

Dell EMC SC Series: Disaster Recovery for Microsoft SQL Server Using VMware Site Recovery Manager

Abstract

This document identifies options available for providing an automated disaster recovery solution for virtualized Microsoft® SQL Server® workloads on Dell EMC™ SC Series storage.

July 2019

Revisions

Date	Description
October 2013	Initial release
July 2016	DSM, Live Volume, technical review
July 2019	Miscellaneous improvements

Acknowledgements

Authors: Doug Bernhardt, 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 © 2013–2019 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. [7/12/2019] [Best Practices] [CML1018]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
1 VMware Site Recovery Manager overview	6
1.1 Architecture.....	6
1.1.1 Active/DR site design	7
1.1.2 Active/active site design	8
1.2 Recovery point objective	8
1.3 Recovery time objective	9
2 Solution components.....	10
2.1 SC Series.....	10
2.2 Snapshots.....	10
2.3 Snapshot profiles.....	10
2.4 Active snapshot	10
2.5 Consistency groups	10
2.6 Replay Manager	10
2.7 Array-based replication.....	10
2.8 Dell Storage Manager.....	11
2.9 Storage Replication Adapter.....	11
3 Storage infrastructure.....	12
3.1 vSphere storage types.....	12
3.2 Virtual machine disk types.....	12
4 Solution design.....	14
4.1 Choosing a snapshot strategy	14
4.1.1 Using Replay Manager with .VMDKs	14
4.1.2 Using Replay Manager with vRDMs.....	15
4.1.3 Using Replay Manager with pRDMs.....	15
4.1.4 Using snapshot profiles	15
4.2 Choosing a replication strategy	15
4.2.1 Asynchronous.....	15
4.2.2 High consistency synchronous.....	16
4.2.3 High availability synchronous	16
4.2.4 Asynchronous active snapshot.....	16

4.2.5 Live Volume	16
4.2.6 Volume replication considerations	17
4.3 Configuring Dell Storage Manager	18
4.4 Configuring recovery in SRM.....	18
4.4.1 Recovery from the active snapshot	18
4.4.2 Recovery from the latest frozen snapshot	19
4.5 Performing a disaster recovery test.....	22
4.6 Performing a disaster recovery.....	22
4.7 Performing a planned migration	23
A Additional resources	24
A.1 Technical support and resources	24
A.2 Referenced or recommended publications.....	24

Executive summary

Data center consolidation by way of x86 virtualization is a trend which has gained tremendous momentum and offers many benefits. One workload type that is generally considered a virtualization candidate is Microsoft® SQL Server®. Although the physical nature of Microsoft SQL Server is transformed once it is virtualized, the necessity for data protection, retention, and recovery remains. This document identifies a variety of options available for providing an automated disaster recovery solution for virtualized SQL Server workloads using Dell EMC™ SC Series, Replay Manager, array-based replication, and VMware® Site Recovery Manager with varying levels of consistency.

1 VMware Site Recovery Manager overview

Site Recovery Manager (SRM) is a disaster recovery testing, execution, and planned migration product for VMware virtualized data centers. It leverages the power of storage replication and virtual machine mobility to provide automated disaster recovery testing and execution as well as planned migrations of virtual machines between active sites. The bundled automation combined with storage replication yields unmatched capabilities to meet RTO and RPO requirements when compared to legacy disaster recovery plans and physical servers.

1.1 Architecture

For most deployments, the Site Recovery Manager infrastructure and resulting architecture is mirrored between two sites. Each site contains storage which replicates between sites, vSphere hosts which provide compute resources for running virtual machines, and lastly the software which is used to manage VMware vSphere®, SRM, and the storage. Each site also contains other infrastructure components such as physical servers, networking, firewalls, authentication, and directory services. Site Recovery Manager 5.0 and newer can accommodate two site designs: the traditional active/DR site design as well as an active/active site design.

1.1.1 Active/DR site design

Traditionally, many disaster recovery plans begin with a single active site and a single DR site. The active site represents the production datacenter. The DR site represents compute, network, and storage capacity where a business could rebuild their IT infrastructure and resume operations. The infrastructure at the DR site remains generally unused until a DR plan is tested or executed.

