VMware vCenter Site Recovery Manager™ 5.0

Evaluation Guide

TECHNICAL WHITE PAPER

Table of Contents

Getting Started

About VMware vCenter Site Recovery Manager

VMware vCenter Site Recovery Manager™ 5.0 (SRM) is an extension to VMware vCenter™ that provides disaster recovery capabilities to VMware customers.

SRM enables integration with array-based replication, as well as the use of a native VMware vSphere®–based replication engine, discovery and management of replicated datastores, automated migration of inventory vCenter environments, automated reprotection, and failback of environments.

SRM servers coordinate the operations of the VMware vCenter Server™ at two sites, so that as virtual machines at one site (the protected site) are shut down, copies of these virtual machines at the other site (the recovery site) start up and, using the data replicated from the protected site, assume responsibility for providing the same services.

Migration of protected inventory and services from one site to the other is controlled by a recovery plan that specifies the order in which virtual machines are shut down and started up, the resource pools to which they are allocated, and the networks they can access. SRM enables the testing of recovery plans, using a temporary copy of the replicated data, in a way that does not disrupt ongoing operations at either site.

SRM runs in conjunction with the VMware vSphere® 5.0 ("vSphere") platform, extending the feature set of the virtual infrastructure platform to provide for rapid business continuity through partial or complete site failures.

About This Guide

The purpose of this guide is to support a self-guided, hands-on evaluation of SRM by IT professionals who are looking to automate their disaster recovery plans with SRM in their vSphere environment.

The *VMware vCenter Site Recovery Manager 5.0 Evaluation Guide* is intended to provide SRM customers and evaluators a guide that walks them through the SRM workflow that must be completed to allow for the successful and automated service failover from the designated SRM protected site to the designated SRM recovery site.

SRM 5.0 introduces a replication engine, independent of traditional array-based replication, called vSphere Replication. vSphere Replication provides a means of duplicating virtual machines between sites and does not require the use of traditional array-based data copying. This guide is designed to illustrate the use of both standard array-based replication and vSphere Replication, although evaluators may choose to use either or both of these methods as part of the assessment, as is appropriate to the requirements of the evaluation.

This guide also provides an overview that includes the considerations and guidance to execute a failback of services from the recovery site back to the site that was originally designated as the SRM protected site. Evaluators can work through the exercises provided in this guide to gain firsthand experience operating the core and new features.

Assumptions

To successfully use this guide, the following is **assumed:**

- VMware ESX®/ESXi™ has been installed on the physical servers designated for this evaluation.
- vCenter Server 5.0 and VMware vSphere® Client™ 5.0 have been installed at each of the SRM protected and recovery sites to manage the ESX hosts.
- A SAN/NFS infrastructure is in place, and set up to replicate designated VMware vSphere® VMFS/NFS datastores between the SRM protected and recovery sites to use array-based replication. This is not a requirement if only vSphere Replication is chosen for evaluation.

- If vSphere Replication is chosen as the replication engine for this assessment, there is no requirement that hardware storage arrays are used. A local disk or even the VMware Virtual Storage Array may be used for the purposes of evaluation, and there is no requirement for an array with a licensed storage-based replication engine.
- The virtual machines that have been selected for protection with array replication for the SRM evaluation have been moved onto the designated replicated datastores. Virtual machines that have not been selected to be array based replication–protected virtual machines for the evaluation should be moved to nonreplicated datastores.
- If vSphere Replication will be evaluated, any virtual machine on any accessible storage may be used. Regardless of this, vSphere Replication–protected virtual machines should not normally reside on a replicated datastore, in order to avoid multiple replications of the same virtual machine.
- Moreover, when using vSphere Replication for evaluation purposes, there is no requirement for multiple physical sites. Customers may choose to base their evaluation on failover between clusters, rather than between sites, to emulate the usage that would occur in production between physical sites.
- If vSphere Replication will be evaluated, a unique database must be provisioned at each site for use by the vSphere Replication management service. This guide will assume that Microsoft SQL Server is being used for the database, and that native SQL permissions will be used for authentication and for access to the database. Database setup and configuration will not be covered in this evaluation guide. Each site must have a separate database configured and reserved for use by the vSphere Replication management service.
- If vSphere Replication will be evaluated, the vCenter Extension vService Dependency must be configured on both vCenter server instances. This is accessible through the vCenter Runtime Settings on each vCenter server, and will be set by configuring the Managed IP Address in the Runtime Settings. For more details, see *http://kb.vmware.com/kb/1008030*.
- The basic installation of the SRM Server in both the protected and recovery sites has been completed. For assistance installing SRM, refer to the VMware vCenter Site Recovery Manager documentation for both administration and installation available at *http://www.vmware.com/support/pubs/srm_pubs.html*.
- Storage Replication Adapters (SRAs) have been installed at protected and recovery sites in case array-based replication is to be used.
- The VMware® Site Recovery Manager™ Plug-In (SRM plug-in) has been installed and enabled on the vSphere Client instances that will be used to access the SRM protected and recovery sites.

