



Compellent Storage Center

VMware Site Recovery Manager 5.x

Best Practices Guide

Dell Compellent Technical Solutions Group
April 2013

THIS BEST PRACTICES GUIDE IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2012 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. Microsoft® and Windows® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.

Table of Contents

1	Preface	1
1.1	Audience	1
1.2	Purpose.....	1
1.3	Customer Support.....	1
2	Introduction.....	2
2.1	Introduction to Site Recovery Manager	2
2.2	What's New in SRM 5.0.....	2
2.3	What's New in SRM 5.1	3
3	Setup Prerequisites	4
3.1	Enterprise Manager	4
3.2	Storage Center	4
3.3	VMware vSphere	4
3.4	Storage Replication Adapter (SRA)	4
4	Site Recovery Manager Architecture	5
4.1	Single Protected Site – Array Based Replication	5
4.2	Dual Protected Site – Array Based Replication.....	6
4.3	Single Protected Site – vSphere Replication.....	7
4.4	Dual Protected Sites – vSphere Replication.....	8
5	Enterprise Manager Configuration.....	9
5.1	Data Collector Configuration.....	9
5.2	Enterprise Manager Logins	9
5.3	Creating dedicated SRA access accounts.....	9
5.4	Saving Restore Points	10
5.5	Validating Restore Points	11
5.6	Automatic Restore Point Saving Schedule	12
5.7	Modifying SRM Settings For Larger Environments	12
6	Configuring Replications	14
6.1	Asynchronous Replications (Supported).....	14
6.2	Synchronous Replications (Supported in Storage Center 6.3 and newer)	14
6.3	Live Volume Replications (Not Supported).....	14

6.4	Data Consistency while Replicating by Replay Schedule	15
6.5	Data Consistency while Replicating the Active Replay	16
6.6	Replication Dependencies and Replication Transfer Time.....	17
6.7	Custom Recovery Tasks	18
7	Site Recovery Manager Configuration	19
7.1	Configuring the Array Managers.....	19
7.2	Creating Array Pairs	21
7.3	Rescanning Array Managers	22
7.4	Creating Protection Groups	23
7.5	Creating Recovery Plans	24
8	Recovery Plan Execution	26
8.1	Testing a Recovery Plan	26
8.2	Running a Recovery Plan	26
9	Reprotect and Failback.....	28
9.1	Overview	28
9.2	Reprotection.....	28
9.3	Failback.....	29
10	Conclusion.....	30
11	Appendix A – Example Scripts	31
12	Appendix B – Additional Resources.....	33
12.1	Compellent Resources	33
12.2	VMware Resources.....	33

Document Revisions

Date	Revision	Author	Comments
8/23/2011	A	Jason Boche	Initial Draft
11/17/2011	B	Jason Boche	Updated for 5.5.4
12/6/2011	C	Jason Boche	Replication sections
3/22/2012	D	Jason Boche	SRA version correction
7/2/2012	E	Jason Boche	Added warning; spelling
10/19/2012	F	Jason Boche	Updated diagrams
10/31/2012	G	Jason Boche	Updated for SRM 5.1
4/15/2013	5.1	Jason Boche	Updated for 6.3 and Sync replication support
7/15/2013	5.1.1	Jason Boche	Corrected two section titles

1 Preface

1.1 Audience

The audience for this document is System Administrators who are responsible for the setup and maintenance of VMware Site Recovery Manager, VMware vSphere, and associated storage. Readers should have a strong working knowledge of SRM, vSphere, Dell Compellent Storage Center and Enterprise Manager.

1.2 Purpose

This document provides an overview of VMware Site Recovery Manager and introduces best practice guidelines for configuring SRM and the SRA when using the Dell Compellent Storage Center.

1.3 Customer Support

Dell Compellent provides live support 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.

2 Introduction

2.1 Introduction to Site Recovery Manager

This document will provide configuration examples, tips, recommended settings, and other storage guidelines a user can follow while integrating VMware Site Recovery Manager with the Compellent Storage Center. This document has been written to answer many frequently asked questions with regard to how VMware interacts with the Site Recovery Manager, as well as basic configuration.

Compellent advises customers to read the Site Recovery Manager documentation provided on the VMware web site before beginning their SRM implementation.

2.2 What's New in SRM 5.0

- **New User Interface (UI)** - Management of the primary and secondary SRM sites is consolidated from two separate interfaces down to one with both sites being visible in one vSphere Client without linked mode.
- **Planned Migration** - SRM can now be used as a tool to gracefully migrate protected virtual machines from the primary to secondary site.
- **Reprotect and Failback** - Once virtual machines are moved from one site to another via planned migration or disaster recovery, the VM re-protection process is automated and includes reverse replication which enables VMs to fail back to the opposite site.
- **vSphere Host-Based Replication (optional)** - A new appliance is introduced which has the ability to provide host based replication for VMs on a per-VM granular basis, abstracting the physical attributes for the storage such as array type and protocol.
- **Faster IP Customization** - Reconfiguring TCP/IP via recovery plan is more efficient and executes faster.
- **New Shadow VM Icons** - Provides better visibility at the secondary site for placeholder VMs.
- **In Guest Scripts** - Script automation can now be generated from within guest VMs themselves.
- **VM Dependency** - 5 Priority Groups and VM dependency relationships within protection groups.
- **Improved Reporting** - Provides increased awareness for historical analysis.
- **IPv6** - Future proof network design.

2.3 What's New in SRM 5.1

- **Reprotect and Failback with vSphere Replication** - Once virtual machines are moved from one site to another via planned migration or disaster recovery, the VM reprotection process is automated and includes reverse replication which enables VMs to fail back to the opposite site. Once an Array Based Replication feature only, this is now supported with vSphere Replication.
- **64-bit SRM Server** - Each SRM Server instance is now developed on 64-bit architecture which paves the way for future scalability enhancements.
- **More Robust VSS Integration with vSphere Replication** - Flushing of application writers for application consistency brings VSS integration parity closer between Array Based and vSphere Replication.
- **Forced Recovery for vSphere Replication** - The Forced Recovery feature is now available to both vSphere and Array Based Replication.
- **vSphere Replication Decoupling** - vSphere Replication was introduced as a feature bundled in SRM 5.0. In SRM 5.1, vSphere Replication is decoupled from SRM and is now included in the vSphere Essentials Plus and above platform bundle.
- **Relaxed Licensing** - VMware added SRM support for the Essentials Plus tier of vSphere licensing.
- **vSphere 5.1 Compatibility** - SRM 5.1 is compatible with vSphere 5.1.

3 Setup Prerequisites

3.1 Enterprise Manager

Compellent Enterprise Manager Version 5.5.4 or greater is required for SRM 5.0 and 6.2.2 or greater is required for SRM 5.1 for the Storage Replication Adapter (SRA) to function. The SRA makes calls directly to the Enterprise Manager Data Collector to manipulate the storage. It is recommended to have the latest version of the Enterprise Manager Data Collector installed to ensure compatibility with SRM 5.0.x and the Compellent SRA.

3.2 Storage Center

It is required to have two Compellent Storage Center (version 5.4 or greater for SRM 5.0, 5.5.3 for SRM 5.1) systems with Remote Data Instant Replay (replication) between the sites licensed and operational. Site Recovery Manager using Array Based Replication requires two Compellent systems replicating between each other. SRM using vSphere Replication can leverage any storage certified for use with vSphere including Dell Compellent Storage Center.

3.3 VMware vSphere

VMware Site Recovery Manager 5. x with vCenter 5.x, ESXi 5.x and/or ESXi/ESX 3.5 or newer are required. Please check the latest Site Recovery Manager Compatibility Matrix for the versions of software required for SRM to function. As of the release of SRM 5.1, vSphere Essentials Plus licensing is required.

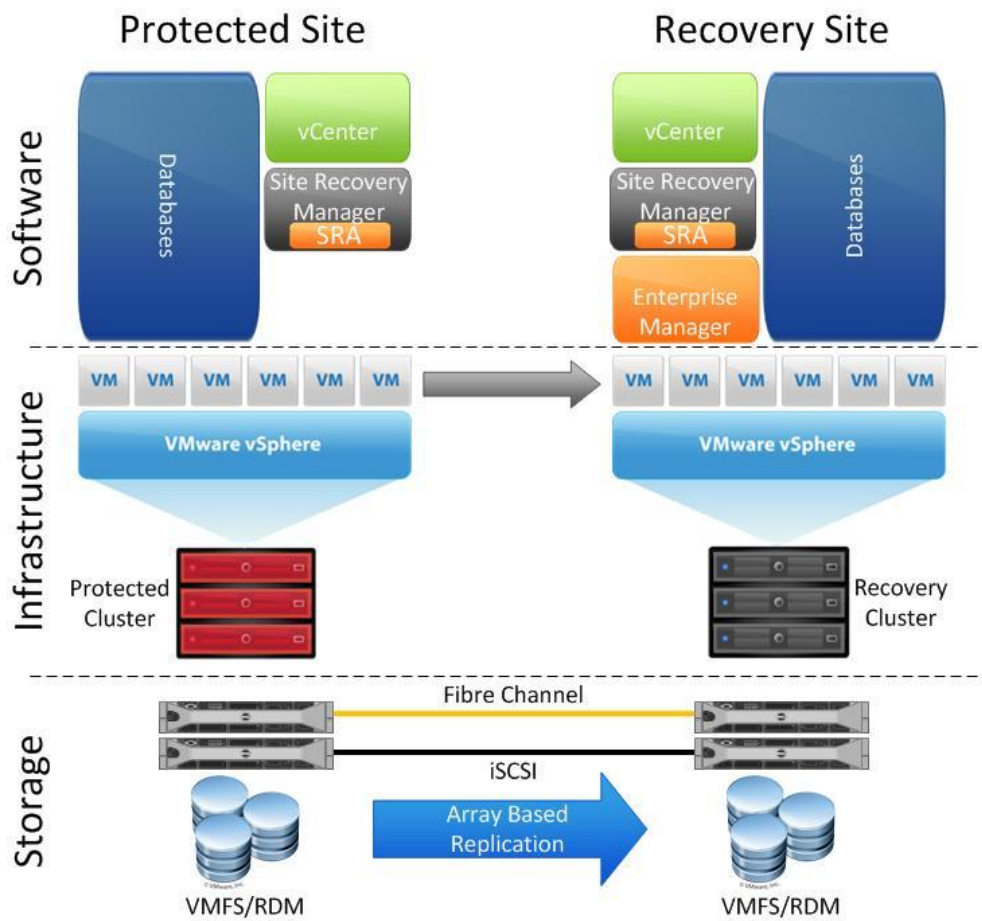
3.4 Storage Replication Adapter (SRA)

The Compellent Storage Replication Adapter (SRA) is required to be running version 5.5.3 or greater for SRM 5.0 and 6.2.2.7 for SRM 5.1.