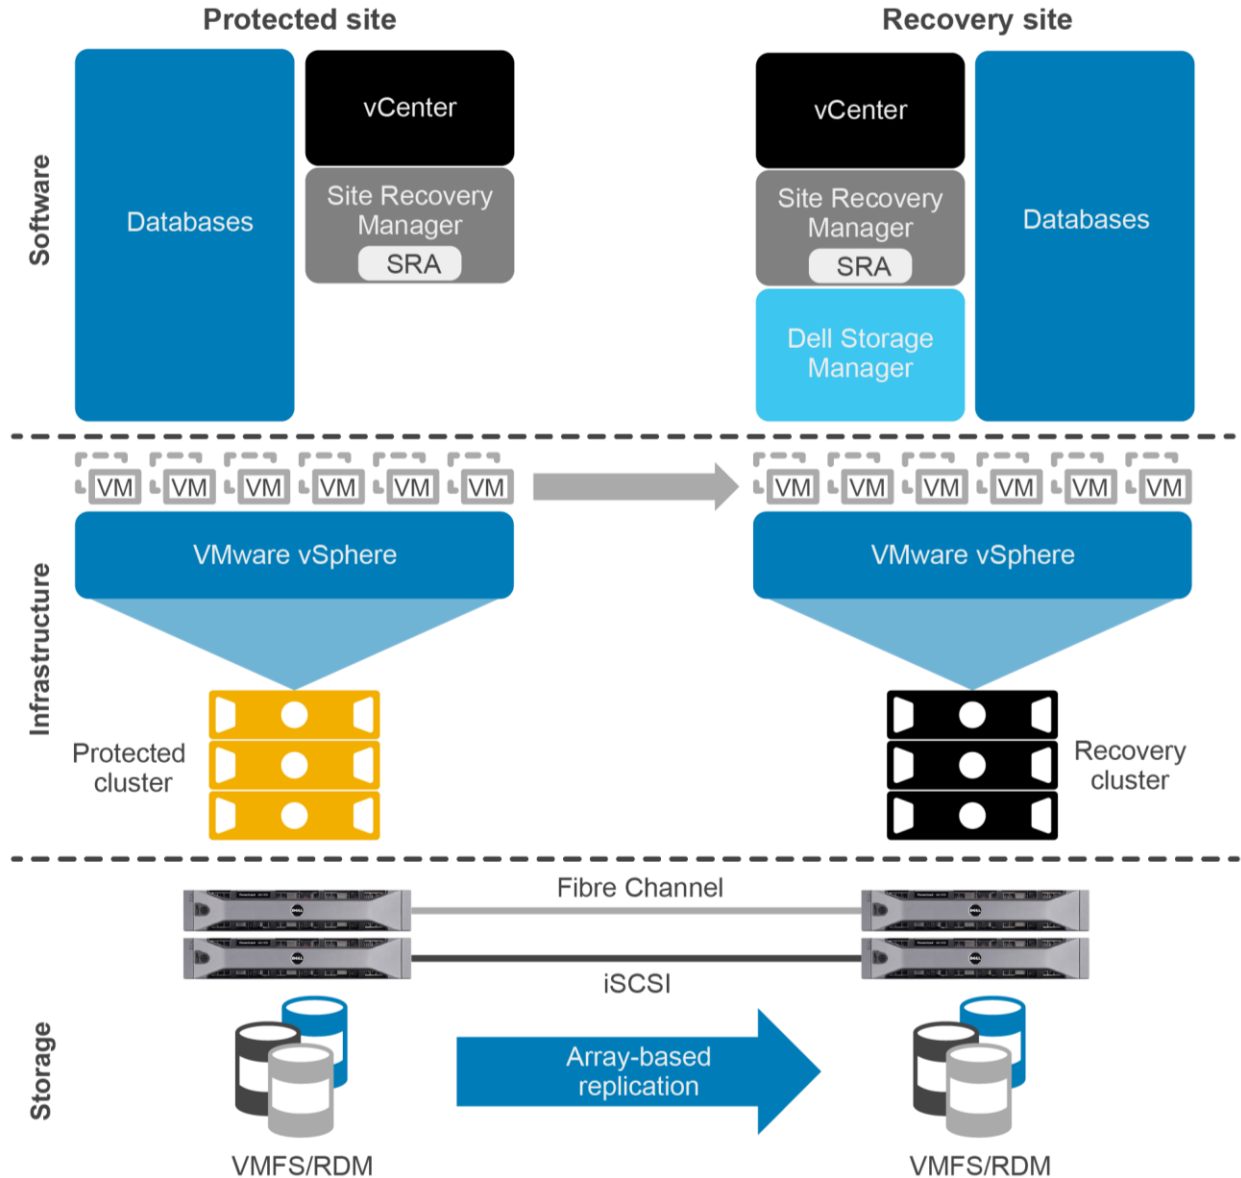


Figure 1 Active/DR site architecture with DSM available at the DR site in the event of a disaster

1.1.2 Active/active site design

Site Recovery Manager also supports a similar design in which two sites exist, but both are actively providing applications and services which are in scope for a comprehensive DR plan. In this design, each site functions as an active site for production applications as well as a recovery location for the other active site.

