

Dell EMC SC Series Best Practices with VMware Site Recovery Manager

Abstract

This document offers best practices for automated disaster recovery of virtualized workloads using Dell EMC™ SC Series arrays, Replay Manager, array-based replication, and VMware® Site Recovery Manager™ with varying levels of consistency.

August 2017

Revisions

Date	Description
August 2011	Initial release
November 2011	Updated for Storage Center OS 5.5.4
December 2011	Updated replication sections
March 2012	SRA version correction
July 2012	Added warning
October 2012	Updated diagrams
October 2012	Updated for SRM 5.1
April 2013	Updated for Storage Center OS 6.3 and sync replication support
July 2013	Corrected two section titles
December 2013	Added selectable replay, QoS, user accounts, and revert to snapshot details
August 2014	Updated Storage Center OS version information
March 2015	Updated for SRM 5.8
July 2016	Updated for DSM 2016 R1, SRM 6.1, and Live Volume support
August 2017	Updated for DSM 2016 R3.11 and SRM 6.5

Acknowledgements

This paper was produced by the following members of the Dell EMC storage engineering team:

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.

© 2011–2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

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

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	5
1 Introduction.....	6
2 Setup prerequisites.....	7
2.1 Storage Replication Adapter.....	7
2.2 Dell Storage Manager.....	7
2.3 SC Series.....	7
2.4 VMware vSphere.....	7
3 Site Recovery Manager architecture.....	8
3.1 Array-based replication: single protected site.....	8
3.2 Array-based replication: dual protected site.....	9
3.3 Array-based replication: Live Volume stretched storage.....	10
3.4 vSphere replication: single protected site.....	12
3.5 vSphere replication: dual protected site.....	13
4 Dell Storage Manager configuration.....	14
4.1 Data Collector configuration.....	14
4.2 DSM logins.....	14
4.3 Creating dedicated SRA access accounts.....	15
4.4 Saving restore points manually.....	15
4.5 Disable automatic restore point saving.....	16
4.6 Validating restore points.....	16
4.7 Modifying SRM settings for larger environments and stretched storage.....	16
5 Configuring replication.....	19
5.1 Asynchronous replication (supported).....	19
5.2 Synchronous replication.....	19
5.3 QoS definitions.....	19
5.4 Live Volume replication.....	19
5.5 Data consistency while replicating the frozen snapshot.....	21
5.6 Data consistency while replicating the active snapshot.....	22
5.7 Replication dependencies and replication transfer time.....	23
5.8 SRM selectable snapshot.....	24
5.9 Using application- and data-consistent frozen snapshots with SRM.....	25
5.10 Custom recovery tasks.....	26

6	Site Recovery Manager configuration	27
6.1	Configuring the array managers	27
6.2	Creating array pairs	29
6.3	Array manager device discovery	29
6.4	Creating protection groups	30
6.5	Creating recovery plans	31
7	Testing a recovery plan	33
8	Reprotect and failback	35
8.1	Reprotection	35
8.2	Failback	36
9	Conclusion	37
A	Example scripts	38
A.1	REST API script: TakeSnapshot.py	38
A.2	CompCU Script: TakeSnapshot.cmd	40
A.3	SC Series command set PowerShell script: TakeSnapshot.ps1	41
A.4	Dell Storage PowerShell SDK script: TakeSnapshot.ps1	41
B	Additional resources	43
B.1	Technical support and resources	43
B.2	VMware support	43

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™ SC Series arrays, Replay Manager, array-based replication, and VMware® Site Recovery Manager™ (SRM) with varying levels of consistency.

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 SC Series solutions. In addition to basic configuration, this document also answers frequently asked questions about VMware interactions with Site Recovery Manager.

Dell EMC advises reading the Site Recovery Manager documentation provided on vmware.com before beginning an SRM implementation.

2 Setup prerequisites

Verify system requirements prior to building or upgrading your environment. To view the current product support matrix, refer to the section, "Storage Replication Adapter for VMware SRM", in the *Dell Storage Manager Administrator's Guide*, available in the Knowledge Center on the SC Series [Customer Portal](#) (login required).

2.1 Storage Replication Adapter

The SC Series Storage Replication Adapter (SRA) must be installed on each SRM server. Live Volume stretched storage support was added in Dell Storage Manager (DSM) 2016 R1 and improved in DSM 2016 R3.11. It is highly recommended to use DSM 2016 R3.11 and SRA version 16.3.10.

2.2 Dell Storage Manager

A DSM deployment is required for SRM functionality. DSM 2016 R1 or newer is required for stretched storage supported by Live Volume in SRM. DSM 2016 R3.11 further improved Live Volume stretched storage support. The Storage Replication Adapter (SRA) makes calls directly to the DSM Data Collector to manipulate the storage in order to carry out SRM requested tasks. To ensure compatibility with SRM and the SRA, refer to the section, "Storage Replication Adapter for VMware SRM", in the *Dell Storage Manager Administrator's Guide*.

2.3 SC Series

SRM and array-based replication requires two SC Series systems licensed with Remote Data Instant Replay (replication) replicating between each other in one or both directions. SRM using vSphere replication can leverage any Dell EMC storage certified for use with vSphere including SC Series storage. For the current product support matrix, refer to the section, "Storage Replication Adapter for VMware SRM", in the *Dell Storage Manager Administrator's Guide*.

2.4 VMware vSphere

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.

3 Site Recovery Manager architecture

This section provides array-based replication architecture for single- and dual-protected sites as well as Live Volume stretched storage. vSphere replication architecture is also included for comparison purposes.

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 need to be protected by SRM. The secondary site exists solely for disaster recovery purposes. The DSM Data Collector server is placed at the disaster recovery site because its availability is required by SRM to execute the recovery plan. Keep this in mind if using the SRM planned migration or reprotect feature because the DSM Data Collector server may no longer be at the recovery site and could be impacted by an unplanned outage at the protected site.

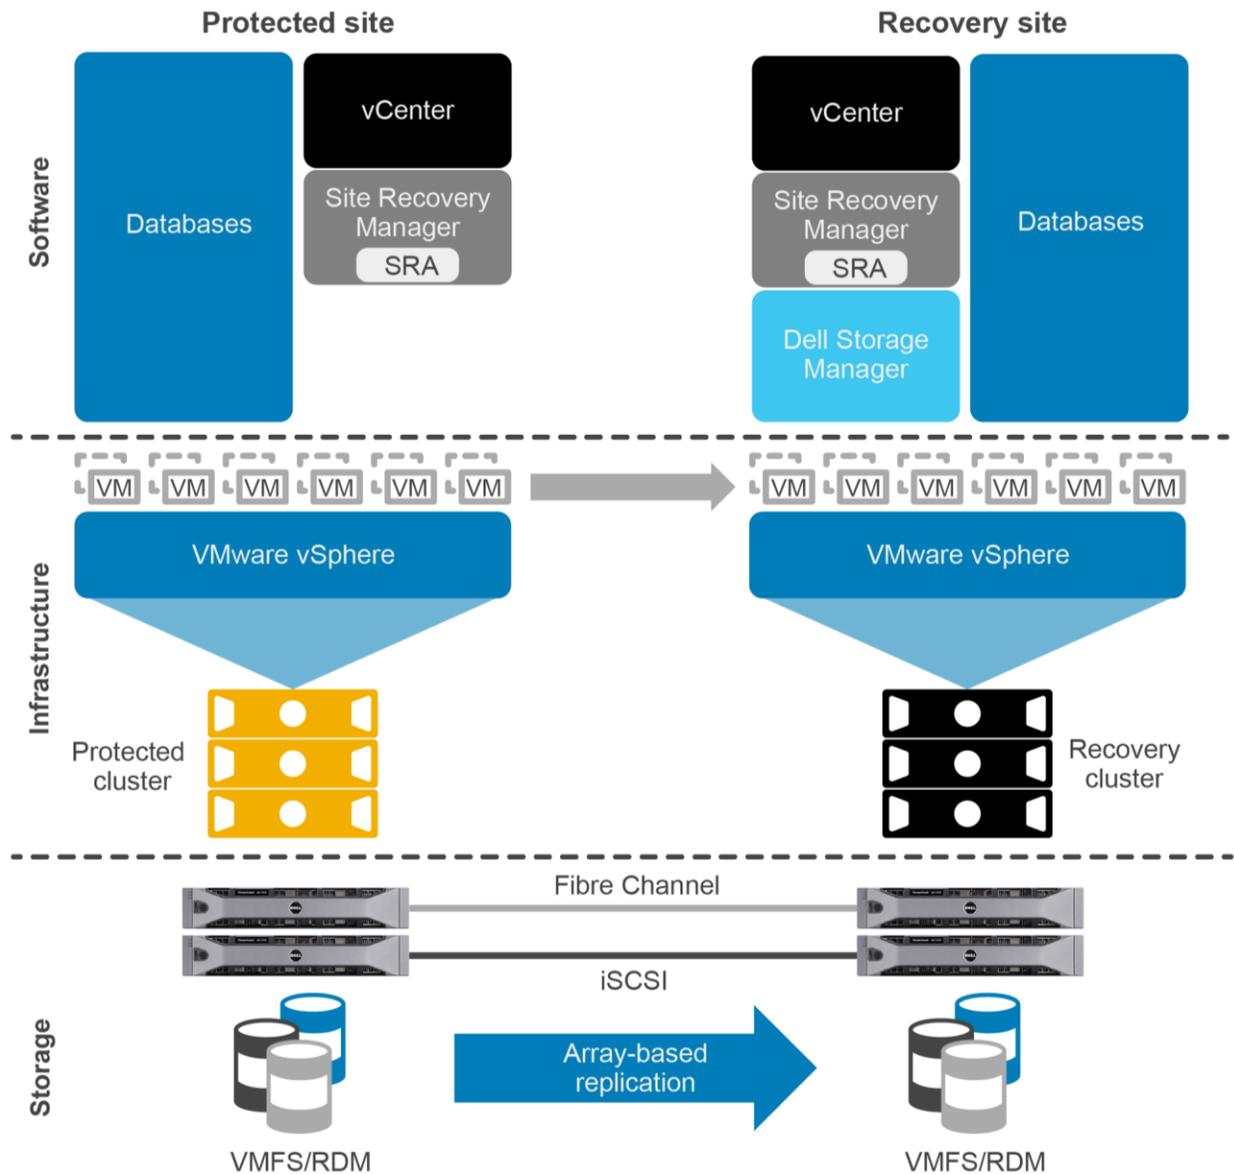


Figure 1 Architecture for a single protected site with array-based replication

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. In this example, each site has virtual machines that must be protected by SRM. Each site replicates its virtual machines to the opposing site where they can be recovered. Both a DSM Data Collector and DSM Remote Data Collector are deployed, which allows one Data Collector to be available for SRM operations in the event of an unplanned outage.

