Technical Report

Deploying VMware vCenter Site Recovery Manager 5 with NetApp FAS/V-Series Storage Systems

Larry Touchette and Julian Cates, NetApp
June 2012 | TR-4064

NetApp Best Practices for SRM5

This document discusses the implementation of VMware® vCenter™ Site Recovery Manager (SRM) Version 5 in an environment using NetApp® FAS storage systems. It provides a conceptual understanding of what is required in a true DR scenario, which is much broader than just failing over the virtual infrastructure and storage environment.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Solution Architecture

1.1 Overview

This document discusses the implementation of VMware vCenter Site Recovery Manager (SRM) Version 5 in an environment using NetApp FAS storage systems. This document intends to provide a conceptual understanding of what is required in a true DR scenario, which is much broader than just failing over the virtual infrastructure and storage environment.

To architect a DR solution, keep the following factors in mind:

- Recovery time objective (RTO): Refers to how quickly a business can recover from a disaster, or specifically, how long it takes to execute the recovery process to make business services available again.
- Recovery point objective (RPO): Refers to how old the recovered data will be once it has been made available, as compared to the time that the disaster occurred.
- Scalability and adaptability in a growing environment.

An ideal solution will have both a low RPO (minutes) and low RTO (minutes to hours). One factor that is often overlooked in a DR solution is the ability to test the DR solution in an efficient manner. In physical environments, DR testing might take many hours or even days and requires that replication between sites be stopped while performing the tests.

1.2 A Traditional Disaster Recovery Scenario

When failing over business operations to recover from a disaster, there are several steps that are manual, lengthy, and complex. Often, custom scripts are written and utilized to simplify some of these processes. However, these processes can affect the real RTO that any DR solution can deliver, and most scripts cannot adapt and update as an environment grows or changes.

Consider the following simplified outline of the flow of a traditional disaster recovery scenario in a virtual environment. Each of these steps might involve several individual tasks.

1. A disaster recovery solution was previously implemented, and replication has been occurring.
2. A disaster occurs that requires a failover to the DR site. This might be a lengthy power outage, which is too long for the business to withstand without failing over, or a more unfortunate disaster, causing the loss of data and/or equipment at the primary site.
3. The DR team takes necessary steps to confirm the disaster and decides to failover business operations to the DR site.
4. Assuming all has been well with the replication of the data, the current state of the DR site and that prior testing had been done to confirm these facts, then:
    a. The replicated storage must be presented to the VMware ESX® hosts at the DR site.
    b. The ESX hosts must be attached to the storage.
    c. The VMs must be added to the inventory of the ESX hosts.
    d. If the DR site is on a network segment different from the primary site, each virtual machine (VM) might need to be reconfigured for the new network.
    e. Make sure that the environment is brought up properly, with certain systems and services being made available in the proper order.
5. After the DR environment is ready, the business can continue in whatever capacity supported by the equipment at the DR site.
6. At some point, the primary site will be made available again, or lost equipment will be replaced.

7. Changes that had been applied to data while the DR site was supporting business will need to be replicated to the primary site. Replication must be reversed to accomplish this.
8. The processes described in step 4 must now be performed again, this time within a controlled outage window, to failover the environment back to the primary site. Depending on how soon after the original disaster event the DR team was able to engage, this process might take nearly as long as recovering from the DR event.
9. After the primary environment is recovered, replication must be established in the original direction from the primary site to the DR site.
10. Testing is done again to make sure that the environment is ready for a future disaster.

As mentioned earlier, a DR process can be lengthy, complex, and prone to human error. These factors carry risk that is amplified by the fact that the process will need to be performed again to recover the operations back to the primary site when it is made available. A DR solution is an important insurance policy for any business. Periodic testing of the DR plan is a must to make sure of its reliability. Due to physical environment limitations and the difficulty of performing DR testing, most environments are only able to do so a few times a year at most, and some not at all in a realistic manner.

1.3 Improved Disaster Recovery with VMware and NetApp

A virtualized environment using VMware vCenter SRM with NetApp storage provides the infrastructure with unique opportunities to implement real working DR processes that are quick and easy to test, consume little additional storage, and significantly reduce the RTO and RPO period.

Site Recovery Manager

One of the most time-consuming parts of a DR failover in a VMware environment is the execution of the steps necessary to inventory, register, reconfigure, and power up VMs at the DR site. VMware has solved these problems with VMware vCenter Site Recovery Manager (SRM). SRM enables separate VMware environments to communicate with each other and negotiate the protection and recovery of VMs. VMware SRM works in conjunction with NetApp SnapMirror® array-based replication technology through the NetApp Storage Replication Adapter (SRA). The NetApp SRA is a small software package that is installed on each VMware SRM server and is available to any customer using VMware SRM.

Figure 1) SRM components.

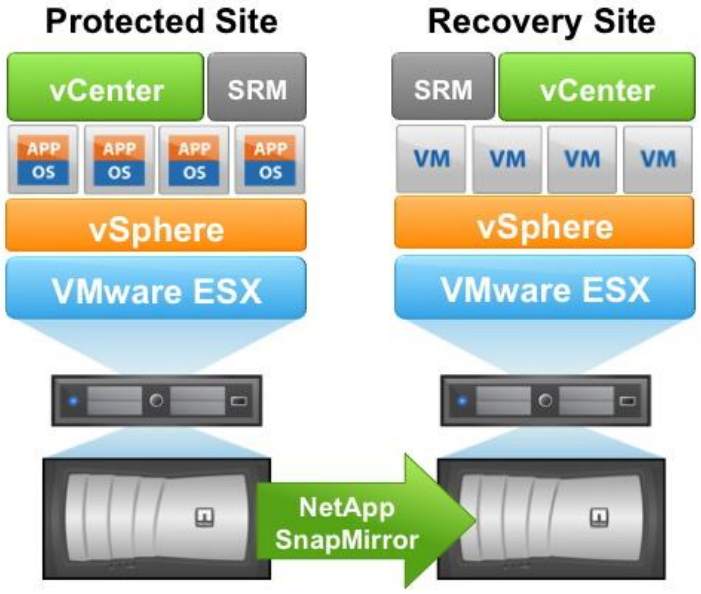

In Version 5 of vCenter Site Recovery Manager, VMware has increased the amount of integration between SRM and the NetApp storage array. SRM 5 introduces the ability to automate the failback process, including the ability to automate the reversal of SnapMirror replication relationships in a new SRM workflow called reprotect. Additionally, SRM now has the capability to trigger SnapMirror relationship synchronization, referred to as a SnapMirror update, during certain workflows. However, keep in mind that SRM does not schedule or manage the periodic update of replication relationships. Replication updates are triggered only during certain SRM operations, and only as directed by the VMware administrator.

Disaster Recovery Test Failover

SRM 5 provides the ability to incorporate storage replication synchronization to the test failover workflow. This allows the VMware administrator to make sure that any changes recently made in the environment, such as the application of patches in the VM guest OS, are replicated to the recovery site and will be present during the test. Storage synchronization is an optional capability in the test failover workflow.

When the VMware administrator runs a test failover operation, SRM will automate the following tasks:

- Optionally trigger the SnapMirror relationships to update the storage at the DR site with any recent changes made at the production site.
- Create NetApp FlexClone® volumes of the FlexVol® volumes on the DR storage appliance.
- Connect the datastores in the FlexClone volumes to the ESX hosts at the DR site.
- Reconfigure the storage settings inside the VMs.
- Connect the VM network adapters to a private test network.
- Optionally reconfigure the VM guest OS network settings as defined for the network at the DR site.
- Power on the VMs in the order defined in the recovery plan.
- Execute any custom commands that have been stored in the recovery plan.

Planned and Unplanned Failover

There are two ways in SRM 5 to perform a real failover. One way is a planned failover, in which proper VM shutdown and storage replication synchronization are incorporated into the process to recover or effectively move the VMs to the recovery site. The planned failover of course requires that the primary protected site is accessible. The other way is an unplanned failover in which the VMs are recovered at the DR site from the last storage replication interval that was able to complete. Depending on the RPO that has been designed into the solution, some amount of data loss can be expected in the unplanned scenario. Both of these types of failover operations are performed in one SRM workflow called failover. Whether the failover is considered planned or unplanned depends on whether or not the VMware administrator has selected the optional storage synchronization capabilities at the time of executing the failover.

When the VMware administrator performs a real failover operation, SRM will automate the following tasks:

- Optionally trigger the SnapMirror replication relationships to update the storage at the DR site with any recent changes made at the production site.
- Shut down the affected VMs at the primary site.
- Trigger storage synchronization in case of a planned failover workflow to make sure the state of the shutdown VMs is replicated to the recovery site.
- Fail over the NetApp SnapMirror relationships.
- Connect the replicated datastores to the ESX hosts at the DR site.
- Optionally power off VMs, such as test/dev instances, at the DR site, freeing compute resources.
- Reconfigure storage settings inside the VMs.

- Connect the VM network adapters to the appropriate recovery site network.
- Optionally reconfigure the VM guest OS network settings as defined for the network at the DR site.
- Power on the VMs in the order defined in the recovery plan.
- Execute any custom commands that have been stored in the recovery plan.

NetApp SnapMirror

SnapMirror provides data replication in an SRM and NetApp environment. Built on NetApp Snapshot™ technology, SnapMirror replication is extremely efficient because it only replicates the 4KB blocks that have been changed or added since the previous update. SnapMirror is easily configured using NetApp OnCommand® System Manager, the Data ONTAP® CLI, or FilerView®.

For cases in which the primary storage is not completely lost, SnapMirror provides an efficient means of resynchronizing the primary and DR sites. SnapMirror can resynchronize the two sites, transferring only changed or new data back to the primary site from the DR site by simply reversing the SnapMirror relationships.

NetApp Deduplication

Additional efficiency is gained when SnapMirror is combined with data deduplication. When deduplication is utilized on the primary storage, only unique data is replicated to the DR site. Additionally, SnapMirror network compression can provide native on wire compression of data sent over the WAN. These technologies result in significant telecommunication and storage capacity savings at the DR site.

Figure 2) NetApp SnapMirror and deduplication.

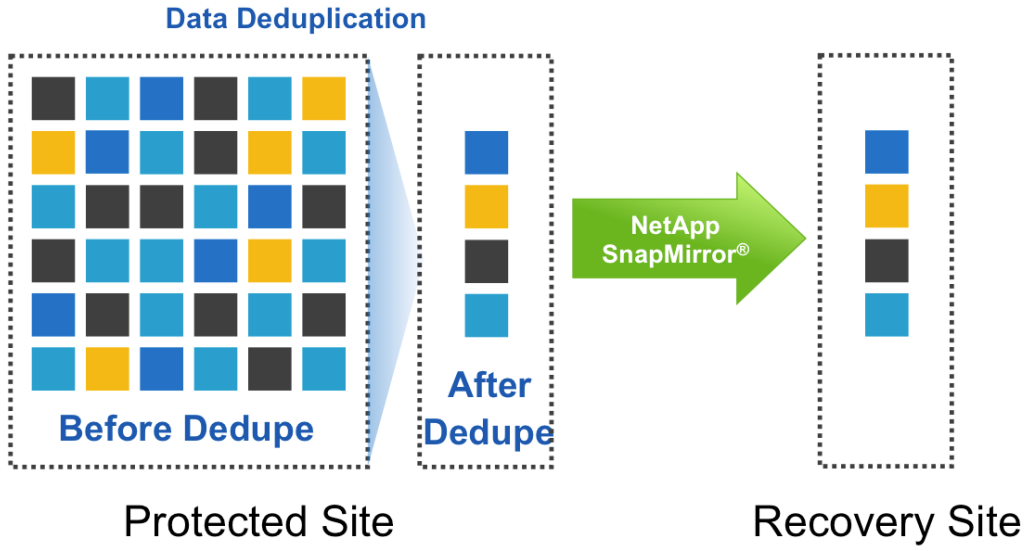

NetApp FlexClone

When FlexClone technology is combined with SnapMirror and SRM, testing the DR solution can be performed quickly and easily, requiring very little additional storage, and without interrupting the replication process. FlexClone quickly creates a read-writable copy of a FlexVol volume. When this functionality is used, an additional copy of the data is not required. For example, for a 10GB LUN, another 10GB LUN isn't required, only the metadata required to define the LUN. FlexClone volumes only store data that is written or changed after the clone was created.

The SRM DR testing component leverages FlexClone functionality to create a copy of the DR data in a matter of seconds, requiring only a small percentage of additional capacity for writes that occur during testing.

Figure 3) Example capacity required for DR testing with FlexClone.

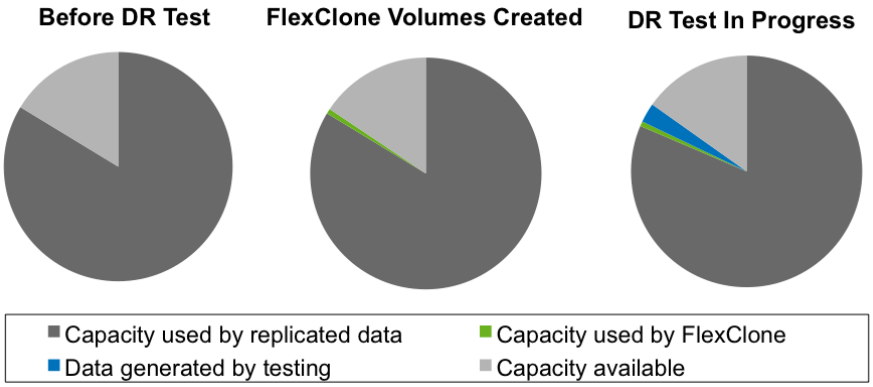

FlexClone volumes share common data blocks with their parent FlexVol volumes but behave as independent volumes. This allows DR testing to be completed without affecting the existing replication processes. Testing of the DR environment can be performed, even for an extended period of time, while replication to the parent FlexVol volume occurs in the background.

NetApp Flash Cache

A disaster recovery scenario is a critical time for any environment. The increase in disk I/O generated by booting many virtual machines can affect the overall performance of the storage array and therefore affect the RTO of the solution. NetApp Flash Cache technology reduces the disk I/O requirements for booting multiple VMs simultaneously, allowing the solution to be deployed with less physical disks, while providing faster boot time and significantly decreasing the recovery time during a DR scenario. Flash Cache is a standard feature with each controller in FAS6240, FAS6280, V6240, and V6280 systems.

Figure 4) NetApp Flash Cache.

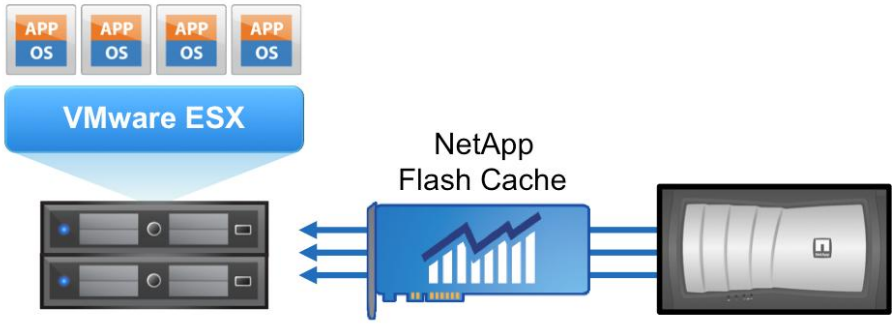

Unified Architecture Flexibility

All NetApp storage systems run the Data ONTAP operating system. This allows for systems with different performance characteristics and different costs to be utilized at the primary and DR sites. For example, depending on the capabilities required, the DR site might contain a lower model storage system, SATA disk vs. FC disk, or the iSCSI protocol vs. FC. A unified architecture, from the low end of the NetApp

storage product line to the high end, also allows for one management and monitoring paradigm to be learned and used by system administrators.

1.4 Architecture

An SRM environment consists of separate vCenter instances. Even though there might be only two instances of vCenter in your environment, SRM does support a shared recovery site model. In the shared recovery site model, multiple vCenter instances can be configured to protect VMs in a single vCenter instance that all the other sites share for recovery resources. Each vCenter instance manages a different set of ESX or ESXi™ hosts. In an SRM environment, the vCenter instance or site in which a VM is currently running is referred to as the protected site for that VM. The site to which the VM's data is replicated is referred to as the recovery site for that VM. When using SRM to manage failover and DR testing, failover and testing occur at the same granularity as the SnapMirror relationship. That is, if you have configured a FlexVol volume as a datastore, all VMs in that datastore will be part of the same SRM protection group and therefore part of the same SRM recovery plan.

A typical SRM environment would consist of the following at each site:

- A number of VMware ESX or ESXi hosts configured in the HA/DRS clusters. Various ESX or ESXi versions are supported, including 3.5U3 (with all patches applied), 4, and 5.
- NetApp FAS or V-Series systems to provide storage for VMFS or NFS datastores.
- VMware vCenter Server.
- Site Recovery Manager Server.
- Microsoft® SQL Server® database.
- Various servers providing infrastructures services such as Active Directory® servers for authentication and DNS servers for name resolution.

Figure 5) Typical SRM environment.

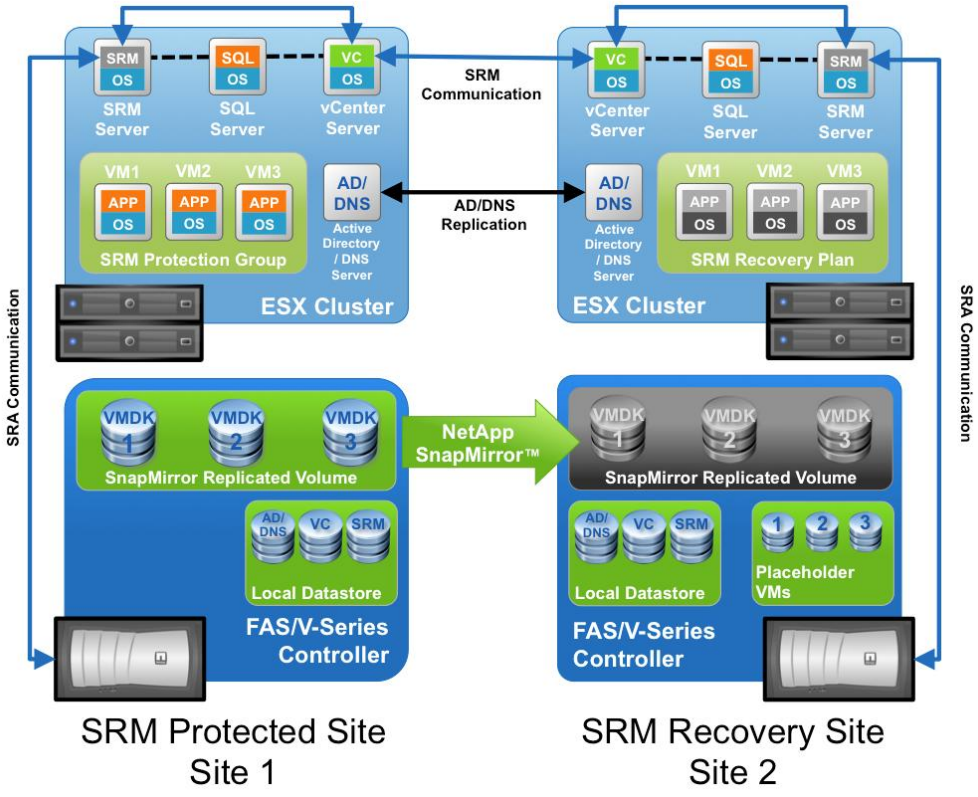

Figure 5 shows VMs that exist at the protected site, site 1, being replicated to the recovery site, site 2. For simplicity, this figure shows replication and protection of VMs going only in one direction from site 1 to site 2. However, replication and protection of VMs can be performed in both directions, with different VMs in different datastores at each site that are configured to be recovered in the opposite site.

In an SRM environment communication does not occur directly between the SRM servers; instead, SRM communication is performed by proxy through the vCenter Server at each site, as shown by the blue arrowed lines. The same is true of communication with the NetApp storage arrays. At no time does the SRM server in site 1 communicate with the FAS/V-Series controller in site 2. If you are working in the SRM interface at site 1 and you are performing some action that requires an operation be performed on the FAS/V-Series controller at site 2, the SRA command is sent by proxy through the vCenter Servers to the SRM server at site 2. The SRM server at site 2 then communicates with the local NetApp controller and sends the response back to the SRM server in site 1, again by proxy back through the vCenter Servers.

It's important that the infrastructure services, such as authentication, name resolution, and VMware licensing, are active and available at both sites.

SnapMirror is used to replicate FlexVol volumes containing VMFS datastores from the primary site to the DR site.

1.5 VM Network Setting Reconfiguration During Failover

Some environments might be able to use the same network IP addresses at both the primary site and the DR site. This is referred to as a stretched VLAN or stretched network setup. Other environments might have a requirement to use different network IP addresses (for example, in different VLANs) at the primary site than what is configured at the DR site. SRM supports both of these scenarios.

SRM 5 has the capability to change the network configuration of a VM as it is recovered. This reconfiguration includes settings such as IP addresses, gateway address, and DNS server settings.

Different network settings, to be applied to individual VMs as they are recovered, can be specified in the properties settings of a VM in the recovery plan. To configure SRM to apply different network setting to multiple VMs without having to edit the properties of each one in the recovery plan, VMware provides a tool called the dr-ip-customizer. Directions for using the utility are provided in section "Customize IP Properties for a Group of Virtual Machines" in Site Recovery Manager Administration Guide. The dr-ip-customizer utility is a tool that takes as input a file containing a comma-separated value (CSV) table of IP settings for multiple VMs and generates a unique customization specification for each VM, and then applies that customization specification to the recovery plan for each VM.

1.6 Using SRM with NFS Storage Connectivity

In NFS environments, the storage network used for connecting to NFS datastores is often a private back-end network that is either physically or logically separated from the VM network. Networks can be separated on different physical switches, on different physical ports on a shared switch infrastructure using VLANs, or on the same ports by using VLAN tagging.

In some cases, multiple private network connections might be configured between an individual ESX host and an individual NetApp controller or vFiler® unit for NFS storage. This practice is sometimes used or recommended to increase the overall available throughput between a single ESX host and a single NetApp controller or vFiler unit. When such a configuration is used, some datastores are accessed on one IP address of a controller, and other datastores are accessed on a different IP on that same controller. This does not increase the throughput to any one particular datastore, but increases the overall throughput between a single ESX host and a single NetApp controller, across multiple datastores. Configuring multiple networks is only done in environments where certain EtherChannel configurations are not possible. For more information about NFS network design for VMware vSphere® on NFS environments, see section 10, "vSphere 5 Storage Networking" in TR-3749: NetApp and VMware

vSphere Storage Best Practices. NetApp recommends using the NetApp Virtual Storage Console (VSC) to provision datastores as described in TR-3749: NetApp Storage Best Practices for VMware vSphere. The VSC will provision NFS datastore mounts using the recommended method of mounting datastores by IP address.

Figure 6 shows an environment where a private storage network has been configured for NFS storage. In this example, the 192.168.10.0 network is used for administration of the NetApp and VMware systems. The 192.168.12.0 and 192.168.13.0 networks are two private networks, used only for NFS storage connectivity.

Figure 6) Private NFS storage network.

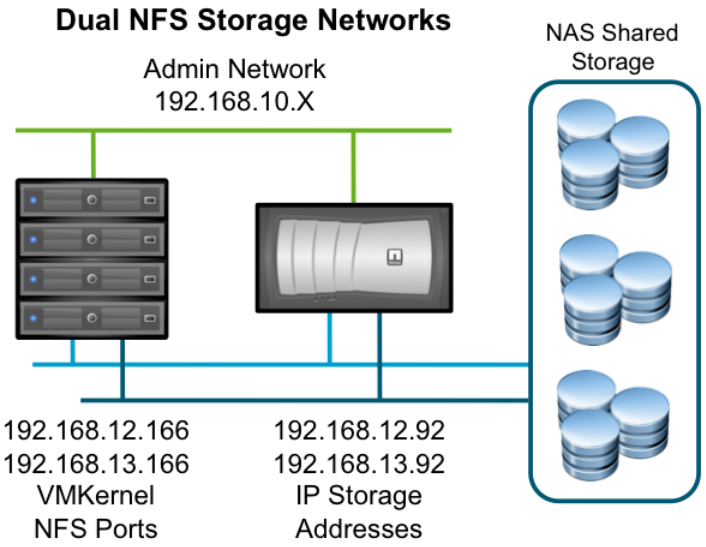

SRM and the NetApp adapter will support connecting to a datastore over a private storage network rather than the network or ports used for system administration, VM access, user access, and so on. The field called NFS IP Addresses, in the NetApp adapter setup screen in the SRM array manager configuration wizard, provides this support. The use of this field is explained in the deployment processes and in the administration guide included in the NetApp adapter download package.