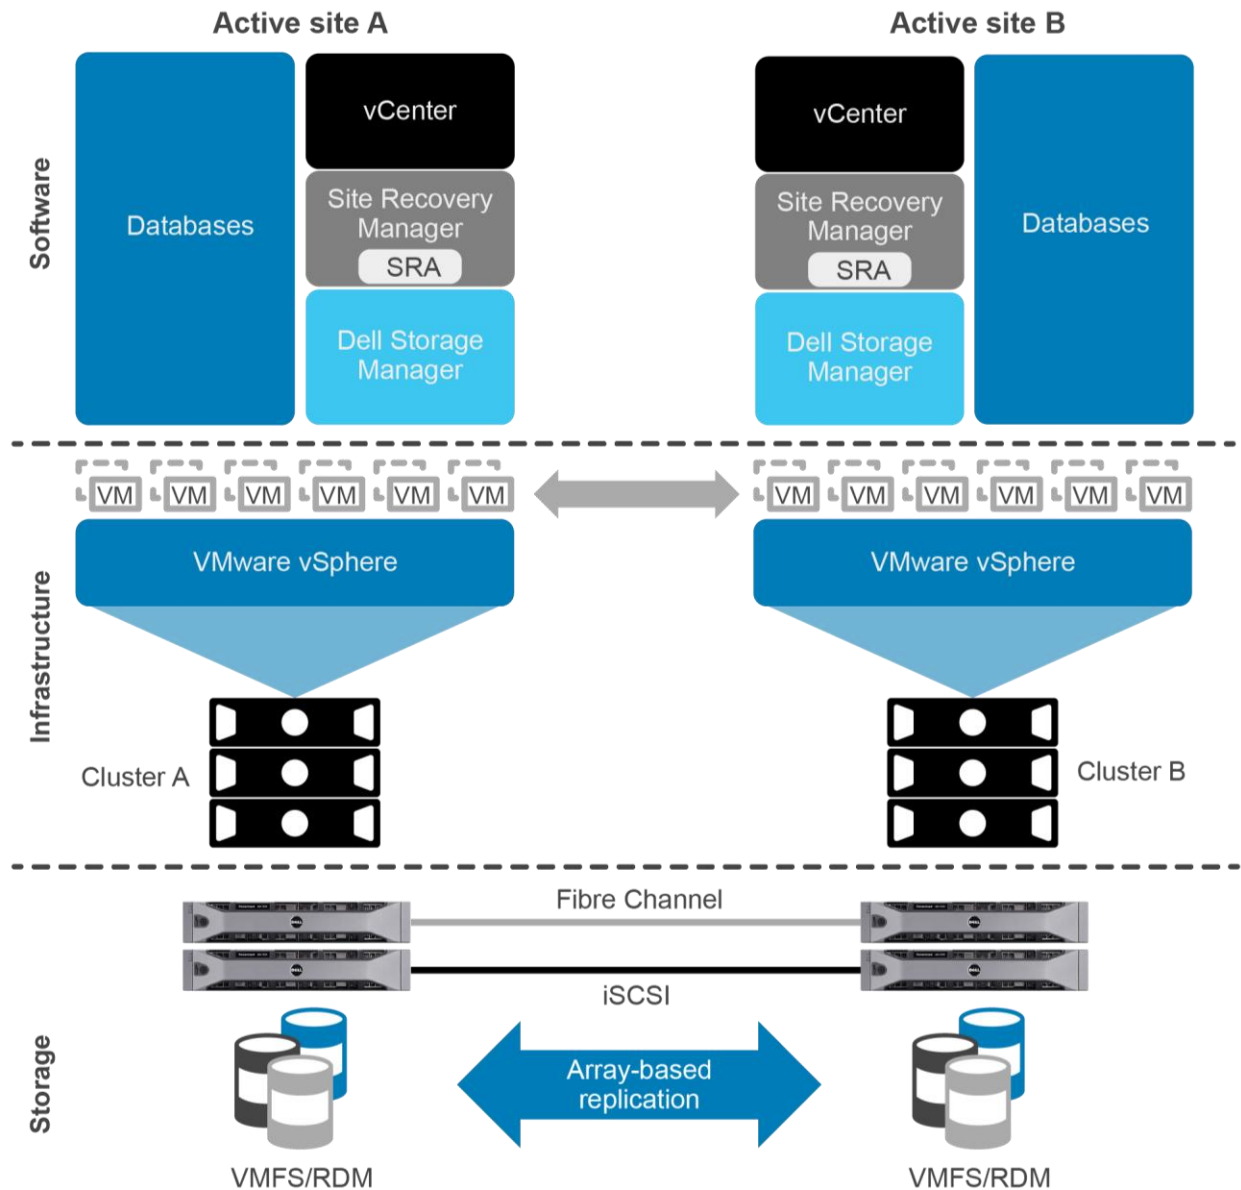


Figure 2 Active/active site architecture with DSM available at both sites in the event of a disaster

1.2 Recovery point objective

Recovery point objective (RPO) is an industry-standard metric which identifies the recovery point or maximum tolerance of data loss when a disaster recovery plan is executed. RPO is defined in a disaster recovery plan itself for a given tier or data set and is subsequently used as a measurement tool to determine the success or failure of an executed plan, whether test or actual. A variety of RPOs may exist for various tiers of applications or data being recovered. RPO is typically measured in terms of hours or minutes. As an example,

a one-hour RPO may be tied to a tier 1 SQL Server application database. This means a maximum of one hour of data may be lost or the executed disaster recovery plan will recover data to a point within one hour or less from the time of the disaster. RPO is improved by increasing the interval at which data is backed up or replicated to the disaster recovery site.

1.3 Recovery time objective

Recovery time objective (RTO) is an industry-standard metric which identifies the maximum allowed recovery time when a disaster recovery plan is executed. RTO is defined in a disaster recovery plan itself for a given tier or data set and is subsequently used as a measurement tool to determine the success or failure of an executed plan, whether test or actual. A variety of RTOs may exist for various tiers of applications or data being recovered. RTO is typically measured in terms of hours or minutes. As an example, a six-hour RTO may be tied to a tier 1 SQL Server application database. This means a maximum of six hours may elapse from the time of the disaster until the time the SQL Server application database is made available again. The starting point for the RTO calculation may vary between organizations but should be clearly defined in the disaster recovery plan. As an example, the RTO calculation could be based on the precise time of the disaster which is common for service providers, or it may be based on an organization's formal declaration of a disaster, rather than the disaster event itself which is the actual starting point of application and data inaccessibility. Declaring a disaster is a process with impacts and as a result the declaration itself consumes measurable amounts of time. The RTO calculation may or may not factor in the time required to make a decision. RTO is generally improved by sound documentation, processes, data integrity, automation, and virtualization.

2 Solution components

The solutions described in this document incorporate various components from SC Series, array-based replication, and VMware Site Recovery Manager. A combination of these components can be leveraged to provide a purpose-built solution meeting the data protection and disaster recovery requirements of the environment.

2.1 SC Series

Dell EMC SC Series storage is a multiprotocol shared storage area network (SAN) designed to provide high availability, performance, automated tiering, and scalability for VMware vSphere virtualized and consolidated environments.

2.2 Snapshots

SC Series storage has the ability to create space-efficient, hardware-based snapshots (replays) of volumes. Blocks of data which are frozen in a snapshot form the basis of data protection mechanisms and cannot be modified. Snapshots can be replicated to remote SC Series arrays through asynchronous or synchronous replication.

2.3 Snapshot profiles

Snapshot profiles (replay profiles) define a schedule by which snapshots will automatically be created throughout a period of time. Snapshot profiles are assigned to each volume which is presented as a VMFS datastore or raw device mapping (RDM) in a vSphere environment. Snapshot profiles have no integration with the Microsoft Volume Shadow Copy Service (VSS).

2.4 Active snapshot

An active snapshot (active replay) contains newly written data or data that has been changed on a volume since the last frozen snapshot was created.