4 Site Recovery Manager Architecture

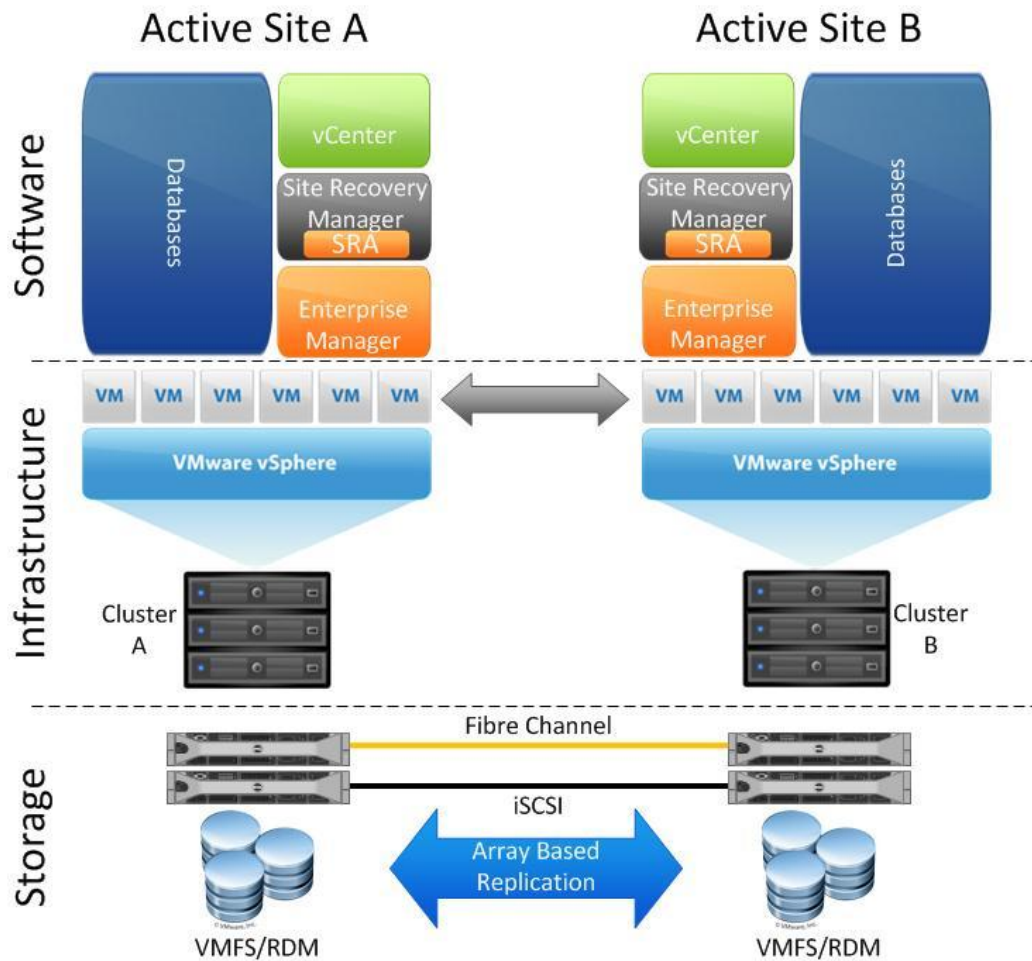
4.1 Single Protected Site – Array Based Replication

This configuration is generally used when the secondary site does not have any virtual machines that need to be protected by SRM. The secondary site functions solely for disaster recovery purposes. The Enterprise Manager Data Collector Server is placed at the disaster recovery site because it is required by SRM to perform recovery functions. An Enterprise Manager Data Collector Server needs to be running at the site opposite of protected virtual machines in the event of a site failure for SRM to function. Keep this in mind if you are planning on using SRM 5.x's planned migration or failback feature.



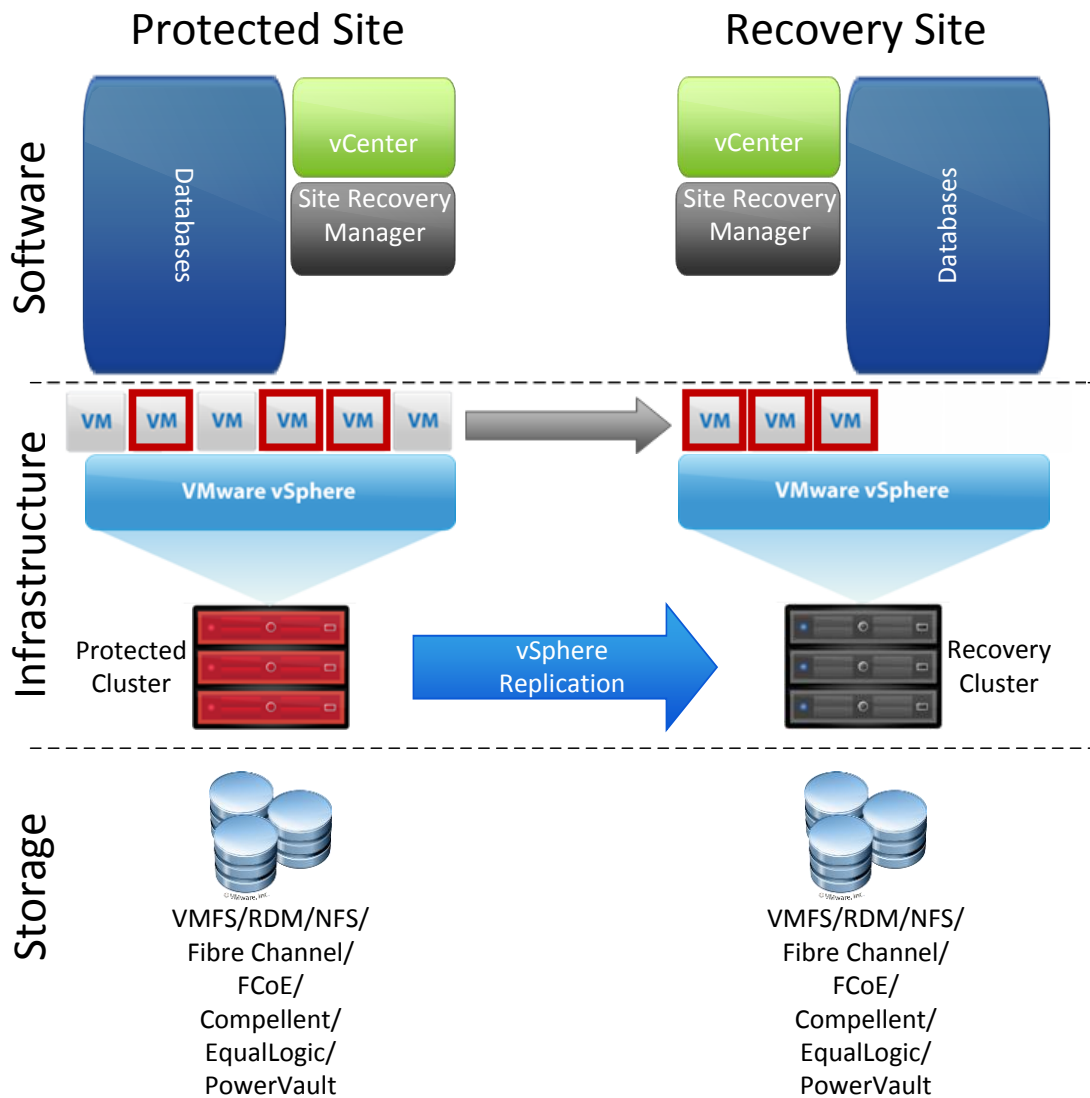
4.2 Dual Protected Site – Array Based Replication

This configuration is generally used when both sites have virtual machines that need to be protected by SRM. This scenario may be commonly used in conjunction with SRM assisted migrations which is a new feature in SRM 5.0. In this example, each site replicates virtual machines to the opposing site in order to protect both sites from a failure or to orchestrate a planned migration of virtual machines. Planned migrations can be performed with just one Enterprise Manager Data Collector server. However, once there are active virtual machines running simultaneously at both the primary and secondary sites, the site with the Enterprise Manager Data Collector will not be adequately protected by SRM. Enterprise Manager Data Collector Servers are placed at each site so that either site can fail.



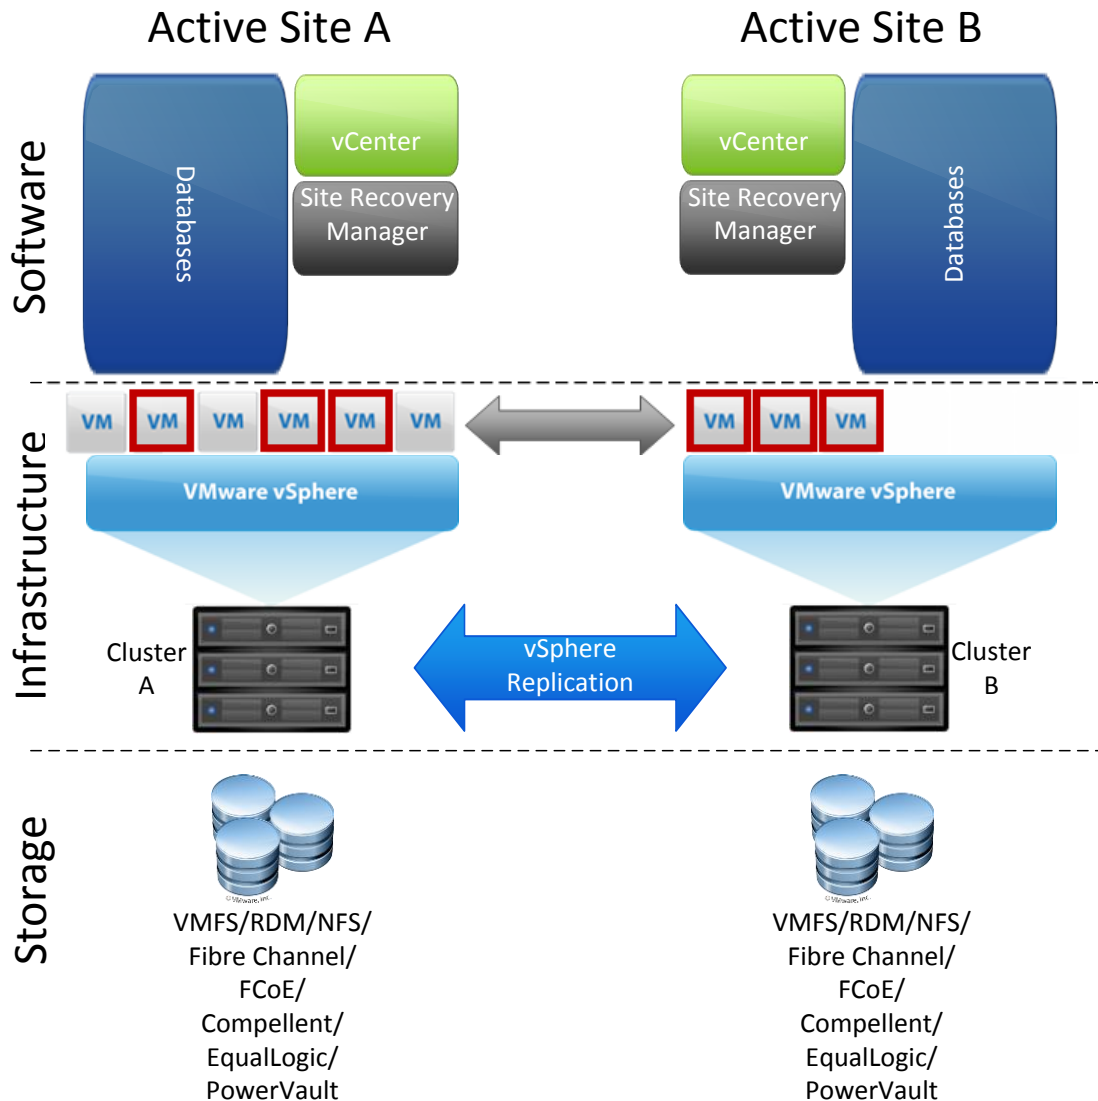
4.3 Single Protected Site – vSphere Replication