1.7 Considerations for Communication Between ESX Hosts During DR Testing

When a DR test is performed, a private test bubble network is created on the ESX host for the VMs; however, this network is not automatically connected to any physical network adapters and therefore does not provide connectivity between the ESX hosts. To allow communication among VMs running on different ESX hosts during DR testing, a physical private network is created between the ESX hosts at the DR site. To make sure the test network is private, the test bubble network can be separated physically or by using VLANs or VLAN tagging. Segregating this network from the production network is required so that as VMs are recovered, they are not placed on the production network with IP addresses that would conflict with actual production systems. When a recovery plan is created in SRM, you can select the test network you have created to be the private network to which VMs will be connected during the test.

1.8 Considerations for Active Directory and Name Resolution Services

The Active Directory (AD) and name resolution (domain name service or DNS) architecture play an important role in successful failover, DR test, and failback scenarios. The AD servers should not be recovered from replicas created by unsupported processes because this could create an update sequence number (USN) rollback scenario. A USN rollback scenario can occur when an AD server having an older version of the AD database is recovered from an unsupported backup method. This might

cause the AD processes on the recovered server to be unable to process authentication requests or other functions. For information regarding supported methods of creating backups of AD servers in virtual server environments, see Microsoft KB 888794.

Providing Active Directory and DNS Services for DR Testing

Remember that the DR testing network is a private network and that you would use the VMware VM console window to access the VMs running in test mode. To perform the tests inside the VMs, such as verifying that an application server is able to connect to a database, you must be able to authenticate to log in to the application server. It's also important that the application server is able to resolve the network address of the database server.

If authentication and name resolution services are required in the private test network to perform testing, you can create clones of the necessary VMs to use in the test network. For example, if you are using Microsoft AD and DNS services to provide user authentication and name resolution services, VMs providing Microsoft AD and DNS services at the DR site may be cloned prior to running the DR test. But before powering on the cloned VMs, reconfigure the VM network connections to connect the cloned VMs only to the private DR test network. These cloned VMs can then be powered on and will be able to provide name resolution and authentication services in the private test network.

The AD server you choose to clone must be one that was configured as a global catalog server. Some applications and AD functions require FSMO roles in the AD forest. To seize the roles on the cloned AD server, after it has been cloned and connected to the private test network, use the procedure described in Microsoft KB 255504. After the roles are seized, it is very important that the clone never be connected to a production VM network. As these servers cannot replicate their databases with the real servers in the production environment, they should be destroyed after DR testing is complete, and new clones created for later DR tests.

Providing Active Directory and DNS Services for Real Failover

Do not use the cloning process described previously in a real DR failover scenario. In a real failover scenario, you must rely on existing AD and DNS servers at the recovery site to provide those services. Just as in the testing scenario, some applications and AD functions require FSMO roles in the AD forest. If the AD server servicing these roles was lost at the protected site, you must seize the roles on an AD server in the recovery site. To do this, use the procedure described in Microsoft KB 255504. However, the five FSMO roles must be seized per the procedure described in Microsoft KB 255504.

1.9 Software Version Requirements

This section describes software and configuration requirements for deploying VMware vCenter Site Recovery Manager 5 with NetApp FAS/V-Series storage arrays.

Managing failover of SnapMirror relationships with VMware vCenter Site Recovery Manager 5 requires the use of NetApp Storage Replication Adapter (SRA) Version 2. NetApp SRA Version 2 can be obtained from the software download section of the NetApp Support site or on the VMware SRM download page. NetApp SRA Version 2 supports SAN and NAS VMware storage protocols, including support for raw device map (RDM) LUNs.

Table 1) VMware vCenter SRM 5 software version requirements.

Software Version	Comments
VMware vCenter Server 5.x	
VMware vCenter Site Recovery Manager 5.x	

Software Version	Comments
ESX or ESXi	Review the Site Recovery Manager Compatibility Matrix on the VMware SRM documentation site and make sure the environment has supported versions of vCenter Server and ESX/ESXi servers.
NetApp FAS or V-Series Storage Replication Adapter 2.x	
NetApp Data ONTAP	Review the supported versions of Data ONTAP operating system in the Site Recovery Manager Storage Partner Compatibility Matrix in the VMware SRM documentation site to make you are running a supported version of Data ONTAP.

1.10 NetApp Software License Requirements

Table 2 lists the software licenses that are required for supporting NetApp storage with SRM 5.

Table 2) NetApp software license requirements.

NetApp Software License Requirements
SnapMirror license
Storage protocol license—NFS, iSCSI, or FC as used in the solution
FlexClone license—required for nondisruptive DR testing

1.11 Infrastructure Requirements

Table 3 lists the infrastructure services required for a complete SRM solution.

Table 3) SRM solution infrastructure requirements.

SRM Solution Infrastructure Requirements
vCenter Server installed at both sites. VMware SRM requires two independent vCenter environments, each managed by its own vCenter Server.
The vSphere client installed at both sites. In SRM 5, both sites can be managed through one vSphere client interface; however, in the event of a disaster, you must have a vSphere client connection.
Site Recovery Manager server installed at both sites.
NetApp Storage Replication Adapter installed on the SRM servers at the primary and DR sites.
NetApp FlexVol volumes and FlexClone technology.
LUNs or NFS exports configured as datastores and connected at the primary site.

1.12 Supported Replication Technologies

Table 4 lists the data replication technologies that NetApp SRA supports.

Table 4) Supported NetApp replication technologies.

Supported NetApp Replication Technologies
Asynchronous volume SnapMirror

Supported NetApp Replication Technologies
Asynchronous qtree SnapMirror
Synchronous SnapMirror

NetApp SRA does not support the replication technologies listed in Table 5.

Table 5) Unsupported NetApp replication technologies.

Unsupported NetApp Replication Technologies
SnapVault®
SyncMirror® in MetroCluster™

Note: SyncMirror is the aggregate-level synchronous mirroring technology used in a NetApp MetroCluster solution. A MetroCluster system may be used as the source or destination of SnapMirror relationships in an SRM environment. However, SRM does not manage MetroCluster failover or the failover of SyncMirror relationships in MetroCluster.

1.13 SRM 5 and NetApp Snapshot Auto-Delete

NetApp Data ONTAP can be configured to automatically remove Snapshot copies in order to preserve capacity in a FlexVol volume. The default setting for this capability will not automatically delete the Snapshot copies that are created by SnapMirror. If SnapMirror Snapshot copies are deleted, then this will prevent the SRA from being able to reverse and resynchronize replication for the affected volume.

To prevent this from happening, make sure the Snapshot auto-delete capability is configured to "try" which means that Data ONTAP will not delete SnapMirror Snapshot copies.

```
snap autodelete <volname> commitment try
```

1.14 SnapMirror and Data ONTAP Version Requirements

Because SRM 5 has the capability to reverse storage replication relationships, attention must be given to the version requirements for the different modes of SnapMirror. To reverse replication, the source and destination storage systems must be running appropriate versions of Data ONTAP.

For asynchronous volume SnapMirror, the destination system must use a Data ONTAP version that is the same as or later than that of the source system, if the source and destination systems belong to different Data ONTAP release families. Data ONTAP 7.3 and Data ONTAP 8.0 are examples of different release families.

There is no version requirement if the source and destination systems belong to the same Data ONTAP release family. Data ONTAP 7.3 and Data ONTAP 7.3.3 are examples of the same release family. For example, SnapMirror volume replication is possible from a source system using Data ONTAP 7.3.2 to a destination system using Data ONTAP 7.3.1.

For synchronous or semi-synchronous volume SnapMirror, the Data ONTAP version must be the same on the source and destination systems.

There is no version requirement for qtree SnapMirror.

1.15 Supported Storage and Replication Layouts

Supported VM Storage and Replication Layouts

Each storage controller or vFiler unit in a NetApp FAS system is considered an array in SRM. SRM supports certain array-to-array (or controller-to-controller) replication layouts. To determine supportability, keep this rule in mind:

A single virtual machine cannot own data (vmdk or RDM) on more than one SRM array (physical NetApp controller or MultiStore® vFiler unit.)

Best Practice
To be able to protect a VM with SRM and the NetApp SRA, all parts of the VM must exist on only one NetApp controller or MultiStore vFiler unit, in both the protected and recovery site.

Relationships where the data (vmdk or RDM) owned by any individual virtual machine exists on only one array (physical controller or vFiler unit) at each site are supported. The SnapMirror relationship layout scenarios shown here in Figure 7, Figure 8, and Figure 9 are supported. Each virtual machine in the replicated volumes owns data on only one array within each site.

Figure 7) Supported storage layout for SRM.

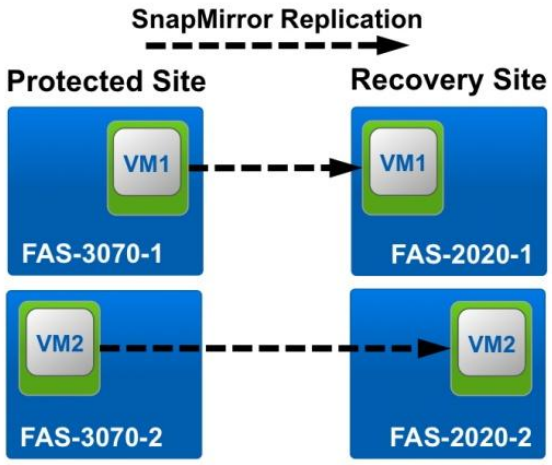

Figure 8) Supported storage layout for SRM.

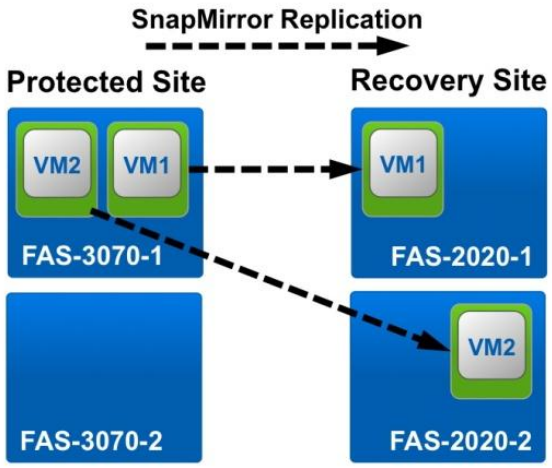

Figure 9) Supported replication layout.

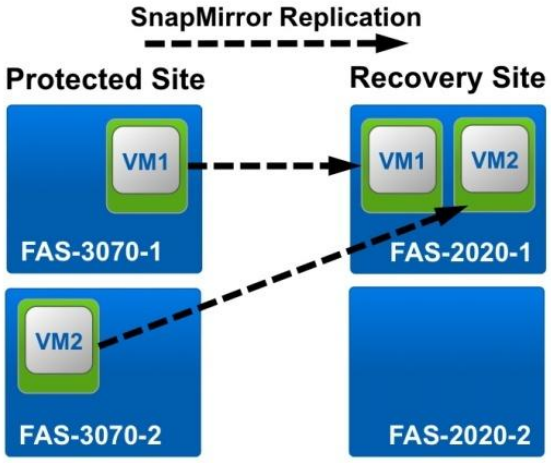

Relationships in which the data (vmdk or RDM) owned by any individual virtual machine exists on multiple arrays (physical controller or vFiler) are not supported. In the examples in Figure 10 and Figure 11, VM5 cannot be configured for protection with SRM because VM5 has data on two arrays.

Figure 10) Unsupported replication layout.

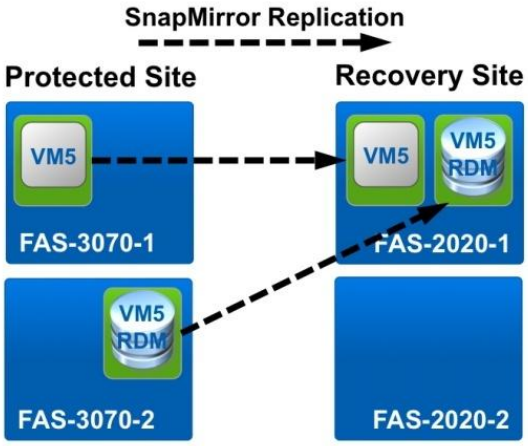

Figure 11) Unsupported replication layout.

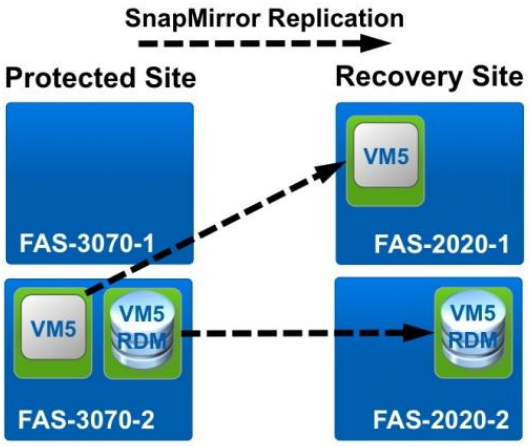

Any replication relationship where an individual NetApp volume or qtree is replicated from one source array to multiple destinations, in the same array or different arrays, is referred to as SnapMirror fan-out and is not supported with SRM. In the example in Figure 12, VM1 cannot be configured for protection in SRM because it is replicated with SnapMirror to two different locations.

Figure 12) Unsupported SnapMirror fan-out replication layout.

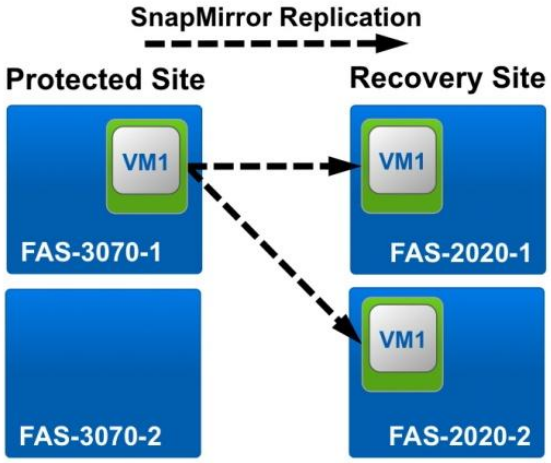

SnapMirror Cascade

SRM 5 does not support cascading of SnapMirror relationships, where a source volume or qtree is replicated to a destination volume, and that destination volume is also replicated by using SnapMirror to another destination volume. In this scenario you cannot use SRM for failover between any sites.

Figure 13) Unsupported SnapMirror cascade replication layout.

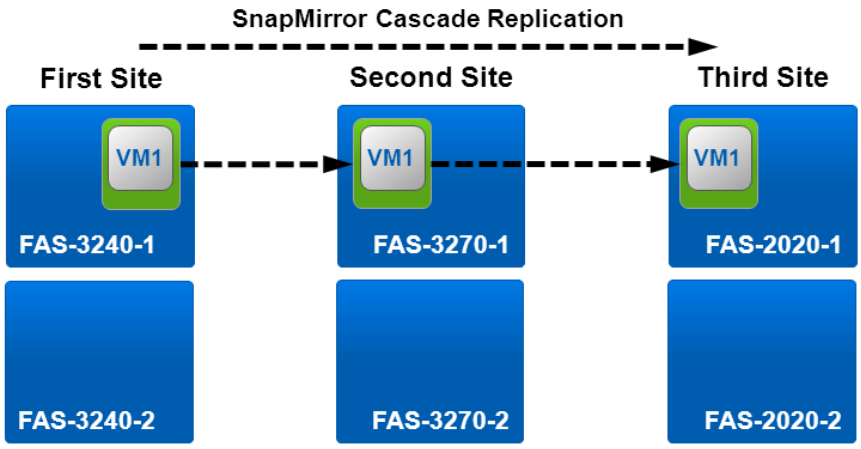

SnapMirror and SnapVault

NetApp SnapVault software enables disk-based backup of enterprise data between NetApp storage systems. SnapVault and SnapMirror can coexist in the same environment; however, SRM supports the failover of only the SnapMirror relationships.

In an environment where SnapVault is used, specific named Snapshot copies are created on the primary storage system. Depending on the configuration implemented, the named Snapshot copies might be created on the primary by a SnapVault schedule or by an application such as NetApp OnCommand™ Unified Manager. The named Snapshot copies created on the primary are then replicated to the SnapMirror destination, and from there they are vaulted to the SnapVault destination.

A source volume or qtree can be in cascade configuration, in which a volume is replicated to a SnapMirror destination in the DR site and from there vaulted to a SnapVault destination, or in a fan-out relationship where one destination is a SnapMirror destination and the other destination is a SnapVault destination. However, SRA does not automatically reconfigure the SnapVault relationship to use the SnapMirror destination volume as the source for the vault when SRM failover or replication reversal occurs.

Best Practice

If using SnapVault and SRM in the same environment, NetApp recommends using a SnapMirror to SnapVault cascade configuration in which SnapVault backups are normally performed from the SnapMirror destination at the DR site. In the event of a disaster that makes the primary site inaccessible, keeping the SnapVault destination at the recovery site allows SnapVault backups to be reconfigured after failover such that SnapVault backups can continue while operating at the recovery site.

In a VMware environment, each datastore has a unique identifier ID (UUID) and each virtual machine has a unique managed object ID (MOID). These IDs are not maintained by SRM during failover or failback. Because datastore UUIDs and virtual machine MOIDs are not maintained during failover by SRM, any applications that depend on these IDs, such as NetApp OnCommand Unified Manager that coordinates SnapVault replication with the vSphere environment, must be reconfigured after SRM failover.

Figure 14) SnapMirror to SnapVault cascade.

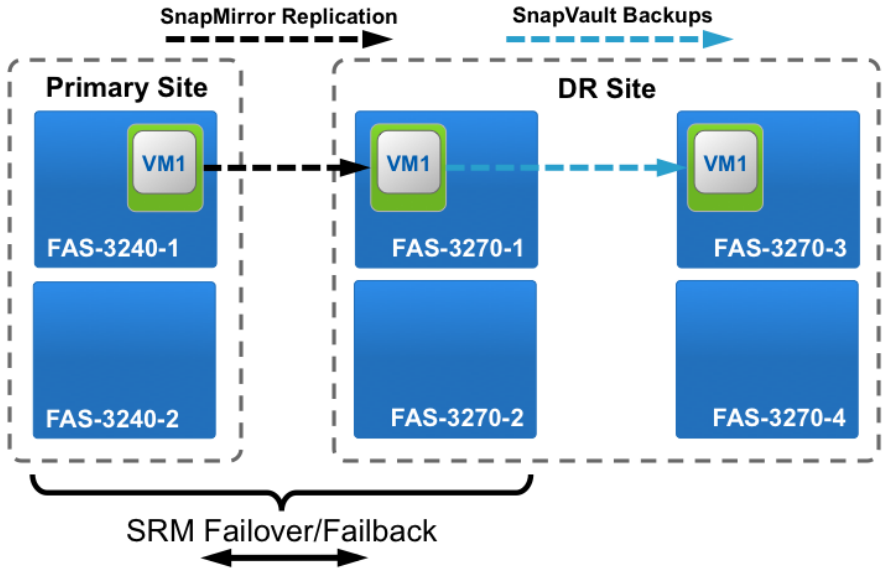

In the configuration shown In a VMware environment, each datastore has a unique identifier ID (UUID) and each virtual machine has a unique managed object ID (MOID). These IDs are not maintained by SRM during failover or failback. Because datastore UUIDs and virtual machine MOIDs are not maintained during failover by SRM, any applications that depend on these IDs, such as NetApp OnCommand Unified Manager that coordinates SnapVault replication with the vSphere environment, must be reconfigured after SRM failover.

Figure 14, a SnapMirror to SnapVault cascade configuration has been implemented. If the SnapVault destination is in the DR site or in a tertiary site, not affected by an outage in the primary site, the environment can be reconfigured to allow backups to continue after failover.

After SRM has been used to reverse SnapMirror replication back to the primary site and the environment has been reconfigured such that SnapVault backups are occurring from what is now the SnapMirror source, the environment will be in a configuration like that shown in Figure 15. This is a SnapMirror SnapVault fan-out configuration.

Figure 15) SnapMirror SnapVault fan-out.

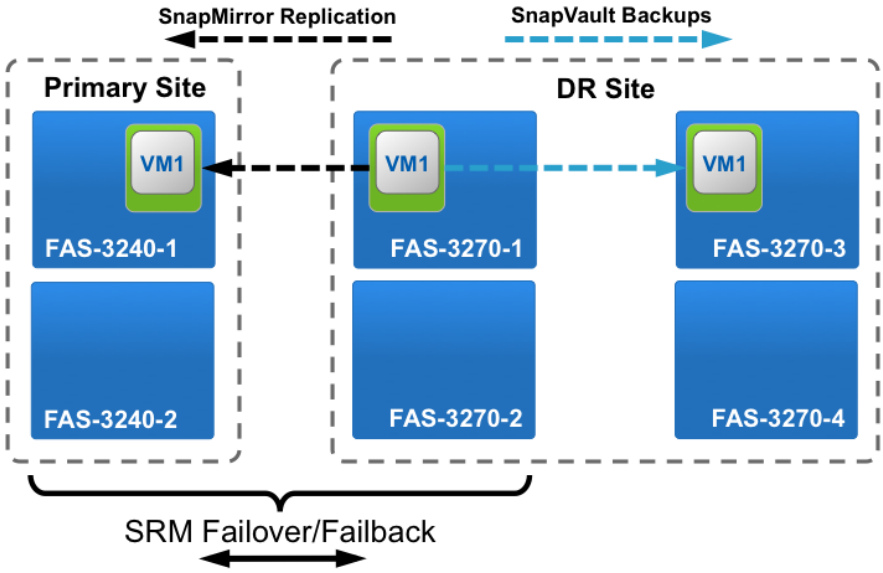

After the SRM failback and a second reversal of the SnapMirror relationships is performed by SRM, the production data, now back at the primary site, will be protected through SnapMirror and SnapVault backups as it was before the failover to the DR site.

Using Qtrees in SRM Environments

NetApp storage systems provide the ability to create a folder called a qtree in the root of a FlexVol volume. A qtree is a means of providing file system quotas for NAS protocol user access and is also the primary storage unit used to define SnapVault backups. Volumes can also be vaulted into qtrees by using SnapVault. SnapMirror can replicate data at the volume level or at the qtree level. Typically, VMware datastores are stored in the root of a volume.

Consider the following caveats when using qtrees in an environment where SRM is used:

- If you have configured qtrees as NFS datastores, you must create an NFS export for each qtree in order for SRM to be able to discover the NFS datastore. If you export only the volume that contains the qtrees, so that there is only one export line for the volume in the /etc/exports file, SRM will not be able to discover the qtree NFS datastores.
- If you are using qtrees in an SRM environment and if you configure multiple qtrees in the same volume, mounting each qtree separately as a different NFS datastore, NetApp recommends using qtree-level SnapMirror with SRM.
- If you are using qtrees in an SRM environment, and if you configure multiple qtrees in the same volume with each qtree containing a LUN that is a different VMFS datastore, NetApp recommends using qtree-level SnapMirror with SRM.

Best Practice

If you are using qtrees, with multiple qtrees in one volume as separate datastores in the VMware environment, the level at which you mirror data should match the level at which you configured the datastores. Meaning, if you are using different qtrees in the same volume to store different datastores, use qtree-level SnapMirror for each qtree. If you are storing each datastore in its own volume, use volume-level SnapMirror for each volume.

- SRM 5 has the capability to failover a datastore from the protected site to the recovery site, reverse that replication, and fail back to the primary site. If you've configured multiple qtrees as individual datastores in the same volume and are using volume-level SnapMirror, SRM will not be able to reverse replication for one of the qtrees in that volume without reversing replication for all the qtrees in that volume.

Configuration Requirement

When multiple qtrees in the same volume are configured containing different datastores, you must make sure these datastores are all included in the same SRM recovery plan so that they are failed over and failed back together.

Volume-Level NFS Datastores or Qtree-Level NFS Datastores

When using NFS, a path on the NFS server is shared to allow NFS clients to mount that path. This is referred to as an NFS export or exported path. The NFS protocol allows mounting a subdirectory under the export path, as well as the export path itself. For example, if you have exported the `/vol/vol1` path, where the `vol1` volume contains multiple qtrees, the NFS protocol allows an NFS client to mount each qtree path directly, such as `/vol/vol1/qtree_name`, without the need for each individual qtree to be exported.

The preceding practice can be problematic in an SRM 5 environment as this means that one set of export security rules, in this case, the security rules for the `/vol/vol1` export, apply to all the datastores in the `vol1` volume. In some cases, SRM failover requires that one datastore be failed over with one set of security rules applied to that datastore export at the recovery site, and a different datastore be failed over with a different set of security rules. This cannot be supported if multiple qtrees are accessed through one shared export path.