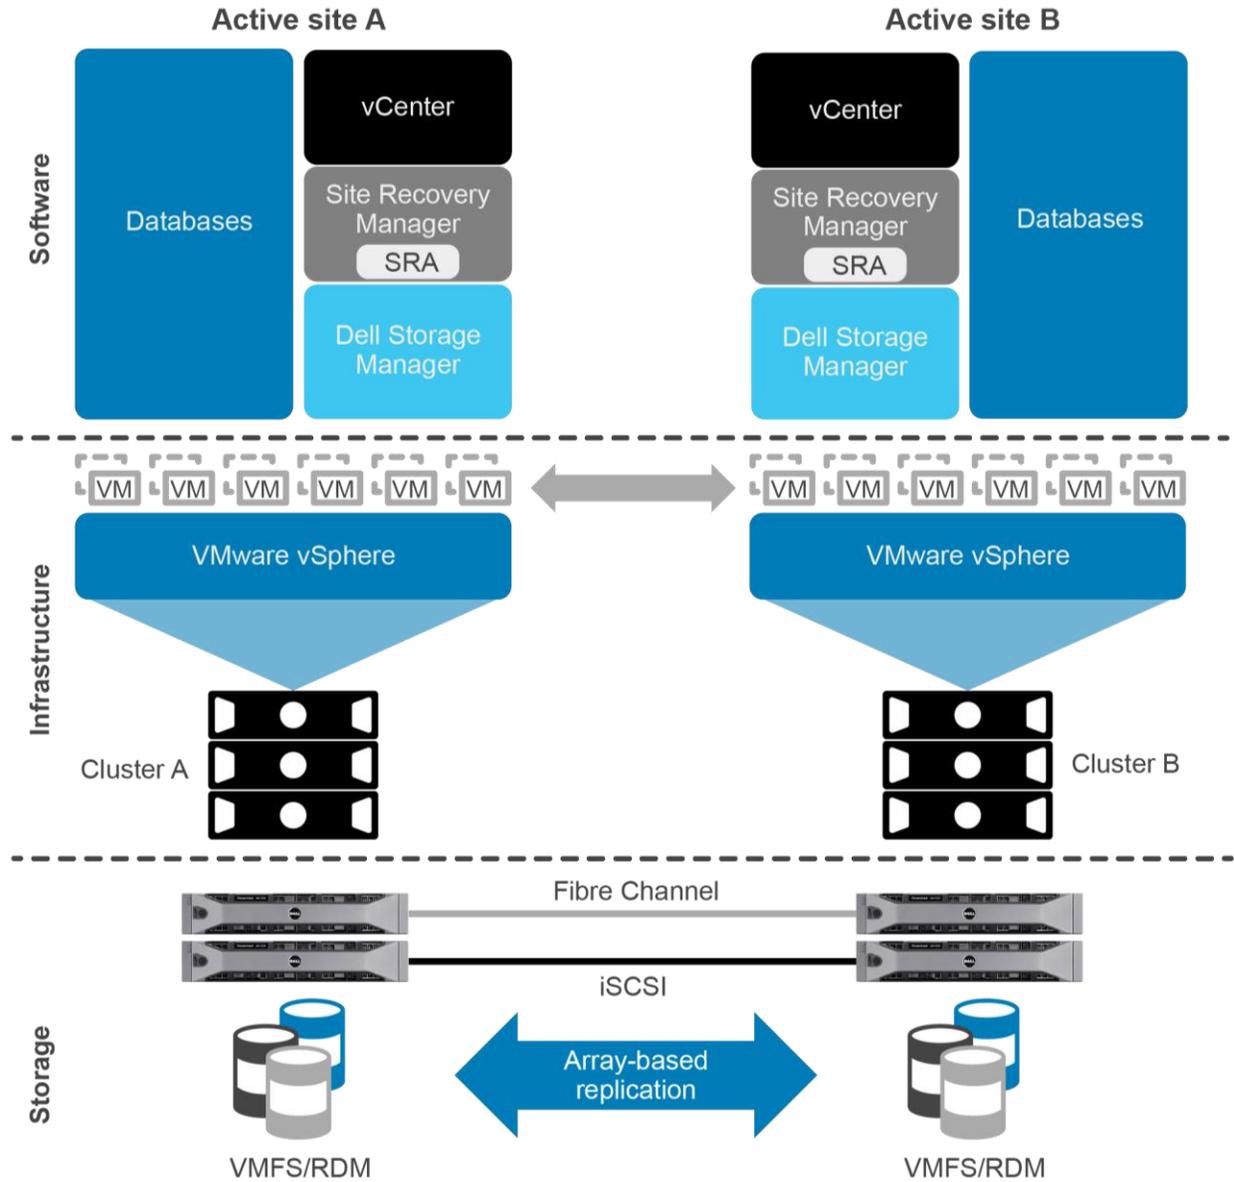


Figure 2 Architecture for a dual protected site with array-based replication

3.3 Array-based replication: Live Volume stretched storage

Live Volume stretched storage typically depicts two active sites where one site may be impacted by an unplanned outage (see Figure 3). When Live Volume is protected by SRM, a Remote Data Collector cannot be used. Because only a single Data Collector can be deployed in this configuration and the Data Collector must be available to execute SRM Recovery Plans, the DSM Data Collector should be deployed at a third site that is available through the network. This ensures that either protected site can be impacted by an unplanned outage and the Data Collector will be available to SRM for recovery. Live Volumes that are configured for automatic failover, automatic role swap, uniform storage presentation, or synchronous high consistency replication cannot be protected by SRM. Live Volume Managed Replications can be used but the Managed Replications are recovered using DSM Activate Disaster Recovery and not by the SRM process.

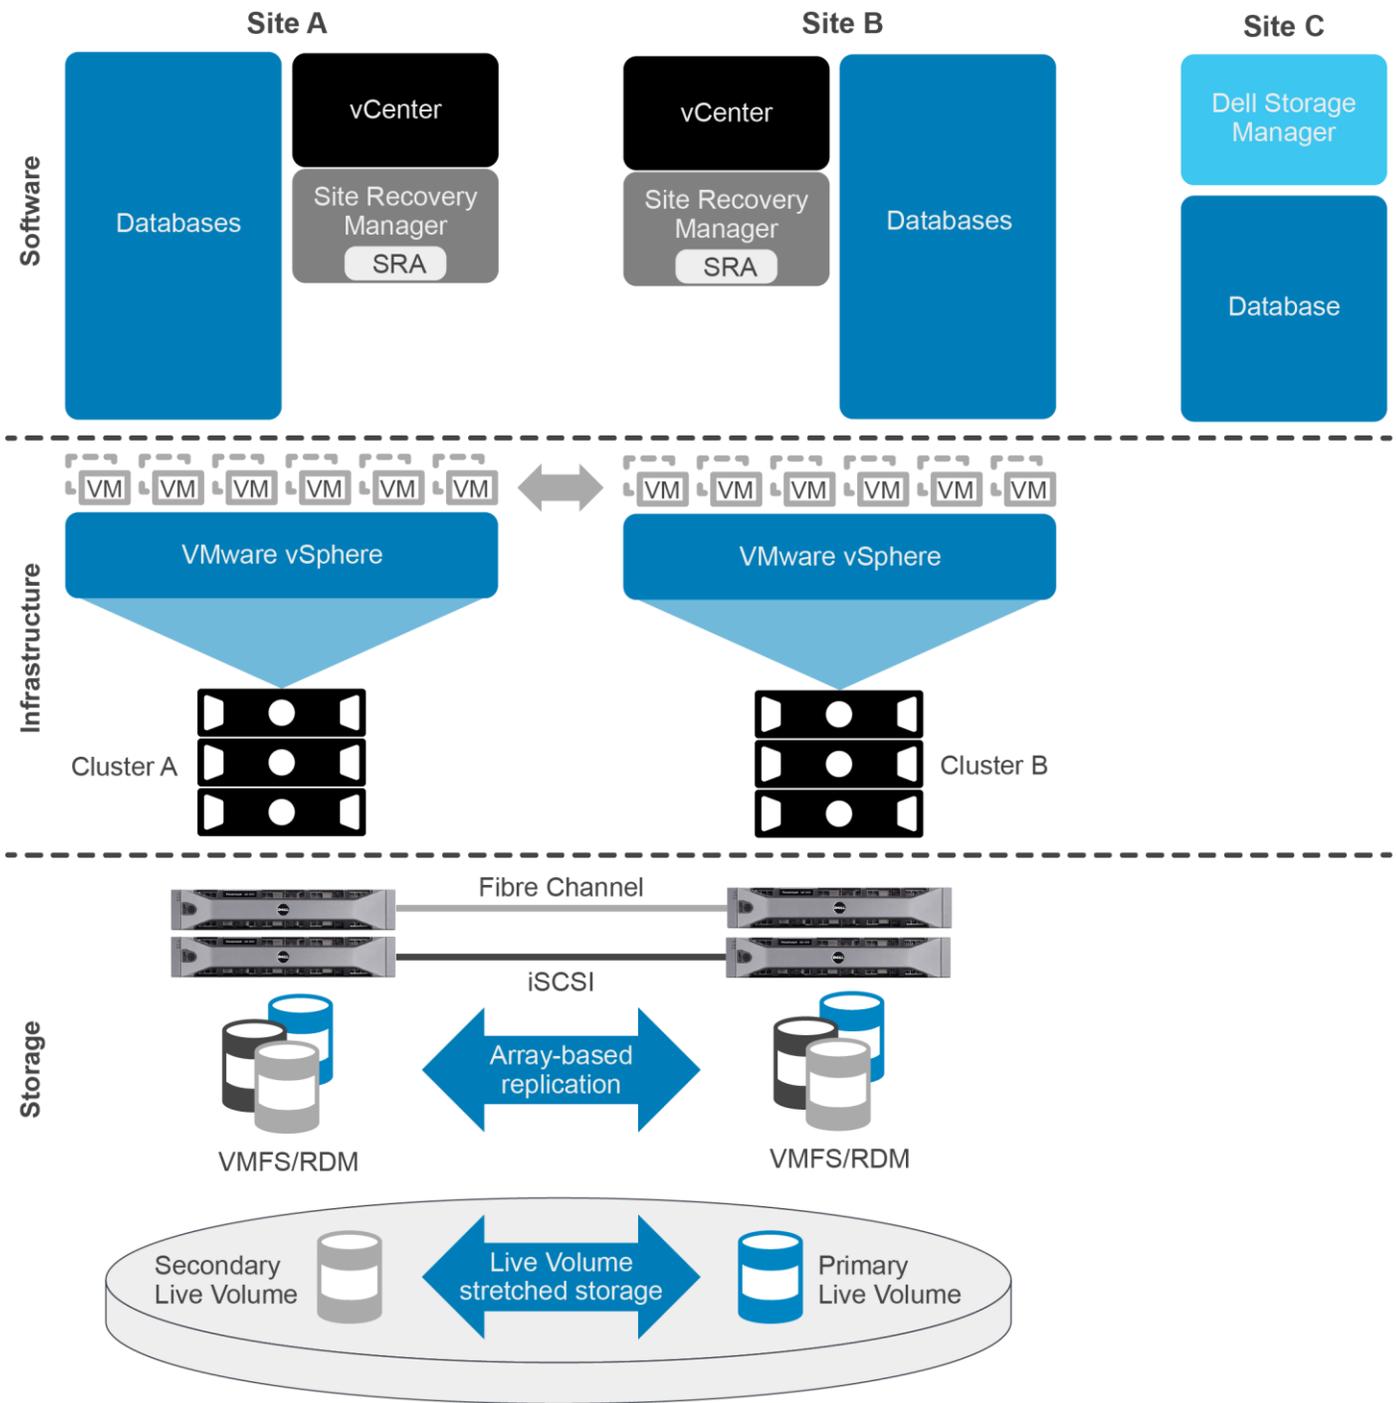


Figure 3 Architecture for sites with Live Volume stretched storage

Note: Refer to the *Dell Storage Manager Release Notes* for the latest information about stretched storage and SRM configuration.

3.4 vSphere replication: single protected site

vSphere replication can be used in addition to or in place of array-based replication. Two of the main advantages of vSphere replication over array-based replication are:

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

vSphere replication, along with other feature support for vSphere replication added in SRM 5.1, makes SRM much more appealing and adaptable as a DR solution for small- to medium-sized organizations with aggressive storage needs or budget constraints.