Although the main focus of this document is SRM integration with Dell Compellent Array Based Replication, it should also be pointed out that as of SRM 5.0, vSphere Replication can be used in addition to or in place of Array Based Replication. vSphere Replication has a few unique advantages over Array Based Replication. Two of the main ones being that a granular selection of individual VMs are replicated instead of entire datastores of VMs and vSphere datastore objects abstract the underlying storage vendor, model, protocol, and type meaning replication can be carried out between different array models and protocols, even 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 businesses with aggressive storage constraints.



4.4 Dual Protected Sites – vSphere Replication

The architectural changes with vSphere Replication are carried into the Active/Active site model. Note in both vSphere Replication architecture diagrams that replication is handled by the vSphere hosts via the vSphere network stack. There is no SRA in this architecture as there is with Array Based Replication. It should also be highlighted that not all components of vSphere Replication are represented in detail here. A deployment of vSphere Replication consists of multiple appliances deployed at each site and on each vSphere host that will be handling the movement of data between sites. Refer to VMware documentation for a detailed look at vSphere Replication.



5 Enterprise Manager Configuration

5.1 Data Collector Configuration

As illustrated in the Architecture section, Enterprise Manager is a critical piece to the SRM infrastructure because the Data Collector processes all of the calls from the Storage Replication Adapter (SRA) and relays them to the Storage Centers to perform the workflow tasks.


Deciding whether or not to use one or two Enterprise Manager Servers depends on whether virtual machines need to be protected in one or multiple sites.

- If protecting virtual machines at a single site, a single Enterprise Manager Data Collector will suffice, and it is highly recommended that it be placed at the recovery site.
- If protecting virtual machines at both sites, it is highly recommended to place Enterprise Manager Data Collectors at each site.

5.2 Enterprise Manager Logins

For SRM to function, the Storage Replication Adapter (SRA) must use Enterprise Manager Login credentials that have rights to both of the Storage Center systems replicating the virtual machine volumes.

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



5.3 Creating dedicated SRA access accounts

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

Following the example above:

- Create an account named "sra-service-acct" on both the protected site array and the recovery site array.
 - This account needs Administrator privileges, so make sure the password assigned is secure.

- For added security, you could create different accounts on both systems with different passwords. For example, on the protected array it could be named "sra-system1" and the secondary system it could be named "sra-system2". The account names and passwords are arbitrary.
- Create a new account within Enterprise Manager named "sra-admin".

User Settings

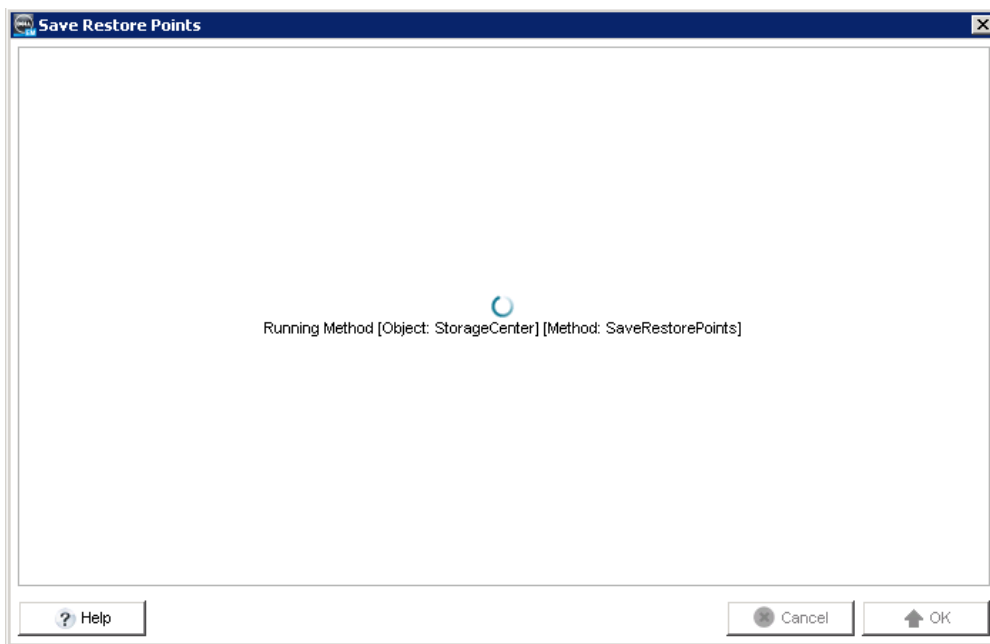
Name	sra-admin
Email Address	sra-admin@techsol.local
Privilege	Administrator
New Password	*****
Confirm Password	*****
Requires Password Change	<input type="checkbox"/> Enabled

The "sra-admin" account used to access Enterprise Manager can now be used for configuring the Storage Center credentials within SRM.

5.4 Saving Restore Points

Saving restore points must be completed for the SRA to be able to query the active replications and should be performed after a major SRM event such as performing a Planned Migration or Disaster Recovery. The process can be initiated in one of two ways:

1. For convenience, it is automatically initiated at the end of the **Create Replication Wizard**:



2. From the **Enterprise Manager Actions** menu:



5.5 Validating Restore Points

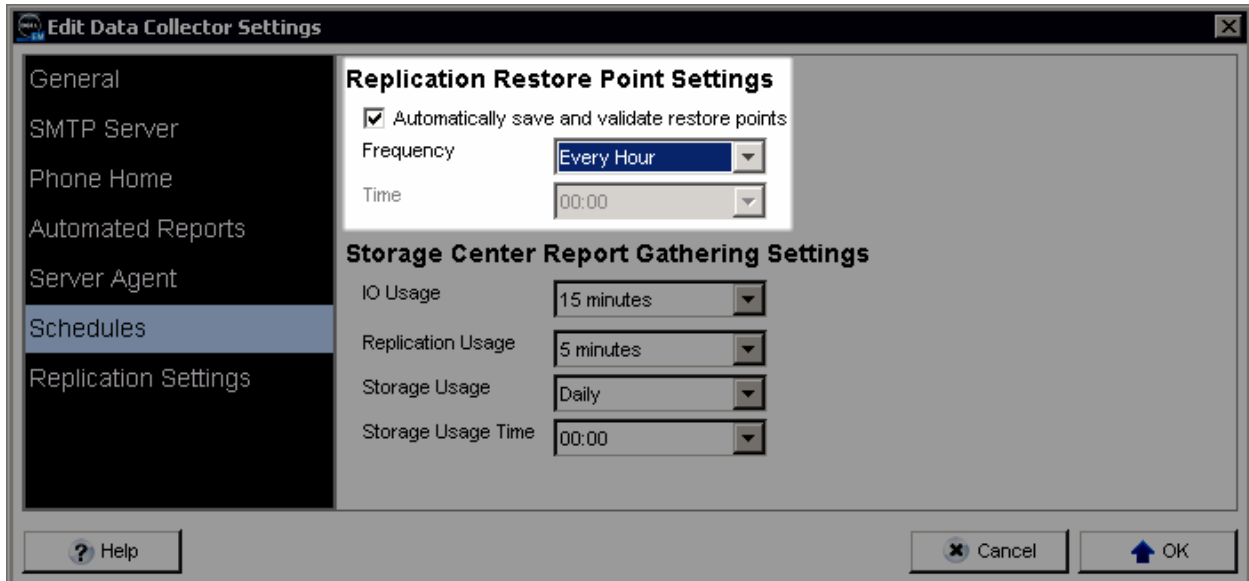
The 'Validate Restore Points' process reconciles the list of saved restore points with the list of replication jobs.

From the **Enterprise Manager Actions** menu:



5.6 Automatic Restore Point Saving Schedule

The **Finish saving Restore Points** screen in the **Save Restore Points Wizard** has the option to save restore points automatically at a selected interval by clicking on the **Set Replication Restore Schedule** link. It is recommended to configure the data collector to save the restore points **hourly**. This helps to ensure that the most current restore points are available for the SRA to query for replication information. If using multiple Enterprise Manager Data Collectors, the first Data Collector is configured as the Primary. The second Data Collector is installed and configured as a Remote Data Collector. Restore points are saved on the Primary Data Collector and replicated to the Remote Data Collector at one minute intervals. Replications must be created and Restore Points saved before the volume can be protected by SRM. Non-replicated volumes won't be discovered as a device by SRM and thus cannot be protected.



5.7 Modifying SRM Settings For Larger Environments

VMware Site Recovery Manager ships with a default configuration which is tuned for a large cross-section of environments. However, each customer environment is unique in terms of architecture, infrastructure, size, and Recovery Time Objective (RTO). Generally speaking, larger or more complex SRM environments may require tuning adjustments within SRM in order for SRM to work properly. VMware KB article 2013167 outlines some of the adjustments that can be made to accommodate such environments.