For detailed information regarding installation, configuration, administration, and usage of vSphere and SRM, refer to the following online documentation:

- vSphere – *http://www.vmware.com/support/pubs/vs_pubs.html*
- SRM – *https://www.vmware.com/support/pubs/srm_pubs.html*

Abbreviations and Terminology

The following disaster recovery (DR), vSphere, and vCenter abbreviations are used throughout this evaluation guide:

ABBREVIATION	DESCRIPTION
ABR	Array-based replication
BC/DR	Business continuity and disaster recovery
VM	Virtual machine on a managed host
VRP	vCenter resource pool
RP	Recovery plan
RPO	Recovery point objective

ABBREVIATION	DESCRIPTION
RTO	Recovery time objective
PG	Protection group
VMFS	Virtual Machine File System
SAN	Storage area network–type datastore shared between managed hosts
VR	vSphere Replication
VRA	vSphere Replication agent
VRMS	vSphere Replication Management Server
VRS	vSphere Replication server
NFS	Network File System

The following DR and SRM **terminology** is used throughout this guide:

DR AND SRM TERMINOLOGY	DESCRIPTION
Array-based replication (ABR)	Replication of virtual machines that is managed and executed by the storage subsystem itself, rather than from inside the virtual machines, the vmkernel or the Service Console.
vSphere Replication	Native software-based replication engine built-in to ESXi 5.0 that can be used to provide replication of virtual machines via SRM.
Logical unit number (LUN)	A single SCSI storage device on the SAN that can be mapped to one or more vSphere hosts.
Failover	Event that occurs when the recovery site takes over operation in place of the protected site after the declaration of a disaster.
Failback	Reversal of failover, returning IT operations to the primary site.
Reprotect	Reversal of direction of replication, and automatic reprotection of protection groups.
Datastore	Storage unit of a managed vSphere host.
Host	vCenter-managed vSphere hosts.
SRM Server	Short form for VMware vCenter Site Recovery Manager™ Server. SRM Server extends vCenter Server to provide disaster recovery capabilities for VMware customers. It enables integration with array-based replication, discovery and management of replicated datastores, and automated migration of VMware inventory from one vCenter server to another.
Protected VM	A VM that is protected by SRM.
Unprotected VM	A VM that is not protected by SRM.
Protected site	The site that initially contains the protected VMs.
Recovery site	The site to which virtual machines will fail over.
Datastore group	Replicated datastores containing complete sets of protected VMs.
Protection group	A group of VMs that will be failed over together to the recovery site during testing or recovery.
Storage Replication Adapter (SRA)	Enables SRM to interact with a storage array replication engine.

DR AND SRM TERMINOLOGY	DESCRIPTION
Placeholder VM	An object found with other VMs in the recovery site vCenter inventory representing a protected site VM that is being replicated to the recovery site. It is represented with an icon showing a lightning bolt.
Recovery point objective	The maximum acceptable amount of data that can be lost during a failure, expressed as a time value. For example, an RPO of four hours indicates that up to four hours worth of data loss are acceptable before a return to an operational state.
Recovery time objective	The maximum acceptable amount of time that a service or services of a datacenter may be nonfunctional during a failure, expressed as a time value. For example, an RTO of 12 hours indicates it is acceptable that up to 12 hours might pass before a service might be restored.
Inventory mappings	Associations between resource pools, virtual machine folders, networks at the protected site and their destination counterparts at the recovery site.
Recovery plan	The complete set of steps needed to recover (or test recovery of) the protected VMs in one or more protection groups.