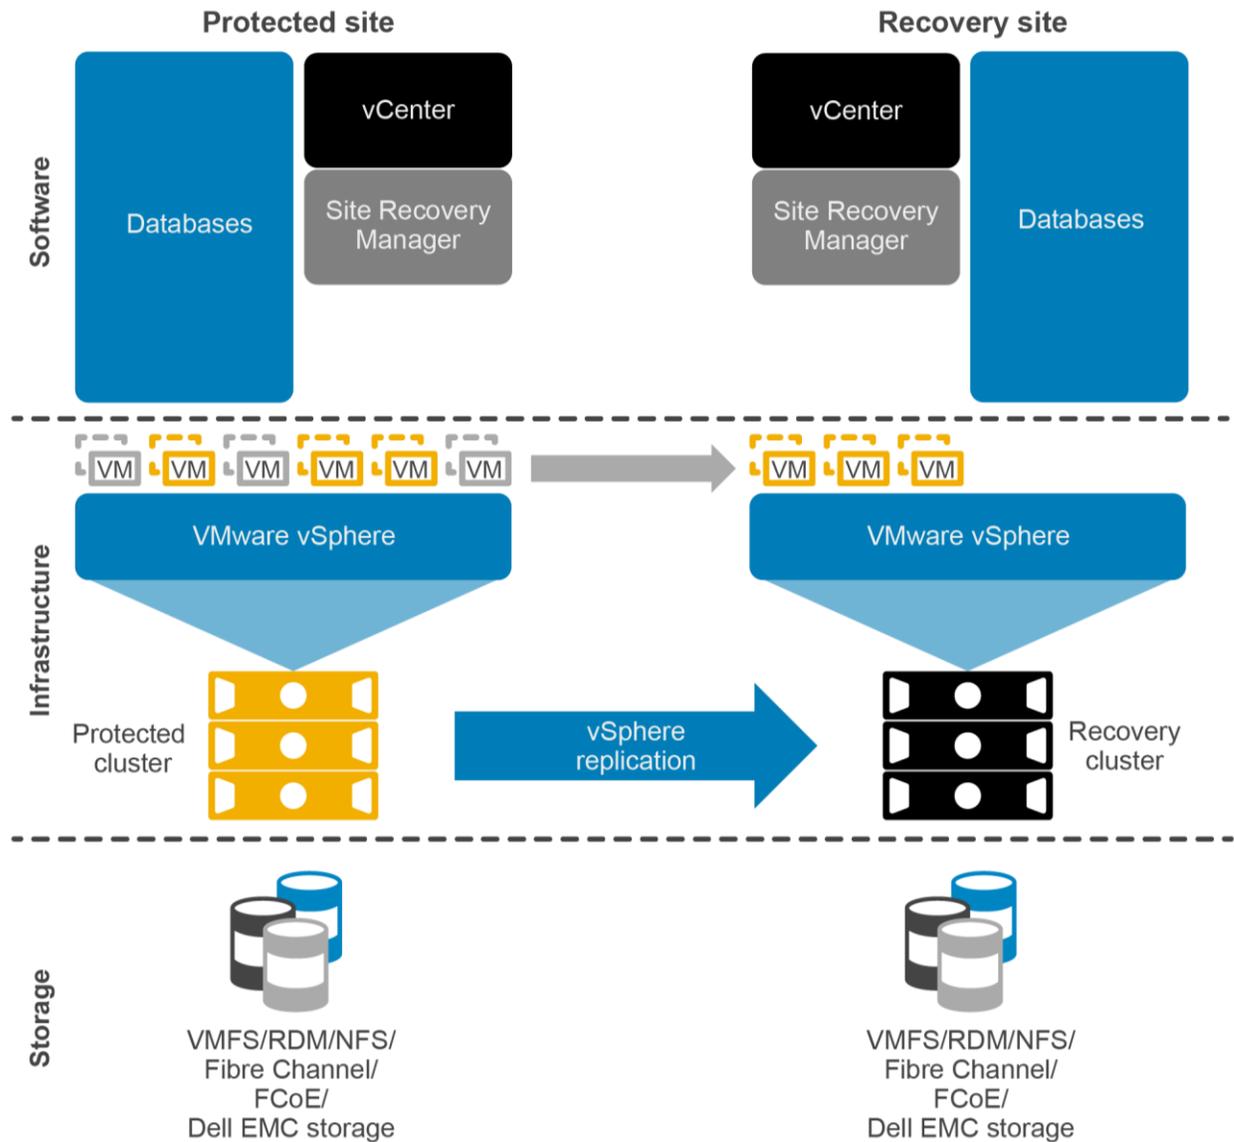


Figure 4 Configuration for a single protected site with vSphere replication

3.5 vSphere replication: dual protected site

The architectural changes with vSphere replication are carried into the active/active site model. In each vSphere replication architecture diagram (Figure 4 and Figure 5), replication is handled by the vSphere hosts using the vSphere network stack. An array-based SRA is not present in this architecture. Note that these figures do not provide representation of 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. Refer to [VMware Documentation](#) for a detailed look at vSphere replication.

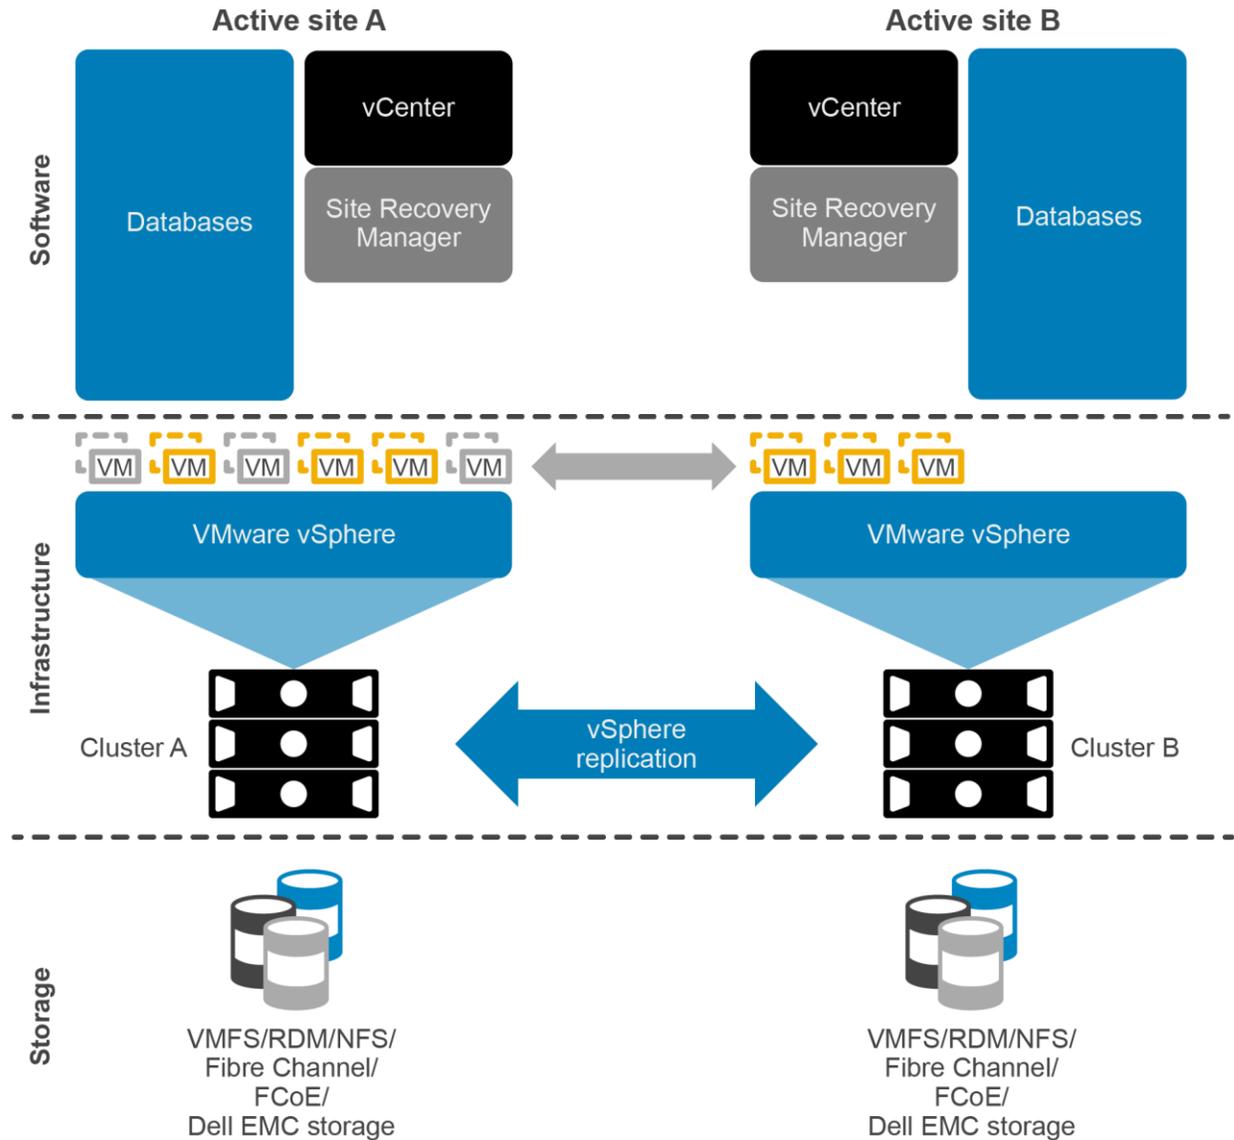


Figure 5 Configuration for a dual protected site with vSphere replication

4 Dell Storage Manager configuration

This section provides best practices for configuring Dell Storage Manager (DSM).

4.1 Data Collector configuration

As illustrated in section 3, DSM is a critical piece to the SRM infrastructure because the Data Collector processes all of the calls from the SRA and relays them to the SC Series arrays to perform the workflow tasks.

When deciding whether to use one or two DSM Data Collector servers, it depends if virtual machines need to be protected in one or multiple sites.

Single site: If protecting virtual machines at a single site, a single Data Collector will suffice; it is required that it be placed at the recovery site.

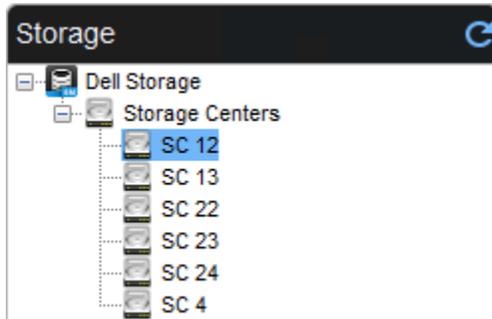
Multiple sites without Live Volume: If protecting virtual machines at both sites and Live Volume is not used, a Data Collector must be available at each site. One site would host the primary Data Collector and the other site would host the Data Collector configured for Remote Data Collector mode.

Multiple sites with Live Volume: If protecting virtual machines at both sites and Live Volume is used, a single Data Collector must be available at a third site.

4.2 DSM logins

For SRM to function, the SRA must use DSM login credentials that have rights to both SC systems replicating the virtual machine volumes.

For example, if SC12 is replicating virtual machine volumes to SC13, the credentials that the SRA uses must have administrator privileges to both systems.



Keep in mind that each Data Collector, whether primary or remote, maintains its own user access database. In a typical active/DR site configuration, a single Data Collector server is configured and deployed, managing both source and destination site arrays. A single set of credentials is needed to register that Data Collector server as an array manager for both sites.

In a typical active/active site configuration, two Data Collector servers are configured and deployed with one managing the source site array and one managing the destination site array. Two sets of credentials are needed to register each respective Data Collector as an array manager in SRM.

4.3 Creating dedicated SRA access accounts

For the SRA to have uninterrupted access to both arrays through the DSM Data Collector, it is recommended to create dedicated accounts for SRM. Using dedicated accounts on each array ensures that service is not disrupted due to a user changing their password.

Following the previous example, creating dedicated accounts requires these steps:

1. Create an account named **srmadmin** on both the protected site array and the recovery site array. This account needs administrator privileges and the password assigned must be secure. For added security, create different accounts on both systems with different passwords. The account names and passwords are arbitrary.
2. Create a new account within DSM named **srmadmin**.

User Settings

User Name	<input type="text" value="srmadmin"/>
Email Address	<input type="text" value="srmadmin@techsol.local"/>
Privilege	<input type="text" value="Administrator"/>
Preferred Language	<input type="text" value="English (US)"/>
New Password	<input type="password" value="*****"/>
Confirm Password	<input type="password" value="*****"/>
Requires Password Change	<input type="checkbox"/> Enabled

The **srmadmin** account used to access DSM can now be used for configuring the SC Series credentials within the SRM Array Manager configuration.

Note: Each DSM Data Collector and SC Series array maintains its own user account database. In an active/active SRM site configuration, two sets of credentials are used to configure the SRM array managers.

4.4 Saving restore points manually

Restore points are queried by the SRA to discover the replication state and available recovery points for data on volumes. Saving the restore points allows the SRA to retrieve the most current information about volumes replicated between sites.

