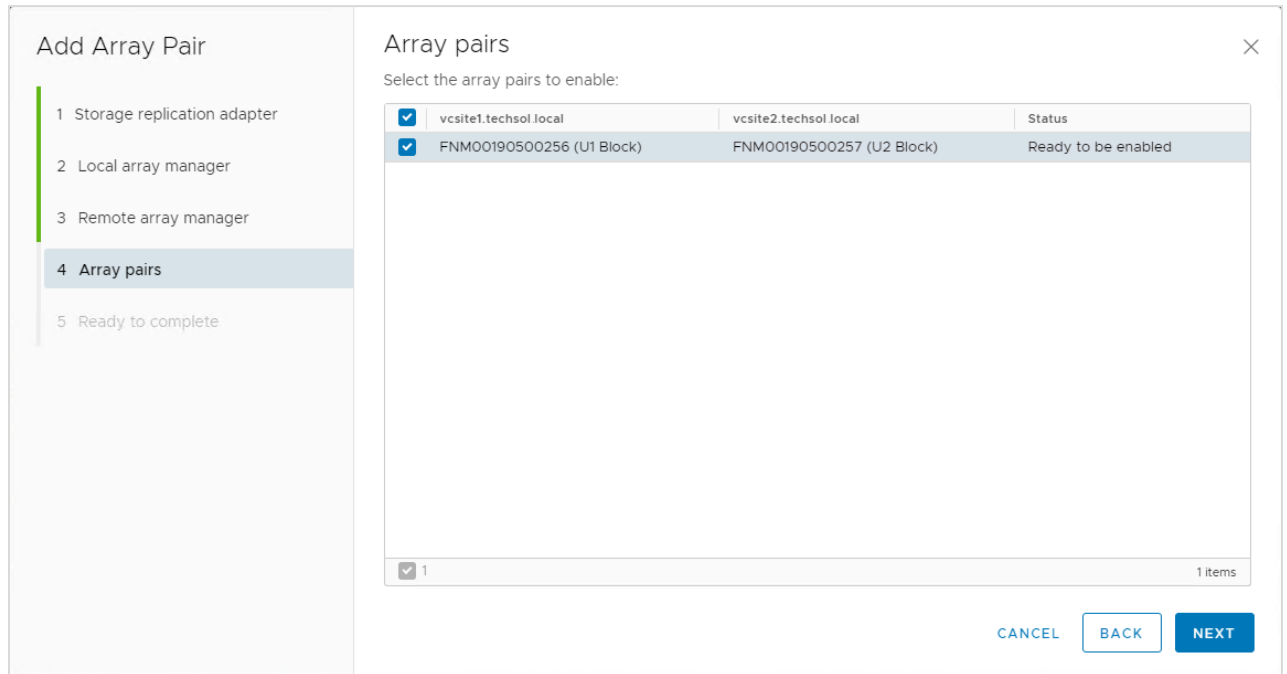


Figure 16 Creating array pairs

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 in the **Array Based Replication > Array Pairs** menu. Run the device discovery on both arrays to ensure a consistent list of devices. Nonreplicated volumes are not discovered nor 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 if vSphere does not use them. For example, replicated volumes could belong to other storage hosts such as physical Microsoft Exchange, SQL Server®, or Oracle® hosts, and file servers.

To obtain the newest array-based replicated device information, select **Discover Devices** as shown in Figure 17 to invoke an SRA query of the Dell EMC Unity system.