What Will Be Covered

This guide provides the following overview of the SRM features and capabilities:

CATEGORY	FEATURES	WHAT WILL BE COVERED	TIME ESTIMATES[1]
SRM Recovery Workflow	Recovery workflow automation	Setting up SRM Recovery Workflow 1. Setting up site-pairing. 2. Setting up array managers for the replicated datastores. 3. Setting up inventory mappings. 4. Setting up protection groups. 5. Setting up recovery plans. 6. Configuring IP customization. 7. Triggering a test recovery.	60 minutes
Deploying vSphere Replication	Deploying vSphere Replication components	1. Deploying VRMSs. 2. Deploying a VRS. 3. Pairing VRMSs. 4. Registering the VRS. 5. Configuring protection/replication for an individual virtual machine.	90 minutes
SRM alarms	Configuring action for an SRM alarm	Configuring action for the Remote Site Down alarm 1. Configuring alarm action to send out notification email.	10 minutes
SRM failover from protected site to recovery site (optional)	Failover	Reading details of failover operations (exercise is optional).	30 minutes

CATEGORY	FEATURES	WHAT WILL BE COVERED	TIME ESTIMATES[1]
SRM failover from recovery site to protected site (optional)	Reprotect/failback	Reading details of reprotect and failback operations (exercise is optional).	90 minutes

1. The real time spent on each exercise is dependent on the specifics of your environment.

Steps

It is highly recommended that you work through the exercises in these sections to experience the SRM features and capabilities firsthand. For failover, reprotection, and failback, you can simply read through the details provided previously in the corresponding sections listed if you decide not to go through the real operations.

After you have successfully installed the vSphere and SRM software components in your environment, you can proceed to perform the evaluation of SRM. For each scenario, you can use the corresponding checklist to ensure that you are following the proper sequence.

Checklist

You can use the following worksheet to organize your evaluation process.

HARDWARE CHECKLIST	
Physical servers	

SOFTWARE CHECKLIST	
ESX/ESXi Server	
vCenter Server (and associated database)	
vSphere Client	
SRM Server (and associated database)	
vSphere Replication database and user ID	
SRM plug-in	
SRA – storage-vendor specific	

EVALUATION EXERCISES	
SRM Recovery Workflow	
vSphere Replication	
SRM Alarms and Site Status Monitoring	
Failover from Protected Site to Recovery Site (optional)	
Failback from Recovery Site to Protected Site (optional)	

Documentation

This guide is intended to provide an overview of the steps required to ensure a successful evaluation of SRM. It is not meant to substitute for product documentation. Refer to the online product documentation for SRM for more detailed information. (See the following links.) You might also consult the online knowledge base if you have any additional questions. If you require further assistance, contact a VMware sales representative or channel partner.

VMware vCenter Site Recovery Manager Resources

Product resources: *http://www.vmware.com/products/srm/resource.html*

Product community: *http://www.vmware.com/products/srm/community.html*

Site Recovery Manager Administration Guide: *http://www.vmware.com/pdf/srm_admin_5_0.pdf*

Product documentation: *http://www.vmware.com/support/pubs/*

Online support: *http://www.vmware.com/support*

Support offerings: *http://www.vmware.com/support/services*

Education services: *http://mylearn1.vmware.com/mgrreg/index.cfm*

Support knowledge base: *http://kb.vmware.com*

VMware Uptime Blog: *http://blogs.vmware.com/uptime/*

VMware Contact Information

For additional information, or to purchase vSphere and VMware vCenter Site Recovery Manager, VMware's global network of solution providers are ready to assist you.

If you would like to contact VMware directly, you can reach a sales representative at 1-877-4VMWARE (650-475-5000 outside North America) or email *sales@vmware.com.* When emailing, include the state, country, and company name from which you are inquiring.

You can also visit *http://www.vmware.com/vmwarestore/*.

Providing Feedback

Your feedback is appreciated on the material included in this guide. In particular, any guidance on the following topics would be extremely helpful:

- How useful was the information in this guide?
- What other specific topics would you like to see covered?
- Overall, how would you rate this guide?

Send your feedback to the following address: *tmdocfeedback@vmware.com*, with "VMware vCenter Site Recovery Manager 5.0 Evaluation Guide" in the subject line. Thank you for your help in making these guides a valuable resource.

How Does Site Recovery Manager Work?