Additionally, if the path that is exported does not match the path that is replicated or the datastore path, it can interfere with the configuration of SRM. For example, suppose the `/vol/vol1` path is exported for NFS, but the datastores actually mounted by the ESX or ESXi hosts use the `/vol/vol1/qtree_name` path. This mismatch between the path that is exported and the path to the datastore can interfere with the ability of the SRA to discover the replicated storage devices.

Best Practice

When multiple qtrees in the same volume are configured as different NFS datastores, create a separate NFS export for each qtree path, rather than creating one export for the volume containing the qtrees.

Mixed FC and iSCSI Environments

VMware SRM and NetApp SRA support the use of the FC protocol at one site and the iSCSI protocol at the other. It does not support having a mix of FC-attached datastores and iSCSI-attached datastores in the same ESX or ESXi host or in different hosts in the same cluster.

NetApp SRA does not support a mixed Asymmetric Logical Unit Addressing (ALUA) configuration. A mixed ALUA configuration is one where a single ESX or ESXi host, or multiple ESX hosts in the same cluster, has some initiators configured in ALUA enabled igroups and the same ESX host or hosts have other initiators configured in ALUA disabled igroups.

An example of an unsupported ALUA configuration on a single ESX host is one where some initiators are used in an ALUA enabled igroup for VMFS LUNs, and different initiators are used in the ALUA disabled igroup for RDM LUNs. This type of configuration is sometimes used to support Microsoft Cluster Services (MSCS) in a VM. This configuration is not supported with SRM because during the SRM failover or test failover, SRM will include all FC and iSCSI initiators in the ESX hosts in the request. A NetApp storage system with Data ONTAP operating in 7-Mode does not support ALUA enabled igroups with iSCSI. It also

does not support presenting the LUNs in FC igroups with ALUA enabled and iSCSI igroups with ALUA disabled at the same time.

An example of an unsupported ALUA configuration in an ESX or ESXi cluster is one where the cluster contains ESX 3.5 hosts that require ALUA disabled and ESX 4 or ESX 5 hosts that have ALUA enabled. As SRM would include all the initiators from all the hosts in one failover/test failover request, the configuration cannot be supported.

Best Practice
SRM and NetApp SRA support having mixed FC and iSCSI protocols between the protected and recovery site. However, each site should be configured with only one protocol, either FC or iSCSI, not both protocols in the same site. If there is a requirement to have both FC and iSCSI protocols configured in the same site, NetApp recommends that some hosts use iSCSI, and other hosts use FC, and that SRM resource mappings be configured such that VMs are configured to failover into one group of hosts or the other.

1.16 MultiStore vFiler Requirements

NetApp MultiStore enables secure logical partitioning of a single NetApp storage controller into virtual controllers called vFiler units. vFiler has the effect of making a single physical storage controller appear to be many logical controllers. vFiler supports IP storage environments including iSCSI VMFS datastores, iSCSI RDM devices, and NFS datastores.

If you are using MultiStore vFiler units as SRM storage arrays, the following are requirements for vFiler support with the NetApp SRA.

Table 6) Requirements for MultiStore vFiler units as SRM storage arrays.

MultiStore vFiler Requirements
Each vFiler unit at the primary site must be added to SRM separately as an array.
Each vFiler unit at the DR site must be added to SRM separately as an array.
The source and destination vFiler units must both be online. This means that SRM cannot be used to manage failover where the MultiStore vFiler DR capability is configured. When a vFiler unit is configured for vFiler DR capability, the destination vFiler unit is in an offline state; this is not supported for an SRM array.
SnapMirror relationships must be defined within the vFiler context, not on the physical hosting array referred to as vFiler0. If you are using NetApp OnCommand Unified Manager (formerly Protection Manager) to create SnapMirror relationships for vFiler units, these relationships cannot be used with SRM, because these relationships are not created in the vFiler context.
In the physical controller, prior to executing a DR test, set the `vfiler.vol_clone_zapi_allow` option to `on`.
Only iSCSI or NFS connected datastores are supported in vFiler, as vFiler does not allow the FC protocol.
In iSCSI environments, the ESX hosts at the DR site must have established iSCSI sessions to the DR site vFiler prior to executing SRM DR test or failover. The NetApp adapter does not start the iSCSI service in the vFiler unit or establish an iSCSI connection to the iSCSI target. You must configure the iSCSI initiator on the ESX hosts at the DR site to establish the iSCSI session.
In an NFS environment, VMkernel port network connectivity to the DR vFiler NFS IP addresses must already exist prior to executing SRM DR test or failover. Make sure that network communication is possible between the VMkernel ports on the ESX hosts and the network interfaces on the NetApp controller, which provide NFS access. For example, make sure that the appropriate VLANs have been configured if using VLAN tagging.

MultiStore vFiler Requirements
A single VM must not have data on multiple vFiler units. The storage layout rules, described in section 1.15, “Supported Storage and Replication Layouts,” apply equally to both physical NetApp arrays and vFiler arrays.

2 Deployment Procedures

2.1 Overview

Configuring SRM 5 to protect virtual machines (VMs) replicated by SnapMirror involves the following steps.

1. Reviewing the configuration requirements.
2. Performing the verification steps to verify that the configuration of the environment is ready to support SRM.
3. Connecting the vCenter sites in the SRM interface.
4. Configuring inventory mappings. This aligns resources such as hosts, cluster, and networks at the protected site with the resources that should be used for recovery in the recovery site.
5. Installing the NetApp FAS/V-Series Storage Replication Adapter (SRA) on each of the SRM servers.
6. Configuring the SRA at each site. This enables the SRM software to communicate with the NetApp storage appliances for issuing SnapMirror failover commands, mapping LUNs to igroups, exporting NFS datastores, and so on.
7. Building SRM protection groups. Protection groups define groups of VMs that will be recovered together.
8. Building SRM recovery plans. Recovery plans identify the startup priority of VMs, timeout settings for waiting for recovered VMs to respond, additional custom commands to execute, and so on.
9. Testing the recovery plan by executing an SRM test failover.

2.2 Verifying Configuration Requirements

There are configuration rules and requirements that must be met for a successful SRM implementation. Review the following requirements to make sure the environment is properly set up prior to configuring the SRA and attempting to perform an SRM test failover.

SAN Environment Verification

The following steps need to be performed only in environments where FC or iSCSI is used.

1. At the protected (source) site, verify that the source FC or iSCSI LUNs are configured in igroups that have the igroup type of “vmware." When using RDM disks with VMware, the igroup type must be set to “vmware”; however, the LUN type may be set to whatever OS type the guest VM requires.

   Note: NetApp SRA for SRM 5 will automatically create igroups as needed at the recovery site, during failover and test failover workflows.

2. At the recovery site (destination), verify that proper storage connectivity exists between the NetApp storage array and the ESX or ESXi hosts. If using MultiStore vFiler units as storage arrays, be sure to verify that storage network connectivity has been configured between the ESX or ESXi hosts and the vFiler unit.

   a. You may do this by checking to see if there is an existing datastore connected to the ESX or ESXi hosts at the recovery site that is already using the storage network. The existence of such a datastore indicates storage network connectivity.

b. If there are no existing datastores connected on the recovery array, you may verify storage network connectivity by issuing the `fc initiator show` or `iscsi initiator show` command on the array to verify connectivity exists between the array and the hosts.

Note: At the recovery site, if using iSCSI, you must configure the ESX or ESXi iSCSI initiators to connect them to the recovery site array. SRM does not configure the connection from an ESX or ESXi iSCSI initiator to the NetApp iSCSI target array.

NAS Environment Verification

The following steps need to be performed only in environments where NFS is used.

1. At the protected site, verify that the NFS datastores are listed in the /etc/exports file and are configured with values in the `rw` (read/write) security field.
   - Example of a `/etc/export` line that can be discovered by SRM:

```
/vol/srm5 -rw=192.168.2.0/24,root=192.168.2.0/24
```

   - Example of a `/etc/export` line that cannot be discovered by SRM:

```
/vol/srm5 -rw,anon=0
```

   Note: Exports made from the CLI, without the `-p` option that adds them to the exports file, will not be discovered by SRM.

   Note: Exports using the default setting of `rw` to all hosts, the `rw` option without any hosts following it, will not be discovered by SRM.

2. At the recovery site (destination), verify that proper storage connectivity exists between the NetApp storage array and the ESX or ESXi hosts. If using MultiStore vFiler units as storage arrays, be sure to verify storage network connectivity has been configured between the ESX or ESXi hosts and the vFiler unit.
   a. You may do this by checking to see if there is an existing datastore connected to the ESX or ESXi hosts at the recovery site that is already using the storage network; the existence of such a datastore indicates storage network connectivity.
   b. If there are no existing datastores connected on the recovery array, you may verify storage network connectivity by issuing the `ping` command on the NetApp array, to ping the vmkernel ports on the ESX or ESXi hosts. Alternatively, you can issue `the vmkping` command on the ESX or ESXi hosts to ping the NetApp controller ports used to serve NFS.

   Note: The `vmkping` command makes the ESX or ESXi host send the ping out a vmkernel interface rather than a management interface.

Replication Verification

The following steps are to be performed in both SAN and NAS environments.

1. Verify all the datastores containing the VMs to be protected with SRM are replicated to the recovery site using SnapMirror. You may do this using NetApp OnCommand System Manager, or CLI.

   Note: SRM does not perform scheduled SnapMirror updates or baseline transfers. Periodic SnapMirror transfers must be managed and scheduled by using NetApp software such as the built-in scheduler in Data ONTAP, OnCommand System Manager, or Protection Manager.

2. NetApp SRA does not support fan-out replication using SnapMirror of a datastore to multiple different destinations. This includes cascaded SnapMirror to SnapMirror relationships. You can use the `snapmirror destinations` command on the source system to display all the destinations for a source. Each source should have only one destination.

   Note: The NetApp SRA ignores SnapVault relationships; it is ok for a source to be replicated with SnapMirror and with SnapVault. However, SnapVault relationships are not reconfigured as SnapMirror relationships are failed over and reversed with SRM.

3. Verify the datastores you want to recover with SRM contain VMs. Datastores must contain VMs or virtual disks owned by VMs to be discovered by SRM. If you want to recover an empty datastore with SRM, you must configure at least one VM in that datastore. This does not need to be a complete VM, and it does not have to be configured with a virtual disk; simply create an empty VM in the datastore.

MultiStore vFiler Units Verification

The following steps are required only if using vFiler units as arrays in SRM.

1. Verify that the `vfiler.vol_clone_zapi_allow` option is set to "`on`" on each physical controller used with SRM.
2. Verify that the SnapMirror relationships are defined within the vFiler context.
   a. On the console of the destination array, switch to the vFiler context.

```
f3170c> vfiler context vfiler5
vfiler5@f3170c>
```

   b. Use the `snapmirror status` command to check the status of the relationship.

```
vfiler1@f3170c> snapmirror status
Source          Destination    State         Lag       Status
vfiler4:volsrc   vfiler5:voldst  Snapmirrored  00:09:29  Idle
```

Note: The SnapMirror relationship should show the name of the destination vFiler unit (in this example, vfiler5) as the relationship destination name and the name of the source vFiler unit (in this example, vfiler4) for the relationship source.

Note: If you are performing SnapMirror from physical controller to vFiler unit, the SnapMirror relationship status should show the host name of the source controller for the relationship source and the name of the destination vFiler unit (in this example, vfiler5) as the relationship destination name.

2.3 Connecting the Protected and Recovery Sites

Each site must have a vCenter Server and SRM server. These may be installed on the same server, but typically the SRM software is installed on a separate VM. The SRM vSphere Client plug-in must be installed and enabled in the vSphere Client you are using.

Table 7) Connecting the protected and recovery site prerequisites.

Connecting the Protected and Recovery Site Prerequisites
The SRM software is installed on the SRM server at each site.
The SRM plug-in is installed in the vSphere Client.
The vCenter Servers at each site are able to communicate by HTTP on port 80.

In SRM 5, the Site Recovery plug-in interface has been changed to allow the management of multiple SRM sites from within one vSphere Client instance. You can perform the following steps from a vSphere Client connected to a vCenter Server in either site.

1. In the Site Recovery plug-in interface in the vSphere Client, select the Sites tab in the lower left and select either of the two sites you want to pair.
2. In the Getting Started tab, click the Configure Connection link to start the Configure Connection wizard.

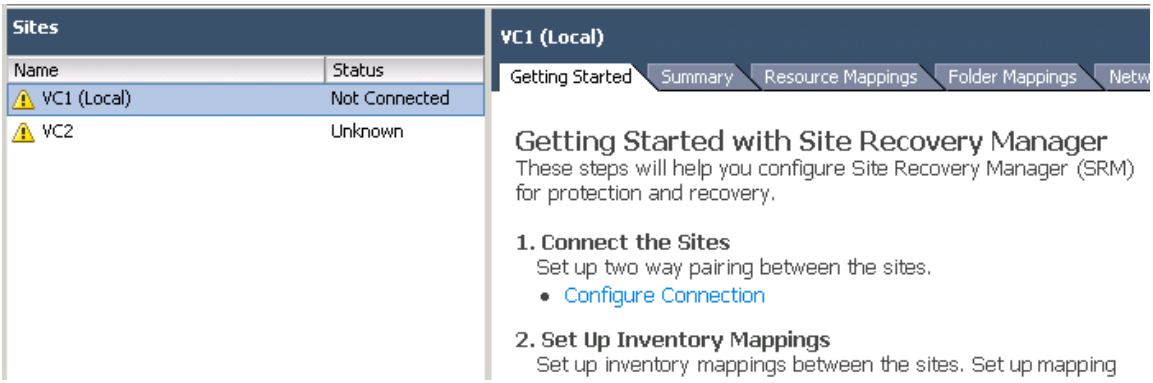

3. Navigate through the Configure Connection wizard, providing the address of the remote vCenter Server and credentials. At the end of the wizard, the Complete Connections screen should show a green check mark next to each of the connection steps. Click Finish to end the wizard.

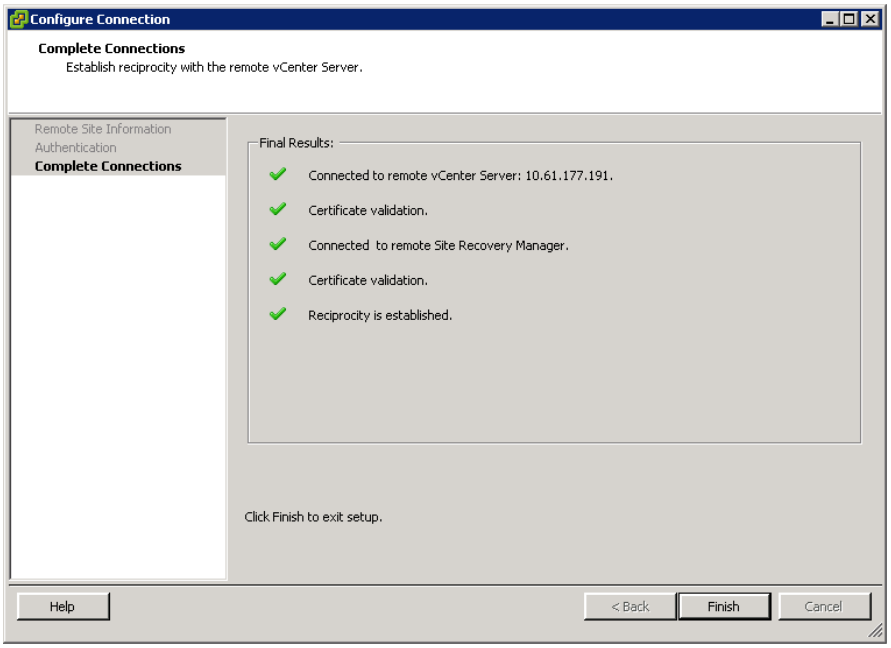

4. Click the Summary tab.

   The SRM interface should now show both sites as being connected.

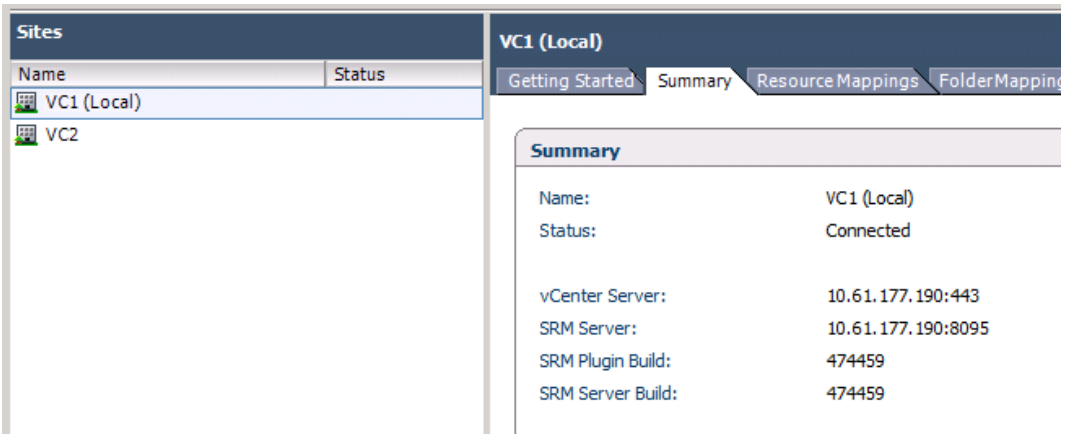

2.4 Configuring Inventory Preferences

The VMware environments at the protected and recovery sites have different sets of resources, such as different VM networks, ESX or ESXi hosts, folders, and so on. In this stage of the configuration, you need to identify a recovery site resource for each corresponding resource at the protected site.

Configure Resource Mappings

1. Select the protected site in the SRM interface and then click the Resource Mappings tab.
2. Expand the Protected Site Resources tree on the left. Select the ESX or ESXi host or cluster where the VMs to be protected are running and click the Configure Mapping link.

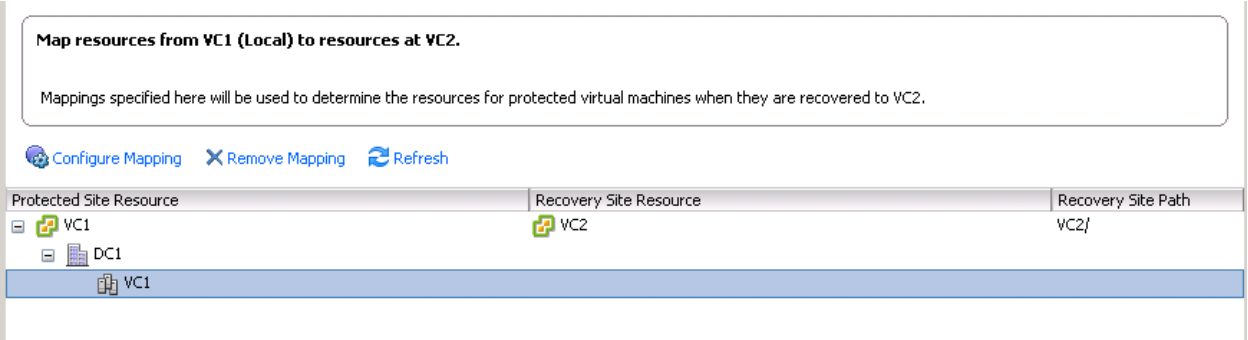

Note: In SRM, resources can be resource pools, ESX or ESXi hosts, or clusters. In this example, a resource mapping is being created between two clusters. You should configure resource mappings between hosts, clusters, resource pools, or folders as appropriate for you environment.

Best Practice

SRM and the NetApp SRA support mixed FC and iSCSI protocols between the protected and recovery site. However, each site should be configured with only one protocol, either FC or iSCSI, not both protocols in the same site. If there is a requirement to have both FC and iSCSI protocols configured in the same site, NetApp recommends that some hosts use iSCSI, and others use FC, and SRM resource mappings be configured such that VMs are configured to failover into one group of hosts or the other.

3. Expand the Recovery Site Resources tree in the dialog that opens. Select the ESX or ESXi host or cluster that you want to use to recover the protected VMs.

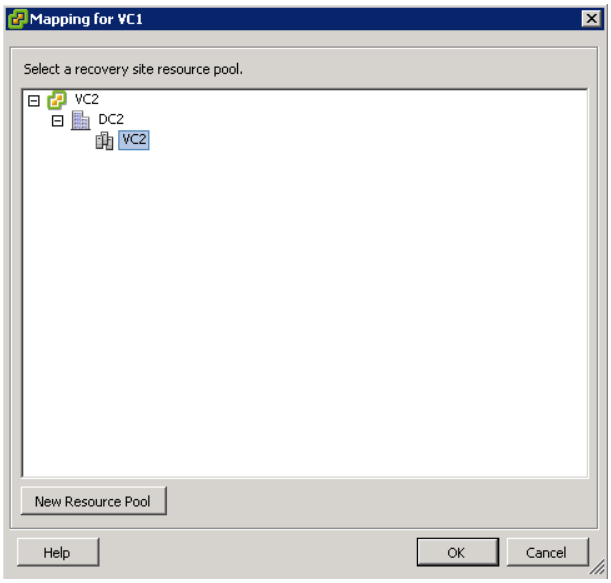

The Resources tab should now show the protected site resource and the corresponding recovery site resource and path.

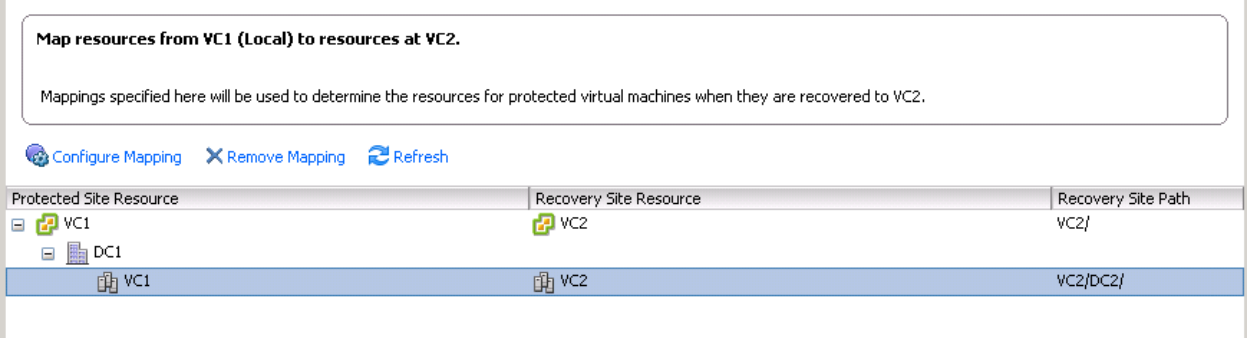

4. Repeat the configuration of resource mappings for other hosts, clusters, or resource pools in the environment.

Configure Folder Mappings

1. Select the protected site in the SRM interface and then select the Folder Mappings tab.
2. Expand the tree of Protected Site Resource folders on the left, select the ESX or ESXi host or cluster where the VMs to be protected are running, and click the Configure Mapping link.

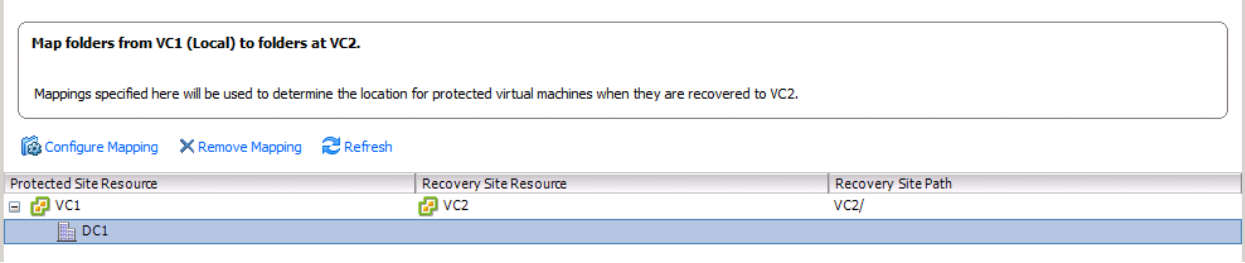

3. Expand the tree of recovery site folders. Select the folder in which you want to place the recovered VM.

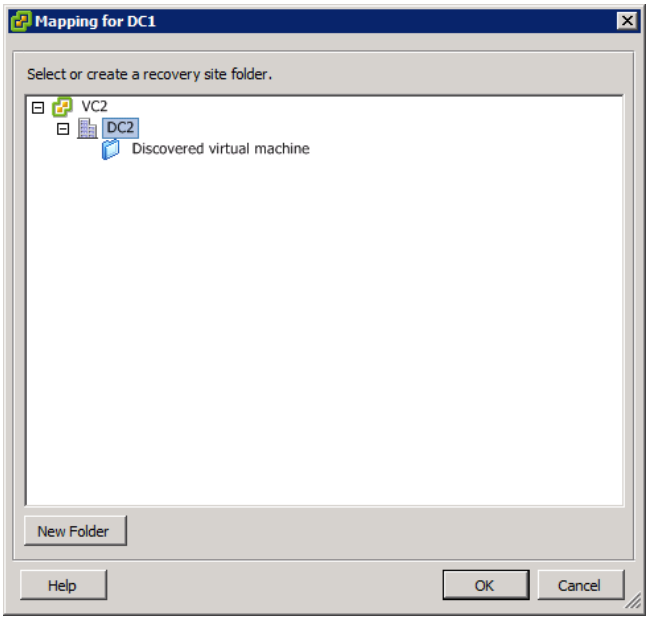

The Folder Mappings tab should now show the protected site resource and the corresponding recovery site resource and path.