Save restore points in the following scenarios:

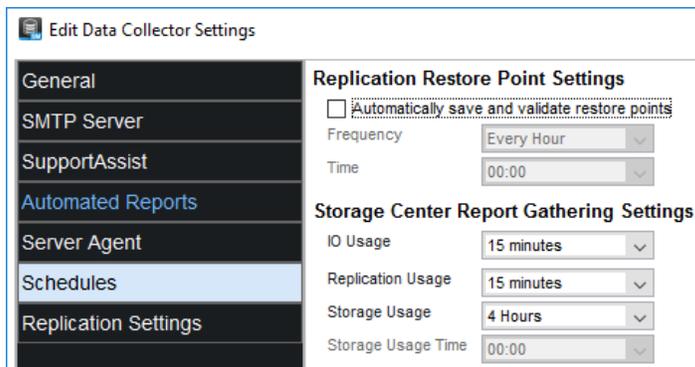
- Prior to a disaster recovery or planned migration recovery plan execution
- After a major SRM event, such as performing a planned migration or disaster recovery

From the DSM **Actions** menu, select **Save Restore Points**.



4.5 Disable automatic restore point saving

As a best practice for all types of replicated storage used with SRM, including Live Volume stretched storage, clear the checkbox labeled **Automatically save and validate restore points** in the Data Collector Settings. This will prevent restore points from becoming degraded after disaster recovery.



4.6 Validating restore points

The validate restore points process reconciles the list of saved restore points with the list of replication jobs and provides an opportunity to clean up restore points that may be orphaned or no longer needed. Select this option in the DSM **Actions** menu.



Validate restore points is performed automatically during a save restore points operation. Therefore, it is not required to perform a subsequent validate restore points operation after a save restore points operation has already been performed.

4.7 Modifying SRM settings for larger environments and stretched storage

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 within SRM (listed in the following bullet points) in order for SRM to work properly. For more information on making adjustments to accommodate such environments, see VMware KB article, [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 or environments with Live Volume stretched storage. Recovery plans with a large number of 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. For environments including Live Volume stretched storage within protection groups, this value should be increased by 90 seconds for each Live Volume managed by SRM.

- `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 re-introduced 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 a large number of datastores and/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.

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

This option specifies the timeout allowed (in seconds) for local stretched devices to be matched to the corresponding remote stretched devices. Increase this value (for example, to 1000 or higher) in the advanced SRM settings. Increasing this value is typically required for environments including Live Volume stretched storage within protection groups.

- `defaultMaxBootAndShutdownOpsPerCluster` – Default: off

This option specifies the maximum number of concurrent power-on operations performed by SRM at the cluster object level. Enable the option globally by specifying a numerical value (such as 32) by modifying the `vmware-dr.xml` file. The option can be added anywhere in the `<config>` section and the Site Recovery Manager Server service should be restarted after making a change.

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

The value can also be configured per cluster by editing the **`srmMaxBootShutdownOps`** in vSphere DRS Advanced Options. This value will override 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. Enable by specifying a numerical value (such as 4) by modifying the `vmware-dr.xml` file. The option can be added anywhere in the `<config>` section and the Site Recovery Manager Server service should be restarted 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:

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

5 Configuring replication

SC Series replication, in coordination with Site Recovery Manager (SRM), can provide a robust and scalable disaster recovery solution. Since each replication method affects recovery differently, choosing the correct method to meet business requirements is important. A brief summary of the different options is provided in this section.

5.1 Asynchronous replication (supported)

In an asynchronous replication, the I/O needs to be committed to and acknowledged by the source system in order for the data to be transferred to the destination in an independent timeframe. There are two different methods to determine when data is transferred to the destination:

By snapshot schedule: The snapshot (replay) schedule dictates how often data is sent to the destination. When each snapshot is taken, the SC Series system determines which blocks have changed since the last snapshot (the delta changes), and then transfers them to the destination. Depending on the rate of change and the bandwidth, it is entirely possible for the replications to fall behind, so it is important to monitor them and verify that the recovery point objective (RPO) can be met.

Replicating the active snapshot: This method transfers data in near real-time, which usually requires more bandwidth than replicating only the delta changes in the snapshots. As each block of data is written on the source volume, it is committed, acknowledged to the host, and then transferred to the destination as fast as possible. Keep in mind that the replications can still fall behind if the rate of change exceeds available bandwidth.

Asynchronous replications usually have more flexible bandwidth requirements making this the most common replication method. Another benefit of asynchronous replication is that the snapshots are transferred to the destination volume, allowing for checkpoints at the source system as well as the destination system.

5.2 Synchronous replication

In a synchronous replication, data is replicated in real-time to the destination. An I/O must be committed on both systems before an acknowledgment is sent back to the host. This limits the type of links that can be used, since they need to be highly available with low latencies. High latencies across the link slow down access times on the source volume.

5.3 QoS definitions

SRM supports replication from a source to a destination and can reverse the direction of replication after recovery and reprotect plans have been invoked. As a best practice, maintain consistent QoS definitions on each SC Series system in a replication pair to provide consistency and prevent unexpected replication latency. Inconsistent QoS definitions between sites may also cause reprotect or failback workflows to fail.

5.4 Live Volume replication

Standard asynchronous replication or synchronous replication (either mode) can be leveraged by VMware vSphere SRM protection groups, recovery plans, and reprotection. Live Volume replications add an abstraction layer to the replication to allow mapping of an abstracted volume derived across two SC Series arrays.

SRM support for stretched storage with Live Volume was added in DSM 2016 R1 and further improved in DSM 2016 R3.11. Supported configurations are asynchronous replication or synchronous high availability replication with non-uniform storage mapping to hosts. For more information on use cases and integrating stretched storage with SRM, see the *Site Recovery Manager Administration* guide at [VMware Documentation](#).

The use of Live Volume stretched storage in protection groups requires asynchronous or synchronous high availability replication and non-uniform storage mapping. Live Volume automatic role swap and automatic failover must be disabled.

The following are not supported in conjunction with SRM: synchronous replication in high consistency mode, uniform storage mapping, managed replication, consistency groups, Live Volume automatic role swap, and Live Volume auto failover.

Using DSM to perform Activate Disaster Recovery of a Live Volume managed replication is a good use case for remote disaster recovery of a Live Volume. However, Live Volume managed replications are not explicitly supported with vSphere SRM. Live Volume managed replication volumes are activated using DSM.

5.5 Data consistency while replicating the frozen snapshot

Figure 6 and the steps that follow describe the consistency states of replications during plan execution while replicating a frozen snapshot.

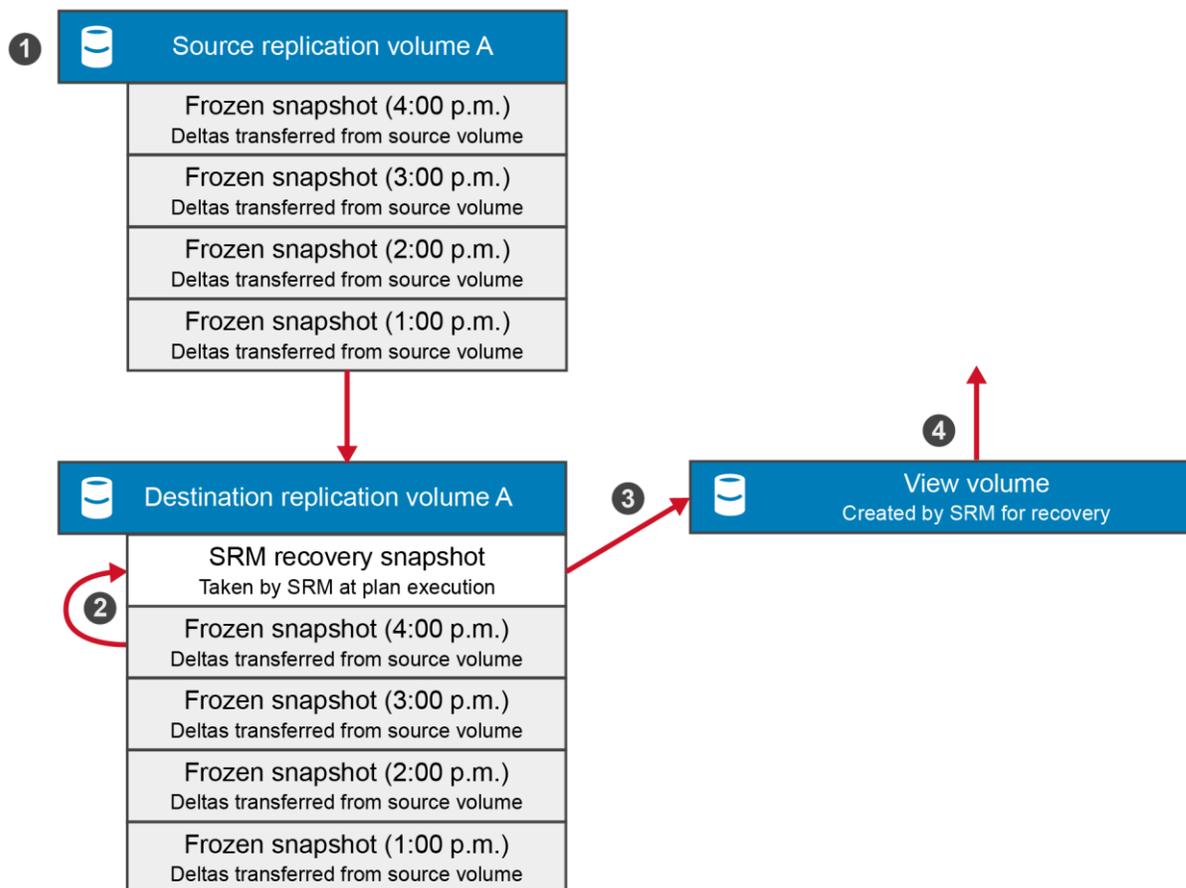


Figure 6 Consistency states of frozen snapshot replications during plan execution

- Once a snapshot is taken of the source volume, the delta changes begin transferring to the destination immediately. The consistency state of the data within this snapshot is dependent on whether or not the application had the awareness to quiesce the data before the snapshot was taken.
- During a recovery plan test, a new snapshot is taken of the destination volume. This is performed per the VMware SRM specification to capture the newest data that has arrived at the DR site. This means that the consistency of the data is dependent on whether or not the previous snapshot was completely transferred. For example, in Figure 6:
 - If the 4:00 p.m. snapshot taken at the primary site was application consistent, but only 75 percent of the snapshot data had been transferred, the data is considered incomplete. It may be necessary to perform a manual recovery in order to present the last complete snapshot to the application (such as the 3:00 p.m. snapshot in Figure 6).
 - If the 4:00 p.m. primary site snapshot was application consistent and performed at the same time as the SRM recovery snapshot, then 100 percent of that snapshot would transfer. As a result, a new snapshot would include all of the 4:00 p.m. snapshot data, and the application consistency of the data would be preserved.
- Once the SRM recovery snapshot has been taken, a view volume is created from that snapshot.
- The view volume is presented to the vSphere host (or hosts) at the DR site for the SRM to begin a test execution of the recovery plan.

5.6 Data consistency while replicating the active snapshot

Figure 7 and the steps that follow describe the consistency states of replications during plan execution while replicating the active snapshot.