- storage.commandTimeout** – **Min:** 0 **Default:** 300 **Max:** 9223372036854775807

This option specifies the timeout allowed (in seconds) for running SRA commands in Array Based Replication related workflows. Recovery Plans with a large number of datastores to manage will fail if the storage related commands take longer than five minutes to complete. Increase this value (i.e. 3600 or higher) in the Advanced SRM Settings.
- storageProvider.hostRescanTimeoutSec** – **Min:** 0 **Default:** 300 **Max:** 9223372036854775807

This option specifies the timeout allowed (in seconds) for host rescan operations 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 (i.e. 600 or higher) in the Advanced SRM Settings.
- storageProvider.hostRescanRepeatCnt** – **Min:** 0 **Default:** 1 **Max:** 9223372036854775807

This option specifies the number of additional host rescans performed 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 (i.e. 2 or higher) in the Advanced SRM Settings.
- defaultMaxBootAndShutdownOpsPerCluster** – **Default:** off

This option specifies the maximum number of concurrent power-on operations performed by SRM at the cluster object level. Enable by specifying a numerical value (i.e. 32) as shown below by modifying the vmware-dr.xml file (or configure per cluster in vCenter | DRS options | srmMaxBootShutdownOps).

```
<config>
<defaultMaxBootAndShutdownOpsPerCluster>32 </defaultMaxBootAndShutdownOpsPerCluster>
</config>
```
- 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 (i.e. 4) as shown below by modifying the vmware-dr.xml file.

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

The vmware-dr.xml file is located in a directory named 'config' which resides within the Site Recovery Manager installation folder which will vary depending on the operating system and SRM version.

i.e.

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

6 Configuring Replications

Storage Center replication in coordination with Site Recovery Manager can provide a robust disaster recovery solution. Since each replication method affects recovery differently, choosing the correct method to meet business requirements is important. Here is a brief summary of the different options.

6.1 Asynchronous Replications (Supported)

- In an asynchronous replication, the I/O needs only be committed and acknowledged to the source system, so the data can be transferred to the destination in a nonconcurrent timeframe. There are two different methods to determine when data is transferred to the destination:
 - **By replay schedule** – The replay schedule dictates how often data is sent to the destination. When each replay is taken, the Storage Center determines which blocks have changed since the last replay (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 to verify that the recovery point objective (RPO) can be met.
 - **Replicating the active replay** – With this method, the data is transferred “near real-time” to the destination, usually requiring more bandwidth than if the system were just replicating the delta changes in the replays. 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 it can”. 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.
- One of the benefits of an asynchronous replication is that the replays are transferred to the destination volume, allowing for “check-points” at the source system as well as the destination system.

6.2 Synchronous Replications (Supported in Storage Center 6.3 and newer)

- The data is replicated real-time to the destination. In a synchronous replication, 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 will slow down access times on the source volume.

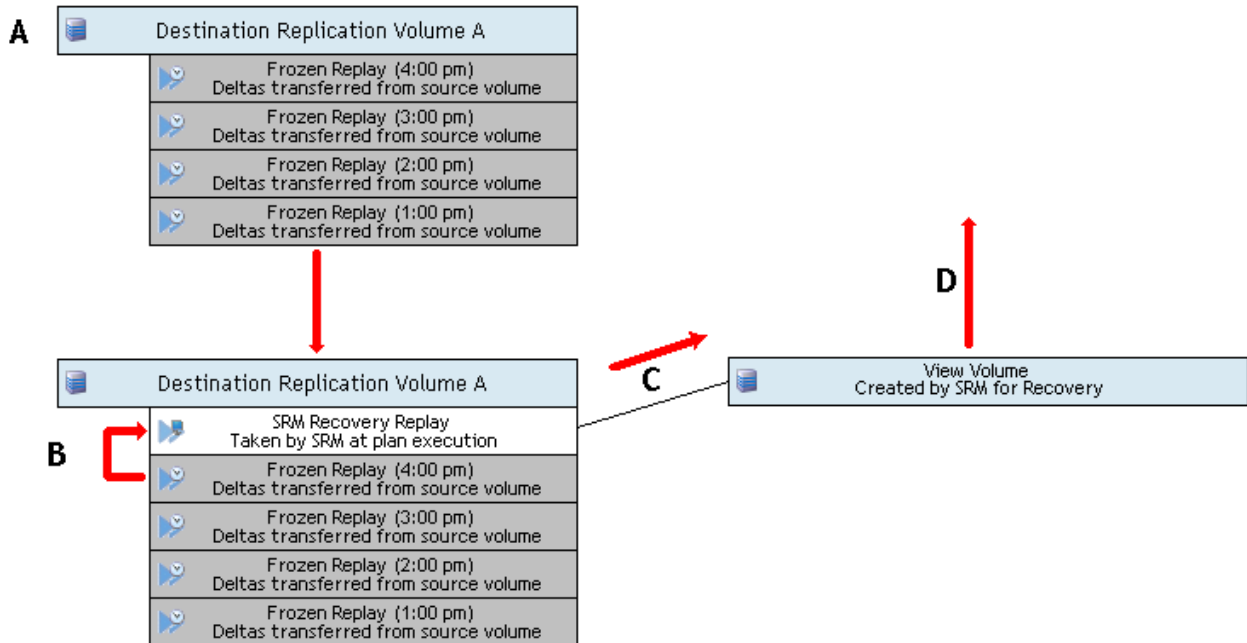
6.3 Live Volume Replications (Not Supported)

- Live Volume replications add an additional abstraction layer to the replication allowing mapping of the same volume through multiple Storage Center systems.
- Live Volume replications are not supported due to the fact that using them with SRM

is mutually exclusive, and because SRM may be confused by the volume being actively mapped at two different sites.

6.4 Data Consistency while Replicating by Replay Schedule

When replicating by replay schedule here are the consistency states of replications during plan execution.



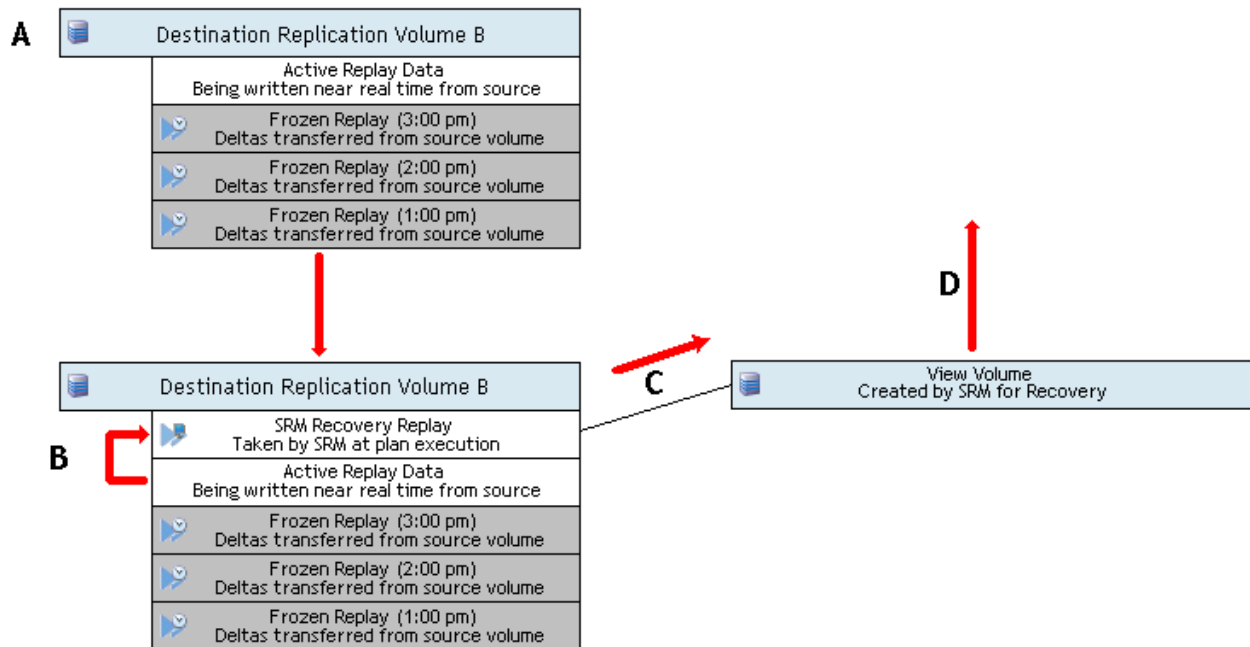
- A. Once a replay is taken of the source volume, the delta changes begin transferring to the destination immediately. The consistency state of the data within this replay is dependent on whether or not the application had the awareness to quiesce the data before the replay was taken.
- B. During a Recovery Plan test, a new replay is taken of the destination volume. This is done per the VMware SRM specification to capture the latest data that has arrived at the DR site. Of course this means that the consistency of the data is dependent on whether or not the previous replay was completely transferred. For example using the figure above:
 - a. If the 4:00 pm replay taken at the primary site was application consistent, but at the time the SRM Recovery Replay was taken, only 75% of that replay's data had been transferred and thus the data is considered incomplete.
 - i. If this scenario is encountered, it may be necessary to perform manual recovery steps in order to present the next latest replay (such as the 3:00pm that was completed and is thus still consistent) back to the application.
 - b. If the 4:00 pm replay taken at the primary site was application consistent, and at the time the SRM Recovery Replay was taken, all 100% of that replay had finished transferring, the resulting newly taken replay will include all of the 4:00pm replay data, and thus the application consistency of the data will be

preserved.

- C. Once the SRM Recovery Replay has been taken, a view volume is created from that Replay.
- D. The View volume is then presented to the ESX(i) host(s) at the DR site for SRM to begin test execution of the Recovery Plan.

6.5 Data Consistency while Replicating the Active Replay

When replicating the active replay, here are the consistency states of replications during plan execution.