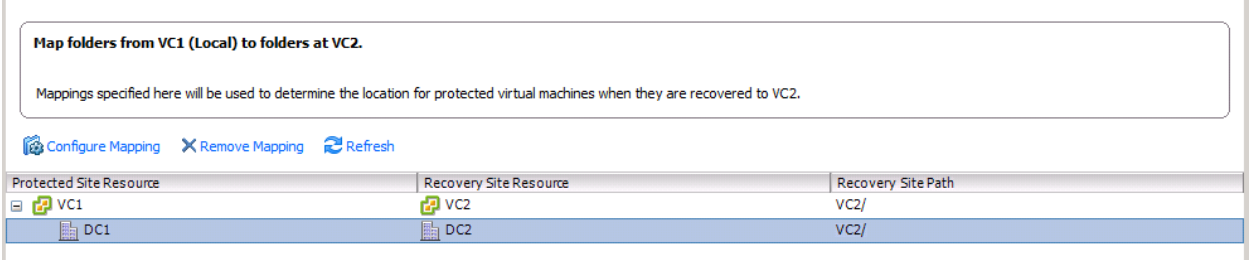

4. Repeat the configuration of folder mappings for the other folders in the environment.

Configure Network Mappings

1. Select the protected site in the SRM interface and then click the Network Mappings tab.
2. Expand the tree of Protected Site Resource networks on the left and select a network at the recovery site. Click the Configure Mapping link.

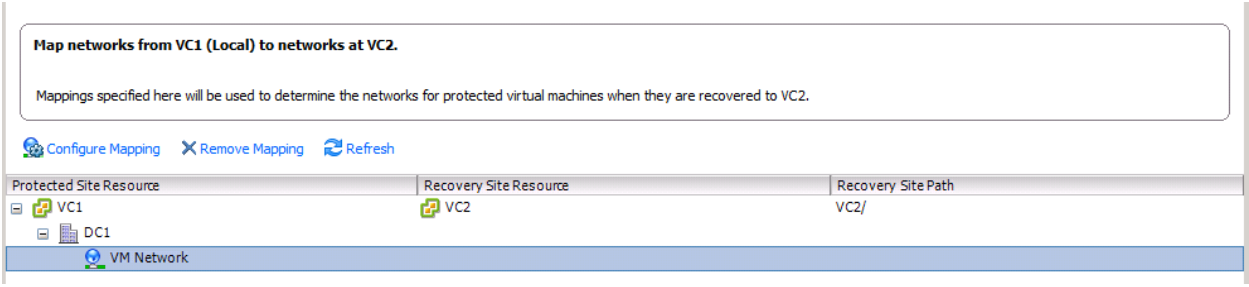

3. Expand the tree of recovery site folders in the dialog that opens. Select the folder in which you want to place the recovered VM.

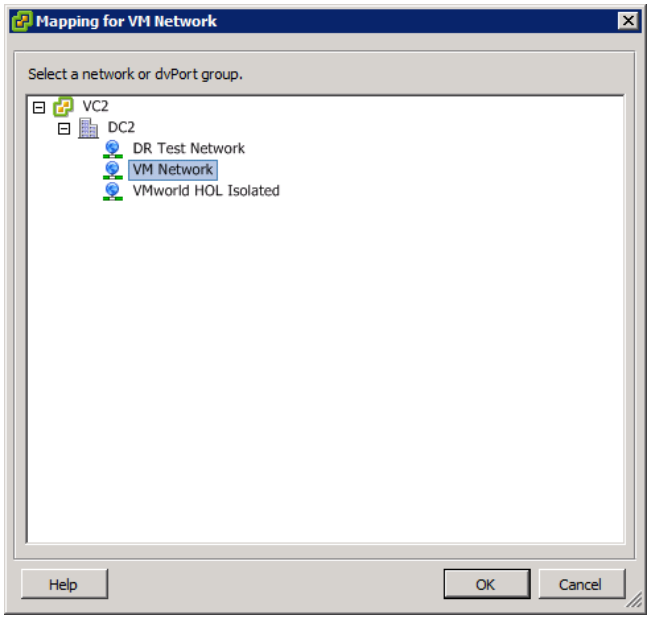

The Folder Mappings tab would now show the protected site resource and the corresponding recovery site resource and path.

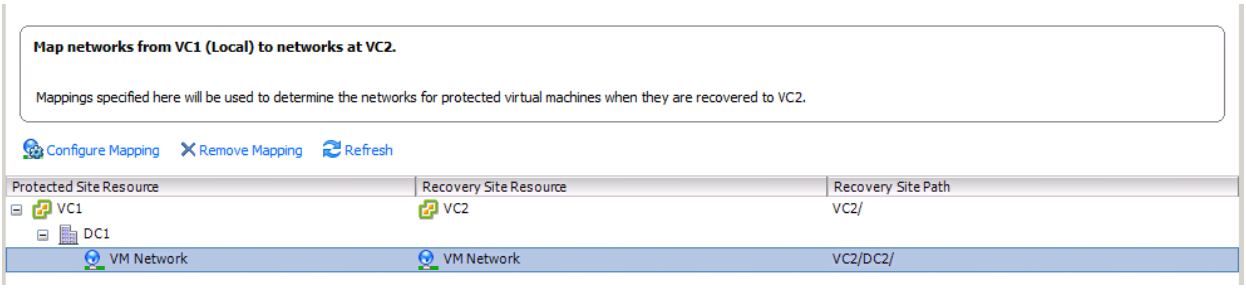

4. Repeat the configuration of network mappings for the other networks in the environment.

Configure Placeholder Datastores

When SRM configures protection of a VM, it creates a placeholder VM with a matching name in the recovery site. This VM does not boot up or consume host resources; it primarily exists to hold a place in the vCenter inventory at the recovery site for the protected VM. You need to specify a datastore to use at the recovery site for storing the placeholder VMs. Placeholder VMs are small, using only a few hundred kilobytes or less, so a large placeholder datastore is not required.

Best Practice

The placeholder VMs are assigned to ESX or ESXi hosts as they are added to the vCenter inventory. The host to which the placeholder VM is assigned is determined by how you configured resource mappings in the environment. The placeholder datastore should be created on a shared storage device, so any host as required can access it. NetApp recommends creating a dedicated small datastore, named srm_placeholder, for example, of a few gigabytes in size depending on the number of VMs in your environment.

If you do not have a datastore in the recovery site to use to store placeholder VMs, create one as follows:

1. Select the protected site in the SRM interface and then select the Placeholder Datastores tab.
2. Click the Configure Placeholder Datastore link.

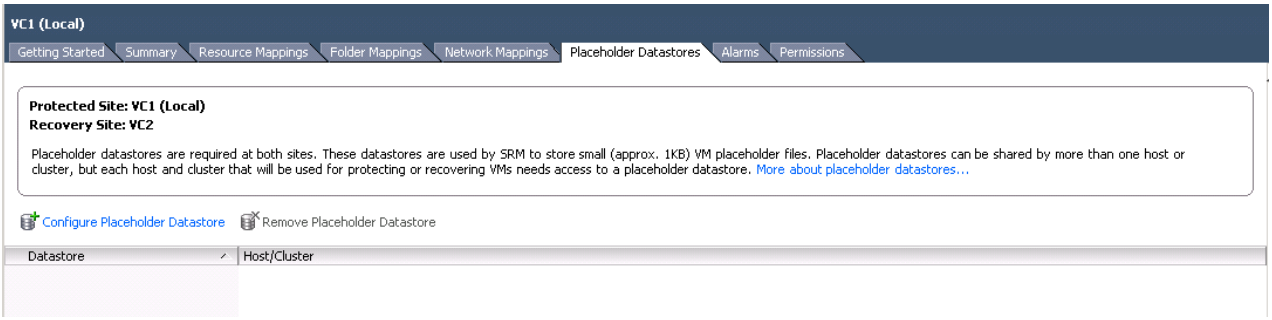

3. Expand the tree of recovery site datastores in the dialog that opens. Select the datastore in which you want SRM to place the placeholder VMs.

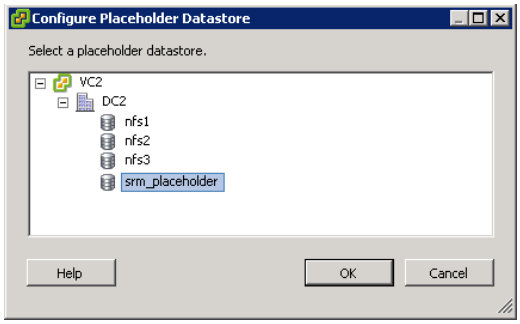

The Placeholder Datastores tab should now show the datastore you selected and the recover site resources that will use that path.

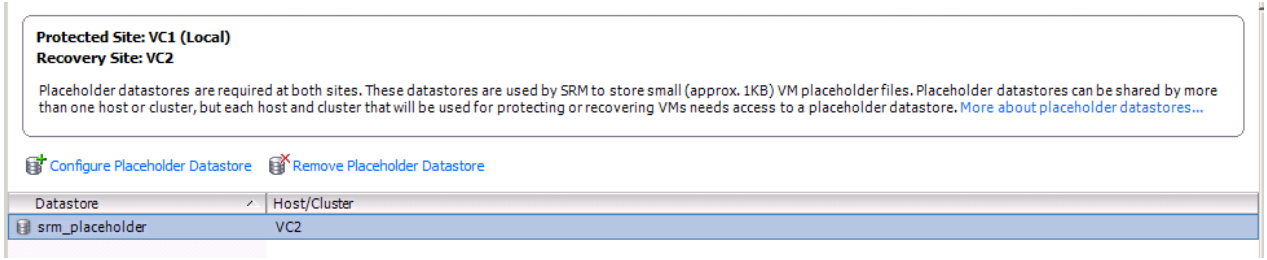

2.5 Install the NetApp Storage Replication Adapter

Table 8) NetApp SRA installation prerequisites.

NetApp SRA Installation Prerequisites
If upgrading from SRM 4 to SRM 5, you must uninstall the previous version of the NetApp adapter before uninstalling SRM 4.
This procedure assumes that VMware vCenter Site Recovery Manager software is installed at each site.

Table 9) NetApp SRA installation requirements.

NetApp SRA Installation Requirements
VMware SRM Version 5 requires NetApp SRA Version 2.x.

As there are no configuration settings entered during the installation of the NetApp SRA, the installation is a simple process of launching the installer, accepting the license agreement, and then proceeding through the install wizard. The SRA will be configured later in the SRM interface.

New Install of the NetApp SRA for SRM 5

Follow these steps when installing the NetApp SRA for SRM 5 on a new SRM server.

1. Launch the NetApp adapter installer on the SRM server at one site.
2. Accept the license agreement in the installer wizard.
3. Proceed through the installer wizard accepting the defaults on all screens.
4. Repeat the process on the other SRM server.

Upgrading from SRM 4 to SRM 5

Previous versions of the NetApp adapter must be uninstalled before installing the NetApp adapter for SRM 5. If you are upgrading from SRM 4 to SRM 5, you must uninstall the previous version of the NetApp SRA prior to uninstalling the SRM software.

Perform the following steps for each SRM server being upgraded, if upgrading the NetApp adapter during upgrade of SRM software from Version 4 to 5.

1. Uninstall the previous version of the NetApp adapter used for SRM 4.
2. Uninstall SRM 4 and install SRM 5 software. Refer to the SRM upgrade procedures described in the Site Recovery Manager Administration available at the SRM documentation site.
3. Launch the new NetApp adapter installer on the SRM server.
4. Accept the license agreement in the installer wizard.
5. Proceed through the installer wizard, accepting the defaults on all screens.
6. Perform the SRM upgrade procedures as described in the Site Recovery Manager Administration available at the SRM documentation site.
7. Repeat the process on the other SRM server.

If the SRM 4 software has already been uninstalled, when you attempt to uninstall the previous version of the NetApp adapter, you will be notified that the adapter cannot be uninstalled because the SRM 4 software does not exist. As a workaround, you will need to manually uninstall the previous adapter. Instructions for manually uninstalling the previous adapter are located in NetApp KB article 2016568.

NetApp SRA Configuration and Storage Discovery

NetApp SRA is configured in the SRM interface. To configure SRA, the source and destination sites must already be connected in SRM, and the SnapMirror relationships must already be configured and replicated.

Table 10) NetApp SRA configuration and storage discovery prerequisites.

NetApp SRA Configuration and Storage Discovery Prerequisites
The SRM software has been installed on the SRM server at each site.
The source and destination vCenter sites have been connected in the SRM interface.
NetApp SRA 2.x has been installed to support SRM 5.
SnapMirror relationships have already been configured and replicated.

Configure the Array Managers

To add a NetApp controller or vFiler unit as an SRM array, perform the following steps. Each controller in a NetApp FAS/V-Series HA pair, or each vFiler unit if using MultiStore, must be added to SRM as an array in the appropriate site. Each array is added only once, in the vCenter site that is local to that array.

1. At the recovery site, check the status of the SnapMirror relationships, noting the name of the source system in the SnapMirror status output.
   a. In the following example, f3170a is the source system host name, and f3170c is the destination system host name. If the name shown as the source in the SnapMirror relationship when viewed at the destination matches the host name of the source system, this SnapMirror relationship will be properly discovered without any special configuration of the SRA. You can proceed to the next step to configure the array manager.

```
f3170c> snapmirror status
```

```
Source           Destination    State        Lag         Status
f3170a:volsrc        f3170c:voldst    Snapmirrored  00:09:29  Idle
```

b. If the source is shown as an IP address, rather than the source system host name, use the ip_hostname_mapping.txt file to configure the SRA to recognize all the SnapMirror relationships as described in the appendix.

```
f3170c> snapmirror status
Source                 Destination    State           Lag       Status
10.72.192.75:volsrc    f3170c:voldst  Snapmirrored   00:09:29   Idle
```

c. If the source is shown as a SnapMirror connection name rather than the source system host name, use the procedure described in the Appendix to configure the SnapMirror relationships to be recognized by the SRA.

```
f3170c> snapmirror status
Source                 Destination    State           Lag       Status
f3170a_conn:volsrc     f3170c:voldst  Snapmirrored    00:09:29  Idle
```

2. In the SRM interface, click the Array Managers tab at the bottom left. Right-click the SRM site and select Add Array Manager.

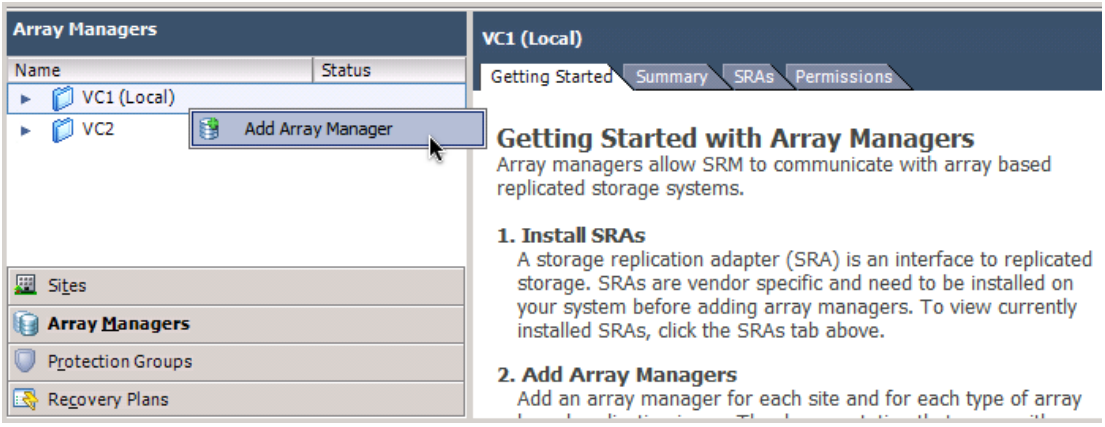

3. In the Add Array Manager screen, enter a name to describe the array, select the SRA type of “NetApp FAS/V-Series Storage Replication Adapter” and click Next.

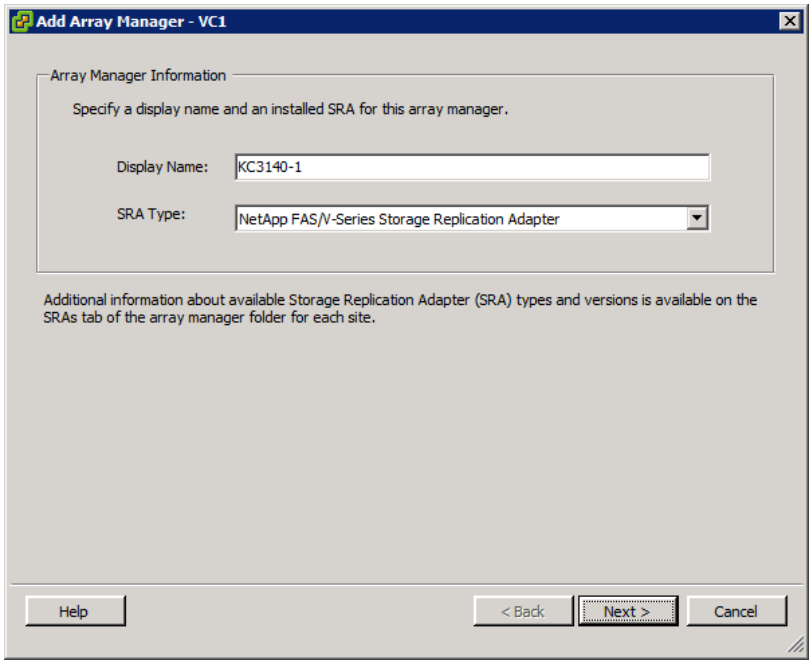

Note: The display name of the array does not have to match the host name of the array. This is a display name used only in the SRM interface.

4. Enter the information required to configure the SRA, according to Table 11.

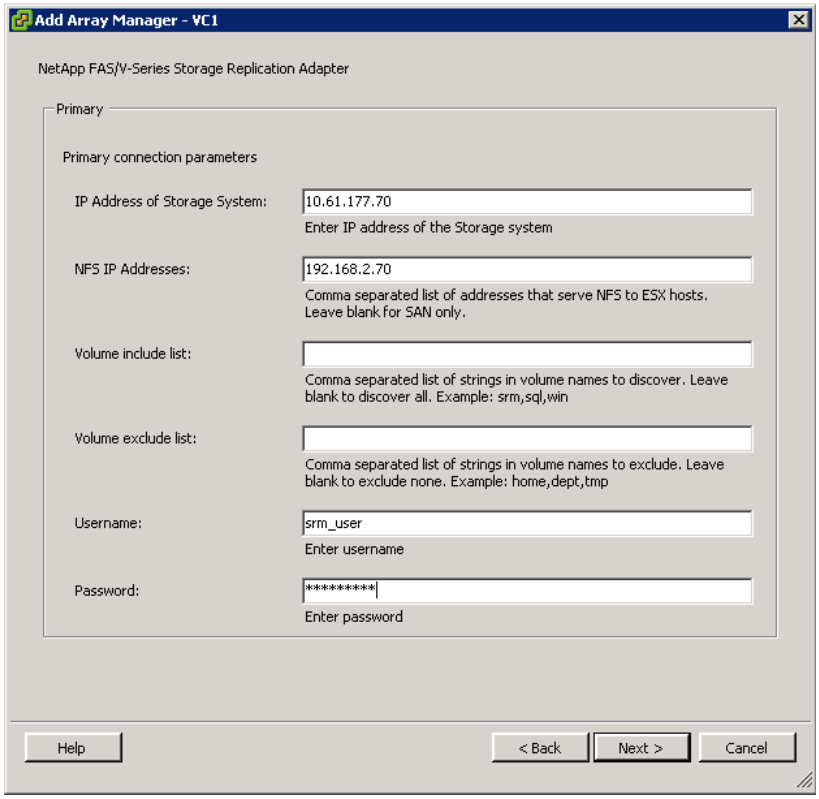

Table 11) Array manager configuration.

In This Field…	Do This…
IP address of storage system	Enter the management address of the controller or vFiler unit.
NFS IP addresses	If using the NFS storage protocol with SRM, enter all the addresses on the controller that are used to serve NFS exports from this controller or vFiler unit to the ESX hosts, separated by commas. If you are not using NFS, enter the same address that was used in the storage system field. *See following note.*
Volume include list	Enter partial names of volumes, or multiple entries separated by commas, that you want to be the only volumes discovered by this array manager. *See following note.*
Volume exclude list	Enter partial names of volumes, or multiple entries separated by commas, that you do not want to be discovered by this array manager. *See following note.*
Username	Enter the credentials required to manage the array. *See following note.*
Password	Enter the credentials required to manage the array. *See following note.*

Note: Regarding NFS IP Addresses field, to support private back-end NFS storage networks, the private address of the storage controller or vFiler unit that is used to serve NFS from that controller to the ESX hosts must be entered in the NFS IP Addresses field. If the controller uses multiple IP addresses to serve NFS data, these IP addresses are entered in the NFS IP Addresses field and are separated by commas. If the address used to administer the controller or vFiler unit is the same address used to serve NFS, you must enter that address in the NFS IP Addresses field as well. If specific IP addresses are not entered in the NFS IP Addresses field to use for NFS datastore connections, the disaster recovery NAS adapter returns all available IP addresses of the storage controller to SRM. When SRM performs a DR test or failover, it will attempt to make NFS mounts to every IP address on the controller, even ones that the ESX hosts would be unable to access from the VMkernel ports such as addresses on other networks, so it is important that the correct addresses are entered in the NFS IP Addresses field.

Note: Regarding the volume include and exclude fields, these fields are used to restrict the volumes that are discovered by this array manager. When storage discovery is performed by SRM, the SRA must return a list of all replicated devices on that array. Devices that are not replicated by design, or replicated devices that are not used by the vCenter site for this array manager, can show up with warnings in the SRM interface. To remove these devices from discovery, so that the SRM interface does not display warnings, use a combination of values in the volume include and exclude lists to control which storage devices are detected by SRM. NetApp recommends using the same values in both the source and destination array manager configuration.

Note: Regarding the Username and Password field, instructions for configuring role-based access controls (RBAC), so that the account used to manage the array can be one with access restricted to only the necessary commands, can be found in NetApp KB article 1013325.

Best Practice

NetApp recommends completely configuring the array manager and SRM first and then testing the SRM functionality using a non-restricted account. After proper functionality of SRM has been verified, change the credentials in the array manager to the RBAC restricted account and refresh the SRA.

5. Click Next. SRM will connect to the SRA, verify the credentials provided, and query the capabilities of the SRA. At this time storage discovery has not yet been performed by SRM. Click Finish to exit the Array Manager Add wizard.

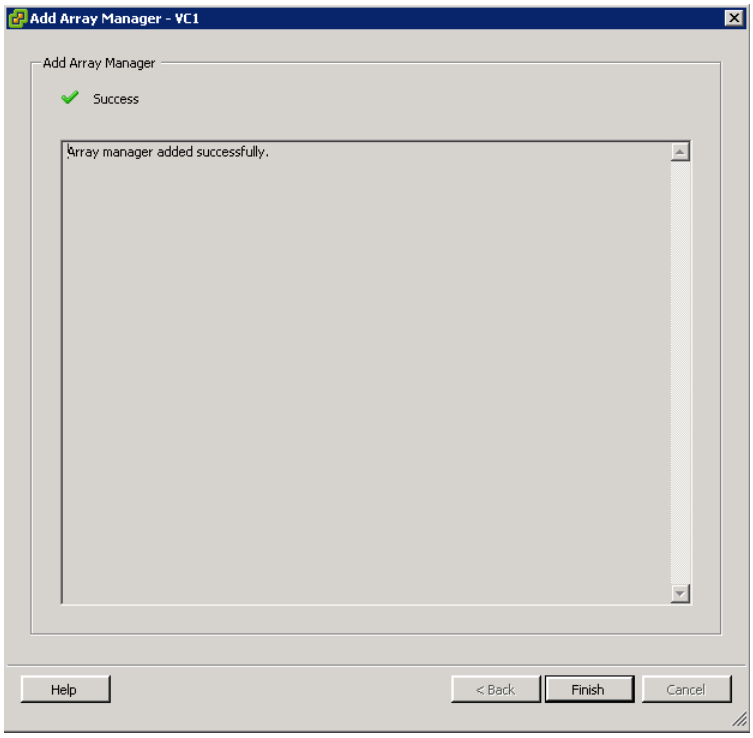

6. Repeat the preceding steps to add the array manager for the controller in the other vCenter site.

Storage Discovery

Perform the following steps to enable the array pairs and perform storage discovery. Each SRM array pair needs to be enabled only from one site or the other.