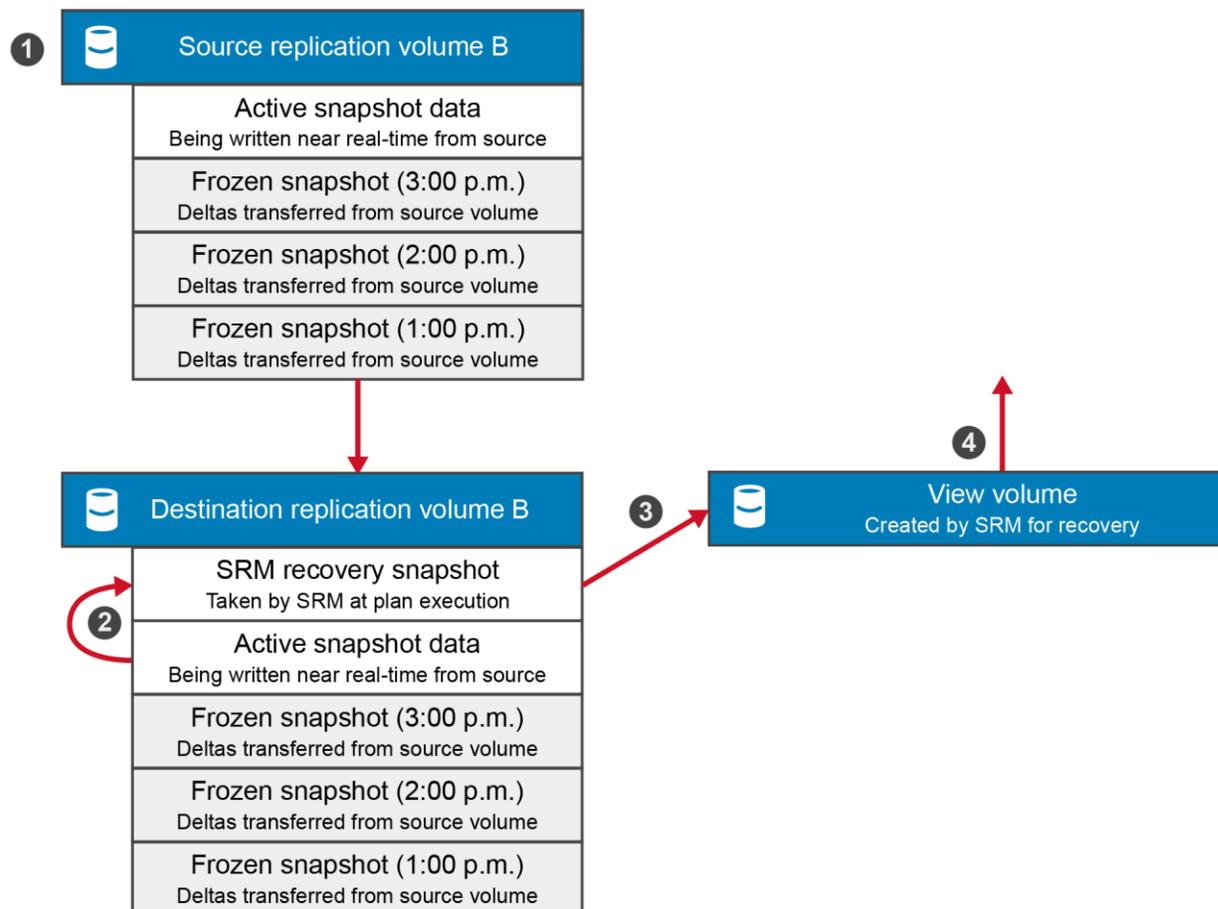


Figure 7 Consistency states of active snapshot replications during plan execution

1. As writes are committed to the source volume, they are almost simultaneously transferred to the destination and stored in the active snapshot. Consistent snapshots can still be taken of the source volume, and the check points are transferred to the destination volume when replicating the active snapshot.
2. During a recovery plan test, a new snapshot is taken of the destination volume that locks in all the data transferred up to that point. Data stored within the active snapshot is most likely crash consistent. For example, in Figure 7:
 - a. Although the 3:00 p.m. primary site snapshot was application consistent at the time of the SRM recovery snapshot, the data must still be deemed crash consistent because it is highly probable that writes into the active snapshot occurred after the 3:00 p.m. snapshot was taken.
 - b. If the application is unable to recover the crash consistent data captured in this snapshot, then manual steps are required in order to present the last known consistent snapshot to the host (such as the 3:00 p.m. frozen snapshot).
 - c. The only time that application data can be considered consistent while replicating the active snapshot is if all I/O has ceased to the source volume, followed by a complete synchronization of the data from the source to the destination.

3. Once the SRM recovery snapshot has been taken, a view volume is created from that snapshot.
4. The view volume is then presented to the vSphere host (or hosts) at the DR site for the SRM to begin a test execution of the recovery plan.

During an actual disaster recovery or planned migration execution, a view volume from a snapshot is not mounted to the remote vSphere hosts. Instead, the destination volume itself is mounted to the remote hosts. This change in behavior from SRM 4.x is to facilitate the reprotect and failback features introduced in SRM 5.x.

5.7 Replication dependencies and replication transfer time

If the application has multiple volumes in a data set, it is important to remember that not all volumes may finish replicating within the same timeframe.

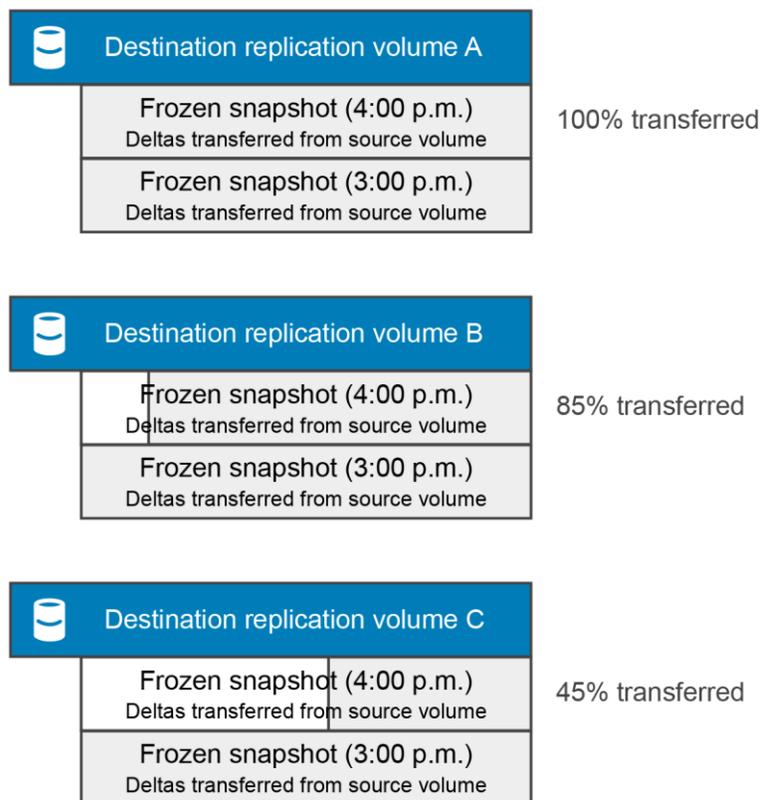


Figure 8 Replication dependencies and transfer times

Due to various factors (such as rate of change, replication network bandwidth, volume size, and replication QoS), it is possible for multiple volumes in a backup set to finish replicating frozen snapshots at different times. In this scenario, SRM does not use incomplete replication snapshots. This means that if the plan is executed before the frozen snapshot replications have completely transferred all of the data (as seen in Figure 8), then the volume A snapshot from 4 p.m. will be mounted, the volume B snapshot from 3 p.m. will be mounted, and the volume C snapshot from 3 p.m. will be mounted. If all three volumes are part of a consistency group, and are in use for a tiered application or database application, then the application volumes presented from different points in time may cause problems with the application. If this happens, manual intervention may be required to present the previous set of consistent snapshots back to the host (for example, the 3 p.m. snapshot from volumes A, B, and C may need to be forced).

5.8 SRM selectable snapshot

SRM selectable snapshot is a feature that is built into DSM. Because multiple methods of replication are supported by SC Series storage, this feature determines whether the active snapshot or last frozen snapshot is used when VMware SRM initiates a failover or test failover. There are four globally applied configuration choices for SRM selectable snapshot:

Always use Active Snapshot (default): Uses the active snapshot (current, unfrozen state of the data transferred to the destination).

Use Active Snapshot If Replicating Active Snapshot: Uses the active snapshot if Replicate Active Snapshot is enabled for the replication, otherwise the last frozen snapshot is used.

Always use Last Frozen Snapshot: Uses the most current frozen snapshot that has been 100 percent transferred to the destination.

Use Restore Point Settings: Uses the pre-configured settings for the restore point of the replication. If Use Active Snapshot is not selected then the last frozen snapshot is used. This option allows granular selection of the SRM Selectable Snapshot configuration on a per-volume-replication basis. The default Use Active Snapshot configuration for each restore point is a cleared checkbox which translates to Use Last Frozen Snapshot for the restore point.

The DSM default is to always use the active snapshot (the current unfrozen state) of the volume for SRM purposes. The SRM selectable snapshot feature integrates only with certain SRM actions; in other actions it is ignored. Table 1 outlines each scenario.

Table 1 When SRM selectable snapshot is honored

SRM action	Recovery type	SRM selectable snapshot configuration honored
Activate recovery plan	Planned migration	No
Activate recovery plan	Disaster recovery	<ul style="list-style-type: none"> • If the protected site is down: Yes • If the protected site is up: No
Test activate recovery plan	N/A	<ul style="list-style-type: none"> • If Replicate recent changes to recovery site check box is cleared in SRM: Yes • If Replicate recent changes to recovery site check box is selected in SRM: No

5.9 Using application- and data-consistent frozen snapshots with SRM

There are a number of methods available for creating a frozen snapshot (replay) on SC Series storage. Once replicated, the snapshot may be used by SRM. While some methods of snapshot creation result in crash-consistent data contained within the snapshot, other methods may be employed that result in application or data consistency within the snapshot. For example, Replay Manager 7.x can be used with vSphere or the vSphere Client plug-in to create a snapshot of a datastore, or the workflow feature can be used to quiesce file systems (if available). Using either of these methods results in a vSphere snapshot with both parent and delta disks frozen, and the snapshot being replicated to the destination site.

Two important facts to recognize are:

- The VM is replicated to the destination site in a vSphere snapshot state and should be dealt with in one way or another to prevent the VM from running continuously over a long time in a vSphere snapshot state.
- The application and data consistency is contained within the frozen parent virtual disk and crash-consistent data is contained within the delta virtual disk.

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

If manually carrying out the previous process on a large scale, this can quickly erode efforts made toward meeting the recovery plan RTO and is not the best use of SRM. In such instances, a more efficient and consistent solution would be to script the snapshot management process using PowerShell and have that process carried out as a pre-power-on, or potentially a post-power-on, step for the VM. Custom recovery tasks are discussed in section 5.10.

5.10 Custom recovery tasks

If the environment requires a custom recovery strategy, both Dell EMC storage and VMware have robust sets of PowerShell cmdlets to customize the recovery steps where needed. The Dell EMC storage cmdlets can control the snapshot selection, view volume creation, volume mappings, and even modify 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.

Add Pre Power On Step ?

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

Name:

Content:

Timeout: minutes seconds

OK Cancel

See appendix A.1A for examples of REST API and PowerShell scripts that can be used within Site Recovery Manager.

6 Site Recovery Manager configuration

This section provides best practices for configuring the Site Recovery Manager.

6.1 Configuring the array managers

To allow SRA to communicate with the DSM Data Collector, array manager configuration can be performed from the Array Managers module. An array manager must be added for each site in the unified interface.

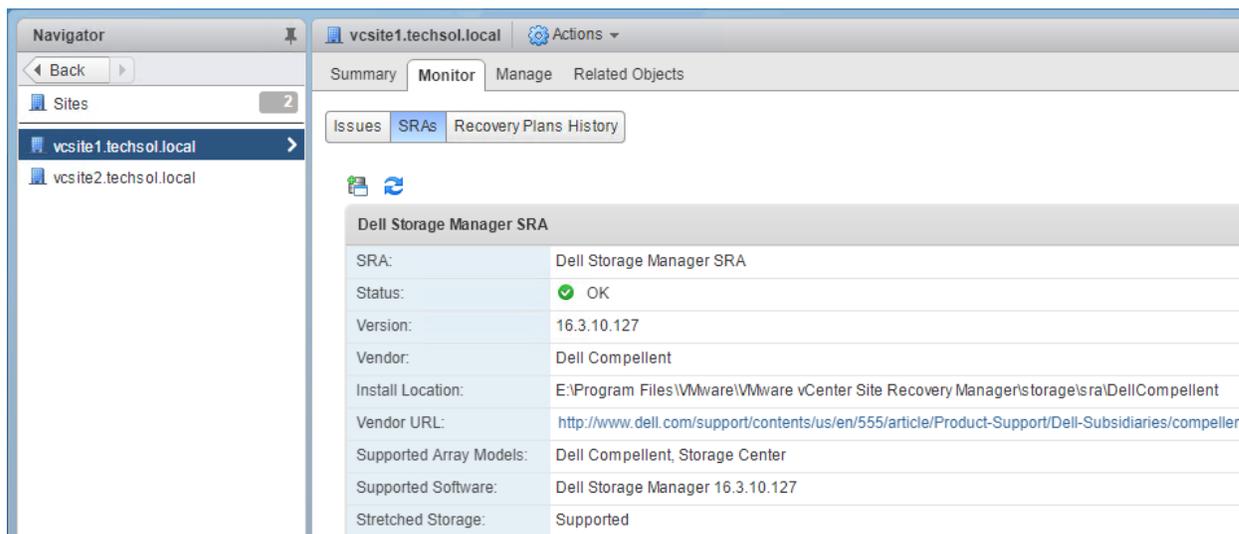


Figure 9 Adding an array manager

The protected site array managers and the recovery site array managers must both be configured for pairing. Depending on the architecture, either a single DSM Data Collector can be added for both sites, or a Data Collector and Remote Data Collector model can be deployed.