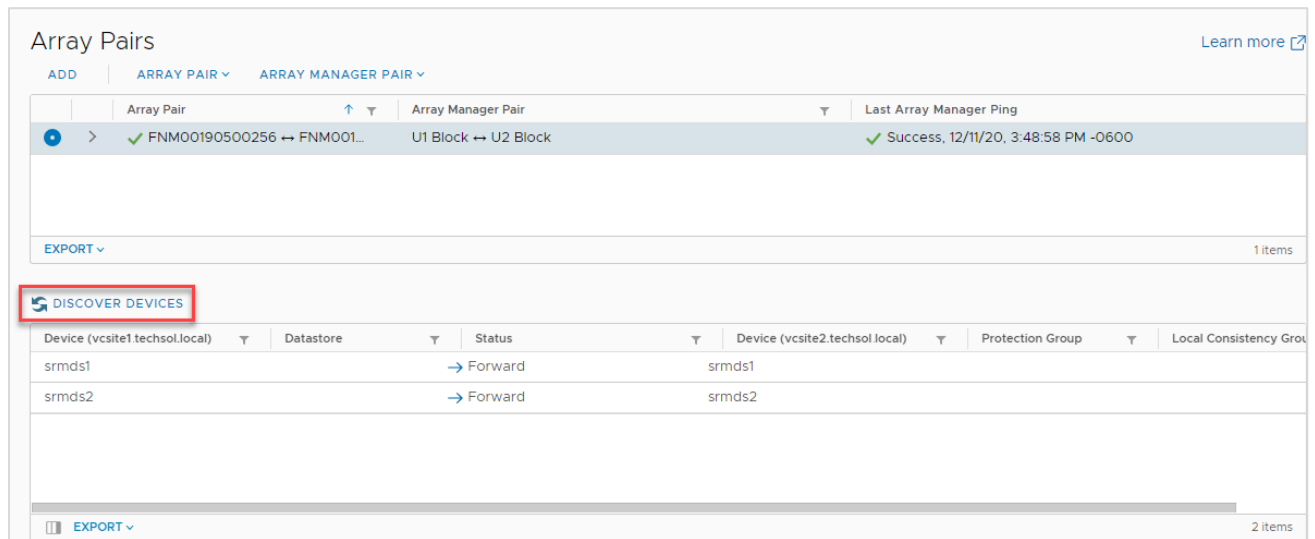


Figure 17 Discover the devices of array pairs.

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. Dell EMC Unity thinly provisions both block and file volumes by default as a space-efficient standard.

Note: The minimum VMFS volume size is 2 GB. However, the minimum Dell EMC Unity volume size is 10 GB. The placeholder datastore should not be used for other types of file storage outside of SRM, or vSphere datastore-capacity alarms may be triggered in the vSphere Client UI.

Typically, only one placeholder datastore per site is required because the disaster recovery and migration processes unregister and re-register the recovered VM 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. When a VM is protected, it is 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 Dell EMC Unity REST API 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 18).

The screenshot displays the 'Exchange Servers' interface in Site Recovery Manager. The 'Recovery Steps' tab is active, showing a plan that is currently paused. The 'Plan status' is 'Waiting for user input' with a progress bar at 1%. The 'Description' states: 'The test has paused at a user prompt. Dismiss the prompt to resume the test.' A 'Prompts' section shows a warning icon and the text '*** Press Any Key To Continue ***' with a 'DISMISS' button. Below this is a table of recovery steps:

Recovery Step	Status	Step Started	Step Completed
1. Prompt: *** Press Any Key To Continue ***	Waiting for user input	Thursday, July 11, 2019 8:59:01 AM	0%
2. Synchronize storage			
3. Restore recovery site hosts from standby			
4. Suspend non-critical VMs at recovery site			
5. Create writable storage snapshot			
6. Configure test networks			
7. Power on priority 1 VMs			
8. Power on priority 2 VMs			
9. Power on priority 3 VMs			
10. Power on priority 4 VMs			
11. Power on priority 5 VMs			

Figure 18 Recovery plan added step prompting to continue

Note: For more information about REST API, see the *Dell EMC Unity REST API Programmer's Guide* and the *Dell EMC Unity REST API Reference Guide* on the [Dell EMC Unity 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 19).