Table 12) Storage discovery requirements.

Storage Discovery Requirements
To be able to protect a VM with SRM and the NetApp SRA, all parts of the VM must exist on only one NetApp controller or MultiStore vFiler unit, in both the protected and recovery site.

1. In the SRM interface, click the Array Managers tab at the bottom left. Select either of the array managers you previously configured in SRM and then click the Array Pairs tab.

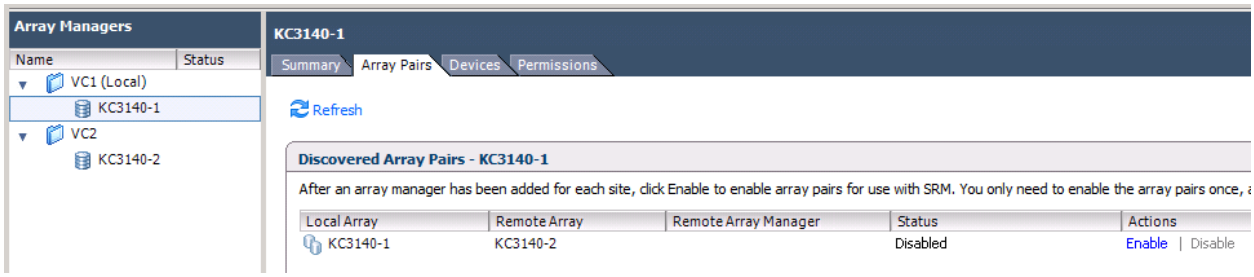

2. Click the enable link on the right to enable this array pair. At this time SRM will run the storage discovery process.
3. When the storage discovery process completes, the SRM interface will show the array pair as Enabled and list the name of the remote array manager in the remote array manager column.

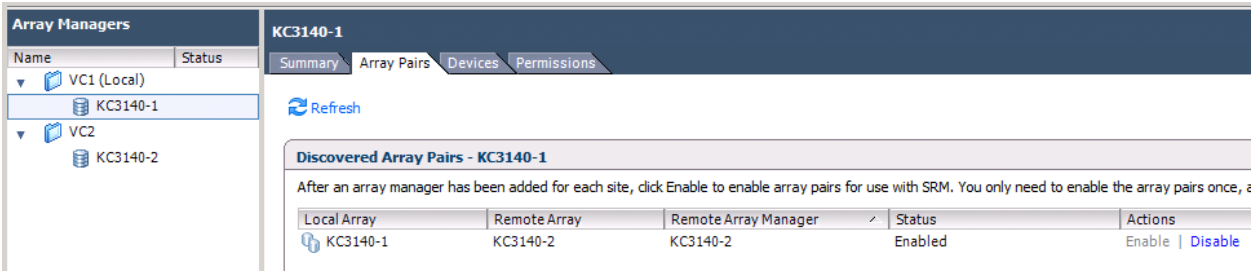

4. Click the Devices tab at the top. The SRM interface will show the devices that were discovered and the current direction of replication.

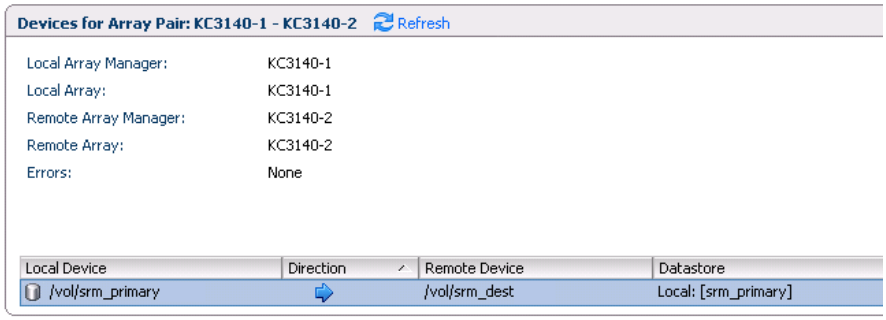

Note: It is import to be aware that in the SRM interface you can view the SRM configuration from either site and to be aware of which array manager you have selected to properly interpret the current replication direction. Notice in the preceding example, with the KC3140-1 array manager selected in site VC1, the SRM interface shows the replication going from the /vol/srm_primary device listed in the Local Device column to the /vol/srm_dest device listed in the Remote Device column. For example, if you were to select the KC3140-2 controller in site VC2, the SRM interface would change to show replication going away from the local site, even though you are still in the same vSphere Client that is connected to site VC1.

5. Repeat the preceding steps to enable the other array pairs in the SRM interface. Remember that each array pair need only be enabled once, from either SRM site.

2.6 Build Protection Groups

Protection groups define VMs and datastores in groups that will be recovered together.

Table 13) Protection group prerequisites.

Protection Group Prerequisites
NetApp SRA has been configured at each SRM site.
The SRM array pair has been enabled, and storage discovery was successfully performed by SRM.

Table 14) Protection group requirements.

Protection Group Requirements
To be able to protect a VM with SRM and the NetApp SRA, you must make sure all parts of the VM exist on only one NetApp controller or MultiStore vFiler unit, in both the protected and recovery site.

In a NetApp environment, replication occurs at the FlexVol volume level when you use volume SnapMirror (VSM). When you use qtree SnapMirror (QSM), replication occurs at the qtree level.

Best Practice

If you are using qtrees, with multiple qtrees in one volume, as separate datastores in the VMware environment, then the level at which you mirror data should match the level at which you configured the datastores. Meaning, if you are using different qtrees in the same volume to store different datastores, then use qtree level SnapMirror for each qtree. If you are storing each datastore in its own volume, then use volume-level SnapMirror for each volume.

To create an SRM protection group, perform the following steps.

1. In the SRM interface, click the Protection Groups tab on the bottom left. Right-click the All Protection Groups folder and select Create Protection Group.

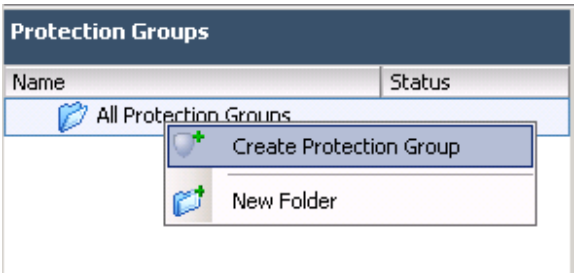

2. Select the protected site. This is the site where the VMs that you want to configure in this protection group currently exist.

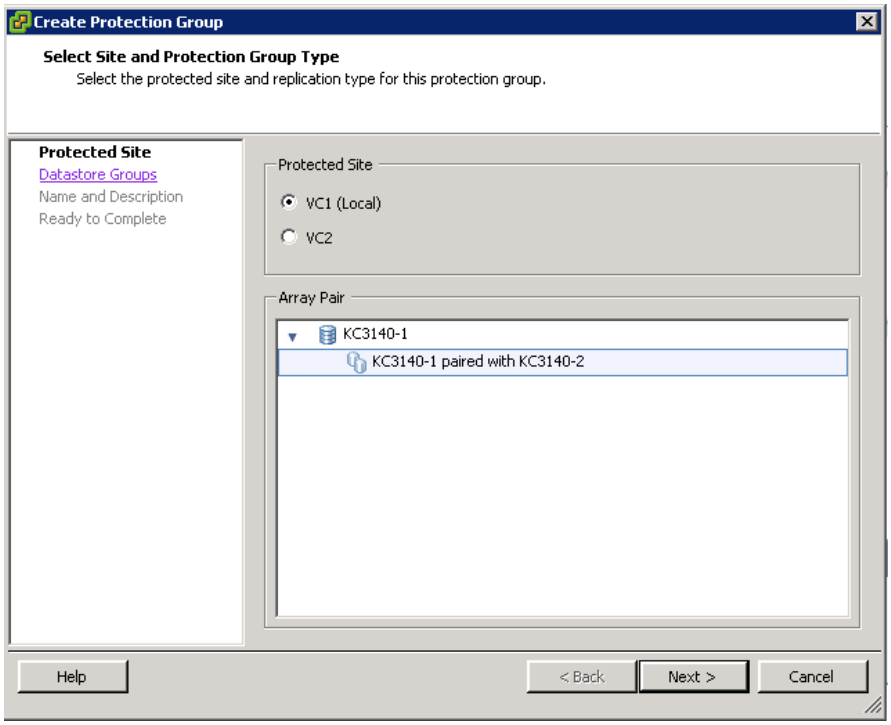

3. Select the datastore group or groups you want to configure in this protection group and then click Next.

Best Practice

When multiple qtrees in the same volume are configured containing different datastores, make sure that these datastores are all included in the same SRM recovery plan, so that they are failed over and failed back together. The best way to do this is by selecting them as datastore groups to be contained in the same protection group.

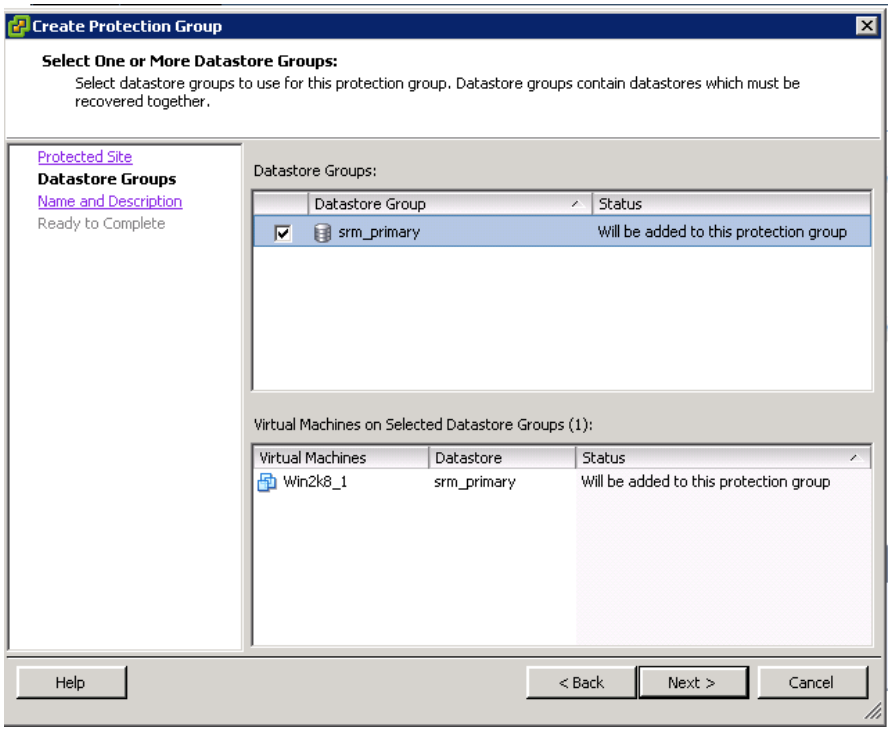

Note: Datastore groups are automatically created by SRM during storage discovery. The vCenter environment is examined by SRM and a minimum set of storage required to recovery each VM is calculated, and then datastores are divided into different datastore groups depending on which VMs are using which datastores.

4. Enter a name and description for this protection group and then click Next.

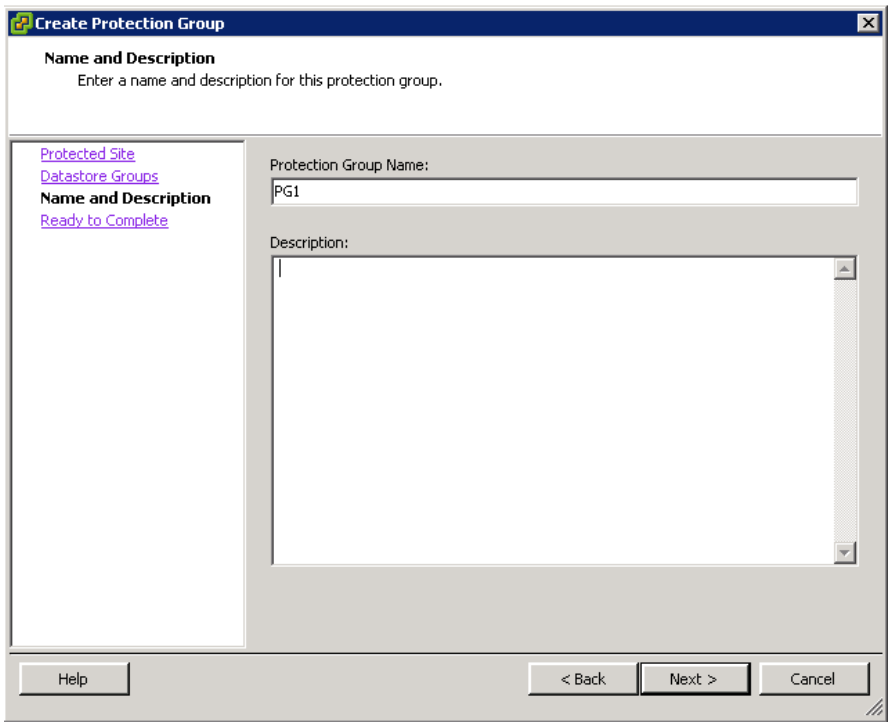

5. Review the final screen of the protection group wizard and click Finish.

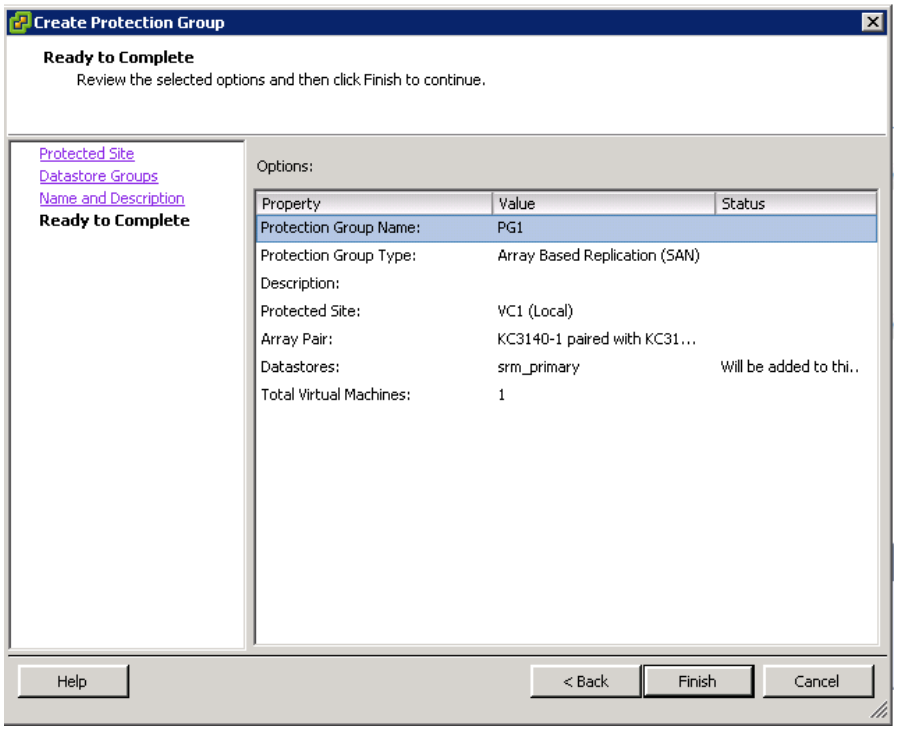

6. Repeat the preceding steps to create any remaining protection groups required in the environment.

2.7 Create Recovery Plans

Recovery plans define which protection groups will be recovered in the same process. Multiple protection groups can be configured in the same recovery plan. Additionally, a single protection group can be included in multiple different recovery plans, if you want more options for the execution of recovery plans. For example, if each protection group is configured in its own recovery plan, you may perform a test failover of that recovery plan that would perform a test only for VMs contained in that plan.

Note: The procedure here assumes the VMs will be recovered at the recovery site having the same network configuration (IP address, subnet mask, gateway address, DNS settings, and so on) that was used at the protected site. If different network settings need to be applied to individual VMs as they are recovered, this can be specified in the properties settings of a VM in the recovery plan. To configure SRM to apply different network settings to multiple VMs, without having to edit separately the properties of each VM in the recovery plan, VMware provides a tool called the dr-ip-customizer. Directions for using the utility are provided in the section "Customize IP Properties for a Group of Virtual Machines" in Site Recovery Manager Administration Guide.

To create a recovery plan, perform the following steps.

1. In the SRM interface, select the Recovery Plans tab on the bottom left, then right-click the All Recovery Plans folder and select Create Recovery Plan.

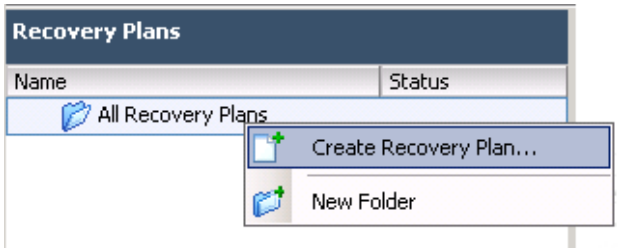

2. Select the recovery site for which you want to create the protection group. You must choose the site to which the datastores are replicated that you want to include in the recovery plan.

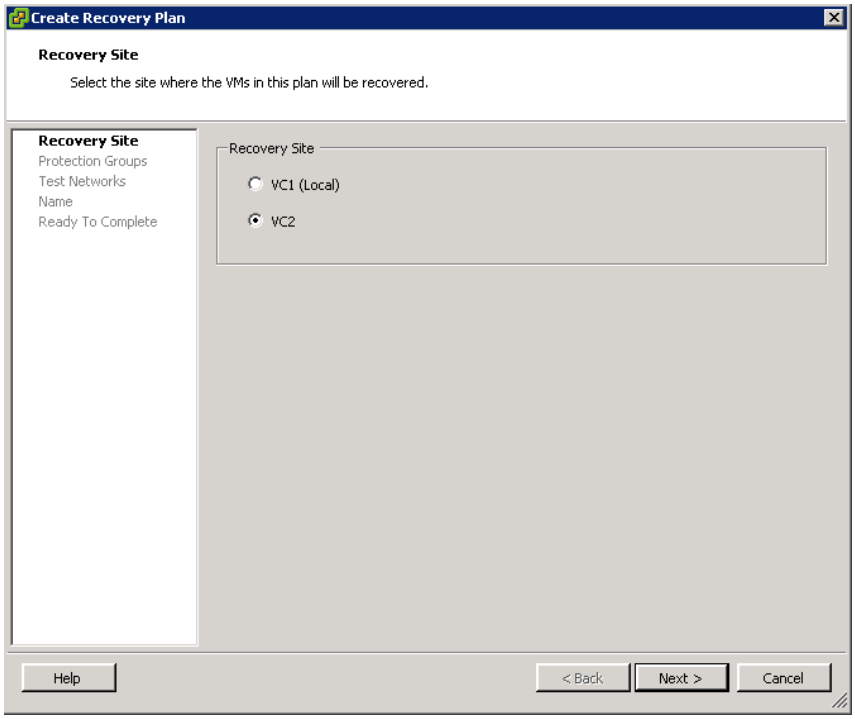

3. Select the protection groups that you want to include in this recovery plan.

Best Practice

When multiple qtrees in the same volume are configured containing different datastores, make sure that these datastores are all included in the same SRM recovery plan, so that they are failed over and failed back together. The best way to do this is by selecting them as datastore groups together to be contained in the same protection group.

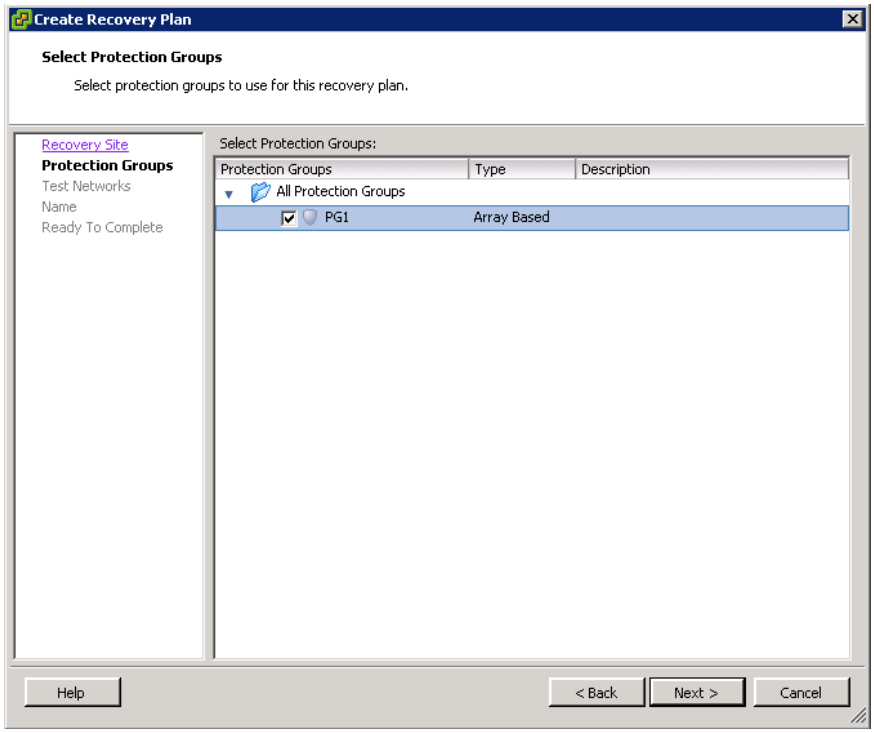

4. Configure a test network to use when running tests of this recovery plan.

Note: If the test network for a recovery network is set to Auto, SRM will automatically create a test bubble network to connect to the VMs to during a DR test. However, if you have created a specific network to use for DR testing, select that network as the test network for use in this recovery plan. If you have configured the environment with a network to use for SRM DR testing, that network will appear in the list of networks in the recovery data center. There is no need to select a test network where your DR test network is listed in the Recovery Network column; you may leave that row set to Auto.