2.5 Consistency groups

A consistency group ties snapshots of multiple volumes together and provides a method of capturing a precise date and time consistent snapshot across all volumes in the group. Snapshot consistency groups have no integration with the Microsoft Volume Shadow Copy Service (VSS) or SC Series array-based replication.

2.6 Replay Manager

Replay Manager creates SC Series snapshots on a scheduled basis with application consistency across volumes. Replay Manager integrates with operating system and application-specific VSS components to provide that application consistency.

2.7 Array-based replication

Array-based replication, licensed as Remote Instant Replay with SC Series, is a common storage feature in which a volume replica is maintained on a remote array. The replica is typically instantiated and kept in sync

automatically through replication at scheduled or continuous intervals. Replication is a significant key to meeting aggressive RTO and RPO in a disaster recovery strategy and serves as the fundamental cornerstone for VMware SRM operations. Various methods of replication exist and will be discussed in further detail.

2.8 Dell Storage Manager

Dell™ Storage Manager is used to manage one or more SC Series arrays and serves a variety of functions in disaster recovery planning, testing, and execution with or without SRM. Among these tasks, the most paramount is the configuration and tracking of replication jobs and saved restore points, and the presentation of volumes during test or execution of predefined recovery plans.

2.9 Storage Replication Adapter

The Storage Replication Adapter (SRA), is a small piece of software provided by Dell EMC storage which is installed on SRM servers at each of the two sites. The SRA interprets a set of storage-related commands from SRM and carries out those commands in conjunction with DSM.

3 Storage infrastructure

Storage is required by vSphere to maintain encapsulated virtual machines and the data each of the VMs contain. vSphere-certified storage is presented to a cluster of vSphere hosts and abstracted by vSphere in a few different ways in order to meet the needs of the VMs. Outside of the disaster recovery context, storage plays major roles in availability, performance, and capacity. However, if and when disaster strikes, storage is needed immediately at the recovery site to recover applications and resume business operations. This section will discuss the storage options available for SQL Server running on a vSphere and SRM infrastructure.

3.1 vSphere storage types

Virtual machines are typically located on one or more types of shared storage which are abstracted as VMFS datastores or in some cases, RDMs. In its current release, VMware SRM supports many of the same storage protocols and storage vendors found on the vSphere HCL including Fibre Channel (FC), iSCSI, and NFS. The key requirement from storage vendors is a VMware-certified SRA. The list of SRA-certified storage vendors and storage types can be found in the [VMware Compatibility Guide](#).

This document focuses on Microsoft SQL Server virtual machines on SC Series which natively supports block storage protocols such as FC, FCoE, and iSCSI.

3.2 Virtual machine disk types

Virtual disks represent drive letters or mount points in the guest operating system and can be presented to a virtual machine in a few different ways. In the majority of use cases, traditional virtual machine disks will be used and each disk is represented by a corresponding .vmdk file on a VMFS datastore. In Windows Disk Management, each .vmdk is abstracted as the physical disk type, VMware Virtual disk SCSI Disk Device. Traditional .vmdk virtual disk types are recommended throughout an environment unless a specific requirement or design decision dictates otherwise.

An RDM is the other virtual machine disk type and there are two varieties of RDM: virtual and physical, notated as vRDM and pRDM, respectively. An RDM presents an entire SC volume to a virtual machine as a disk. Outside of in-guest clustering use cases, RDMs are only presented to a single virtual machine as opposed to being shared by multiple virtual machines. An RDM is also formatted by the guest operating system using a native file system as opposed to the vSphere VMFS file system.

Since traditional virtual disks and RDMs are abstracted as physical disks, either disk type may be carved up into one or more partitions inside the guest operating system to logically isolate data by drive letter or mount point. In addition, all disk types may be expanded or grown, providing the guest operating system supports the feature. Because of abstraction and virtualization, the guest operating system is not aware of its virtual disk type, whether it is a .vmdk or an RDM. From a vSphere perspective, the major difference between a .vmdk and an RDM has already been identified in how an SC volume is presented and abstracted. When comparing the two available RDM types, virtual and physical, the differentiator from an operational perspective is that a vRDM can be included in a vSphere snapshot while a pRDM cannot. This is important to take into consideration if vSphere snapshots are intended to be leveraged as part of an application- or data-consistent data protection and recovery mechanism.

One additional type of storage available for a virtual machine would be SAN or NAS mapped directly to the guest operating system itself instead of being presented through the vSphere storage stack. An example of this would be in-guest iSCSI where an SC volume is presented directly to the IQN of the operating system built-in software iSCSI initiator. While this storage configuration would function, there is very little hypervisor

Additional resources

or application integration available due to lack of visibility to this disk by the hypervisor. The focus of this paper is traditional .vmdk and RDM virtual disk types for Microsoft SQL Server virtual machines.

4 Solution design

SRM can be configured to use either vSphere Replication (VR) or array-based replication. This paper will only cover using SRM with array-based replication, specifically SC Series replication. Before creating the recovery plan in SRM, the type of replication, the method used to create snapshots, and the frequency for which snapshots are taken need to be selected. For virtual machines running SQL Server, it is critical to create snapshots that SRM can use to reliably recover databases. Replay Manager leverages VSS to provide the application and data consistency that ensures consistent, reliable database recovery. Once a replication and snapshot strategy is chosen, DSM and SRM can be configured to facilitate a failover that meets the RPO and RTO requirements.

4.1 Choosing a snapshot strategy

When creating snapshots for volumes containing SQL Server data, it is recommended to use Replay Manager. This provides the most reliable snapshots from which to recover SQL Server data. Replay Manager leverages the VSS to ensure application consistency. If Replay Manager cannot be used, it is highly recommended to use a consistency group when protecting SQL Server data stored across multiple volumes.