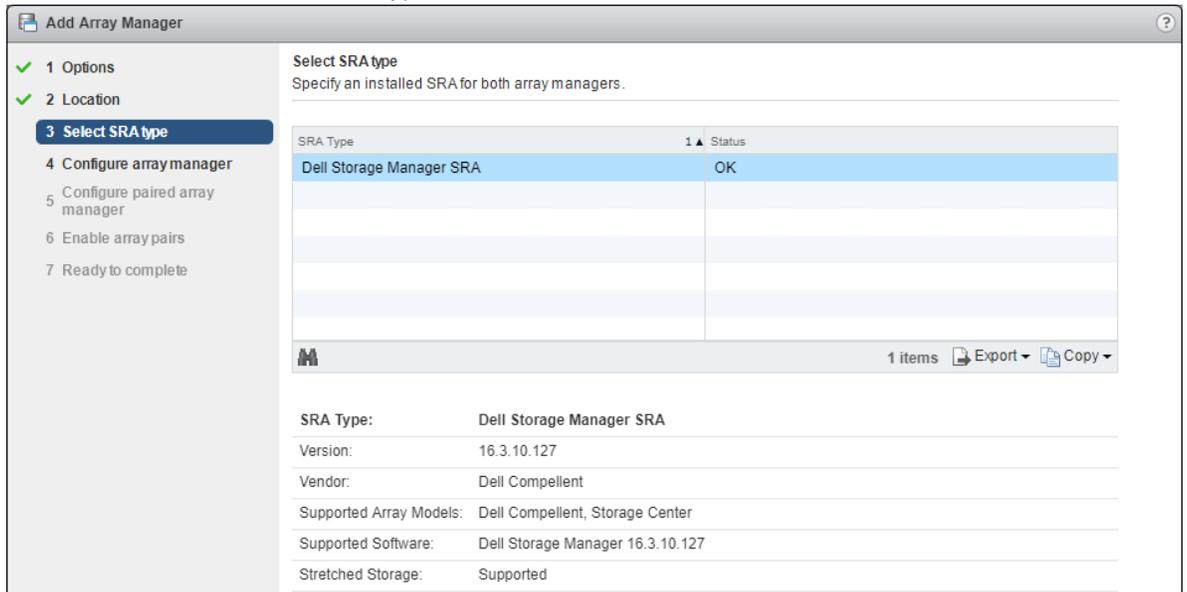
Single Data Collector: Specify the same single Data Collector for both the protected site array manager and the recovery site array manager. This Data Collector would typically reside in the recovery site. If Live Volume stretched storage is protected by SRM, this Data Collector would be located at a third site.

Multiple Data Collectors: A Data Collector and a Remote Data Collector are deployed between the two sites. Specify the Data Collector for the Site A array manager, and the Remote Data Collector for the Site B array manager. Or specify the Remote Data Collector for the Site A array manager, and the Data Collector for the Site B array manager.

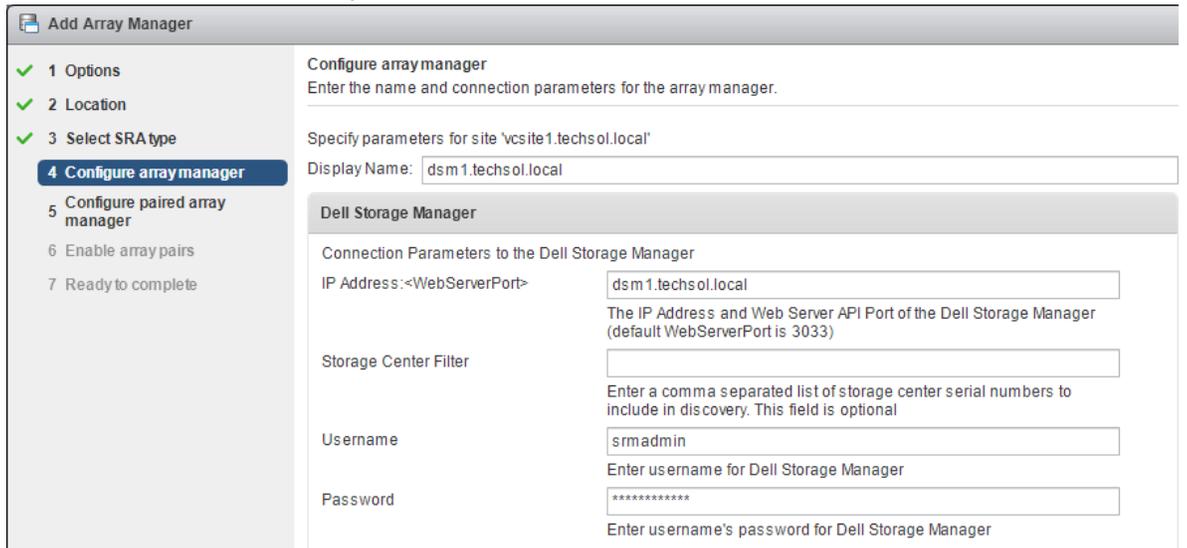
Note: The Remote Data Collector is not supported for Live Volume stretched storage protected by SRM. A single Data Collector must be used at a third site. This Data Collector must be available when disaster recovery or planned migration Recovery Plans are executed.

Perform the following steps to configure the array manager:

1. Choose the installed DSM SRA type.

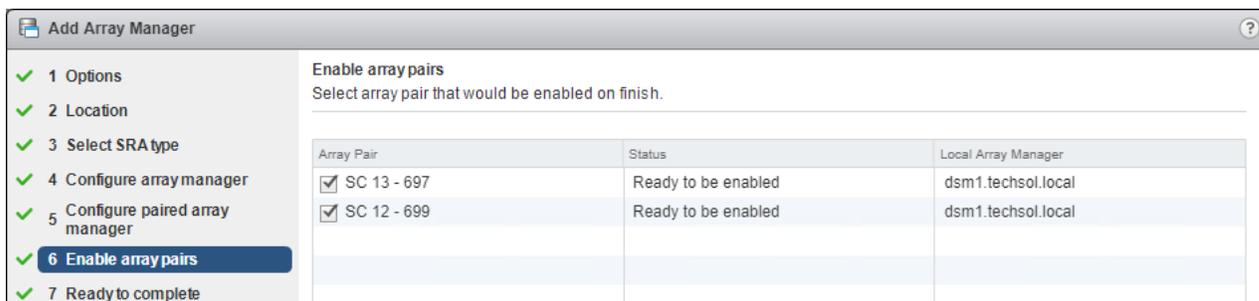


2. Provide the DSM connection parameters.



6.2 Creating array pairs

Once an array manager has been added to each of the two sites in SRM, the arrays need to be paired so that replicated volumes can be discovered by SRM as eligible devices. In older versions of SRM, pairing was an action that was performed after the initial installation of SRM. However, as of SRM 5.8, pairing can be performed as part of the process of adding array managers to sites. Examples of both methods are shown on the following screens.



Add Array Manager

1 Options
2 Location
3 Select SRA type
4 Configure array manager
5 Configure paired array manager
6 Enable array pairs
7 Ready to complete

Enable array pairs
Select array pair that would be enabled on finish.

Array Pair	Status	Local Array Manager
<input checked="" type="checkbox"/> SC 13 - 697	Ready to be enabled	dsm1.techsol.local
<input checked="" type="checkbox"/> SC 12 - 699	Ready to be enabled	dsm1.techsol.local

Array Pairs | Permissions

Array pairs must be enabled for use with Site Recovery Manager. You may enable the array pairs from either the protected or recovery site. Details for the selected array pair are shown below.

Local Array	Remote Array	Status	Local Array Manager	Remote Array Manager
SC 13	697	✓ Enabled	dsm1.techsol.local	dsm1.techsol.local
SC 12	699	✓ Enabled	dsm1.techsol.local	dsm1.techsol.local

Arrays cannot be unpaired while downstream SRM dependencies, such as protection groups, exist.

6.3 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 can be found in the **Array Based Replication** module > **Manage** tab. Run the device discovery on both arrays to ensure a consistent list of devices. Non-replicated volumes will not be 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, all SC Series replicated volumes are discovered as devices in SRM, even if they are not for use by vSphere (for example, replicated volumes belonging to other storage hosts such as physical Microsoft® Exchange®, SQL Server®, Oracle®, and file servers).

Select Discover Devices as shown on the following screen to invoke an SRA query to the DSM Data Collector to obtain the newest array-based replicated device information.

Local Array	Remote Array	Status	Local Array Manager	Remote Array Manager
SC 13	697	✓ Enabled	dsm1.techsol.local	dsm1.techsol.local
SC 12	699	✓ Enabled	dsm1.techsol.local	dsm1.techsol.local

Local Device	Datastore	Status	Remote Device	Protection Group	Local Consistency
srmds3	Local: [srmds3]	→ Outgoing Replication	Repl of srmds3		
srmds4	Local: [srmds4]	→ Outgoing Replication	Repl of srmds4		
Repl of srmds1	Remote: [srmds1]	← Incoming Replication	srmds1		
Repl of srmds2	Remote: [srmds2]	← Incoming Replication	srmds2		

6.4 Creating protection groups

If not completed already, create a small VMFS datastore at the disaster recovery site as a placeholder for VM configuration files. For each virtual machine protected, SRM creates a shadow VM at the opposite site serving as a placeholder for required processor, memory, and network capacity in a disaster recovery or planned migration scenario.

Although this datastore only needs to be large enough to hold the configuration files for all the recoverable virtual machines, creating a standard-sized 500 GB datastore will suffice. SC Series Dynamic Capacity thinly provisions the volume making this a suitable standard which is space efficient.

In most cases, only one placeholder datastore per site is required 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 because VMware SRM only places transient data on this volume that can be easily regenerated within the UI.

With the placeholder datastore ready, protection groups can be created. Replicated or Live Volume stretched storage 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. Manually moving files that belong to a virtual machine off of a datastore is not supported with SRM; the VM will not be 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. Refer to the VMware *Site Recovery Manager Administration Guide*. Dell Storage SC Series arrays offer dynamic block architecture and Data Progression which provides automated sub-LUN tiering for virtual machines without interfering with SRM protection groups.

6.5 Creating recovery plans

When testing or running recovery plans, SRM does not have built-in mechanisms to determine whether the replication volumes are fully synced before the storage is prepared for recovery. In other words, there may be in-flight data that is actively replicated to the secondary site influencing the outcome of the recovery. This will be true if you configure replication to also replicate the active snapshot (replay).

As a best practice, check **Replicate recent changes to recovery site** when executing a test plan to help ensure that all data is successfully replicated to the secondary site. During an actual disaster recovery cutover, this option may or may not be available. For planned migrations using SRM, this step is required in order to proceed.