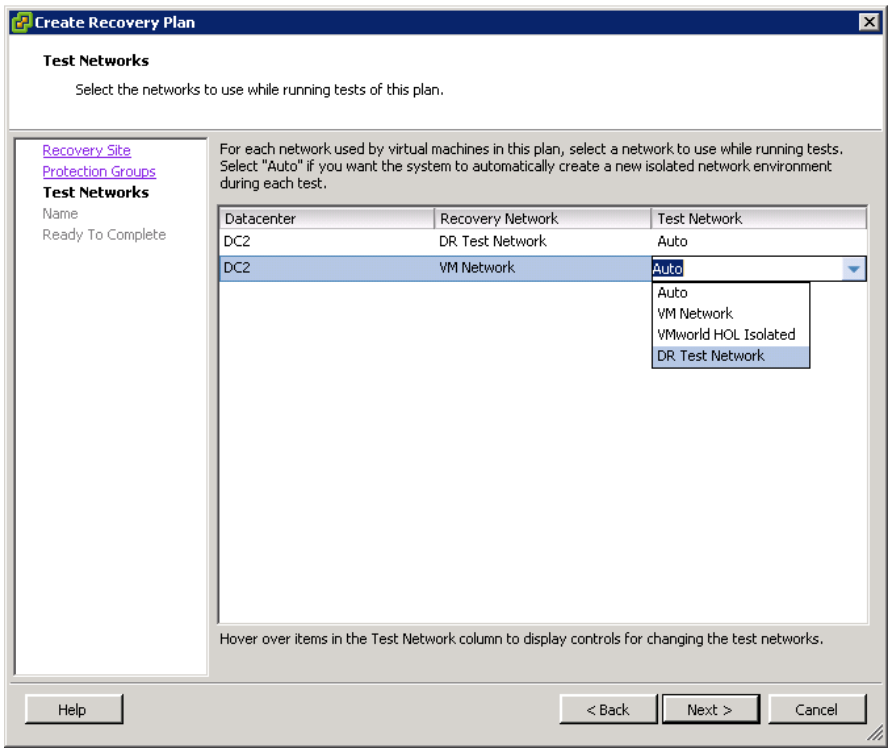

5. Enter a name and description for this recovery plan and click Next.

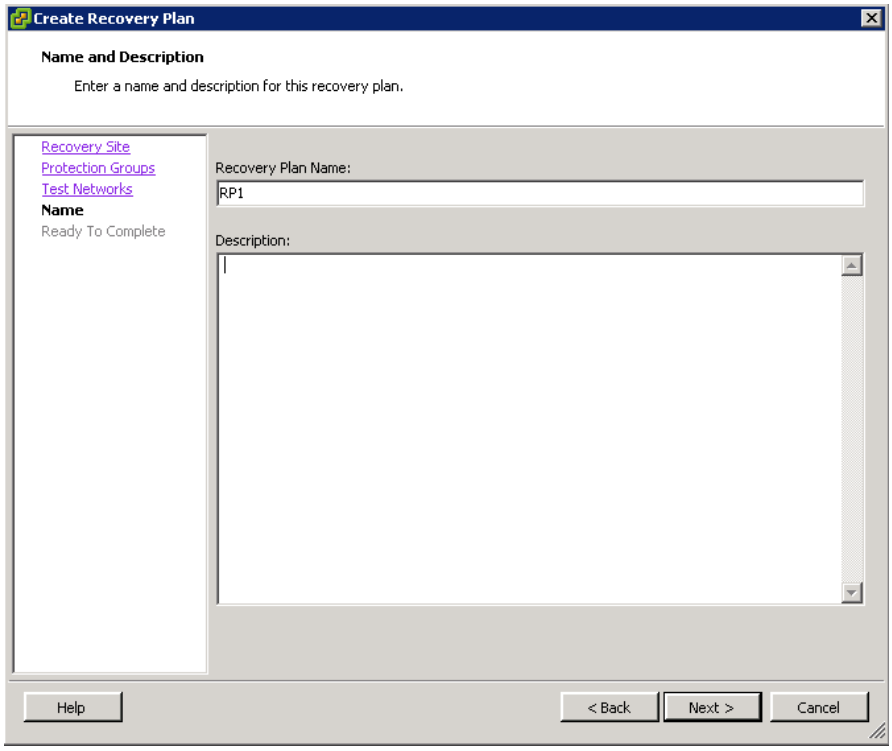

6. Review the final screen of the recovery plan wizard and click Finish.

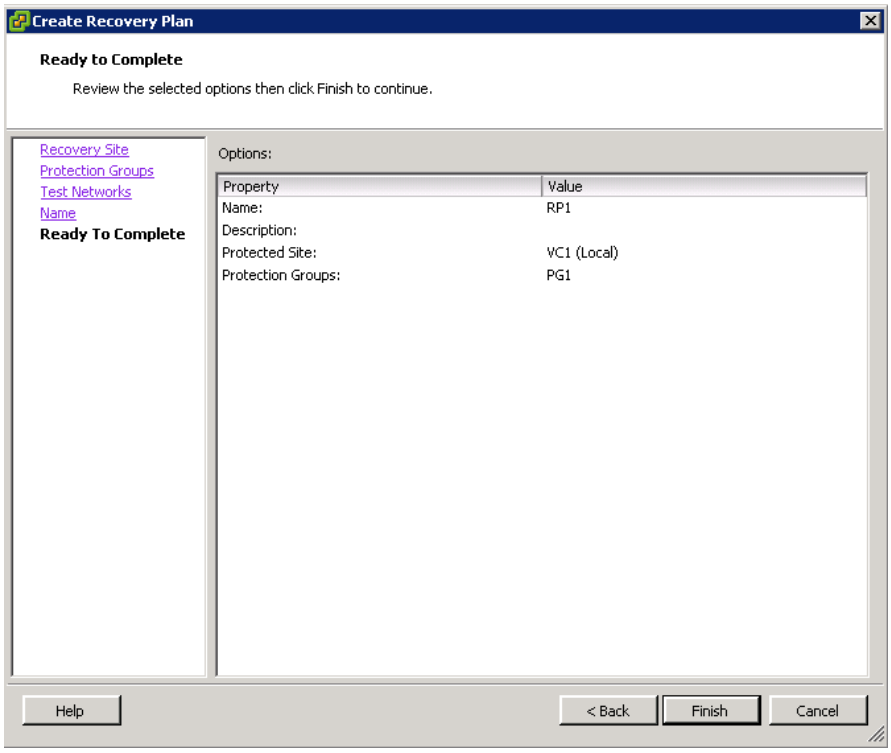

7. Repeat the preceding steps to create any remaining recovery plans required in the environment.

3 Operational Procedures

3.1 Perform a Test Failover

Disaster recovery testing with SRM is an operational procedure that should be performed as your business requirements dictate. When a test failover is executed, SRM uses NetApp SRA to create a FlexClone volume of the replicated volume at the recovery site. The datastore in this volume is then mounted to the hosts at the recovery site and the VMs configured and powered on. During a test, the VMs are connected to the DR test bubble network, rather than the public virtual machine network, or to a preconfigured test network as defined in the recovery plan properties.

To understand how SRM performs a test failover nondisruptively, it is important to understand how the VMs are connected to the network both before and during a test failover operation. Figure 16 shows the vSphere Client map of the network at the recovery site before the failover test. Note the virtual machines that are not running are the placeholder virtual machines created by SRM, and that there is an Active Directory/DNS server at this site called Site2DC.

Figure 16) VM to network map view in vCenter prior to failover test.

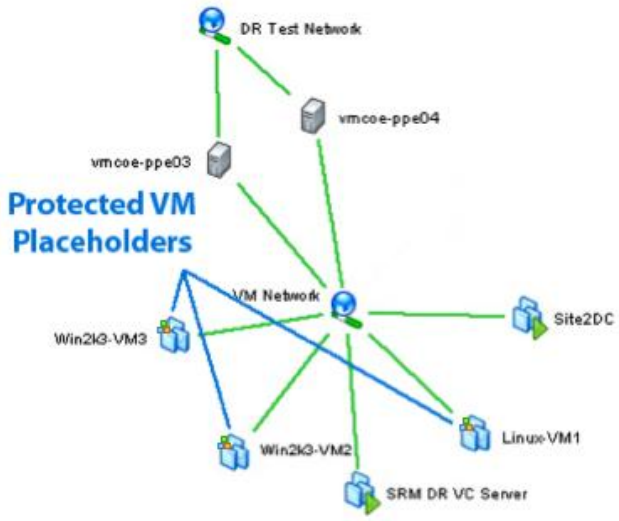

To perform an SRM failover, test execute the following steps.

1. (Optional) If you have preconfigured a private DR test network and require AD/DNS services in the testing environment, you can clone an AD/DNS server at the recovery site prior to running the DR test. Before powering on the clone, make sure to reconfigure the network settings of the cloned AD server to connect it only to the DR test network, as shown in Figure 17.

Figure 17) Cloned AD/DNS server in DR test network.

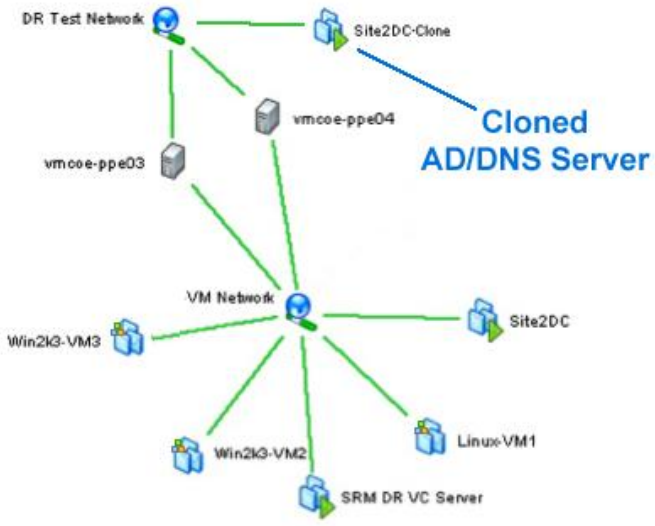

Note: The cloned AD server must be one that was configured as a global catalog server. Some applications and AD functions require FSMO roles in the Active Directory forest. To seize the roles on the cloned AD server in the private test network, use the procedure described in Microsoft KB 255504. The five roles are schema master, domain naming master, RID master, PDC emulator, and infrastructure master. After the roles are seized, it is very important to make sure that the clone never be connected to a production VM network and that it is destroyed after DR testing is complete.

2. In the SRM interface, click the Recovery Plans tab at the bottom left. Select the recovery plan you want to test and click the Recovery Steps tab.

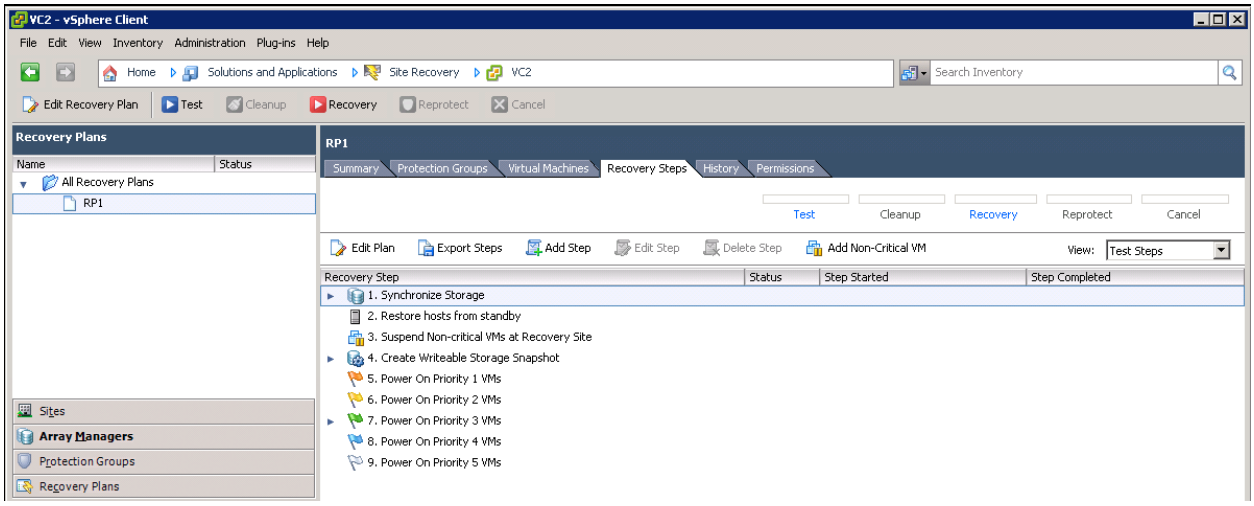

Note: You do not have to click the Recovery Steps tab; however, it is helpful to see the steps as they are executed.

3. Select the Test link in the upper right to begin a test failover.
4. Click Test to launch a test of the selected recovery plan.

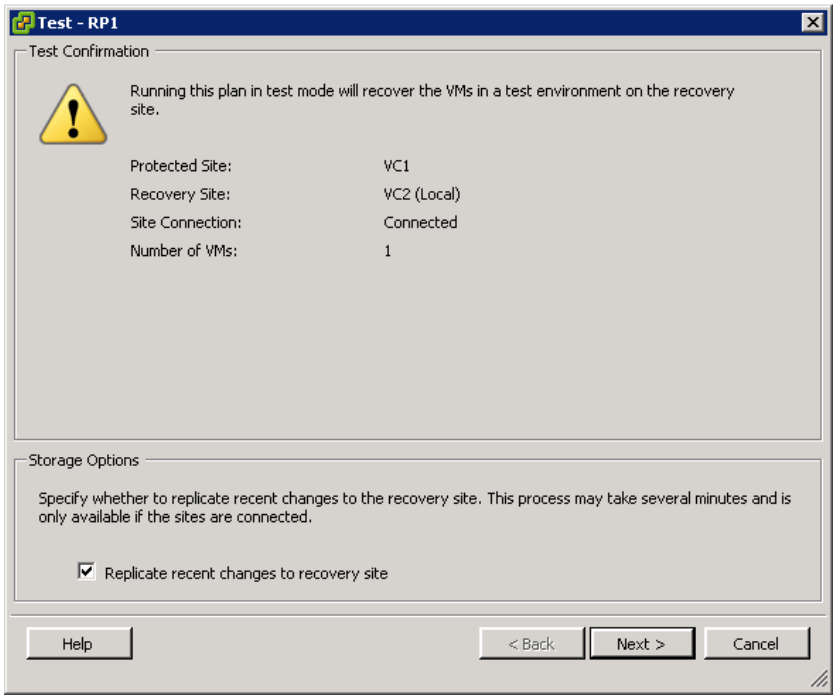

Note: If you select the “Replicate recent changes to the recovery site” checkbox, SRM will use the NetApp SRA to perform a one-time SnapMirror update for each relationship involved in the recovery plan.

5. Click Next to execute the test to proceed to the next screen of the wizard. Review the information about the recovery plan and click Start to perform the test.

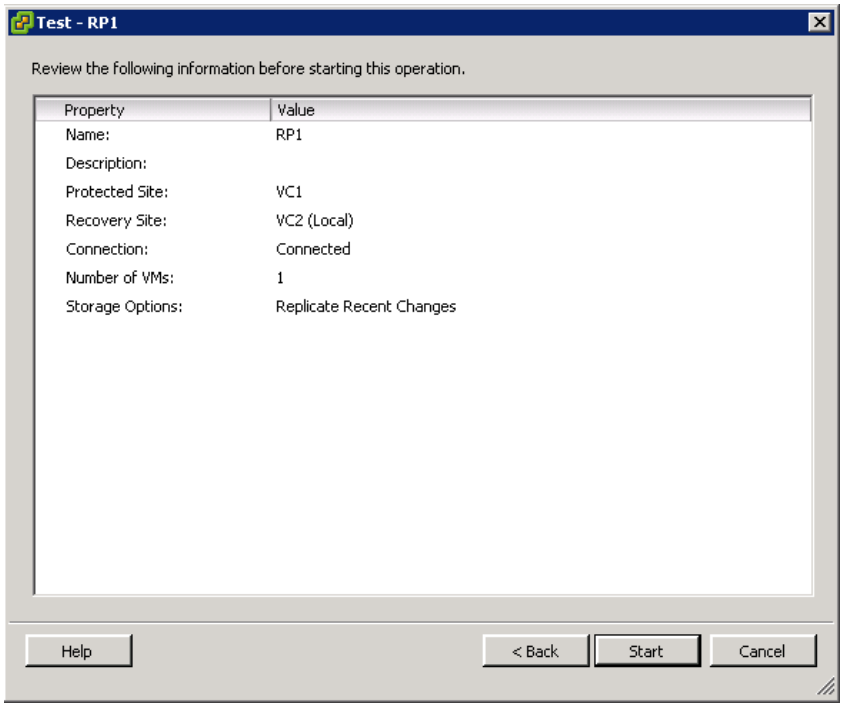

6. Review the progress of the recovery plan by clicking the Recovery Steps tab.

Note: If you selected the box to replicate recent changes to the recovery site, SRM will remain in the synchronize storage step until the replication process is complete. SRM uses the SRA to periodically check the status of the replication relationship; once it is complete, the recovery plan will proceed to the next step.

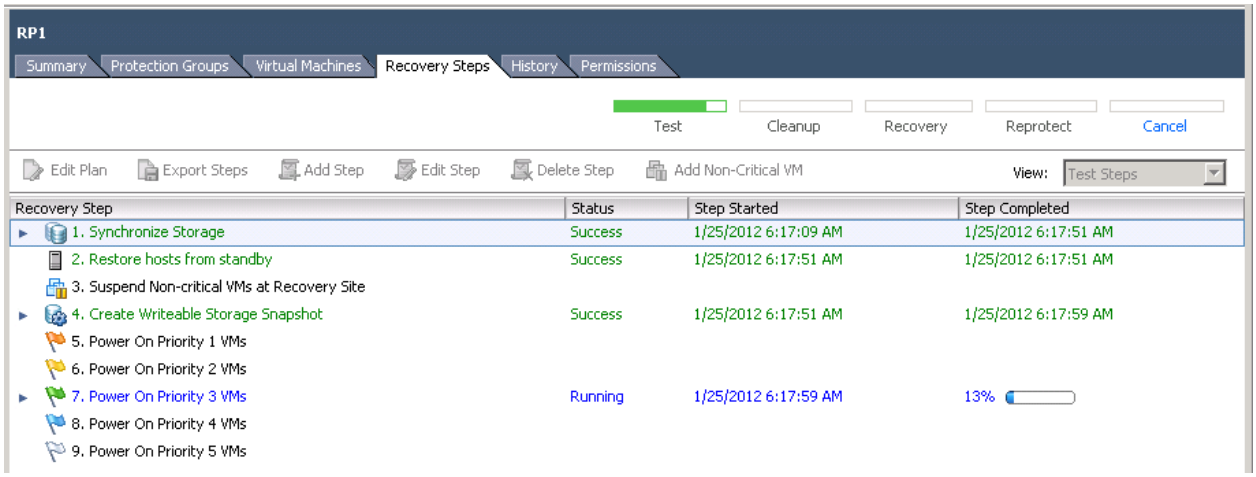

Figure 18) VM to network map in vCenter client while in DR test mode.

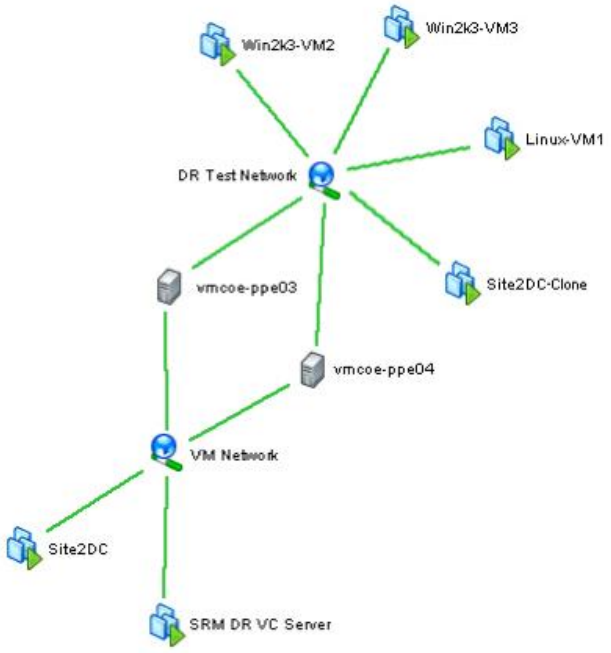

Note: After all the virtual machines have been recovered and the test run of the recovery plan has finished, the network map in the vSphere Client will change as shown in Figure 18. The recovered VMs (Win2k3-VM2, Win2k3-VM3, and Linux-VM1) have been connected to the bubble network; the VM named Site2DC-Clone is a clone of an AD/DNS server that was connected to the DR test network to provide authentication and name resolution in the test environment.

7. After the recovery plan has completed, the VMs will be running in the test network. You may proceed with your DR testing by connecting to the console of the appropriate VMs and perform any tests necessary inside the test network.

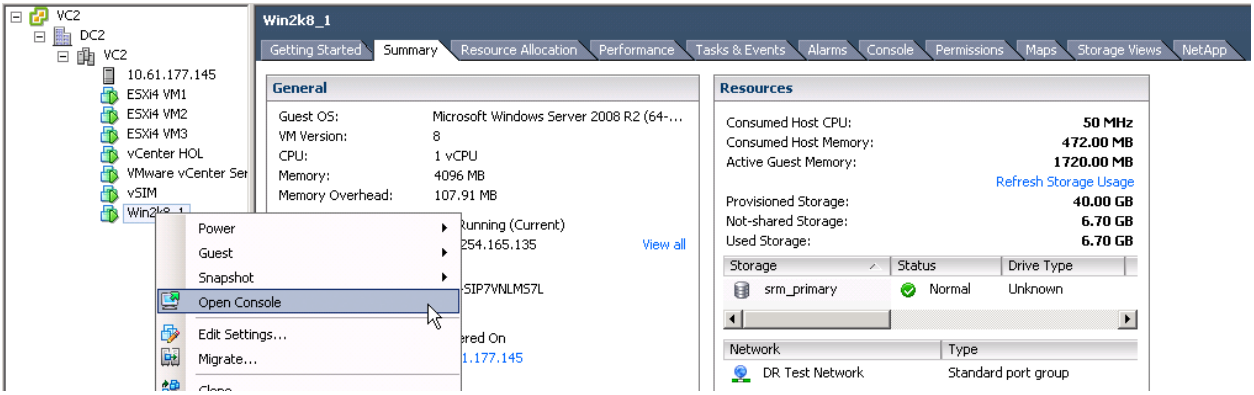

Note: The tests you perform will vary depending on what you are using the VMs for in your environment. For example, if your VMs support a Web application, then inside a test VM, open a Web browser and attempt to access your application in the private test network.

Figure 19) Datastores mounted on different NFS IP addresses.

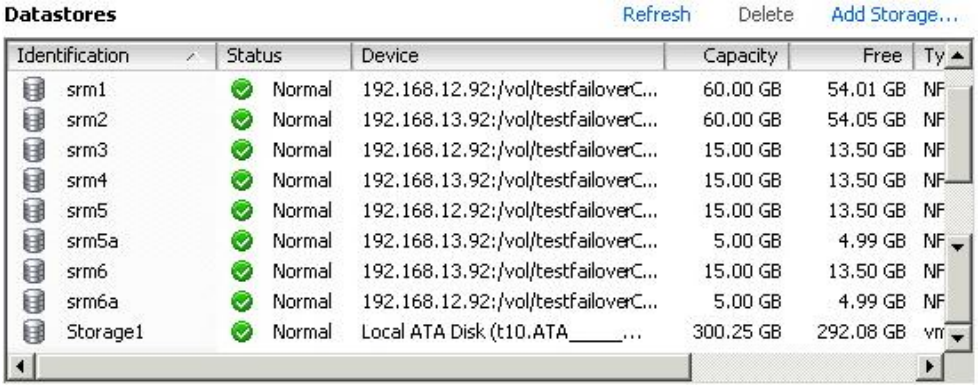

Note: When NFS datastores are being mounted by SRM for a test failover, SRM uses each IP address that was provided in the NFS IP Addresses field when the SRA was configured. The first recovered datastore is mounted at the first NFS IP address provided. All ESXi hosts mounting that datastore use this same IP to mount it. Then, the second datastore being recovered is mounted at the second NFS IP address, likewise by all ESXi hosts accessing the datastore. Figure 19 is an example of multiple NFS datastores being mounted on IP addresses 1921.68.12.92 and 192.168.13.92 during an SRM failover operation. SRM continues mounting datastores on all ESXi hosts (which ESX hosts is determined by SRM inventory mappings), alternating between each NFS IP address provided as each datastore is recovered. SRM does not currently provide a mechanism to specify which IP address should be used for each datastore. Some environments might have three, four, or more IP addresses on the NetApp array to enter into the NFS IP Addresses field.