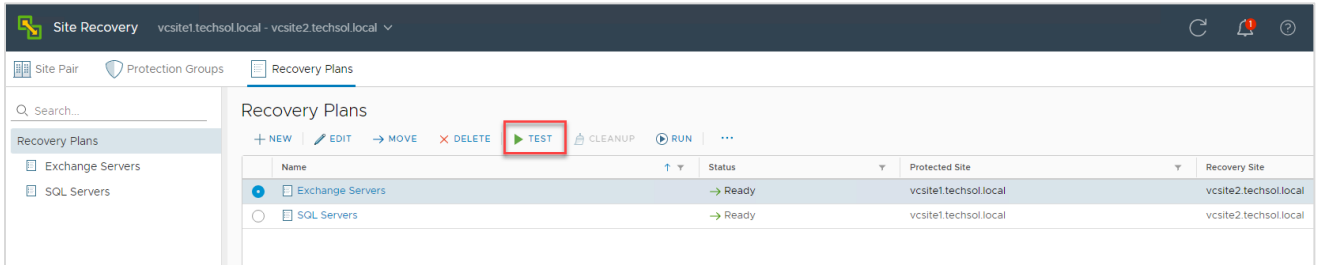


Figure 19 Testing a recovery plan

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 20).

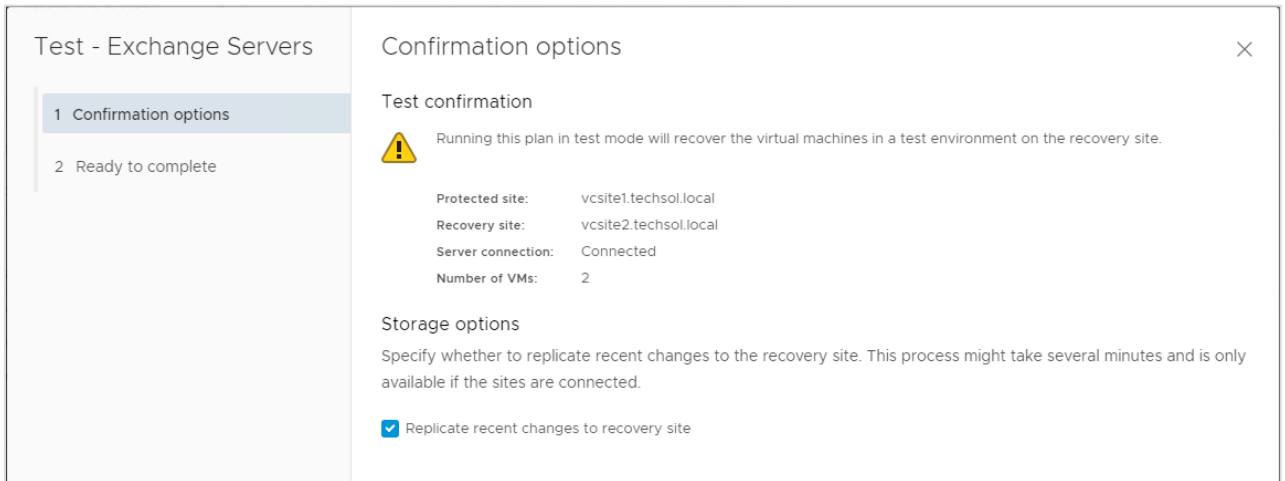
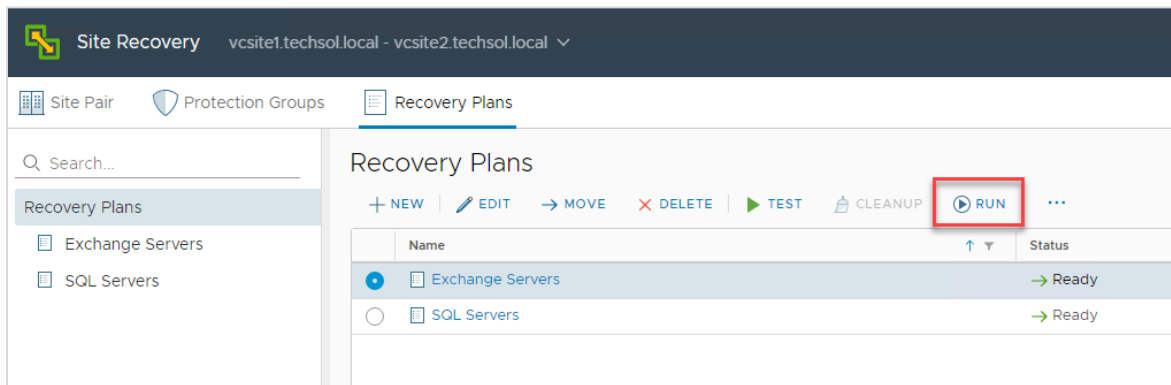