SRM provides business continuity and disaster recovery protection for virtual environments. Protection can extend from individual replicated virtual machines or datastores to an entire virtual site. The virtualization of the datacenter by VMware offers advantages that can be applied to business continuity and disaster recovery, including the following:

- The entire state of a virtual machine (memory, disk images, I/O, and device state) is encapsulated. Encapsulation enables the state of a virtual machine to be saved to a file. Saving the state of a virtual machine to a file allows the transfer of an entire virtual machine to another host.
- Hardware independence eliminates the need for a complete replication of hardware at the recovery site. Hardware running ESX at one site can provide business continuity and disaster recovery protection for hardware running ESX at another site. This eliminates the cost of purchasing and maintaining a system that sits idle until disaster strikes.
- Hardware independence allows an image of the system at the protected site to boot from the disk at the recovery site in minutes, instead of days.

SRM uses replication between a protected site and a recovery site. The workflow that is built into SRM automatically discovers which datastores or VMs are configured for replication between the protected and recovery sites. SRM can be configured to support bidirectional protection between two sites.

SRM provides protection for the operating systems and applications encapsulated by the virtual machines running on ESX. An SRM server must be installed at the protected site and at the recovery site. The protected and recovery sites must each be managed by their own vCenter server. The SRM Server uses the extensibility of vCenter Server to provide the following:

- Access control
- Authorization
- Event-triggered alarms

Site Recovery Manager Prerequisites

SRM has the following prerequisites:

- A vCenter server installed at the protected site
- A vCenter server installed at the recovery site
- Preconfigured array-based replication between the protected site and the recovery site (if array-based replication will be used)
- Network configuration that allows TCP connectivity between SRM servers and vCenter servers
- An Oracle or Microsoft SQL Server database that uses ODBC for connectivity in the protected site and in the recovery site
- An SRM license installed with a sufficient number of per–virtual machine licenses to cover the systems protected in the evaluation

Site Recovery Manager Configuration and Protection Workflow

The following workflows accomplish setup and configuration for the protected and recovery sites. The SRM user interface is installed as a plug-in into the vSphere Client. SRM uses the vSphere Client as the user interface (UI). The SRM UI is accessed by clicking the **Site Recovery** icon on the vSphere Client Home page (found in the Solutions and Applications menu) and is used for the setup of the SRM workflows, recovery plan testing, and services failover from the protected site to the recovery site.

It is important to complete the workflows in the order they are presented in this guide.

The **recovery site configuration** workflow involves the following activities:

- The user installs the SRM Server.
- The user installs the SRA.
- The user installs the SRM plug-in into the vSphere Client.

The **protected site configuration** workflow involves the following activities:

- The user installs the SRM Server.
- The user installs the SRA.
- If a different vSphere Client is used to access the protected and recovery sites, the user installs the SRM plug-in into the vSphere Client. Otherwise, this activity can be skipped.
- With SRM 5.0, all administration might be done from one instance of the vSphere Client without requiring that activities be performed in different clients connected to the protected and recovery sites.
- The user configures SRM for pairing sites, arrays, and scans for available SRAs.
- Inventory, such as networks, resource pools, and folders, is mapped from the protected site to appropriate and correlated inventory containers at the recovery site.
- SRM identifies available arrays and replicated datastores and determines the datastore groups.

The **protection** workflow involves the following activities:

- Using the inventory mapping interface, the user maps the networks, resource pools, and virtual machine folders in the protected site to their counterparts in the recovery site.
- The user creates protection groups from the datastore groups discovered by SRM.
- For each protected virtual machine, the user can override default values.
- The user creates a recovery plan.
- SRM creates the recovery plan steps.
- Optionally, the user has the ability to customize the recovery plan.

Failover and Testing Workflow

SRM automates many of the tasks required at failover. With the push of one button, SRM does the following:

- Shuts down the protected virtual machines if there is connectivity between sites and they are online.
- Synchronizes any final data changes between sites.
- Suspends data replication and Read/Write enables the replica devices.
- Rescans the ESX servers at the recovery site to find iSCSI and Fibre Channel (FC) devices and mounts replicas of NFS volumes (NFS mounts do not require that the host be scanned to be located).
- Registers the replicated virtual machines.
- Suspends nonessential virtual machines (if specified) at the recovery site, to free up resources for the protected virtual machines being failed over.
- Completes power-up of replicated protected virtual machines in accordance with the recovery plan.
- Provides a report of failover results.
- Offers the user the option to choose to Reprotect the environment.
- Has Reprotect communicate with the SRA to reverse the direction of replication of storage arrays, ensuring that protection groups will then be replicated from the recovery site, where the virtual machines will be running, back to the original primary site.
- Offers the user the option to choose to failback the environment.
- Has Failback execute the same recovery plan that was used to migrate or fail the environment over to the recovery site.
- Since the replication is now reversed, it ensures that a subsequent run of the same failover recovery plan will migrate the environment back to its original location at the first site.