Regardless of the method used to create snapshots, be sure to take snapshots often enough to meet RPO requirements. For example, if the RPO is 60 minutes, snapshots should be taken at least every 60 minutes. If the link between sites is slow or unreliable, snapshots might need to be taken more often.

The recovery of virtual machines by SRM can vary depending on the method used to create snapshots. When choosing a snapshot mechanism, be sure to consider the recovery implications. To help simplify the recovery process, choose the same snapshot mechanism for all snapshots on a given volume. PowerShell scripts written to automate a recovery plan will be less complex if they do not have to determine the type of snapshot used for the recovery.

4.1.1 Using Replay Manager with .VMDKs

Either one of the VMware backup extensions can be used to create application consistent snapshots of SQL Server data stored on .vmdks. Consider the following when using Replay Manager with .vmdks:

- When using either VMware extension:
 - If a virtual SCSI controller hosts more than 7 virtual disks in a virtual machine, vSphere snapshots will work, but it will not provide application consistency.
 - The vSphere snapshot will fail if any VSS-aware application data, like a SQL Server database, are stored on a disk type other than a .vmdk or a vRDM.
- When using the **VMware Datastores** extension:
 - Ensure all datastores used by protected SQL Server instances are included in the backup set.
 - If the vSphere snapshot for one or more virtual machines fails, the job will still succeed.
- When using the **VMware Virtual Machines** extension:
 - Select the SQL Server virtual machines for the backup set. Replay Manager will automatically build a list of volumes to include based on the VMs selected.
 - If the vSphere snapshot for one or more virtual machines fails, the entire job fails.

4.1.2 Using Replay Manager with vRDMs

The **VMware Virtual Machines** extension can be used to create application consistent snapshots of SQL Server data stored on vRDMs. This backup set will back up all virtual disks (.vmdks) and vRDMs used by the selected virtual machines. If there are other virtual machines using any of the same datastores, it is recommended to include those virtual machines in the same backup set as well.

4.1.3 Using Replay Manager with pRDMs

To create application consistent snapshots of SQL Server data stored on pRDM, use the **SQL Databases** backup extension. This extension is used inside of the guest to take snapshots of the volumes containing SQL Server databases. Consider the following when using Replay Manager with pRDMs:

- Use the **SQL Databases** backup extension to take snapshots of the database volumes on pRDMs. For other types of volumes, use the appropriate Replay Manager backup extension or use an SC Series snapshot profile.
- With respect to VMware, the SQL Databases backup extension works with pRDMs only. When using the SQL Databases backup extension, Replay Manager cannot create application consistent snapshots of databases that are stored on either .vmdks or vRDMs. If pRDMs are used for one database, it is highly recommended to use pRDMs for all databases, including the system databases.

4.1.4 Using snapshot profiles

Snapshot profiles can also be used to create snapshots directly on SC volumes. However, those snapshots will not be application consistent. SQL Server databases may not be recoverable from those snapshots. The odds of having recovery issues are higher in environments where database files for a given database are placed on multiples volumes, with write-intensive databases being particularly vulnerable. The risk of recovery problems can be reduced by using a consistency group to take a snapshot of all database volumes at the same point in time.

4.2 Choosing a replication strategy

Replicating data enables greater efficiency in disaster recovery, yielding major improvements in RPO and RTO. In other words, replicating virtual machines to a DR site means that during a DR test or actual DR execution, VMs can simply be powered on at the DR site as opposed to rebuilding or restoring operating systems and their applications, which is where much of the RTO improvement is introduced.

VMware SRM leverages the replication of storage in its automated recovery workflow. At the time of this publication, two replication approaches are supported: vSphere host-based replication, which is bundled with vSphere 5.1+ and replicates at the virtual machine level, or storage vendor array-based replication which occurs at the datastore or volume level. This paper focuses on SC Series replication and the integrated solutions that are available with it.

SC Series supports a variety of replication configuration options. Between arrays, replication can be carried out through FC, iSCSI, or a combination of both.

4.2.1 Asynchronous

Asynchronous replication occurs on a per-volume basis and the replication interval is typically based on the volume's snapshot schedule or the frequency at which snapshots are created for a volume. When asynchronous replication is configured for a volume, each time a snapshot completes it is replicated to the target system. As an example, if a snapshot is created on a volume every hour, asynchronous replication is

instantiated hourly for that volume. The length of time the replication will take depends on the size of the snapshot, which is dictated by the rate of change on the volume, and the bandwidth between the source and target systems. In the given example, assuming the replication bandwidth is sufficient, a one-hour RPO is established for the applications and data on the volume. A majority of customers implement asynchronous replication because it provides a good balance between cost and acceptable RPO.

4.2.2 High consistency synchronous

High consistency synchronous replication is also configured on a per-volume basis and the data on each volume is replicated to a remote system. While data replicated asynchronously is sent in batch intervals, synchronous replicates written data immediately when it is written at the source. In effect, each incoming write I/O becomes a dual-write process whereby data is written at the destination system and the source system. After the write is committed in both locations, a write acknowledgement is sent to the operating system and application where the write originated. With this in mind, synchronous replication provides a clear benefit over asynchronous: zero data loss in the event of a disaster, providing the initial replication is not out of date and replication is not paused. However, synchronous replication is more expensive in terms of application latency and site connectivity costs. A highly reliable replication link between sites is needed to accommodate the bulk of data being replicated while minimizing application latency at the same time.

Although synchronous replication is constantly replicating data as it is ingested at the source, it should be noted that as of Storage Center OS (SCOS) 6.3, frozen snapshots are also replicated to the remote site automatically. This provides flexible recovery options in that the data at the remote site can be recovered either from the active snapshot or from a frozen snapshot.

4.2.3 High availability synchronous

High availability synchronous replication is the same as high consistency synchronous replication except that replication is allowed to fall behind in favor of source site application availability during periods of extreme site to site latency or an unplanned remote site connectivity outage. Because replication is allowed to fall behind, data loss could occur in the event of a disaster at that time.