8. Once testing is complete, discard the test environment by clicking the Cleanup link in the upper right.

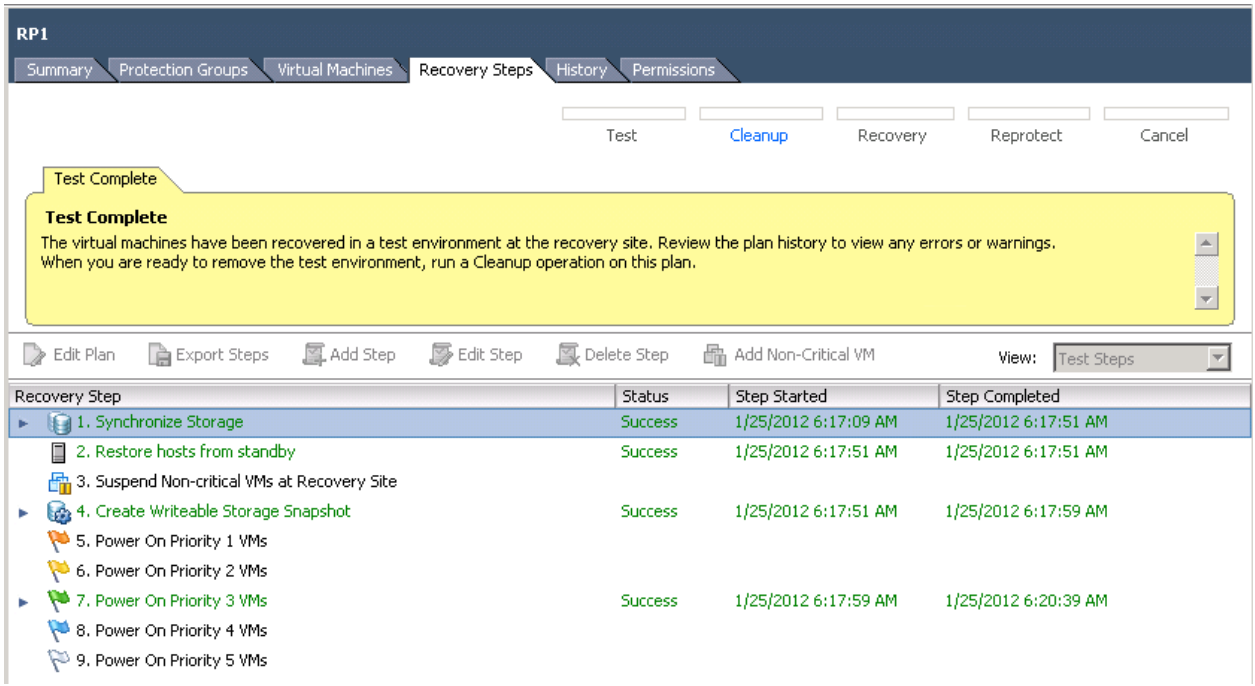

9. Click Next in the cleanup wizard.

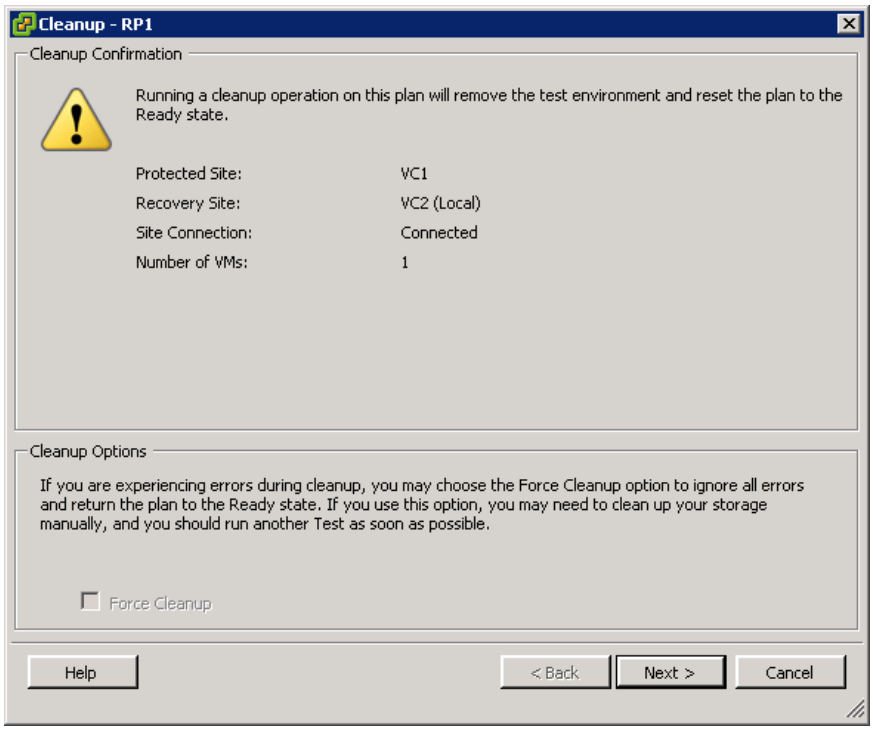

10. Review the details of the cleanup process and click Start. SRM will shut down the test VMs, disconnect the datastores, and use the NetApp SRA to remove the FlexClone volumes from the array.

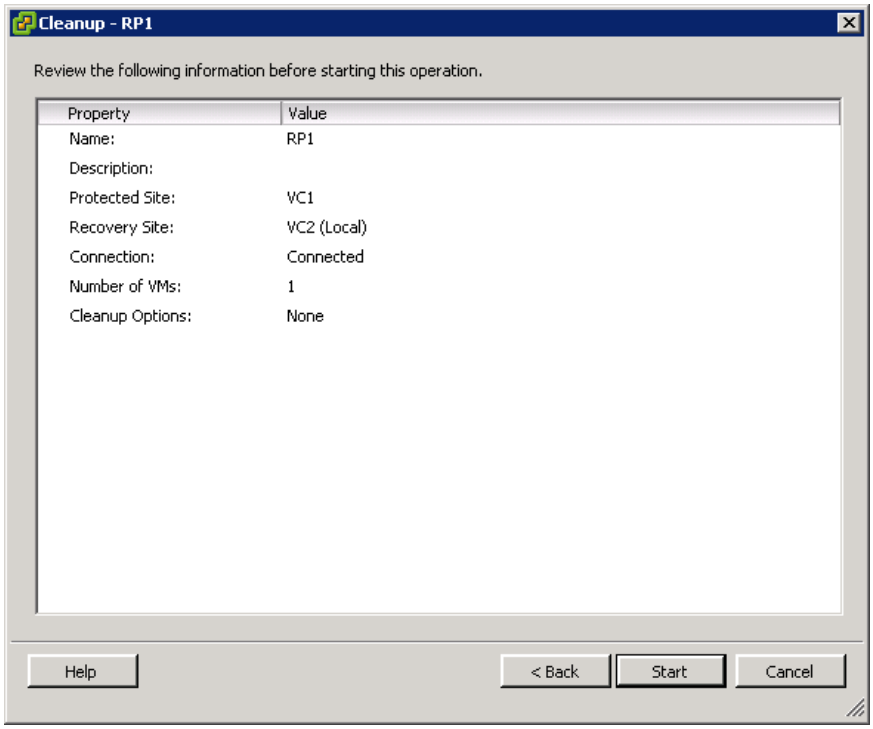

11. (Optional) If you cloned an AD/DNS server at the DR site, it can now be shut down, removed from inventory, and deleted. A new clone should be created if rerunning DR tests after some time, as the

information in the AD/DNS server might become outdated before the next failover test. Again, make sure that the cloned AD server is never connected to a public VM network.

3.2 Perform a Planned or Unplanned Failover

SRM provides two methods for performing a real failover; both methods are performed in the SRM recovery workflow:

- Planned failover, to be used in the event that you want to perform a disruptive migration of VMs from the protected site to the recovery site. This workflow includes the shutdown of the VMs at the protected site, update of storage replication, and startup of VMs at the recovery site.
- Disaster recovery (unplanned failover), to be used in the event the protected site has become unavailable, or the storage for a particular set of VMs at the protected site has become unavailable. In this workflow, SRM will attempt to shut down the VMs at the protected site if they are accessible and then start up the VMs at the recovery site.

Note: SRM always attempts to access the protected site to shut down the VMs at that site that are in the recovery plan being executed for a real failover. The vCenter Server and SRM server at that site must be accessible for this to be successful. If your vCenter or SRM server at the protected site is down, but you are not sure if the protected VMs are still operational, you should verify this before executing an SRM recovery plan.

Table 15) Planned or unplanned failover prerequisites.

Planned or Unplanned Failover Prerequisites
The SRM and NetApp environment has been set up properly and all SRM recovery plans have been tested successfully.
During recovery plan testing, the functionality of the applications and services inside the test environment was verified.
Connectivity to the DR network is available from wherever the DR operations will be performed.
Systems administrators have access to a workstation or server desktop session from which to administer the DR site and perform the failover.
Systems administrators have all appropriate credentials, accounts, passwords, and so on required to access the systems.
Certain infrastructure servers already exist in the DR site and are accessible. These systems provide basic services necessary for the administrators to work in the environment and execute the recovery plan, for example: • The DR site vCenter and SRM servers • DR site Active Directory services • DR site DNS services • DR site VMware license services • System time synchronization (The systems at the recovery site have time synchronization configured to the same source or a source that was in sync with the primary site.)
All required NetApp volumes are being replicated to the DR site using SnapMirror.
The SnapMirror operations have been monitored and are up to date with respect to the designed RPO.
Required capacity exists on the DR NetApp controller. This refers to capacity required to support day-to-day operations that will now occur in the DR environment.
All DR ESX hosts already have iSCSI/FC sessions established with the DR storage arrays (physical arrays or vFiler units).

Planned or Unplanned Failover Prerequisites
Plans have been made for providing users access to the applications and services at the DR site.
Some method exists to isolate the failed primary network from the DR site. This might be necessary if the event causing the disaster were temporary or intermittent in nature, such as an extended power outage. If power should be reapplied to the primary site, restarting systems and services might conflict with the recovered operations that are now running at the DR site.

To execute an SRM failover, perform the following steps.

1. In the SRM interface, click the Recovery Plans tab in the bottom left and then select the recovery plan you want to execute.

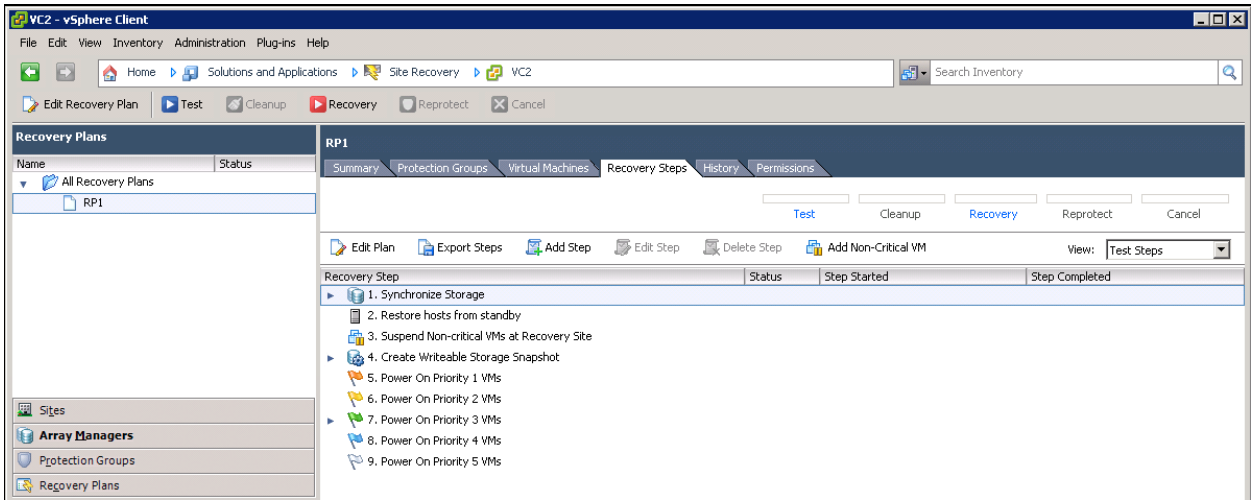

Note: When accessing the SRM interface, you might be prompted to enter administrative credentials for SRM to connect to the other site. If in an unplanned failover scenario, this authentication is expected to fail if the vCenter or SRM server at the other site is unavailable.

2. Click the Recovery link in the upper right to begin a failover.

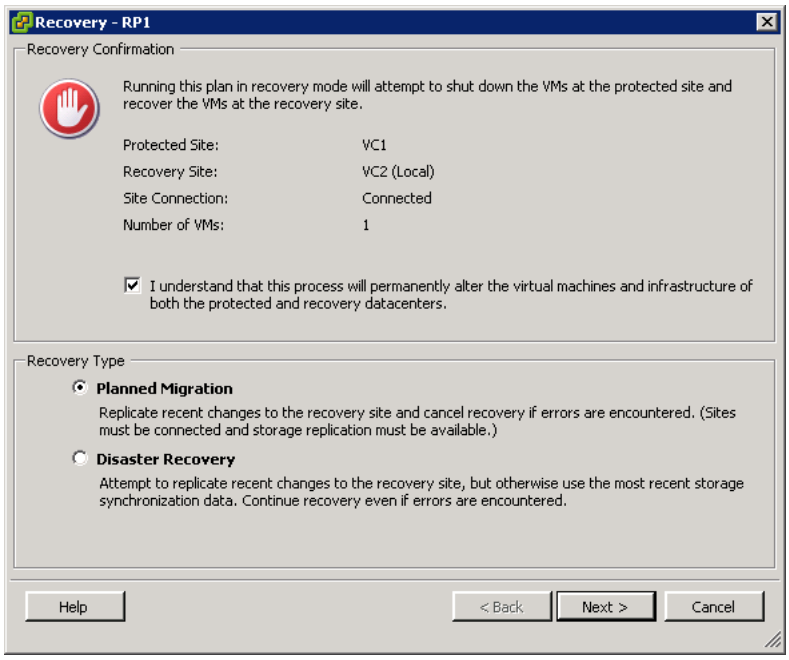

Note: To perform a planned failover, select the Planned Migration button under Recovery Type.

Note: To perform an unplanned failover to recover from a disaster, select the Disaster Recovery button under Recovery Type. Before executing the recovery plan in unplanned failover mode, the administrative team should take predetermined steps to verify that a disaster has actually occurred and it is necessary to run the plan. After the disaster is confirmed, the team can continue with execution of the recovery plan.

3. Click Next and review the details of the recovery plan execution on the next screen and then click Start to execute the recovery plan.

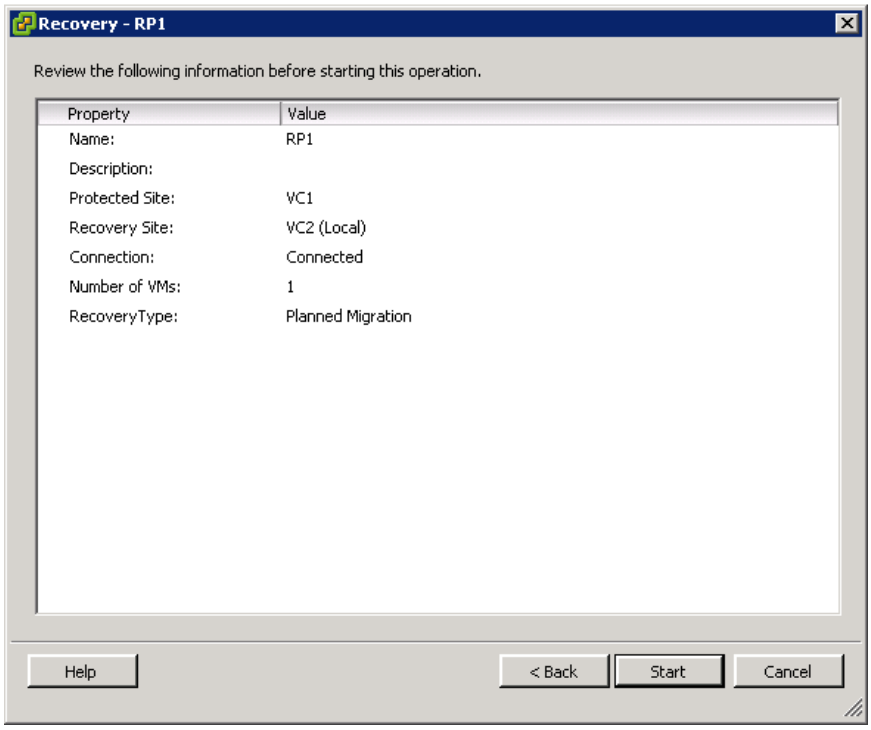

4. Click the Recovery Steps tab to view the progress of the recovery plan as it is running. SRM breaks the SnapMirror relationship, making the DR site volumes writable, and mounts the datastores to the ESX hosts. The virtual machines are recovered and begin booting.

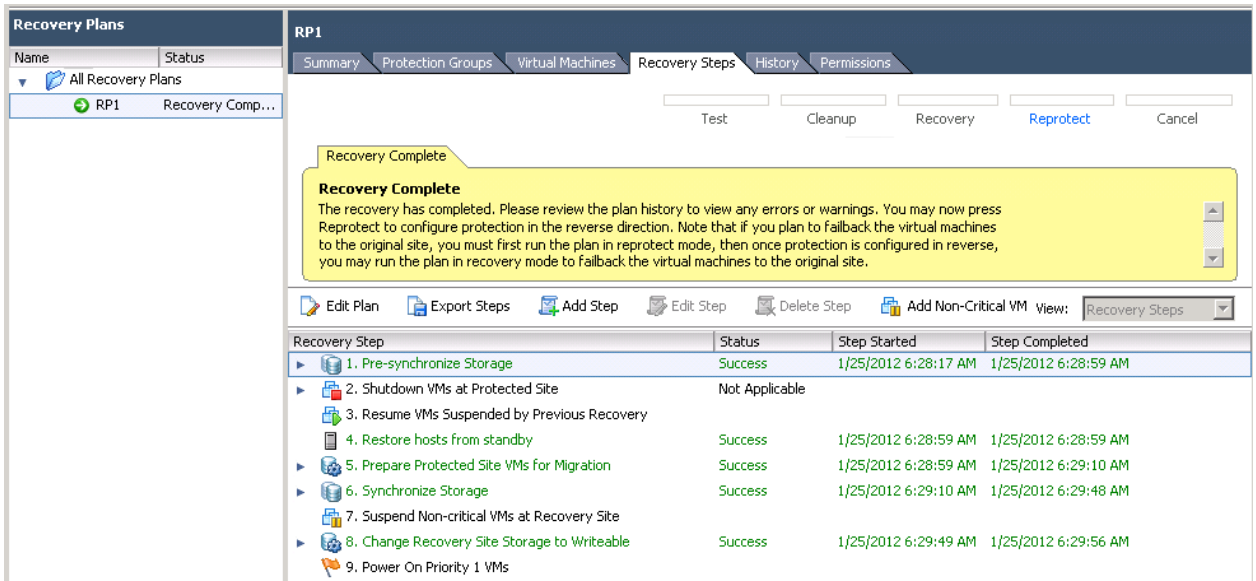

Note: If you selected Planned Migration, SRM will perform a storage replication update. SRM will remain in the storage synchronization steps until the replication process is complete. SRM uses the SRA to periodically check the status of the replication relationship; once it is complete, the recovery plan will proceed to the next step.

Note: If Disaster Recovery failover is selected, SRM will attempt to shut down the VMs at the protected site; however, this is expected to fail if the vCenter Server, SRM server, or VMs at that site are inaccessible.

Note: During a real failover, SRM will mount NFS datastores in the same way that it mounts datastores during a failover test by alternating datastore mounts on IP addresses that were entered into the NFS IP Addresses field in the Array Setup dialog.

Note: It is important to make sure that a server with the necessary Active Directory FSMO roles is available at the DR site. For information regarding transferring roles or seizing these roles, refer to Microsoft KB 255504.

5. Connect to some of the recovered VMs and verify the recovery, or notify applications administrators that their applications have been restored and should be verified.
6. Continue with the steps necessary to complete a disaster recovery for business-critical applications and services.
7. If possible, isolate the primary site to prevent any conflicts if the infrastructure services should be reestablished there without warning.

3.3 Reversing Replication and SRM Failback

There is no workflow in SRM 5 called “failback." Instead a complete failback is performed by using a combination of two different workflows in SRM: the reprotect workflow and the recovery workflow.

To automate failback with SRM 5, first execute the reprotect workflow. During the reprotect workflow, SRM will use the SRA to automate the reversal of the SnapMirror relationships and reverse the roles of the two sites in the recovery plan; what was the recovery site will become the protected site, and what was the protected site will become the recovery site for the VMs and datastores in the recovery plan. Then to complete the failback, execute the recovery workflow just as you would for a planned failover event, except now the failover is performed in the other direction back to the original protected site.

As shown in Figure 20, after an initial failover event in SRM 5, further failback and failover operations are a cycle of the recovery, reprotect, and then recovery and reprotect workflows in SRM to manage failover and failback between two environments. Failover testing should be performed before recovery back to the primary site, just as you would do so before you would depend on the processes for failover in a DR scenario.

Figure 20) Cycle of failover failback between two sites.

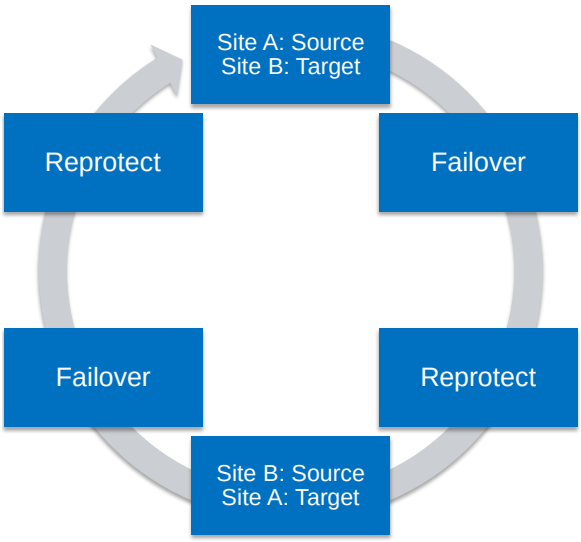

Table 16) Reprotect operation prerequisites.

Reprotect Operation Prerequisites
If the failover was of an unplanned nature, storage at the original protected site might need to be brought back online.
ESX or ESXi hosts must be powered on at the original protected site. If the failover was of an unplanned nature, care should be taken to make sure that the site is fenced from the network as ESX hosts might still have datastore connections and might be configured to automatically start booting VMs.
If the failover was of an unplanned nature, then the vCenter Server and SRM server at the original protected site might need to be powered back on.

To perform the reprotect operation, follow these steps.

Reversing replication might require certain versions of Data ONTAP depending on the direction of replication and the mode of SnapMirror used. For more information about Data ONTAP version requirements for SnapMirror, see section 1.14, “SnapMirror and Data ONTAP Version Requirements.”

1. In the SRM interface, click the Sites tab in the lower left and select either of the two sites between which you will be performing the reprotect.
2. Click the Summary tab. The SRM interface should show the site status as connected.

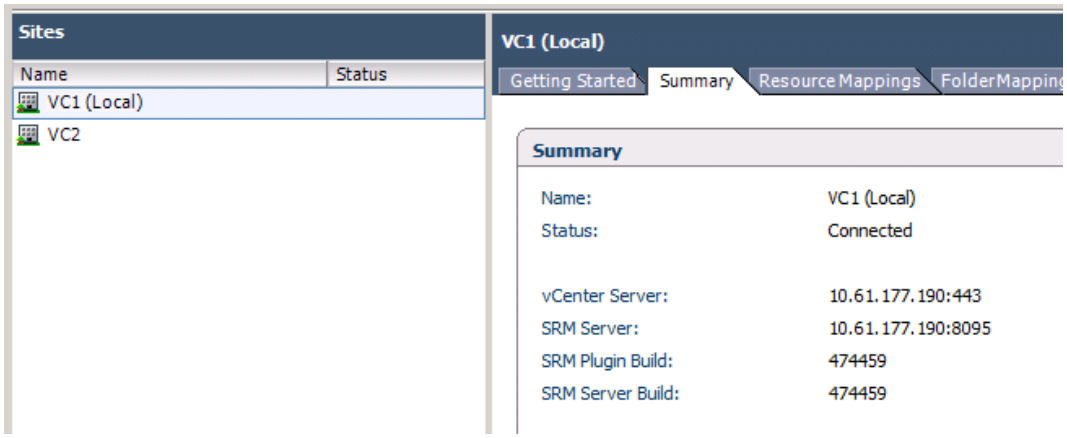

Note: When accessing the SRM interface, you might be prompted to enter administrative credentials for SRM to connect to the other site.

3. In the SRM interface, click the Recovery Plans tab in the bottom left and select the recovery plan you want to execute.

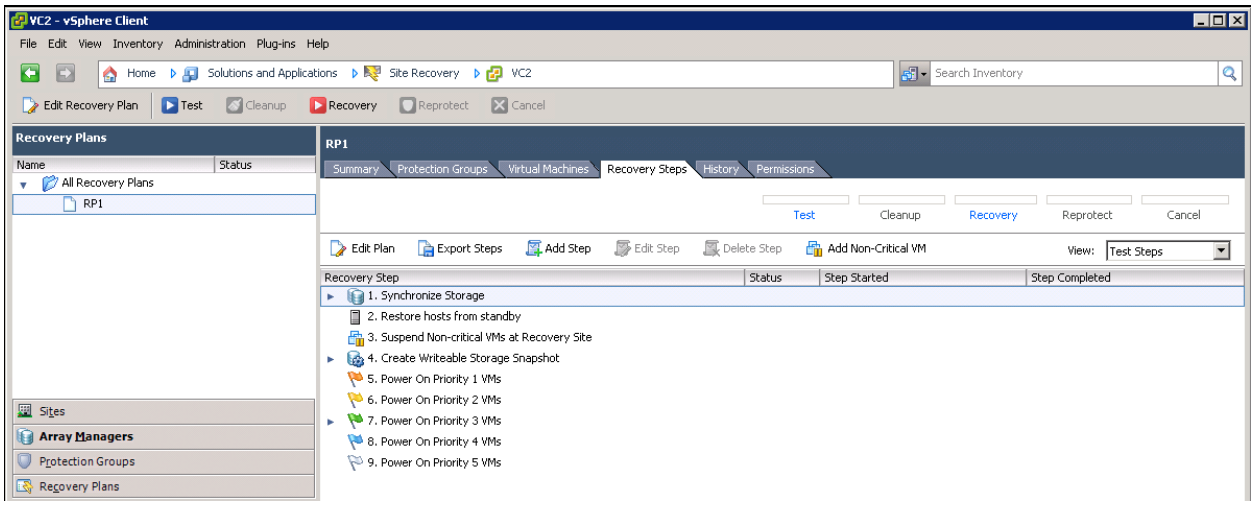

4. Select the Recovery link in the upper right to begin the reprotect operation.
5. Select the box indicating that you understand this operation cannot be undone and then click Next.