SRM does not require production system downtime to run tests. This means you can test often to ensure that you are protected in case of a disaster. For testing, SRM performs the following tasks:

- Creates a test environment, including network and storage infrastructure, that is isolated from the production environment.
- Rescans the ESX servers at the recovery site to find iSCSI and FC devices, and mounts replicas of NFS volumes (NFS mounts do not require that the host be scanned to be located).
- Registers the replicated virtual machines.
- Suspends nonessential virtual machines (if specified) at the recovery site to free up resources for the protected virtual machines being failed over.
- Completes power-up of replicated protected virtual machines in accordance with the recovery plan.
- Automatically deletes temporary files and resets storage configuration in preparation for a failover or next scheduled SRM test.
- Provides a report of test results.

Multiple vCenter servers can be joined together using vCenter Linked Mode to enable them to be managed using a single vSphere Client connection. However, with SRM 5.0, there is no mandatory requirement for use of vCenter Linked Mode to see and manage the SRM environments at both sites. SRM 5.0 includes the ability to see and manage all information that is important for configuration or management of SRM for both protected and recovery sites with or without the use of vCenter Linked Mode. vCenter Linked Mode may still be used, and is recommended for use, because it greatly simplifies license management and allows for easier management of the vSphere environment above and beyond the use of SRM. vCenter Linked Mode, moreover, will gracefully migrate SRM and other VMware licenses between sites, because the environments are moved using SRM. Refer to the *VMware vSphere Basic System Administration Guide* for more information about vSphere Linked Mode.

Site Planning and Preparation

Site planning and preparation at the **protected site** involves the following:

- Identify which virtual machines will be designated as protected virtual machines.
- Note that in this guide virtual machines protected by array-based replication are labeled **ATestWK1** through **ATestWK5** and vSphere Replication protected virtual machines are labeled **VTest1** through **VTest3.**

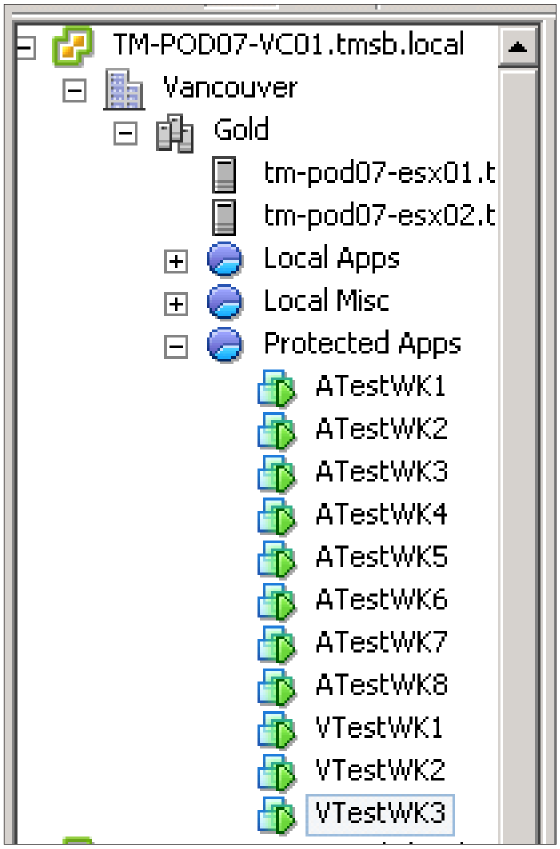

- TM-POD07-VC01.tmsb.local
  - Vancouver
    - Gold
      - tm-pod07-esx01.t
      - tm-pod07-esx02.t
      - Local Apps
      - Local Misc
      - Protected Apps
        - ATestWK1
        - ATestWK2
        - ATestWK3
        - ATestWK4
        - ATestWK5
        - ATestWK6
        - ATestWK7
        - ATestWK8
        - VTestWK1
        - VTestWK2
        - VTestWK3

Figure 1. Protected Virtual Machines in This Guide

- Identify which virtual machines will be designated as unprotected virtual machines (for example, VMs to be protected with vSphere Replication, Active Directory servers, DNS servers, print servers, SRM servers, vCenter servers, and so on).
- Determine which datastores are to hold the array-protected virtual machines (ATestWK1 through ATestWK5). If existing datastores will be used for the protected virtual machines, identify which datastores need to be configured for replication, otherwise provision the required number of new datastores to host the protected virtual machines. Work with your storage team to ensure all the datastores that will host protect virtual machines are configured for replication. Refer to the SRA configuration guide for details on supported replication configurations and the storage replication documentation for details on configuring the replication.