4.2.4 Asynchronous active snapshot

SC Series asynchronous replication also has the ability to replicate a volume's active snapshot. An active snapshot contains newly written data or data that has been changed on a volume since the last frozen snapshot was created. When the replication of active snapshot data is enabled on a per-volume basis, the data is replicated to the remote system as it is written to the source volume on an immediate but best-effort basis. Replicating the active snapshot is often called *semi-synchronous*. It is similar to synchronous replication except there is not a guarantee that data was written to the remote system before it is acknowledged on the source system. The two types are similar as long as the remote site bandwidth connectivity can support the data rate of change being replicated. However, if it does not, it is allowed to immediately fall behind and the result is no application latency penalty or predictable data loss, should a disaster occur at that time. Replicating the active snapshot is a feature of asynchronous replication, meaning frozen snapshots will be replicated as well.

4.2.5 Live Volume

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

SRM version 6.1 support for stretched storage with Live Volume has been added in DSM 2016 R1. 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 SRM Administration documentation available at [VMware Documentation](#).

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

With respect to the Live Volume managed replication feature, disaster recovery at a remote site, independent of the site or sites containing the Live Volume, is a good use case. However, Live Volume managed replications are not explicitly supported with VMware SRM. Live Volume managed replication volumes are activated using DSM.

Replay Manager is not compatible with Live Volume. Since automatic failover is not supported when using Live Volume with SRM, regular replication is often a better choice when protecting SQL Server virtual machines using SRM. SRM with replication provides quick recovery times without losing the protection benefits of application consistent snapshots provided by Replay Manager.

4.2.6 Volume replication considerations

Once a replication methodology is chosen, configuring replication for a single volume is straightforward. However, there are a few points that need to be raised to be sure they are covered. The first is that all of the data on a single volume is replicated. In a vSphere consolidated environment, it is likely that many virtual machines will reside on the datastore being replicated and therefore all of those VMs will be replicated to the remote system whether that was the original intent or not. In situations where only one or a handful of VMs on a datastore needs to be protected using replication, those VMs must be separated from the non-essential VMs using a Storage vMotion or cold migration process. This will ensure only the essential VMs in scope of the disaster recovery plan are being replicated which results in the best use of the replication bandwidth.

In cases where a SQL Server VM utilizes more than one virtual disk spanning datastores or perhaps uses RDMS, attention must be given to ensure consistency between the volumes in use. From a snapshot perspective, all of the volumes in use by SQL Server should be on a consistent snapshot schedule ensuring all volumes have precisely matching point-in-time snapshots. This can be handled automatically using Replay Manager or snapshot consistency groups if not using Replay Manager.

From a replication perspective, adequate sizing of the transport must be provided. This will help ensure snapshots are replicated to the remote system in a timely fashion. In turn, this will also aid in maintaining replication consistency because there is not currently an SC Series mechanism to guarantee this. For example, if a SQL Server utilizes six SC volumes with consistent snapshots and the consistent snapshots of those volumes are asynchronously replicated to a remote system as part of a DR plan, it is likely that some volumes will complete their replication interval before the others, based on the rate of change of the data being replicated. If a disaster were to occur after all volumes have completed their replication cycle, this is not a problem. However, if a disaster were to occur in the middle of a replication cycle, the completed replication of volume snapshots would be inconsistent at the DR site. The result could be the availability of one-hour RPO data for three of the volumes and two-hour RPO for the other three volumes. Recovering SQL Server databases with inconsistent recovery points will almost always be a problem.

4.3 Configuring Dell Storage Manager

Dell Storage Manager is a required component for SRM. It must be up and available in the recovery site in order for SRM to be able to carry out the automated failover workflow when the primary site goes down. In an active/active site configuration, a DSM Data Collector is required at both sites, with the primary Data Collector in one site and a remote Data Collector in the other. The primary or remote Data Collector can be at either site. In the interest of application locale and responsiveness, it is recommended to have the primary Data Collector at the location where the majority of the SC Series administration is performed.

DSM also contains a configuration setting that controls how SRM recovers volumes. This setting can be configured using the DSM client. Volumes can be recovered by either using the active snapshot or the latest frozen snapshot. To modify this setting, do the following:

1. Start the DSM client.
2. In the top-right area of the screen, click **Edit Data Collector Settings**.
3. On the left side of the **Edit Data Collector Settings** box, click **Replication Settings**.
4. Below **VMware SRM Settings**, there is a drop-down list for **SRM Selectable Snapshot**. Select one of the following options:
 - **Always use Active Snapshot** (default): Volumes will be recovered using the active snapshot.
 - **Use Active Snapshot if Replicating Active Snapshot**: Volumes will be recovered using the active snapshot if replication is configured to replicate the active snapshot. If the active snapshot is not being replicated, the latest frozen snapshot is used.
 - **Always use Last Frozen Snapshot**: Volumes will always be recovered using the latest frozen snapshot, even if the active snapshot is being replicated.
 - **Use Restore Point Settings**: Volumes will be recovered using the method defined in the restore point. For example, this option would allow database volumes to be recovered using the latest frozen snapshot, and all other volumes using the active snapshot.
5. Click **OK**.

Note that this setting will be ignored in the following scenarios where the most current data is replicated as part of the recovery plan workflow:

- When performing a planned migration
- When performing a disaster recovery while the primary site is up
- When performing a test recovery when the **Replicate recent changes to recovery site** option is selected

4.4 Configuring recovery in SRM

When recovering virtual machines running SQL Server, it is critical to understand the implications of the two methods that SRM can use to recover volumes. SRM can either use the active snapshot or the latest frozen snapshot when recovering volumes. If SQL Server databases cannot be recovered from the volumes created by SRM, manual intervention will be required to successfully complete the recovery.