- A. As writes are committed to the source volume, they are near simultaneously transferred to the destination and stored in the Active Replay. (See figure above) Keep in mind that consistent Replays can still be taken of the source volume, and the check points will be transferred to the destination volume when replicating the active replay.
- B. During a Recovery Plan test, a new replay is taken of the destination volume which “locks in” all the data that has been transferred up to that point in time. Data stored within the Active Replay will most likely be crash consistent. For example, using the figure above:
 - a. Although the 3:00 pm replay taken at the primary site was application consistent, at the time the SRM Recovery Replay was taken, the data must still be deemed crash consistent because it is highly probable that writes into the Active Replay occurred after the 3:00 pm replay was taken.
 - b. If the application is unable to recover the crash consistent data captured in this replay, then manual steps would be required to present the last known consistent replay (such as the 3:00pm frozen replay) back to the host.
 - c. The only time that application data can be considered consistent while replicating the active replay, is if ALL I/O has ceased to the source volume,

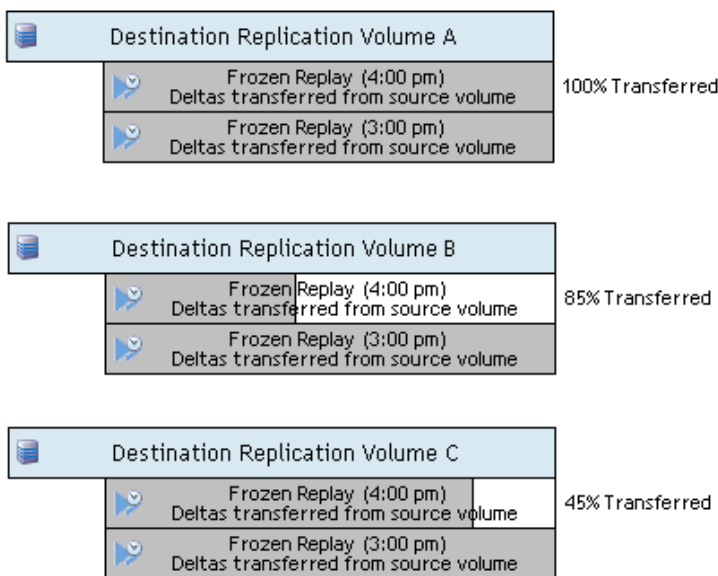
followed by a complete synchronization of the data from the source to the destination.

- C. Once the SRM Recovery Replay has been taken, a view volume is created from that Replay.
- D. The View volume is then presented to the ESX(i) host(s) at the DR site for SRM to begin test execution of the Recovery Plan.

During an actual Disaster Recovery or Planned Migration execution, a View volume from a Replay 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 in SRM 5.x.

6.6 Replication Dependencies and Replication Transfer Time

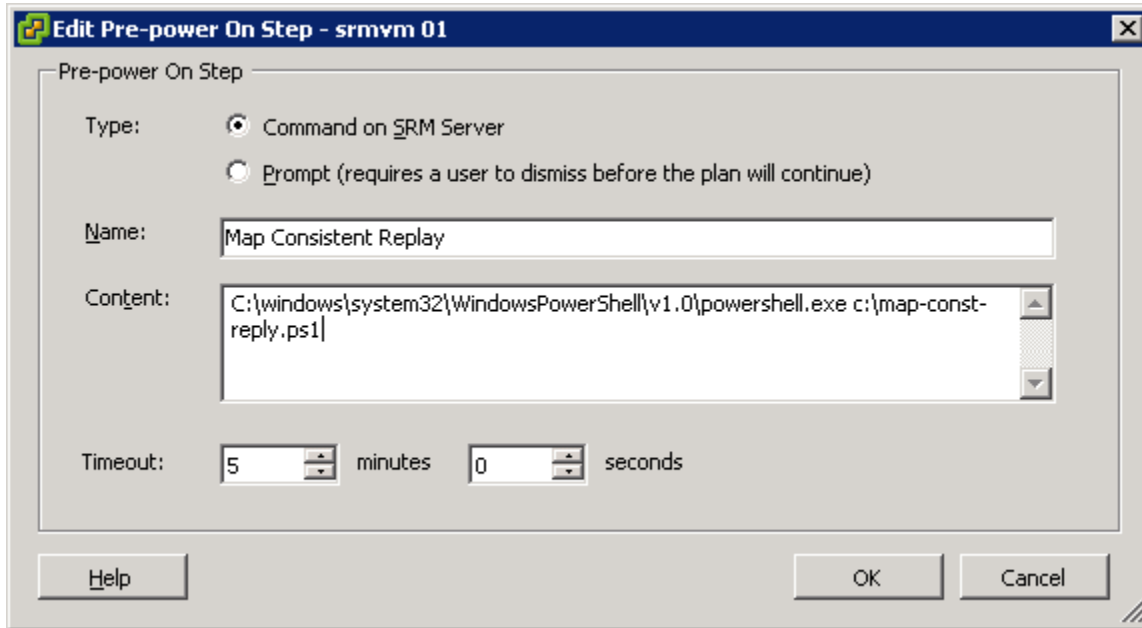
If the application has multiple volumes as part of its data set, it is important to remember that not all volumes may finish replicating within the same timeframe.



Due to various factors such as rate of change, link bandwidth, volume size, and replication QOS, it is possible for multiple volumes in a backup set to finish replicating at different times. This means that if the plan is executed before the replications have completely transferred all of the data (as seen in figure 12), Volume A may have application consistent data, while Volumes B and C will have partially transferred data and be considered crash consistent. Depending on the application, partially transferred volumes will most likely cause problems with the application. If this happens, manual intervention may be required to present the previous set of replays back to the host (for example, the 3:00 pm replay from each volume may need to be used instead of the 4:00 pm).

6.7 Custom Recovery Tasks

If the environment has applications that require custom recovery strategies to avoid any of the situations mentioned above, both Dell Compellent and VMware have a robust set of PowerShell cmdlets in which to customize the recovery steps where needed. The Storage Center cmdlets can control which replay are selected, view volume creation, volume mappings, and even modifying replications. Within the same script, the VMware cmdlets can rescan HBA's, manipulate vDisks, add virtual machines to inventory, and most every other conceivable task required for recovery.

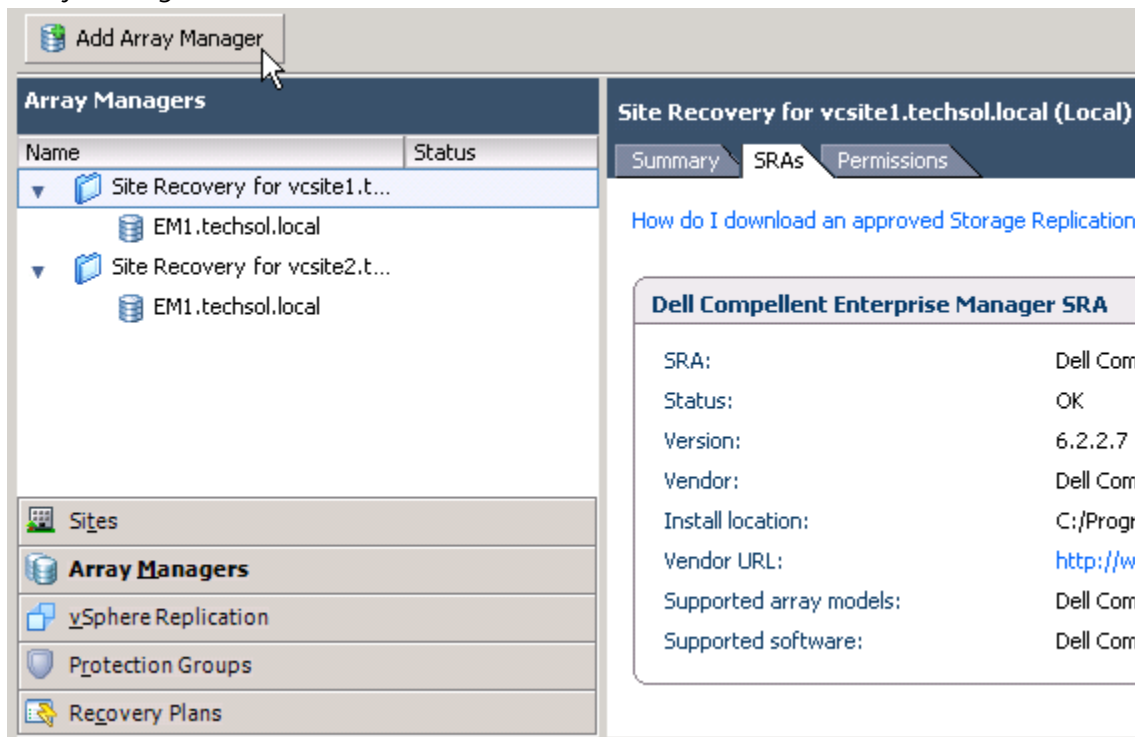


See Appendix A for examples of CompCU and PowerShell scripts that can be used within Site Recovery Manager.

7 Site Recovery Manager Configuration

7.1 Configuring the Array Managers

Configuring the array managers so the Storage Replication Adapter can communicate with the Enterprise Manager Data Collector is performed from the **Array Managers** module. An Array Manager must be added for each site in the unified interface.



The **Protected Site Array Managers** and the **Recovery Site Array Managers** must both be configured so that they can be paired. Depending on the architecture, a single Enterprise Manager Data Collector can be added for both sites, or an Enterprise Manager Data Collector and Remote Data Collector model can be deployed.

- Single Enterprise Manager Data Collector
 - The Protected Site Array Manager and the Recovery Site Array Manager should both specify the Data Collector at the recovery site.
- Multiple Enterprise Manager Data Collectors
 - The Primary Site Array Manager should specify the Data Collector at the primary site, while the Secondary Site Array Manager should specify the Remote Data Collector at the secondary site.

For example, in a single data collector configuration, the **Protection Site Array Manager** should specify the recovery site data collector. Likewise, the **Recovery Site Array Manager**

should specify the same data collector residing within at the same physical site.

Add Array Manager - Site Recovery for tsvc01.techsol.local

Array Manager Information

Specify a display name and an installed SRA for this array manager.

Display Name:

SRA Type:

Additional information about available Storage Replication Adapter (SRA) types and versions is available on the SRAs tab of the array manager folder for each site.