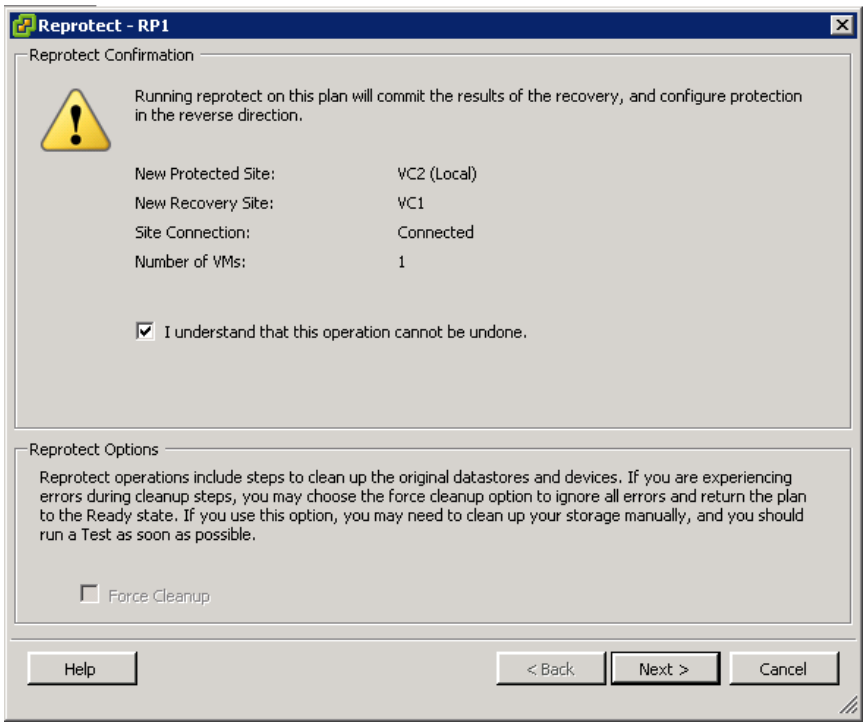

6. Review the details of the reprotect operation and click Start.

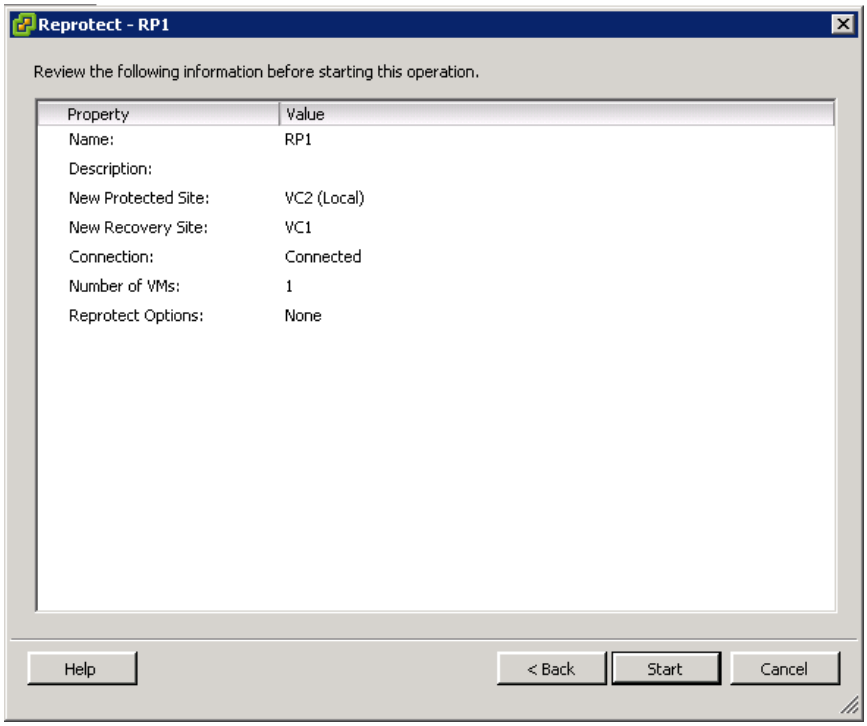

7. Click the Recovery Steps tab to review the progress of the reprotect operation.

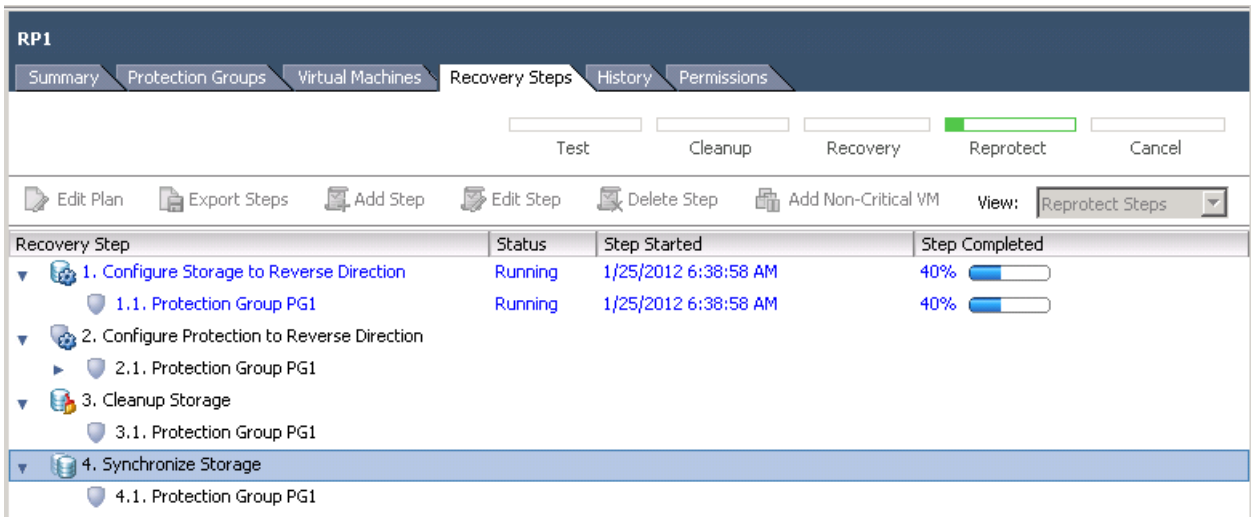

8. After the reprotect operation is complete, click the Array Managers tab at the bottom left and select the array manager for the array at the original protected site.
9. Click the Devices tab. The direction column should show that replication has been reversed for all datastores in the reversed recovery plan.

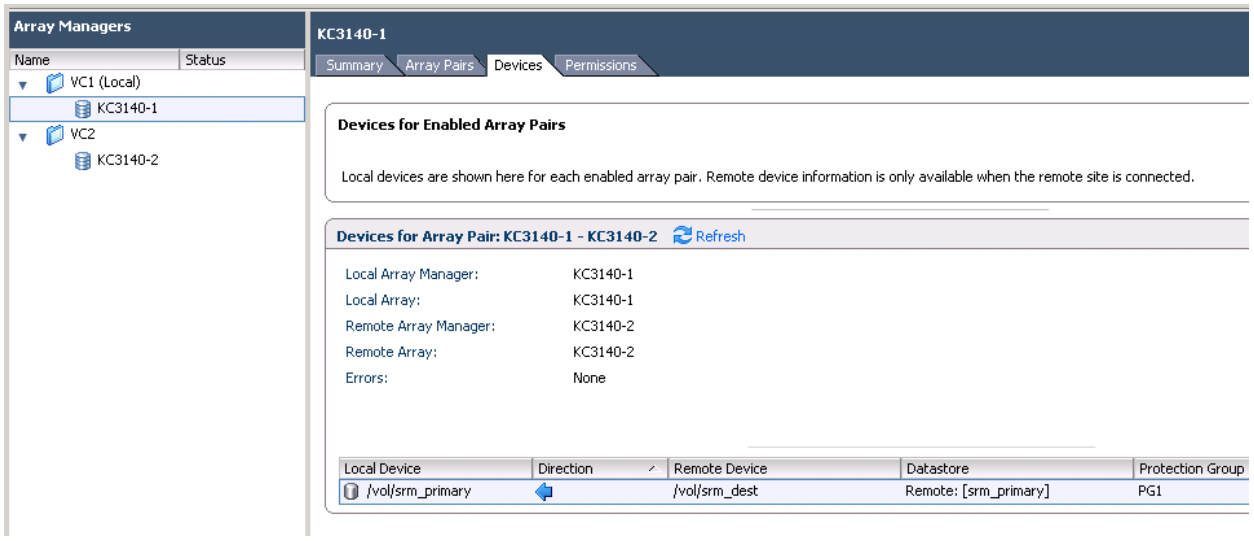

Note: If you are using SnapMirror network compression or a custom SnapMirror TCP window size setting, and you want these settings to be used for replication in the reversed direction, you must apply these two settings manually at the original protected site in the SnapMirror configuration file on the NetApp array.

Snapshot Cleanup after SRM Reprotect

SnapMirror maintains Snapshot copies at the volume level in the source and destination volume to maintain replication relationships and to be able to reverse and resynchronize relationships. When SRA reverses replication, it does not remove the Snapshot copies that were used to maintain the replication in the previous direction. In environments using SRA 2.0, these Snapshot copies must be removed manually after running the reprotect workflow. These Snapshot copies will not interrupt or prevent reversed replication and do not have to be removed immediately after the reprotect workflow is executed; however, they should be removed soon as they will lock some amount of capacity in the volume.

In the example below, the reprotect operation was performed and replication now occurs from controllerB to controllerA.

To remove Snapshot copies that were used to maintain replication in the opposite direction, follow these steps for each volume or qtree affected by the reprotect operation.

1. After the reprotect operation has completed successfully, connect to the console of current destination controller or vFiler unit. In this example, controllerA is the current destination after the reprotect operation.
2. On the current destination, list the Snapshot copies for the volume affected (if you are using qtree SnapMirror, the Snapshot copies occur at the volume level).

```
controllerA> snap list nfs02
Volume nfs02
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 15 18:46  controllerA(0135018611)_nfs02.3
  0% ( 0%)    0% ( 0%)  May 15 18:44  controllerA(0135018611)_nfs02.2
  4% ( 4%)    1% ( 1%)  May 06 13:16  controllerB(0118051885)_nfs02.6 (snapmirror)
  4% ( 0%)    1% ( 0%)  May 06 13:14  controllerB(0118051885)_nfs02.5
```

Note: In the example above, the Snapshot copies named `controllerB(0118051885)_nfs02.n` were used for replication in the previous direction prior to failover. The `(snapmirror)` text is a tag that indicates the Snapshot copy is still locked by the previous relationship.

3. Remove the (`snapmirror`) lock on current destination volume using the `snapmirror release` command.

```
current_destination_array> snapmirror release vol_name current_source_array:vol_name
```

Example:

```
controllerA> snapmirror release nfs02 controllerB:nfs02
```

4. The `snap list` command will now show the (`snapmirror`) lock removed.

```
controllerA> snap list nfs02
Volume nfs02
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 15 18:46  controllerA(0135018611)_nfs02.3
  0% ( 0%)    0% ( 0%)  May 15 18:44  controllerA(0135018611)_nfs02.2
  4% ( 4%)    1% ( 1%)  May 06 13:16  controllerB(0118051885)_nfs02.6
  4% ( 0%)    1% ( 0%)  May 06 13:14  controllerB(0118051885)_nfs02.5
```

Note: In the preceding example, `current_source_array` is the name of the controller or vFiler unit that is the source of the SnapMirror relationship following the reprotect operation. This command unlocks the Snapshot copies (on the array where the release command runs) that were used to replicate data before replication was reversed.

5. Connect to the console of the current source controller or vFiler unit. In this example controllerB is the current source after the reprotect operation.
6. List the Snapshot copies on the volume affected.

```
current_source_array> snap list vol_name
```

Example:

```
controllerB*> snap list nfs02
Volume nfs02
working...

  %/used       %/total  date          name
```

```
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 15 18:46  controllerA(0135018611)_nfs02.3 (snapmirror)
  4% ( 4%)    1% ( 1%)  May 06 13:16  controllerB(0118051885)_nfs02.6 (snapmirror)
  4% ( 0%)    1% ( 0%)  May 06 13:14  controllerB(0118051885)_nfs02.5
```

Note: SnapMirror created Snapshot copies include the name of the destination controller. You can safely delete all the Snapshot copies that contain the name of the current source controller from the affected volume on the current destination array, as SnapMirror no longer uses these Snapshot copies.

7. For each volume or qtree affected by the reprotect operation, delete only the Snapshot copies that contain the name of the current source controller or vFiler unit.

```
current_source_array> snap delete vol_name snapshot_name
```

Example:

```
controllerB*> snap delete nfs02 controllerB(0118051885)_nfs02.6
Tue May 15 18:57:48 EDT [wafl.snap.delete:info]: Snapshot copy controllerB(0118051885)_nfs02.6 on
volume nfs02 NetApp was deleted by the Data ONTAP function snapcmd_delete. The unique ID for this
Snapshot copy is (151, 1604672).

controllerB*> snap delete nfs02 controllerB(0118051885)_nfs02.5
Tue May 15 18:58:00 EDT [wafl.snap.delete:info]: Snapshot copy controllerB(0118051885)_nfs02.5 on
volume nfs02 NetApp was deleted by the Data ONTAP function snapcmd_delete. The unique ID for this
Snapshot copy is (150, 1604652).

controllerB*> snap list nfs02
Volume nfs02
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 15 18:46  controllerA(0135018611)_nfs02.3 (snapmirror)
```

8. The next scheduled SnapMirror update will automatically propagate the Snapshot copy deletion to the current destination. Alteratively, you can perform an update of the SnapMirror relationship using the `snapmirror update` command on the destination controller.

Example:

```
controllerA> snapmirror update nfs02
Transfer started.
Monitor progress with 'snapmirror status' or the snapmirror log.

controllerA> snapmirror status
Snapmirror is on.
Source                  Destination             State          Lag        Status
controllerB:nfs02   controllerA:nfs02   Snapmirrored   00:00:24   Idle
```

Continuing SRM Failback

To fail back to the original protected site and reestablish replication in the original direction, do as follows:

1. Perform a test failover of the recovery plan following the process described in section 3.1, “Perform a Test Failover.”
2. Schedule an outage during which you can perform a planned failover to migrate operations back to the original protected site by using the process described in section 3.2, “Perform a Planned or Unplanned Failover.”
3. Run the reprotect process again to reestablish replication in the original direction.

Appendix

Configuring SRM and NetApp SRA to Use HTTP/SSL

NetApp SRA communicates with the NetApp controllers and vFiler units by using Data ONTAP API calls. The SRA never communicates with the controller by SSH or telnet. By default, all connections made by the SRA are made by using HTTP. If using physical NetApp controllers as arrays in SRM, you can configure the NetApp SRA to use SSL (HTTPS) to connect to the arrays.

Note: NetApp MultiStore vFiler units do not support SSL (HTTPS) connections.

This section describes the procedure to enable SSL communication between the SRA and a NetApp storage controller. The NetApp adapter is developed in Perl. To use SSL with the NetApp adapter requires that two additional modules be installed into the instance of Perl that is distributed with SRM.

Secure communication using SSL is a global option set per SRA instance. Once enabled on the SRA, it must be enabled on each storage controller with which SRA communicates. This procedure must be performed on each SRM server that is required to use SSL to communicate with a NetApp controller.

To enable SSL (HTTPS) in SRM and the NetApp SRA, perform the following steps.

1. In the SRM server, open a command window.
2. Change the working directory to the location of the Perl binaries included with the SRM distribution.

```
cd <system_drive>\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\external\perl-5.8.8\bin
```

Note: The preceding is the default location for the Perl installation in SRM. If SRM is installed in a different location in your environment, change to the appropriate directory.

3. Launch the Perl package manager.

```
ppm
```

4. From the PPM > prompt, add the Bribes package repository.

```
ppm> rep add Bribes http://www.bribes.org/perl/ppm
```

PPM indicates that the new repository has been installed.

```
Repositories:
[1] ActiveState Package Repository
[2] Bribes
```

5. Install the “Tree-DAG_Node” Perl module.

```
ppm> install “Tree-DAG_Node”
```

PPM indicates that the module has been installed.

```
Successfully installed Tree-DAG_Node version 1.05 in ActivePerl 5.8.8.817.
```

6. Install the “Net-SSLeay” Perl module.

```
ppm> install “Net-SSLeay”
```

PPM indicates that the module has been installed.

```
Successfully installed Net-SSLeay version 1.41 in ActivePerl 5.8.8.817.
```

7. Type `exit` to leave PPM and then close the command window.

```
exit
```

8. Open a console connection to the NetApp controller and, if not already configured, enable SSL on the NetApp storage controller.

```
secureadmin setup ssl
```

```
secureadmin enable ssl
```

9. Verify that SSL is turned on.

```
options ssl.enable
```

10. Using a text editor on the SRM server, edit the SRA configuration file and change the default option “ssl = off” to “ssl = on." The default location for the SRA configuration file in SRM 5 is:

```
<install drive>:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\
storage\sra\ONTAP\ontap_config.txt
```

11. Verify that secure communication between the SRA and the storage controller works by refreshing the device list for the array pair within the SRM management interface.

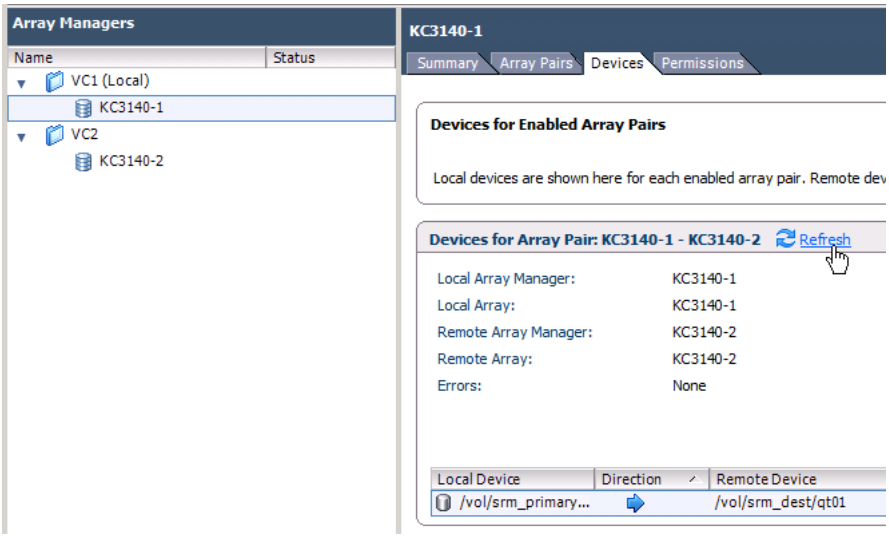

12. To verify SSL is being used, see the Data ONTAP audit logs located in the NetApp controller in the `/etc/log/auditlog` folder. API calls from the SRA are logged as using 'https', whereas non-SSL API calls use HTTP.

SnapMirror Definition Using IP or FQDN

It is possible to define SnapMirror relationships by using the host name, IP address, or fully qualified domain name (FQDN) of the source controller. Typically, a host name definition is required for use with SRM. However, there might be cases where this is not possible, such as when directing SnapMirror to use a specific interface for replication. To resolve issues with SRM configuration when using IP or FQDN, do the following.

1. Where possible, edit the snapmirror.conf file on the destination controller and change the relationship definition so the source name in the definition matches the host name of the source controller.

   Where this is not possible, SRA provides a means to map IP addresses and FQDNs to a host name through the use of the ip_hostname_mapping.txt file.

2. Edit the ip_hostname_mapping.txt file, which is located in the <install drive>:\Program Files(x86)\VMware\VMware vCenter Site Recovery Manager\storage\sra\ONTAP directory, so that its entries map a host name to its corresponding IP address or FQDN.

   When using FQDNs, the file should take on the following format:

```
hostA = <hostA FQDN>
hostB = <hostB FQDN>
```

   Where <hostA FQDN> is the FQDN of that controller, for example:

fas3070-01 = fas3070-01.netapp.com
fas2020-01 = fas2020-01.netapp.com

When using IP addresses, the file should take on the following format:

```
hostA = <hostA IP>
hostB = <hostB IP>
```

Where <hostA IP> is the FQDN of that controller, for example:

```
fas3070-01 = 10.10.10.10
fas2020-01 = 10.10.10.20
```

3. Once the entries have been made in the ip_hostname_mapping.txt file, an option must be set in the ontap_config.txt file, located in the same directory to enable the IP to host name mapping. Open the file and edit the following line to enable the option:

```
use_ip_for_snapmirror_relation = on
```

Note: It is important to note that entries in the ip_hostname_mapping.txt file and the use_ip_for_snapmirror_relation option are global options and apply for all relationships between the two controllers. All relationships defined for a given source-destination controller pair must be defined in the same way. It is not supported to mix short host name, FQDN, and IP address definitions.

SnapMirror Definition Using SnapMirror Connection Names

It is also possible to define a SnapMirror relationship by using a connection name, which is an alias that can be created in a snapmirror.conf file to refer to one or two possible paths by a single name. To work properly with SRM, the connection name in the destination snapmirror.conf file must be made to match the host name of the source controller.

In this example, volume nfs01 is replicated by using SnapMirror from controller fas3070-01 to controller fas2020-01 by using a connection name of "fas3070-01_conn" to allow the use of SnapMirror multipathing.

SnapMirror status is shown in the SnapMirror destination as follows:

```
fas2020-01> snapmirror status
Source             Destination        State          Lag        Status
fas3070-01_conn:nfs01   fas2020-01:nfs01   Snapmirrored   00:07:41   Idle
```

In the `snapmirror status` command output earlier, the source controller is identified as "fas3070-01_conn," which is a connection name specified in the snapmirror.conf file on the destination controller as shown in the following example. This snapmirror.conf file configuration will not work with SRM.

Destination controller snapmirror.conf file that does not support SRM:

```
fas3070-01_conn=multi(10.10.10.20,10.10.10.10)(192.168.10.20,192.168.10.10)
fas3070-01_conn:nfs01 fas2020-01:nfs01 - 0 22 * *
```

To be supported with SRM, the connection name must be changed to match the host name of the source controller as in the following example.

Destination controller snapmirror.conf file that supports SRM:

```
fas3070-01=multi(10.10.10.20,10.10.10.10)(192.168.10.20,192.168.10.10)
fas3070-01:nfs01 fas2020-01:nfs01 - 0 22 * *
```

In the preceding example, the connection name entry and the source controller name on the SnapMirror relationship have been changed from "`fas3070-01_conn`" to "`fas3070-01`" to match the source controller host name.

When the SRA reverses this relationship during the SRM reprotect workflow, SRA will build a connection name for the reversed relationship in the other controller.

Note: Using this configuration to support SRM implies that all SnapMirror relationships between this source controller and this destination controller will use the multipath connection as specified by the connection name that is defined. If you have other SnapMirror relationships between these two controllers that are not intended to be used with VMware SRM, you may configure a different connection name for those relationships to use.

The SnapMirror process will reread the snapmirror.conf file every minute. After the file has been changed and SnapMirror has read the new file, the `snapmirror status` command output should automatically change to show the host name as the source controller name.

SnapMirror status as shown on the destination with connection name matching source host name:

```
fas2020-01> snapmirror status
Source              Destination         State          Lag        Status
fas3070-01:nfs01    fas2020-01:nfs01    Snapmirrored   00:07:41   Idle
```

You can run a `snapmirror update` command for the relationship to confirm the update will work and then SRM will be able to detect the relationship.

Version History

Date	Version	Description	Author
May, 2012	1.0	Initial release	Larry Touchette, Julian Cates
June, 2012	1.1	Clarified co-existence of SRM and SnapVault in same environment. Updated procedures for Snapshot cleanup after SRM reprotect. Added mention of Snapshot auto-delete configuration.	Larry Touchette, Julian Cates

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com