4.4.1 Recovery from the active snapshot

Recovering from the active snapshot can provide the lowest RPO, as it will contain the latest view of the volume at the target site. For volumes using synchronous replication, this is the recommended recovery method. When using asynchronous replication, recovering volumes from the active snapshot is not always reliable. If SRM recovery fails using the active snapshot, the recovery will need to be completed manually

using the latest frozen snapshot. Consider the following before using the active snapshot with asynchronous replication:

- Writes are queued up to be replicated in write order. However, if replication gets behind, it can consolidate multiple writes to the same logical block address (LBA) so that only the latest version of the LBA is sent. This type of write consolidation can prevent the successful recovery of SQL Server databases. The risk of this type of recovery problem is low when there is sufficient bandwidth between data centers to prevent replication from falling behind. To eliminate this risk, recover from a frozen snapshot.
- Since each volume is replicated independently, it is likely that the active snapshots of the transaction log and data volumes will be at different points in time at the target site. While the SQL Server crash recovery mechanism is very good, there is a risk that the database recovery will fail if the data and transaction log volumes are too far out of sync with each other. Recovering from frozen snapshots will help minimize this risk. However, when using frozen snapshots, there is a slight risk that the latest frozen snapshots on a given set of volumes won't be from the same point in time. This risk is low if there is sufficient bandwidth between sites. This risk can be eliminated by placing all database files for a given database on the same volume. Be sure to consider the implications of putting all database files on the same volume.

4.4.2 Recovery from the latest frozen snapshot

For volumes replicated asynchronously, using the latest frozen snapshot is the recommended recovery method. In particular, this method is ideal for database volumes when combined with application consistent snapshots created by Replay Manager. The recovery procedure will vary based on how the snapshot was created.

Since each volume is replicated independently, there is a risk that the latest frozen snapshots for a given set of replicated volumes will not be from the same point in time, even if snapshots are taken at the same time on the source volumes. This risk is low if there is sufficient bandwidth between sites. This risk can be eliminated by placing all database files for a given database on the same volume. Be sure to consider the implications of putting all database files on the same volume.

4.4.2.1 Using snapshots created by the VMware backup extensions in Replay Manager

When recovering from snapshots created by the VMware backup extensions, the virtual machine will need to be rolled back to the VMware snapshot created by Replay Manager.

For manual recovery, configure the recovery plan for the virtual machine to leave the virtual machine powered off. After SRM recovery is complete, use the vSphere Snapshot Manager to revert the virtual machine back to the snapshot created by Replay Manager. Once that has been done, power the virtual machine on.

For automated recovery, configure the recovery plan to power the virtual machine on. Create a recovery step to run the following PowerShell cmdlets before the virtual machine is powered on:

```
# Assign Variables

$vCenterDnsName = "<vCenter DNS Name>"
$VmName         = "<Virtual Machine Name>"

# Load PowerCLI

Add-PSSnapin VMware.VimAutomation.Core
```

```

# Connect to vCenter

Connect-VIServer -Server $vCenterDnsName

# Get the virtual machine

$Vm = Get-VM -Name $VmName

# Get the latest Replay Manager snapshot ( there should be only one )

$VmSnapshot = Get-Snapshot -VM $Vm `
    | Where-Object { $_.Description -like "*Replay Manager*" } `
    | Sort-Object -Property Created `
    | Select-Object -Last 1

# Revert the VM to the snapshot

Set-VM -VM $Vm -Snapshot $VmSnapshot -Confirm:$false

# Remove the snapshot

Remove-Snapshot $VmSnapshot -Confirm:$false

# Disconnect from vCenter

Disconnect-VIServer -Server $vCenterDnsName

```

Note: This is an example of the PowerShell commands required to rollback a virtual machine to a VMware snapshot. This is not a production-ready script.

4.4.2.2 Using snapshots created by the SQL databases backup extension in Replay Manager

Volumes recovered from snapshots of pRDMs created by Replay Manager will not be in a usable state after the virtual machine is powered on by SRM. Attributes placed on the volume by VSS when the snapshot was created will need to be removed.

For manual recovery, configure the recovery plan to power on the virtual machine. After SRM recovery is complete, run the following PowerShell cmdlets, on each pRDM recovered from VSS snapshots:

```

# Assign Variables

$DiskSerialNumber = "<disk serial number>"

# Load the Storage Center PowerShell Command Set

If (!( Get-PSSnapin | Where-Object { $_.Name -eq
"Compellent.StorageCenter.PSSnapin" } ))
{
    Add-PSSnapin -Name "Compellent.StorageCenter.PSSnapin" | Out-Null
}

```

```

}

# Reset the volume

Set-CMLDiskDevice -SerialNumber $DiskSerialNumber -ReadOnly:$False
Set-CMLDiskDevice -SerialNumber $DiskSerialNumber -ResetSnapshotInfo
Set-CMLDiskDevice -SerialNumber $DiskSerialNumber -Online

```

Once all of the database volumes have been cleaned up and brought online, start the SQL Server service as well as any other services used with SQL Server (like the SQL Server Agent).