Help < Back Next > Cancel

Add Array Manager - Site Recovery for tsvc01.techsol.local

Dell Compellent Enterprise Manager SRA

Compellent Enterprise Manager

Connection Parameters to the Dell Compellent Enterprise Manager

IP Address: <WebServerPort >:
The IP Address and Web Server API Port of the Dell Compellent Enterprise Manager (default WebServerPort is 3033)

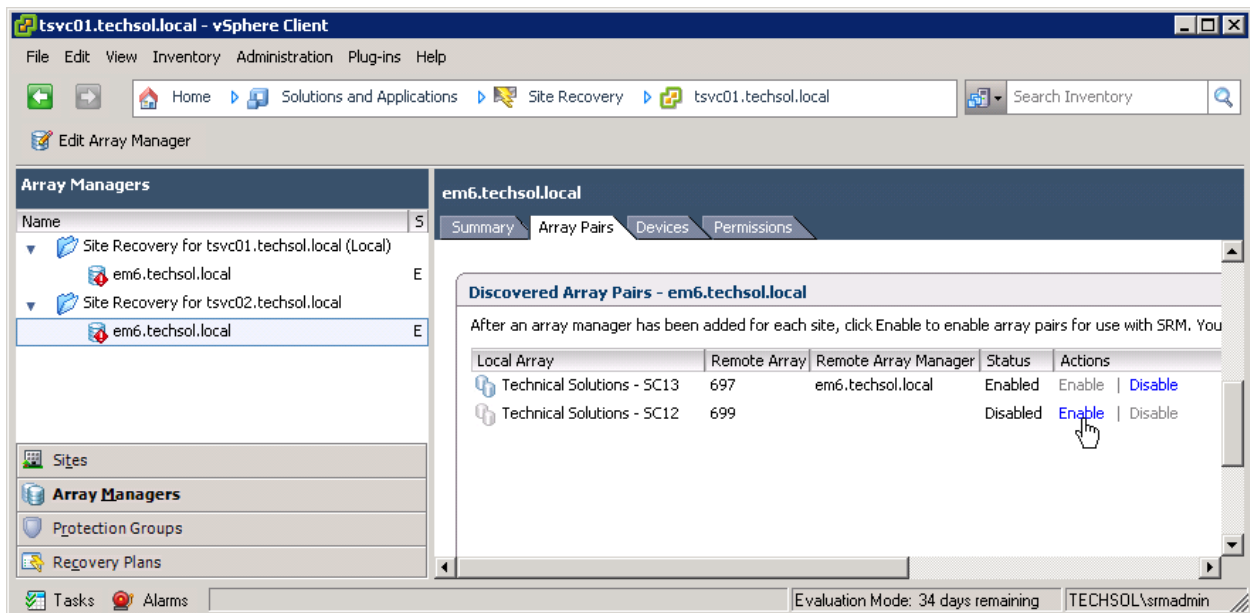
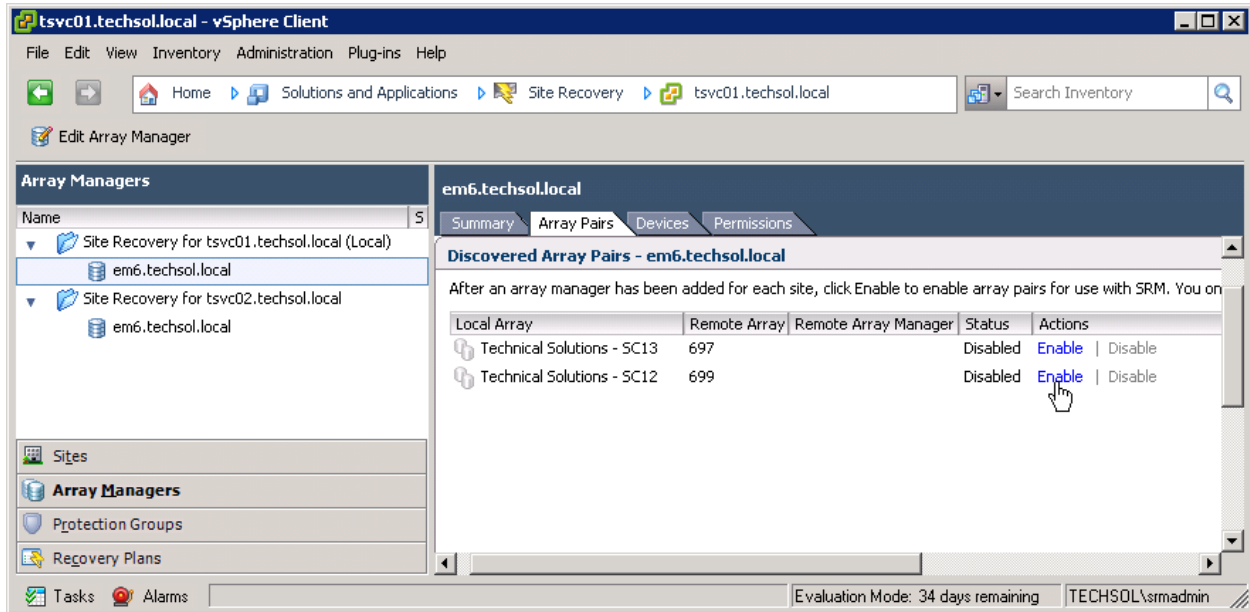
Username:
Enter username for Enterprise Manager

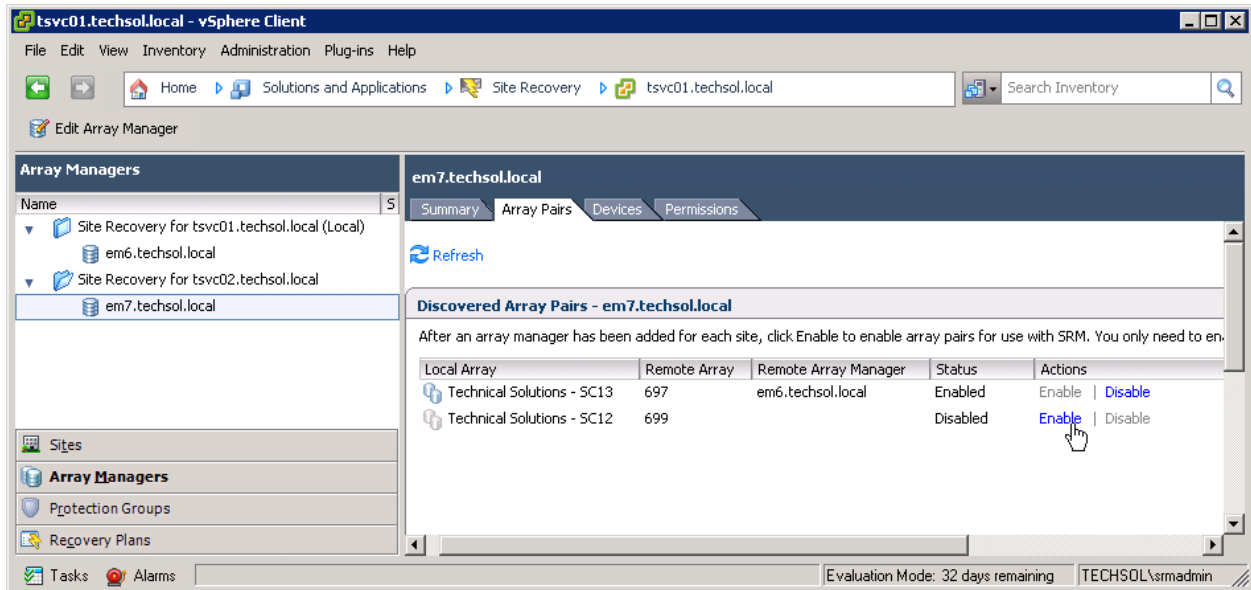
Password:
Enter username's password for Enterprise Manager

Help < Back Next > Cancel

7.2 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 devices. Pairing is an action that is typically only performed after the initial installation of SRM. Once the arrays are paired, they cannot be unpaired while downstream dependencies such as Protection Groups exist.

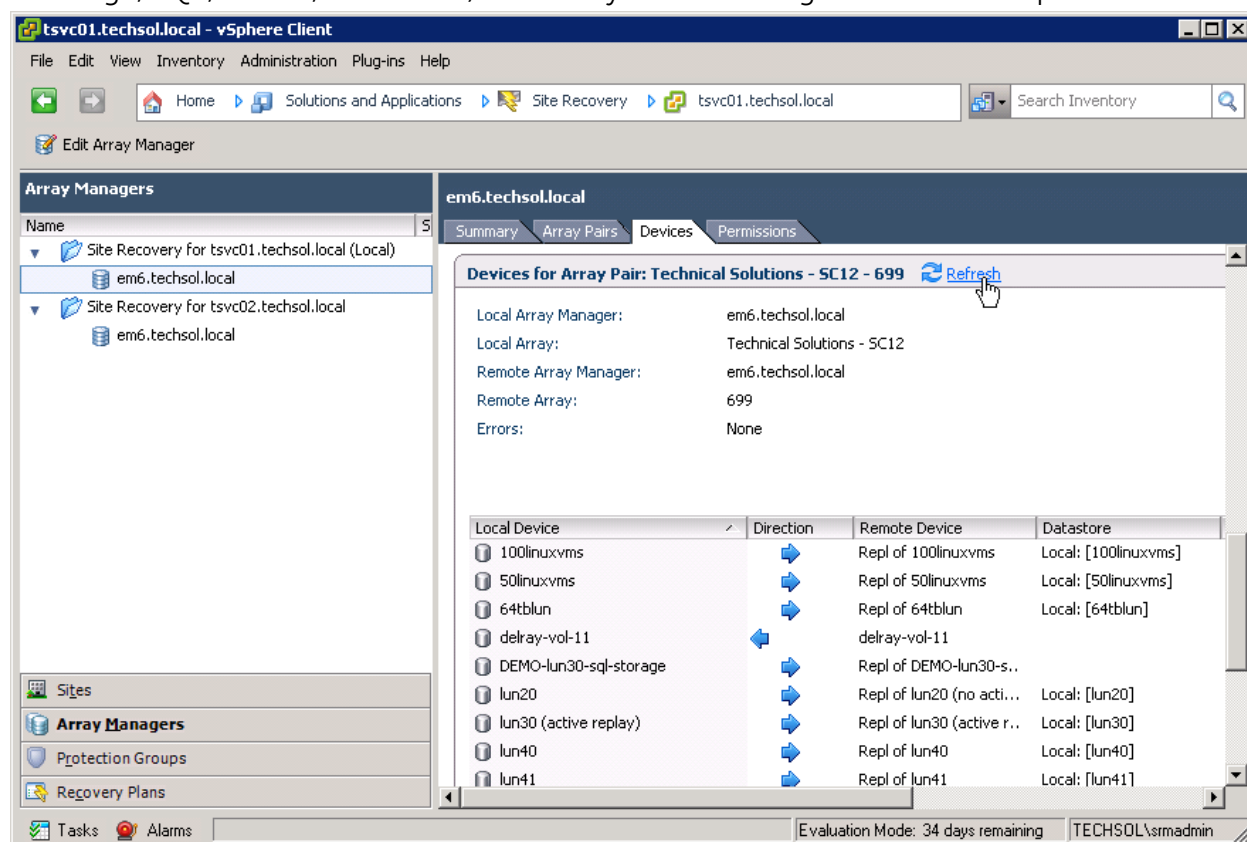




7.3 Rescanning Array Managers

Whenever a new virtual machine datastore or replicated Storage Center volume is added to the environment, the arrays should be rescanned within SRM in addition to rescanning for new LUNs within the ESXi hosts. The Refresh link can be found on the Devices tab in the Array Managers module. Both Array Manager pairs should be refreshed to provide a consistent list of devices. Non-replicated volumes won't be discovered and displayed as devices in SRM. Keep this in mind as a troubleshooting tip if your datastores aren't showing up in SRM. Conversely, all Storage Center replicated volumes will be discovered as devices in SRM, even if they are not for use by vSphere (ie. replicated volumes belonging to physical

Exchange, SQL, Oracle, file servers). This is by VMware design and an SRA requirement.