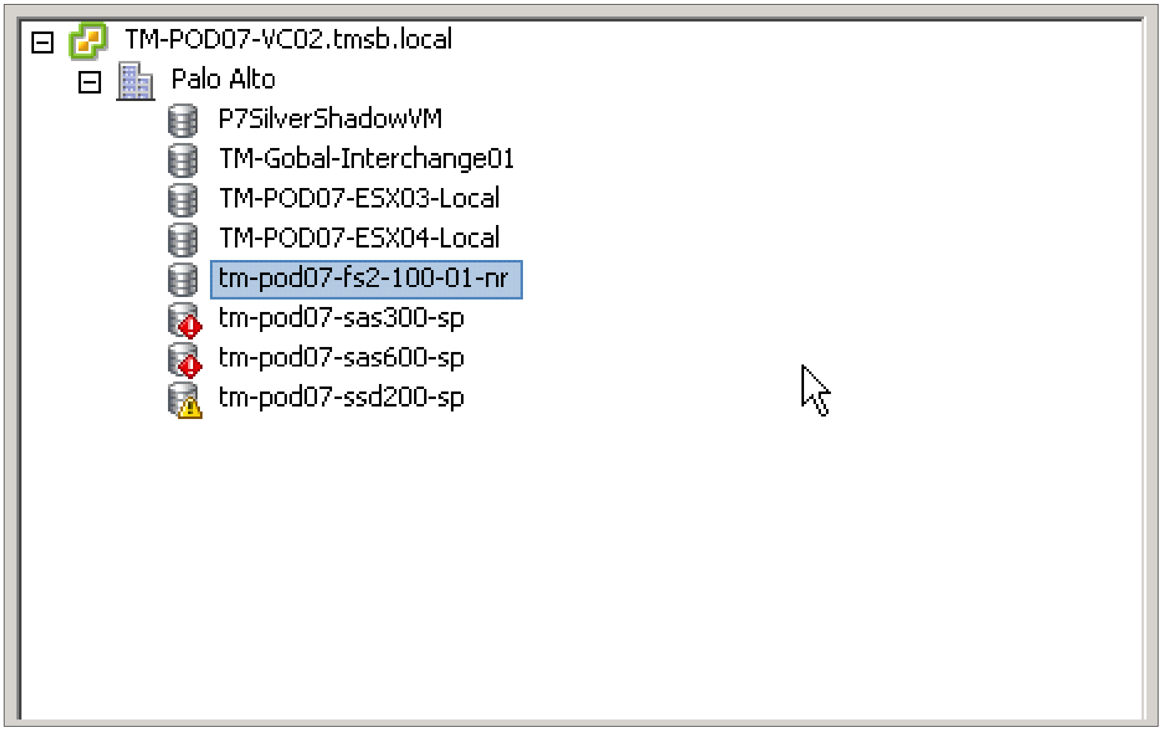

TM-POD07-VC02.tmsb.local
Palo Alto
P7SilverShadowVM
TM-Gobal-Interchange01
TM-POD07-ESX03-Local
TM-POD07-ESX04-Local
tm-pod07-fs2-100-01-nr
tm-pod07-sas300-sp
tm-pod07-sas600-sp
tm-pod07-ssd200-sp

Figure 2. Sample Datastores (Different Storage Types Will Be Used in This Evaluation Guide – Local, Shared, Replicated, and Nonreplicated)

- Move all the designated protected virtual machines onto the replicated datastores. vSphere Storage vMotion can be used to complete the relocation of the protected virtual machines with zero service downtime. If possible, ensure that there are only protected virtual machines on the datastores that are being replicated from the protected site to the recovery site.
- If vSphere Replication will be used, VMs that will be protected with this mechanism (VTest1 through VTest3) may reside on any datastore visible to the vSphere cluster at the protected site.
- If array-based replication will be used in conjunction with vSphere Replication, ensure that VMs to be protected by vSphere Replication reside on nonreplicated datastores that are not protected by array-based replication. This will ensure easier management, and that VR-protected VMs are not accidentally registered at the recovery site as part of an SRA-based protection group. This will also minimize disk space requirements, because the protected VMs will then not be replicated twice.