Figure 20 Replicate recent changes to recovery site during test plans

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

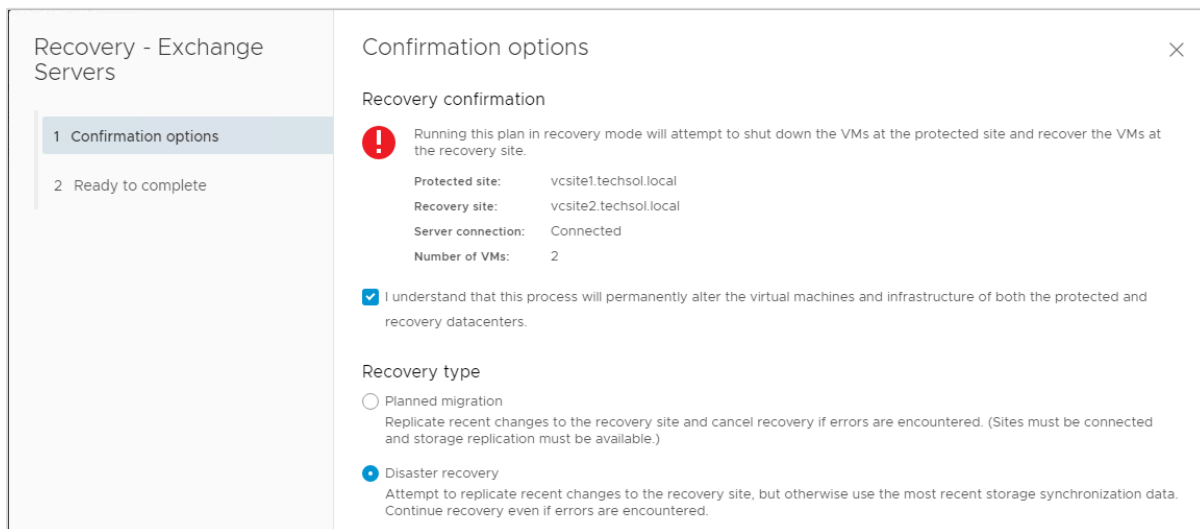
When running a planned migration or disaster recovery plan (as opposed to running a test), keep in mind this procedure is disruptive. It results in VMs being powered off at the primary site, replication mirrors being broken, and VMs being recovered at the secondary site.

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.

Exchange Servers | EDIT | MOVE | DELETE | TEST | CLEANUP | RUN | ... [Learn more](#)

Summary | **Recovery Steps** | Issues | History | Permissions | Protection Groups | Virtual Machines

EXPORT STEPS | TEST | CLEANUP | RUN | **REPROTECT** | CANCEL

Plan status: ✔ Recovery complete

Description: 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 plan to failback the virtual machines to the original site, you must first run the plan in reprotect mode, then once protection is configured in reverse, you can run the plan in recovery mode to failback the virtual machines to the original site.

View: Recovery Steps ▾

Recovery Step	Status	Step Started	Step Completed
1. Pre-synchronize storage	✔ Success	Monday, June 29, 2020 8:00:35 PM	Monday, June 29, 2020 8:01:09 PM
2. Shut down VMs at protected site	✔ Success	Monday, June 29, 2020 8:01:09 PM	Monday, June 29, 2020 8:01:23 PM
3. Resume VMs suspended by previous recovery			
4. Restore recovery site hosts from standby	✔ Success	Monday, June 29, 2020 8:01:23 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 8:01:39 PM
7. Synchronize storage	✔ Success	Monday, June 29, 2020 8:01:39 PM	Monday, June 29, 2020 8:02:13 PM
8. Suspend non-critical VMs at recovery site			
9. Change recovery site storage to writable	✔ Success	Monday, June 29, 2020 8:02:13 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 8:02:29 PM	Monday, June 29, 2020 8:02:32 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 completed in a series of automated steps (see Figure 21).