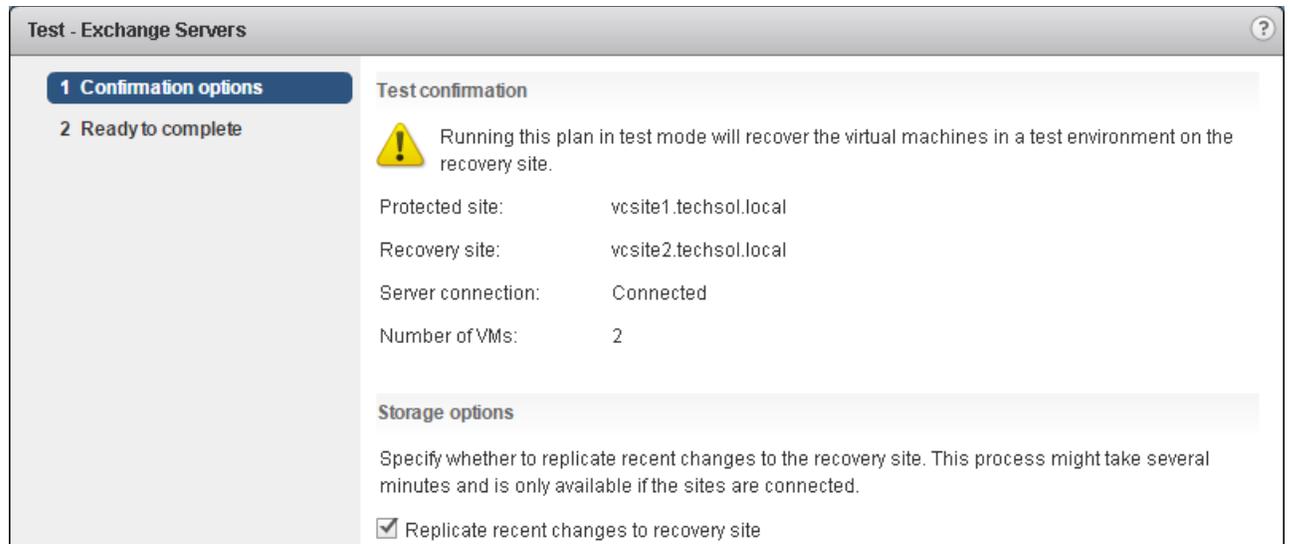


Figure 10 Replicate recent changes to recovery site during test plans

Another best practice to further automate DR failover or planned migration may be to add prompts or SRM server-side commands to the recovery plan to help ensure that all the data has been replicated before the subsequent storage section is executed.

For example, integrate an SC Series REST API or PowerShell script into the recovery plan to take current snapshots of all the volumes and make sure the most recent data has been replicated (see appendix A for examples). When the recovery plan executes, it pauses. However, recovery plan execution will not pause at a command on an SRM server step.

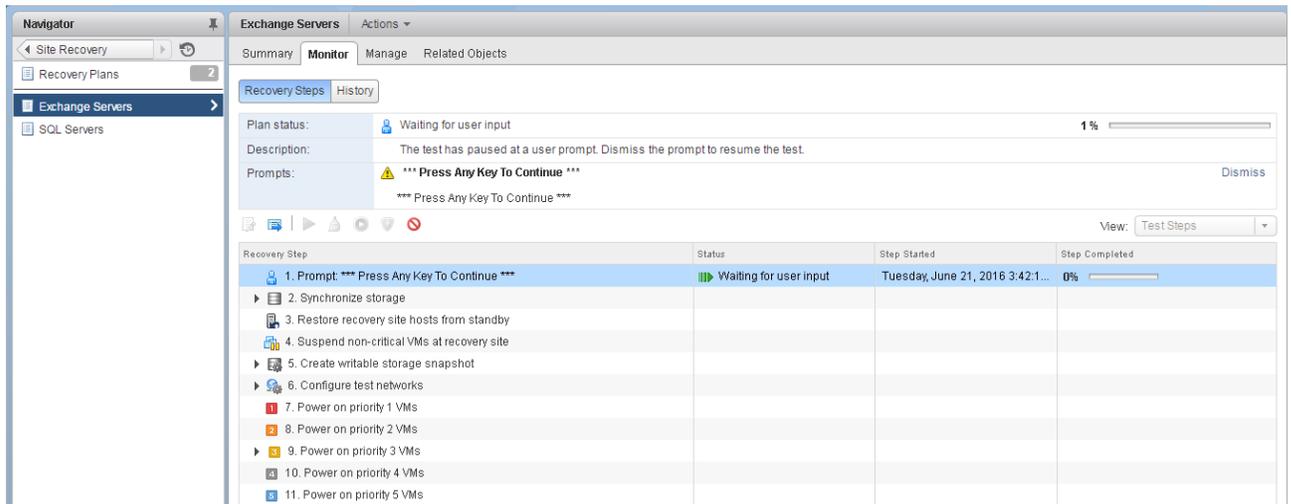
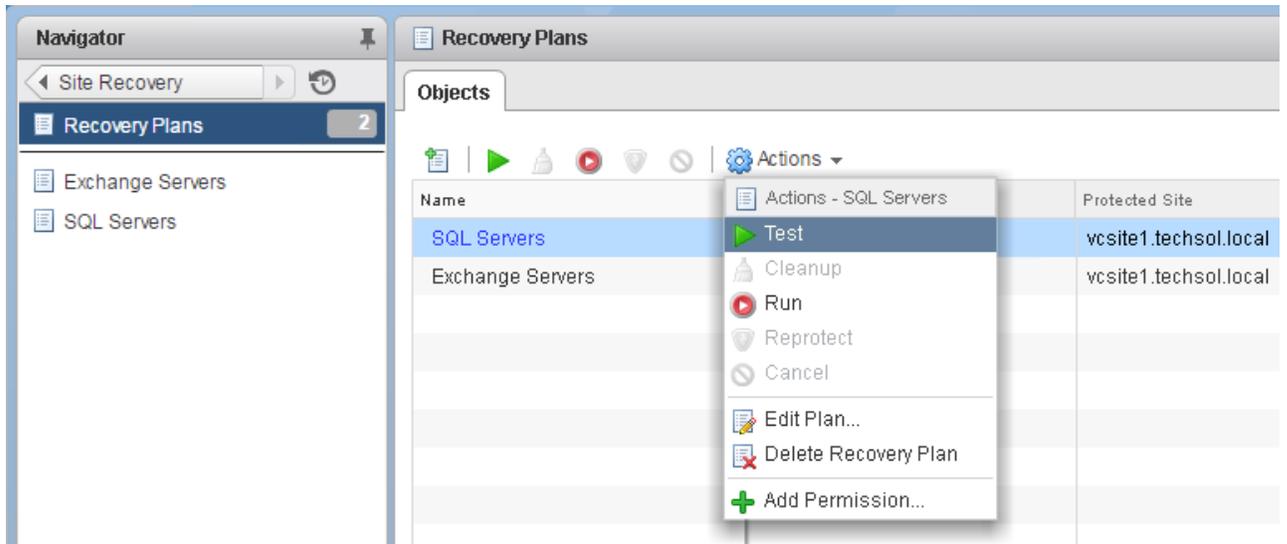


Figure 11 Recovery plan prompting to continue

7 Testing a recovery plan

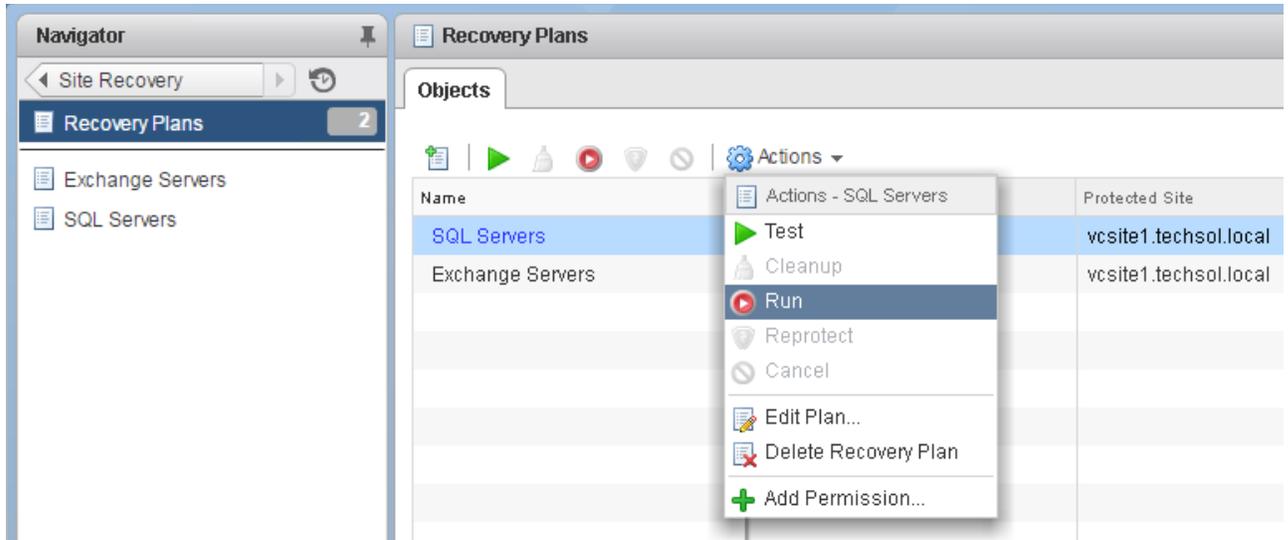
Testing the recovery plan is not disruptive to the storage replications, production volumes, and VMs because the test recoveries use SC Series view volumes created from snapshots (replays) when running the recovery plan tests. This means that when testing a recovery plan, any tests, changes, or updates can be performed on the recovered virtual machines because they will be discarded when the test recovery plan cleanup takes place. While the test plan is executing, production virtual machines and replication continues to run normally without interruption.

To test a disaster recovery plan, right-click the recovery plan, and select **Test**.

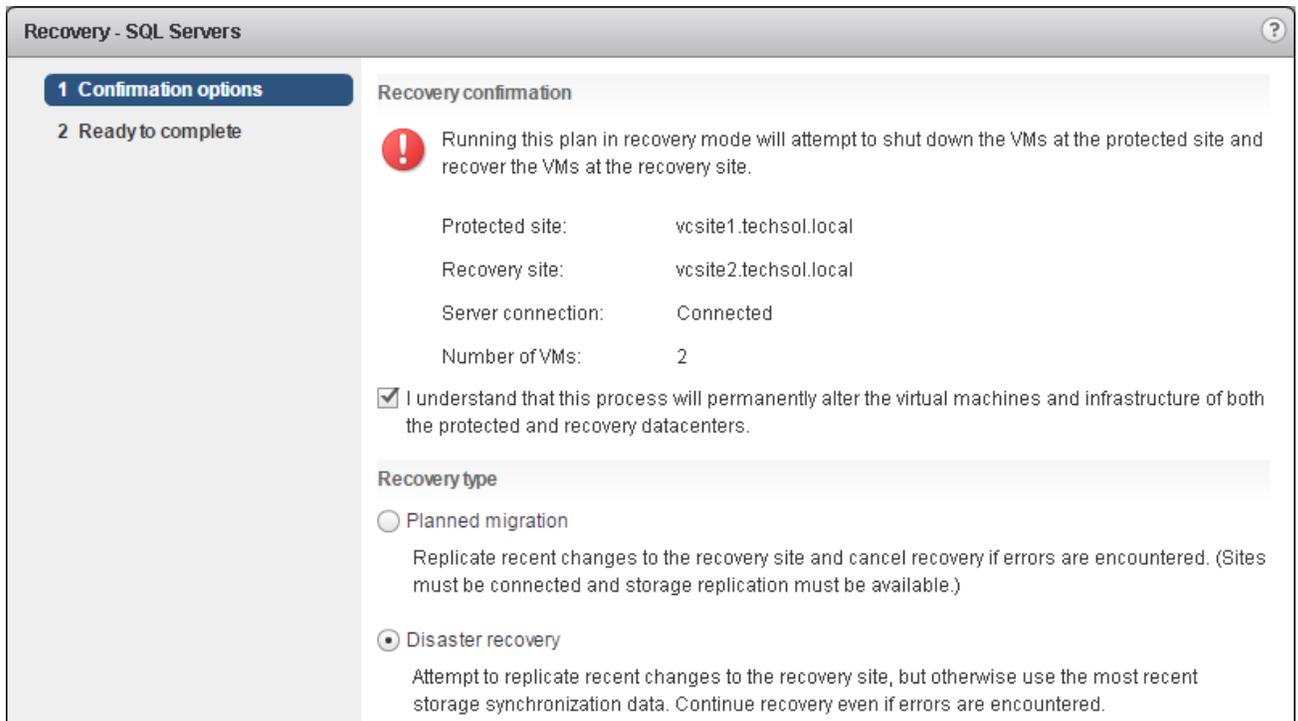


When choosing to run a planned migration or disaster recovery plan (as opposed to running a test), keep in mind this procedure is disruptive and results in virtual machines being powered off at the primary site, replication mirrors being broken, and virtual machines being recovered at the secondary site.

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



Acknowledge the safety precaution message to execute a live plan.



8 Reprotect and failback

After virtual machines are migrated from one site to another using either 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 virtual machines at the recovery site. Today, SRM automates the reprotect process and prepares the virtual machines for failback.