By clicking on the Refresh link on this screen, the SRA will re-query the Enterprise Manager Data Collector to obtain the new replicated virtual machine volume information.

7.4 Creating Protection Groups

Before creating protection groups, it is recommended that a small VMFS datastore be created at the disaster recovery site to hold the Placeholder VM configuration files. For each virtual machine protected, SRM will create a "Shadow VM" at the opposite site serving as a placeholder for required CPU, 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 and should not be thought of as unreasonable because Storage Center dynamic capacity will thinly provision the volume.

In most cases, only one "Placeholder" datastore per site should be required because the disaster recovery and migration processes will unregister and reregister the recovered virtual machine with the .vmx file on the recovered volume. Also important to note is that the Placeholder volume does not need to be replicated. VMware SRM does not place any data on this volume which cannot be easily regenerated within the UI.


After the Placeholder datastore is created, creating protection groups follows the same general process from previous versions of SRM. Replicated datastore volumes are the foundation which protection groups are built upon. Immediately after the protection group is created, virtual machines residing on the datastore or datastores in that protection group are protected. Legacy rules apply in that once a VM is protected, it is essentially “pinned” to the datastore or datastores its .vmx and .vmdk files reside on. Moving files belonging to a virtual machine is not supported with SRM and will result in the VM no longer being protected or replicated from its original datastore or datastores. From a design and operational standpoint, this means that automated Storage DRS (SDRS) and Storage vMotion cannot be used with SRM protected VMs. Dell Compellent Storage Center arrays offer Dynamic Block Architecture and Data Progression which provides automated sub-LUN tiering for virtual machines without interfering with SRM protection groups.

Do not enable VMware Storage IO Control (SIOC) on datastores protected by Site Recovery Manager. SIOC can prevent datastore unmounts leading to the failure of a Planned Migration. For more information, see Storage IO Control and SRM Planned Migration at <http://blogs.vmware.com/vsphere/2012/06/sioc-and-srm.html> as well as VMware KB articles 2008507, 1037393, and 2004605

7.5 Creating Recovery Plans

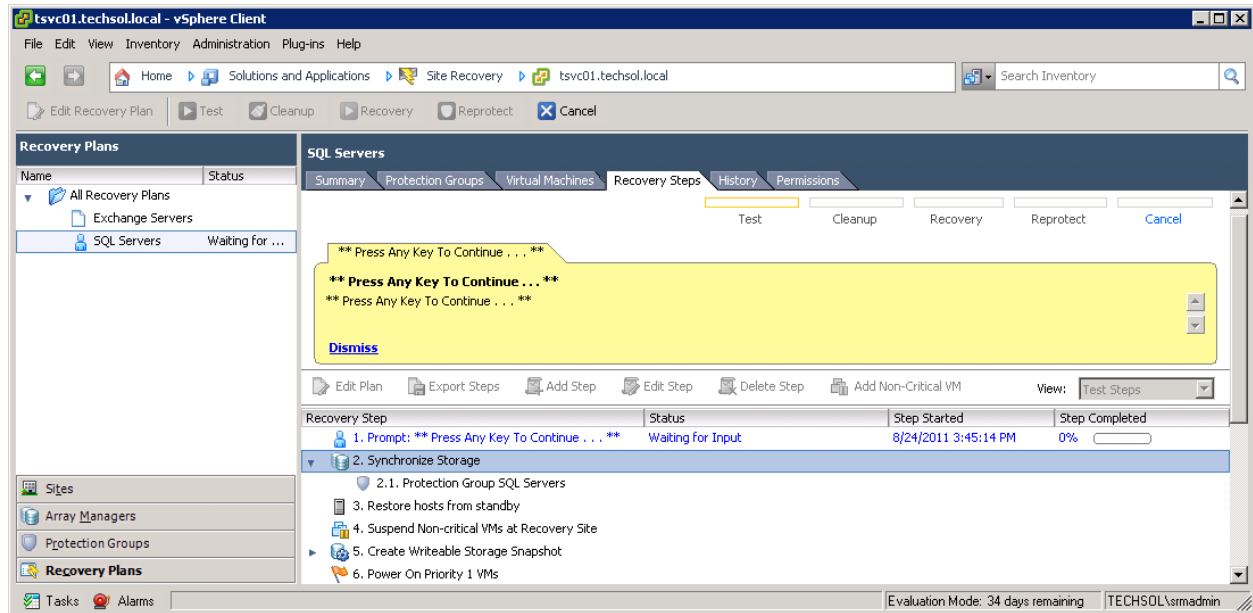
When testing or running recovery plans, SRM has no built-in mechanisms to determine whether or not the replication volumes are fully synced before the storage is prepared for the recovery. In other words, there could still be in-flight data actively being replicated to the secondary site that may influence the outcome of the recovery. This will be particularly true if you configure replication to also replicate the Active Replay.

To help ensure that all data has successfully been replicated to the secondary site, it is recommended that you check the box Replicate recent changes to recovery site when executing a test plan. During an actual disaster recovery cutover, this option may or may not be available. For planned migrations using SRM, this step is required to complete successfully in order to proceed.



In addition, consider adding Prompts and/or SRM server-side Commands to the recovery plan to help ensure all data has been replicated before the subsequent Storage section is executed.

For example, the Compellent CompCU utility or Compellent PowerShell scripts could be integrated into the recovery plan to take current replays of all the volumes to make sure the most recent data has been replicated. An example CompCU script can be found in Appendix A of this document. When the recovery plan executes, it will wait at an added **Pause** step. However, recovery plan execution will not pause at a **Command on SRM Server** step.

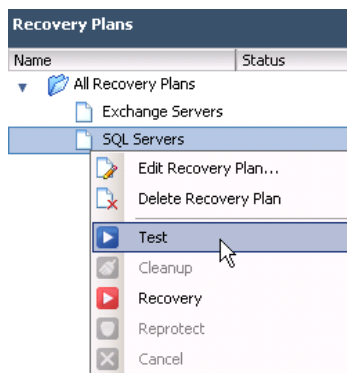


8 Recovery Plan Execution

8.1 Testing a Recovery Plan

Testing the recovery plan is non-disruptive to the storage replications and production volumes and VMs because the test recoveries use Storage Center View Volumes created from Replays to run 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 later 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.

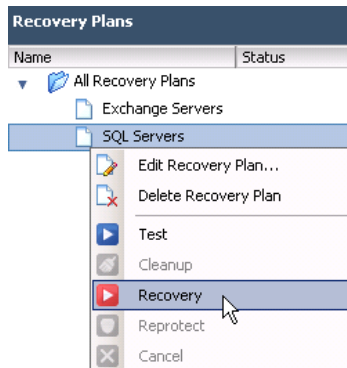
To test a disaster recovery plan, highlight the recovery plan to be tested, right click, and select the Test menu item:



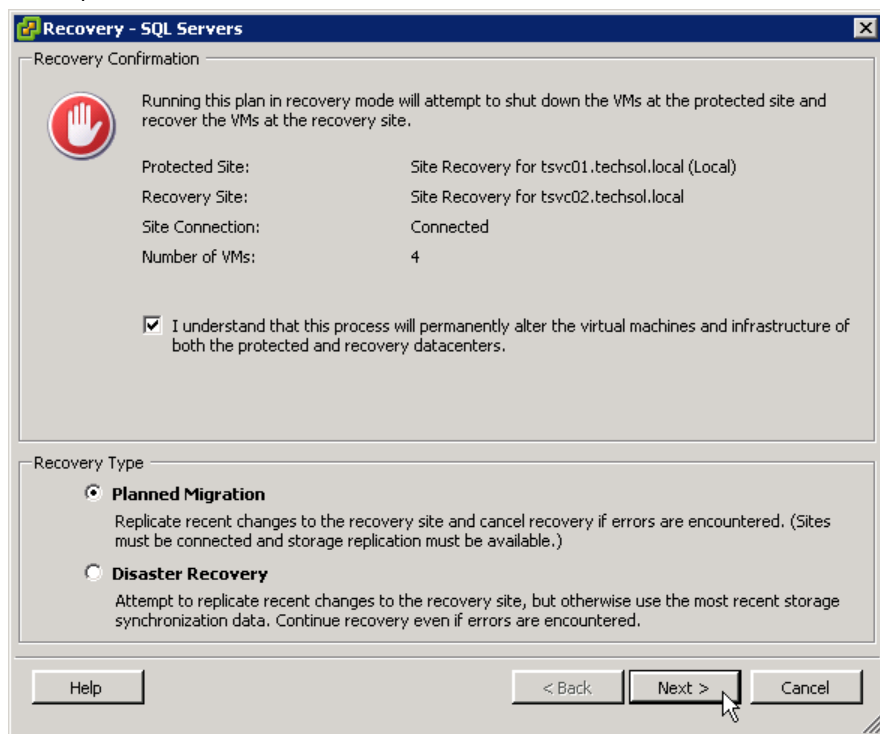
8.2 Running a Recovery Plan

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 will result 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 the requirement to execute a planned migration, highlight the appropriate recovery plan, right click, and choose the **Recovery** option:



As a safety precaution, a warning message will appear and must be acknowledged to execute a live plan.