Site planning and preparation at the **recovery site** involves the following:

- Ensure that you have sufficient resources (in other words, CPU, memory and network) at the recovery site for the recovered virtual machines to utilize.

Exercise 1. Site Configuration and Recovery Workflow Setup

SRM Recovery Workflow	Recovery workflow automation	Set up Recovery Workflow 1. Set up site-pairing. 2. Set up array managers for the replicated datastore. 3. Set up inventory mappings. 4. Set up the protection group. 5. Set up the recovery plan. 6. Configure IP customization. 7. Trigger a test recovery.	60 minutes

Step 1: Set up connection pairing

To set up connection pairing:

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.

NOTE: The recovery site must be the replication target of arrays managed by the SRA at the protected site.

3. Click the **Site Recovery icon** on the vSphere Client Home page under Solutions and Applications. An authentication window might pop up with regards to SSL certificates. You should choose to install the certificate and ignore certificate warnings.

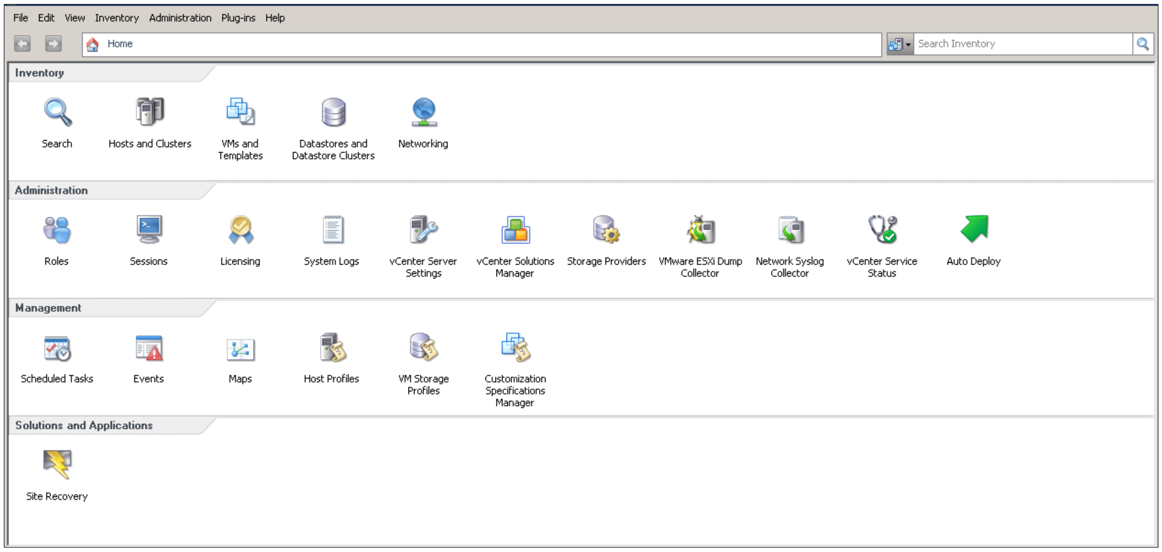

Figure 3. vSphere Client Home Page – Site Recovery Icon

4. In the **Commands** area of the **Summary** window, click **Configure Connection** to begin pairing the protected and recovery sites.