For automated recovery, configure the recovery plan to power the virtual machine on. Create a recovery step to run the following PowerShell cmdlets after the virtual machine is powered on:

```

# Assign Variables

$SQLServerServiceName = "MSSQLSERVER"
$SQLAgentServiceName = "SQLSERVERAGENT"

# Load the Storage Center PowerShell Command Set

If (!(Get-PSSnapin | Where-Object { $_.Name -eq "Compellent.StorageCenter.PSSnapin" } ))
{
    Add-PSSnapin -Name "Compellent.StorageCenter.PSSnapin" | Out-Null
}

# Get all offline disks (assumes offline disks are pRDMs recovered from VSS
snapshots)

$DiskArray = Get-CMLDiskDevice | Where-Object { $_.Status -eq "Offline" }

# Reset each disk

ForEach ( $Disk in $DiskArray )
{
    Set-CMLDiskDevice -SerialNumber $Disk.SerialNumber -ReadOnly:$False
    Set-CMLDiskDevice -SerialNumber $Disk.SerialNumber -ResetSnapshotInfo
    Set-CMLDiskDevice -SerialNumber $Disk.SerialNumber -Online
}

# Start SQL Server services

Start-Service -Name $SQLServerServiceName
Start-Service -Name $SQLAgentServiceName

```

Note: This is an example of the PowerShell commands required to recover pRDMs from snapshots created by Replay Manager. This is not a production-ready script.

4.5 Performing a disaster recovery test

SRM provides the ability to test the recovery process without performing an actual failover. Virtual machines can be brought online in an isolated environment at the disaster recovery site. Because the test environment is isolated, recovered virtual machines can have the same name or IP address as production virtual machines. The production environment is not impacted by the test.

To perform a disaster recovery test failover, perform the following:

1. In the vSphere web client, open the **SRM** plug-in.
2. On the left side of the client, double-click **Recovery Plans**.
3. Under the **Objects** tab, select the recovery plan that needs to be tested.
4. Under the **Actions** drop-down menu (or the green arrow icon), select **Test**.

To clean up after the test is complete, perform the following:

1. In the vSphere web client, open the **SRM** plug-in.
2. On the left side of the client, double-click **Recovery Plans**.
3. Under the **Objects** tab, select the recovery plan that was tested and needs to be cleaned up.
4. Under the **Actions** drop-down menu (or the broom icon), select **Cleanup**.

When performing a disaster recovery test, consider the following:

- By default, the **Replicate recent changes to recovery site** option will be checked in the **Storage Options** pane of the **Test** dialog box. If this option is selected, the test will create new crash consistent snapshots at the primary site and recover the volumes using those snapshots once they have been replicated to the recovery site. For an accurate disaster recovery test, uncheck this option. Snapshots created by Replay Manager will not be used if this option is selected.
- SC Series storage will not be accessible from inside the recovered virtual machines. The test environment is isolated from the network.

4.6 Performing a disaster recovery

When a disaster is declared, SRM can be used to efficiently move virtual machines from the primary site to the disaster recovery site. A fully automated recovery plan can greatly reduce the recovery time while also preventing costly mistakes.

To perform a disaster recovery failover, perform the following:

1. In the vSphere web client, open the **SRM** plug-in.
2. On the left side of the client, double-click **Recovery Plans**.
3. Under the **Objects** tab, select the recovery plan that needs to be run.
4. Under the **Actions** drop-down menu (or the red circle with white arrow icon), select **Run**.
5. In the **Recovery Confirmation** pane, check the box **I understand that this process will permanently alter the virtual machines and infrastructure of both the protected and recovery datacenters**.
6. In the **Recovery Type** pane, select the **Disaster Recovery** option .
7. Click **Next**.
8. Click **Start**.
9. If there are multiple recovery plans, repeat the process for each recovery plan.

If the primary site is still available, SRM will effectively perform a planned migration. SRM will power off the virtual machines at the primary site, create new crash consistent snapshots, and recover the volumes from those snapshots once they have been replicated to the recovery site. This will provide a failover with no data loss. Since the recovery from crash consistent snapshots is different than the recovery from Replay Manager snapshots, any recovery plans using Replay Manager snapshots also need to accommodate crash consistent snapshots.

After SRM completes the recovery, additional setup or reconfiguration may be needed for systems that cannot be managed by SRM. For example, things like tape backups or Replay Manager may need to be reconfigured in order to function correctly at the recovery site.

4.7 Performing a planned migration

SRM can also be used to move virtual machines from one site to another. This type of recovery functions the same as a disaster recovery. However, both sites must be online and available for the migration to succeed.

To perform a planned migration, perform the following:

1. In the vSphere web client, open the **SRM** plug-in.
2. On the left side of the client, double-click **Recovery Plans**.
3. Under the **Objects** tab, select the recovery plan that needs to be run.
4. Under the **Actions** drop-down menu (or the red circle with white arrow icon), select **Run**.
5. In the **Recovery Confirmation** pane, check the box **I understand that this process will permanently alter the virtual machines and infrastructure of both the protected and recovery datacenters**.
6. In the **Recovery Type** pane, select the **Planned Migration** option.
7. Click **Next**.
8. Click **Start**.
9. If there are multiple recovery plans, repeat the process for each recovery plan.

A planned migration will power off the virtual machines at the primary site, create new crash consistent snapshots, and recover the volumes from those snapshots once they have been replicated to the recovery site. This will provide a failover with no data loss. Since the recovery from crash consistent snapshots is different than the recovery from Replay Manager snapshots, any recovery plans using Replay Manager snapshots also need to accommodate crash consistent snapshots.

After SRM completes the recovery, additional setup or reconfiguration may be needed for systems that cannot be managed by SRM. For example, tape backups or Replay Manager may need to be reconfigured in order to function correctly at the recovery site.

A Additional resources

A.1 Technical support and resources

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

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

A.2 Referenced or recommended publications

Dell EMC publications:

- Dell EMC SC Series Arrays and Microsoft SQL Server:
http://en.community.dell.com/techcenter/extras/m/white_papers/20438053
- Dell EMC SC Series Best Practices with VMware Site Recovery Manager:
http://en.community.dell.com/techcenter/extras/m/white_papers/20438016

VMware publications:

- SQL Server on VMware – Availability and Recovery Options
<https://communities.vmware.com/docs/DOC-13270>
- SQL Server on VMware – Best Practices Guide:
<https://communities.vmware.com/docs/DOC-13249>