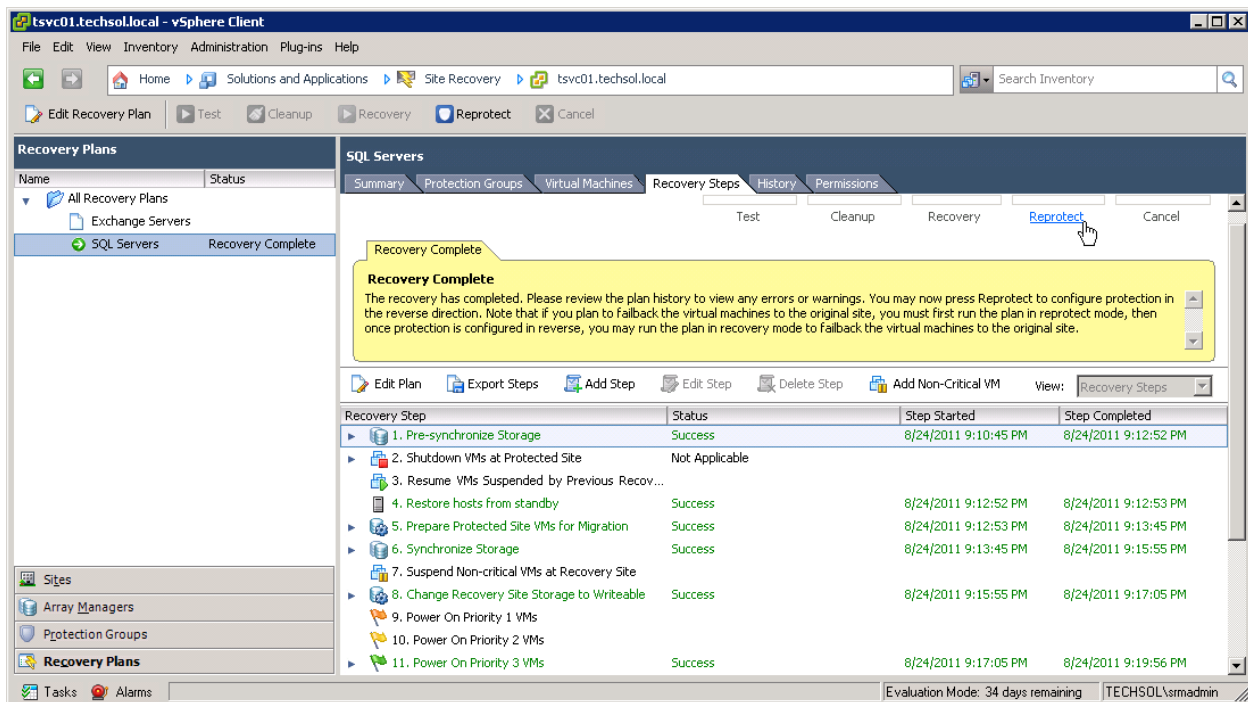
9 Reprotect and Failback

9.1 Overview

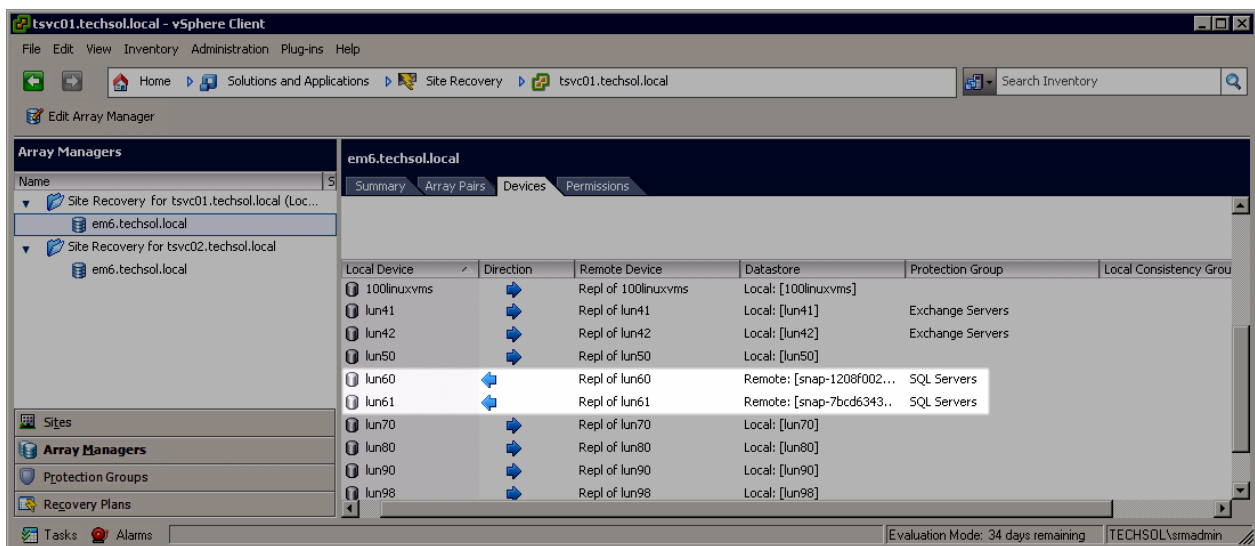
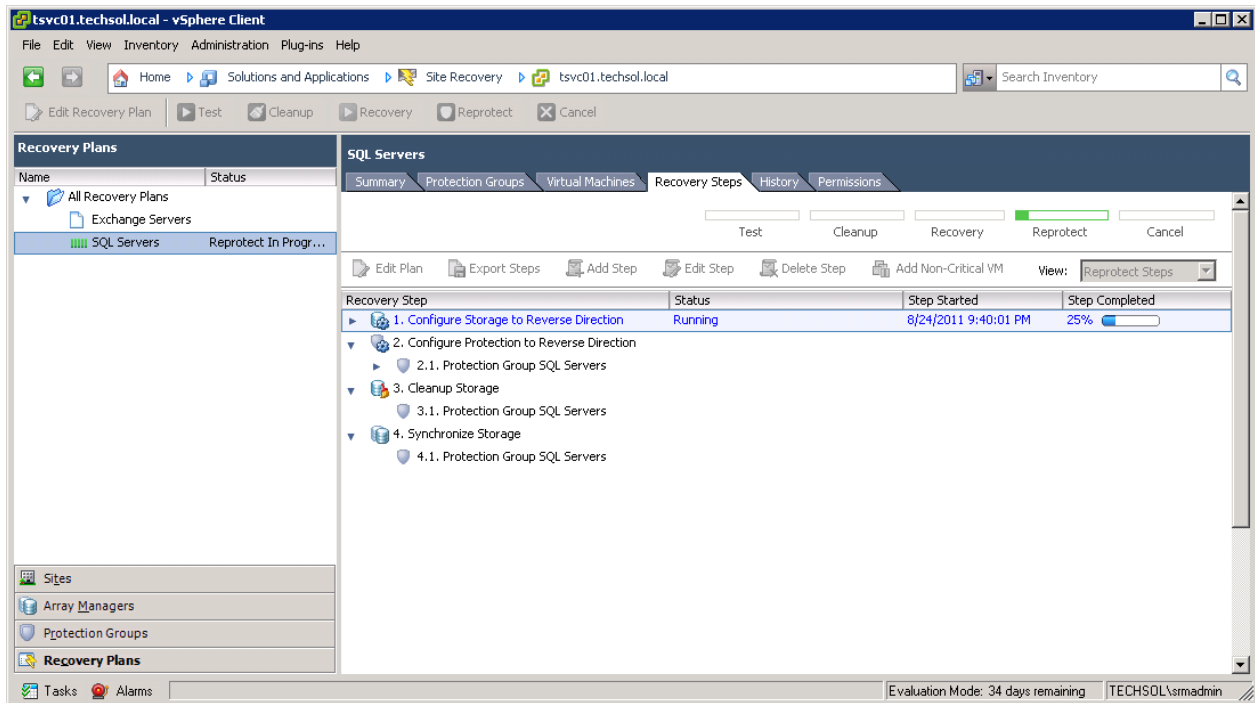
After virtual machines are migrated from one site to another using either the Disaster Recovery or Planned Migration features in SRM, they are in an active running state on the network at the alternate site. However, at this point they are vulnerable to a site failure with no SRM protection. This was true in previous versions of SRM and is true today in SRM 5.x. Previous versions of SRM required a manual re-protection of the virtual machines at the recovery site. SRM 5.x automates the reprotect process and prepares the virtual machines for failback.

9.2 Reprotection

Once protected virtual machines are migrated or disaster recovery failed over to the secondary site, they are unprotected. Immediately following the migration of a protected group, SRM 5.x automates the ability to reprotect the virtual machines. It does this by performing a series of steps.



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, which was originally set up at the primary site is migrated to the secondary site. Placeholder VMs, which were originally set up at the secondary site, are now created at the opposite site (which can now be considered the new recovery site) on its respective placeholder datastore.



9.3 Failback

Failback is no more than an SRM term which describes the ability to perform a subsequent Disaster Recovery or Planned Migration after a recovery and reprotect have already been successfully performed. The benefit that Failback brings in SRM 5.x is the automated ability to move back and forth between sites with minimal effort. This can facilitate a number of use cases including the ability to perform production processing live in the disaster recovery site, resource balancing, and improved disaster recovery infrastructure ROI.

10 Conclusion

VMware vSphere, SRM 5.x and Dell Compellent Storage Center 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 datacenter.

11 Appendix A – Example Scripts

CompCU Script: TakeReplay.cmd

Description: This is an example of a script which can be folded into an SRM Recovery Plan. The script leverages the Compellent Command Utility (CompCU) to take replays of the source replication system volumes to make sure that the most current replay 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 script will connect to a Storage Center with an IP address of "192.168.1.10" with a username of "Admin" and a password of "mmm" to take a replay of "Volume_Name_x" with a replay expiration set to 60 minutes. The ^ symbols are used in this script for line continuation and readability, but could be excluded if the entire command is placed on one line.

The Compellent Command Utility download, and its associated documentation, can be found in the Compellent Knowledge Center. See Appendix B.

PowerShell Script: TakeReplay.ps1

Description: This is an example of the same script from above which can be folded into an SRM Recovery Plan, but written in PowerShell. The script leverages the Compellent's Storage Center Command Set Shell to take replays of the source replication system volumes in an effort to make sure that the most current replay 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 "Replay w/ 1 day retention" -Connection  
$SCConnection
```

This PowerShell script will connect to a Storage Center with a host name of "sc12.techsol.local" with a username of "sradmin" and a password of "mmm" to take a replay of "lun40" with a replay expiration set to 1 day. It is easiest to run this PowerShell script from the Compellent Storage Center Command Set Shell which automatically loads Compellent's Compellent.StorageCenter.PSSnapin snap-in. However, this script can be run from any PowerShell prompt provided the Compellent.StorageCenter.PSSnapin snap-in is manually loaded or loaded as part of the PowerShell profile.

12 Appendix B – Additional Resources

12.1 Compellent Resources

- Compellent Home Page
 - <http://www.compellent.com>
- Compellent Knowledge Center
 - <http://kc.compellent.com>

12.2 VMware Resources

- VMware Home Page
 - <http://www.vmware.com>
- VMware Knowledge Base
 - <http://kb.vmware.com>
- VMware Technology Network
 - <http://communities.vmware.com/>
- VMware Documentation
 - <http://www.vmware.com/support/pubs/>