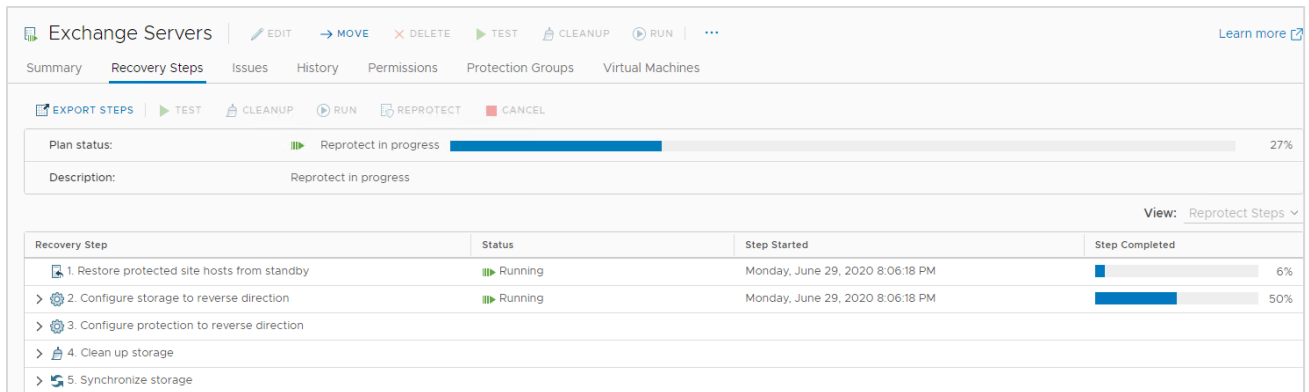


Figure 21 Reprotect workflow of the Exchange Servers protection group

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 22).

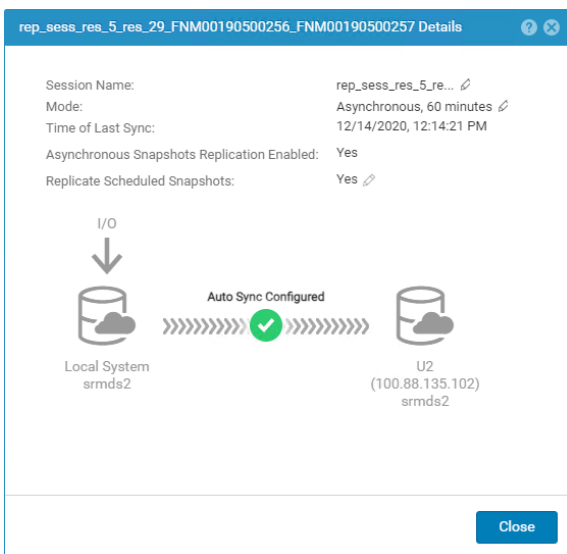


Figure 22 Replication direction and session status can be monitored in Unisphere.

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, perform resource balancing, and improve the ROI of the disaster-recovery infrastructure.

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

9 Conclusion

VMware vSphere, Site Recovery Manager, and Dell EMC Unity storage 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](#) is focused on meeting customer needs with proven services and support.

[Storage and data protection technical documents and videos](#) provide expertise that helps to ensure customer success with Dell EMC storage and data protection products.

The [Dell EMC unity Info Hub](#) provides detailed documentation about how to install, configure, and manage Unity systems.

See the following Dell EMC Unity snapshot and replication-related resources:

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

A.2 VMware support

For VMware support, see the following resources:

- [VMware.com](#)
- [VMware support](#)
- [Education and training](#)
- [Online documentation](#)
- [VMware communities](#)