8.1 Reprotection

Once protected virtual machines are migrated, or disaster recovery failed over to the secondary site, the VMs are unprotected. Following the migration of a protected group, SRM offers the ability to automate the reprotection of the virtual machines. The reprotection is carried out in a series of automated steps.

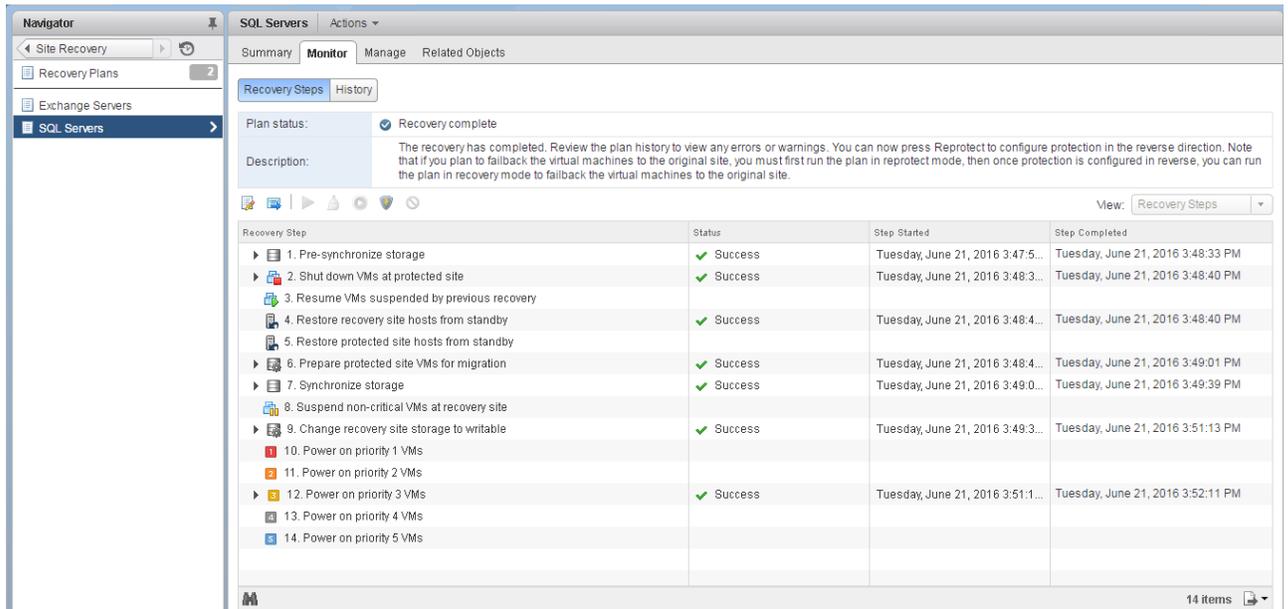
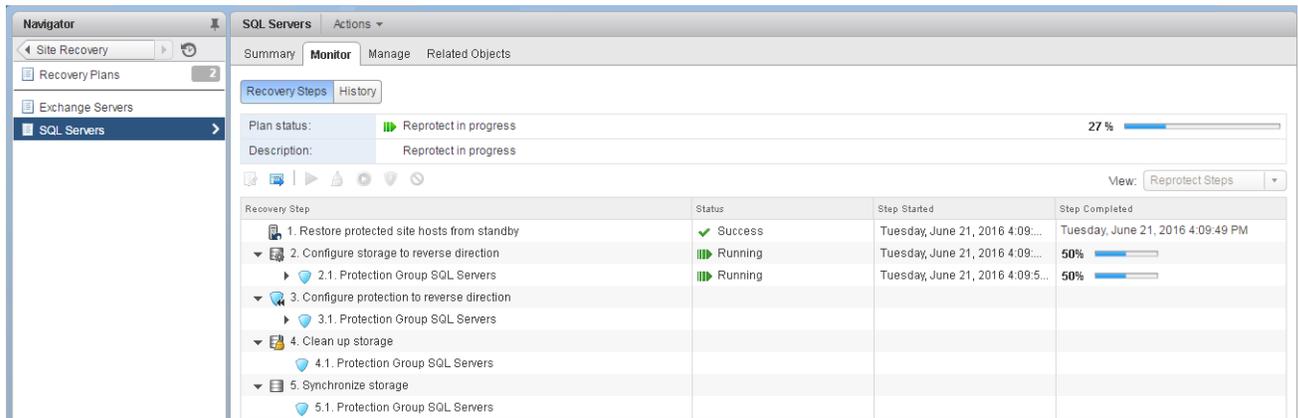


Figure 12 Review the recovery plan history

During a reprotect, SRM commands the SRA to reverse storage replication for each of the datastores/volumes in the protection group in the opposite direction. The protection group originally set up at the primary site is migrated to the secondary site. Placeholder VMs originally set up at the secondary site are now created at the opposite site (the new recovery site) on its respective placeholder datastore.



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 facilitates a number of use cases including the ability to run production applications at the disaster recovery site, resource balancing, and improved disaster recovery infrastructure ROI.

9 Conclusion

VMware vSphere, Site Recovery Manager, and Dell EMC SC Series arrays combine to provide a highly available business platform for automated disaster recovery with the best possible RTO and RPO, as well as planned migrations for your virtualized data center.

A Example scripts

A.1 REST API script: TakeSnapshot.py

This is an example REST API script that can be folded into an SRM recovery plan. It leverages the Dell RESTful API to take a snapshot (replay) of the source replication system volume to make sure that the most current snapshot is replicated to the DR site.

```
#
# main.py
#
# import modules into Python script
import requests, json, http, httplib, urllib, urllib2
import os, sys, subprocess, math
import math, time
import logging
from simplejson import scanner

# setup logging to scapi.log
logging.basicConfig(level=logging.DEBUG, filename='scapi.log',
format='[%(asctime)s] %(levelname)s %(message)s')

if __name__ == '__main__':
    # define env incl. DSM IP addr, port & login credentials
    DSM_ip = '192.168.1.10'          # IP address of DSM instance
    DSM_port = '3033'              # Default port of DSM instance
    DSM_id = 'Admin'               # Login credentials for DSM
    DSM_pass = 'mmm'               # Password
    verify_cert = False           # Default = False
    apiversion = '2.0'            # Default = 2.0

    # disable warnings from requests module
    if not verify_cert:
        requests.packages.urllib3.disable_warnings()

    # define base URL for DSM REST API interface
    baseURL = 'https://%s:%s/api/rest/' % (DSM_ip, DSM_port)

    # define HTTP content headers
    header = {}
    header['Content-Type'] = 'application/json; charset=utf-8'
    header['Accept'] = 'application/json'
    header['x-dell-api-version'] = apiversion

    # define the connection session
    connection = requests.Session()
    connection.auth = (DSM_id, DSM_pass)
```



```
# create a replay object from Volume_Name_x volume object
payload = {}
payload['Description'] = 'Replay of Volume_Name_x'
payload['ExpireTime'] = '60' # in minutes, 0 = Never Expire
REST = '/StorageCenter/ScVolume/%s/CreateReplay' %
volList['Volume_Name_x']['instanceId']
completeURL = '%s%s' % (baseUrl, REST if REST[0] != '/' else REST[1:])
json_data = connection.post(completeURL, data=json.dumps(payload,
ensure_ascii=False).encode('utf-8'), headers=header, verify=verify_cert)
#stdout = json.loads(json_data.text)
#print stdout

# logout from DSM instance
payload = {}
REST = '/ApiConnection/Logout'
completeURL = '%s%s' % (baseUrl, REST if REST[0] != '/' else REST[1:])
print connection.post(completeURL, data=json.dumps(payload,
ensure_ascii=False).encode('utf-8'), headers=header, verify=verify_cert)
```

This REST API script connects to an SC Series system with an IP address **192.168.1.10**, a username **Admin**, and a password **mmm**. It then takes a snapshot of **Volume_Name_x** with a snapshot expiration set to **60** minutes.

A.2 CompCU Script: TakeSnapshot.cmd

This is an example CompCU script that can be folded into an SRM recovery plan. It leverages the SC Series command utility (CompCU) and takes a snapshot of the source replication system volume to make sure that the most current snapshot is replicated to the DR site.

```
"C:\Program Files\Java\jre6\bin\java.exe" ^
-jar c:\scripts\compcu.jar ^
-host 192.168.1.10 ^
-user Admin ^
-password mmm ^
-c "replay create -volume 'Volume_Name_1' -expire 60"

"C:\Program Files\Java\jre6\bin\java.exe" ^
-jar c:\scripts\compcu.jar ^
-host 192.168.1.10 ^
-user Admin ^
-password mmm ^
-c "replay create -volume 'Volume_Name_2' -expire 60"
```

This CompCU script connects to an SC Series system with an IP address **192.168.1.10**, a username **Admin**, and a password **mmm**. It then takes a snapshot of **Volume_Name_x** with a snapshot expiration set to **60** minutes. The carrot (^) symbols in this script are used for line continuation and readability, but could be excluded if the entire command is placed on one line.

A.3 SC Series command set PowerShell script: TakeSnapshot.ps1

This script leverages the SC Series command set to take a snapshot of the source replication system volume in an effort to make sure that the most current snapshot is replicated to the DR site.

```
$SCHostname = "sc12.techsol.local"
$SCUsername = "srmadmin"
$SCPassword = ConvertTo-SecureString "mmm" -AsPlainText -Force
$SCConnection = Get-SCConnection -HostName $SCHostname -User $SCUsername -
Password
$SCPassword
New-SCReplay (Get-SCVolume -Name "lun40" -Connection $SCConnection) -
MinutesToLive 1440 - Description "Snapshot w/ 1 day retention" -Connection
$SCConnection
```

This PowerShell script will connect to an SC Series system with a host name **sc12.techsol.local**, a username **srmadmin**, and a password **mmm** to take a snapshot of **lun40** with a snapshot expiration set to 1 day. Run this PowerShell script from the SC Series command set shell to automatically load the **Compellent.StorageCenter.PSSnapin** snap-in. This script can also be run from any PowerShell prompt provided the **Compellent.StorageCenter.PSSnapin** snap-in is manually loaded or loaded as part of the PowerShell profile.

A.4 Dell Storage PowerShell SDK script: TakeSnapshot.ps1

This script leverages the Dell Storage PowerShell SDK to take a snapshot of the source replication system volume in an effort to make sure that the most current snapshot is replicated to the DR site.

```
# Import the module for the Dell Storage PowerShell SDK
Import-Module "C:\PS_SDK\DellStorage.ApiCommandSet.psd1"

# Assign variables
$EmHostName = "em1.techsol.local"
$EmUserName = "srmadmin"

# Prompt for the password
$EmPassword = Read-Host -AsSecureString `
               -Prompt "Please enter the password for $EmUserName"

# Create the connection
$Connection = Connect-DellApiConnection -HostName $EmHostName `
                                           -User $EmUserName `
                                           -Password $EmPassword

# Assign variables
$ScName = "SC 12"
$VolumeName = "lun40"
$VolumeFolderPath = "VMware/Demo/SRM 6.x/"
$SnapshotDescription = "Created by PowerShell SDK"
```

Example scripts

```
# Get the volume
$Volume = Get-DellScVolume -Connection $Connection `
                        -ScName $ScName `
                        -VolumeFolderPath $VolumeFolderPath `
                        -Name $VolumeName

# Create a Snapshot that will expire in 1 day
$OneDayInMinutes = 1 * 24 * 60 # 1 day * 24 hours/day * 60 minutes/hour
$Snapshot = New-DellScVolumeReplay -Connection $Connection `
                                -Instance $Volume `
                                -Description $SnapshotDescription `
                                -ExpireTime $OneDayInMinutes
```

This PowerShell script will connect to a DSM which is managing an SC Series system named **sc12.techsol.local** with a username **srmadmin** and prompt for a password to take a snapshot of **lun40** with a snapshot expiration set to 1 day. This script can be run from any PowerShell prompt provided the **DellStorage.ApiCommandSet.psd1** module exists in the specified folder location or has already been imported.

B Additional resources

B.1 Technical support and resources

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

[Dell TechCenter](https://delltechcenter.com) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware, and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC storage platforms.

B.2 VMware support

For VMware support, see the following resources:

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