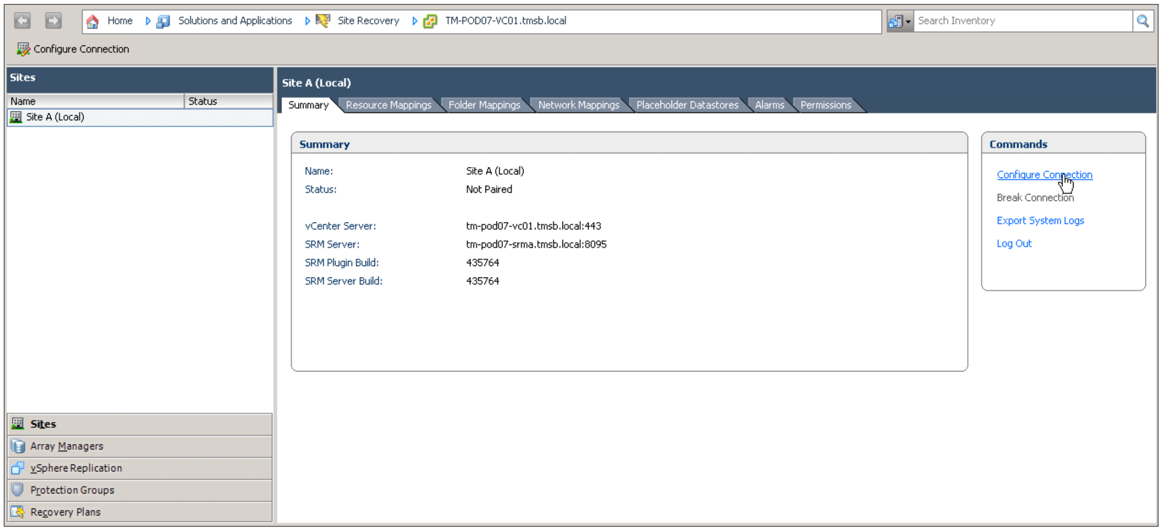

Name: Site A (Local)
Status: Not Paired

vCenter Server: tm-pod07-vc01.tmsb.local:443
SRM Server: tm-pod07-srma.tmsb.local:8095
SRM Plugin Build: 435764
SRM Server Build: 435764

Figure 4. Configure Connection Pairing of Protected and Recovery Sites

5. On the **Remote Site Information** page, type the host name or IP address of the vCenter server at the recovery site and click **Next.** Accept any certificates to proceed.

NOTE: If you are using credential-based authentication, you must supply exactly the same information here that you entered when installing the SRM Server. If you entered an IP address for that step, enter it again here. If you entered a host name or fully qualified domain name for that step, enter it here in exactly the same way.

As a general practice, it is recommended that fully qualified names be used in all scenarios for all name, address, hostname, and other fields, and to ensure that DNS resolution is reliable and consistent for all systems. This includes forward, reverse, short, and FQDN DNS resolution for all systems.

NOTE: Note what format is being used at this step, whether fully qualified, hostname, or IP address. This will be important for future steps.

Port 80 is provided as the default to use for the initial connection to the remote site. After the initial HTTP connection is made, the two sites establish an SSL connection over port 80 to use for subsequent connections.

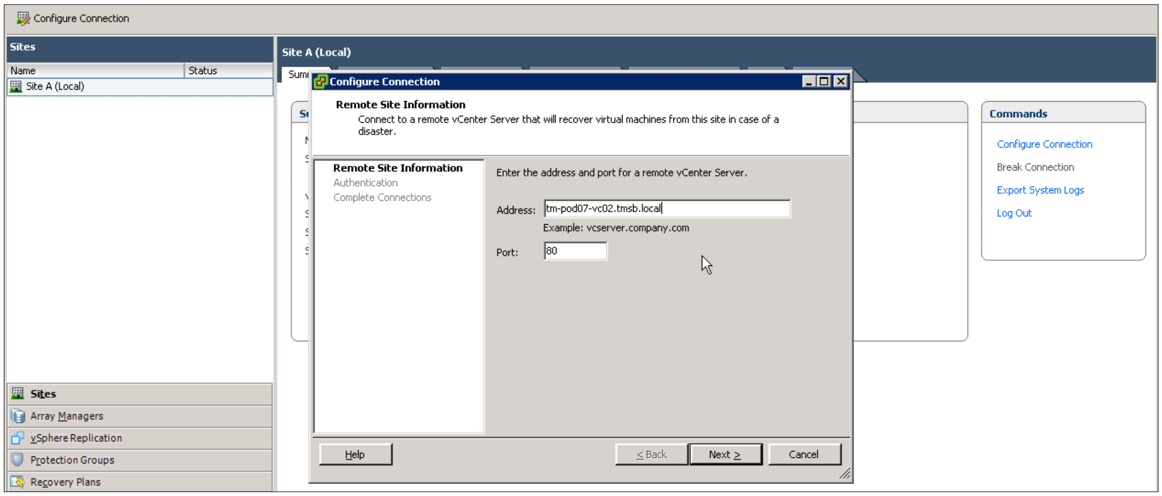

Remote Site Information
Connect to a remote vCenter Server that will recover virtual machines from this site in case of a disaster.

Enter the address and port for a remote vCenter Server.

Address: tm-pod07-vc02.tmsb.local
Example: vcserver.company.com

Port: 80

Figure 5. Enter Remote Site Information

6. On the **vCenter Server Authentication page,** provide the appropriate vCenter administrator credentials (user name and password) for the remote site and click **Next.** If you are using credential-based authentication, you must supply exactly the same information here that you entered when installing the SRM Server.

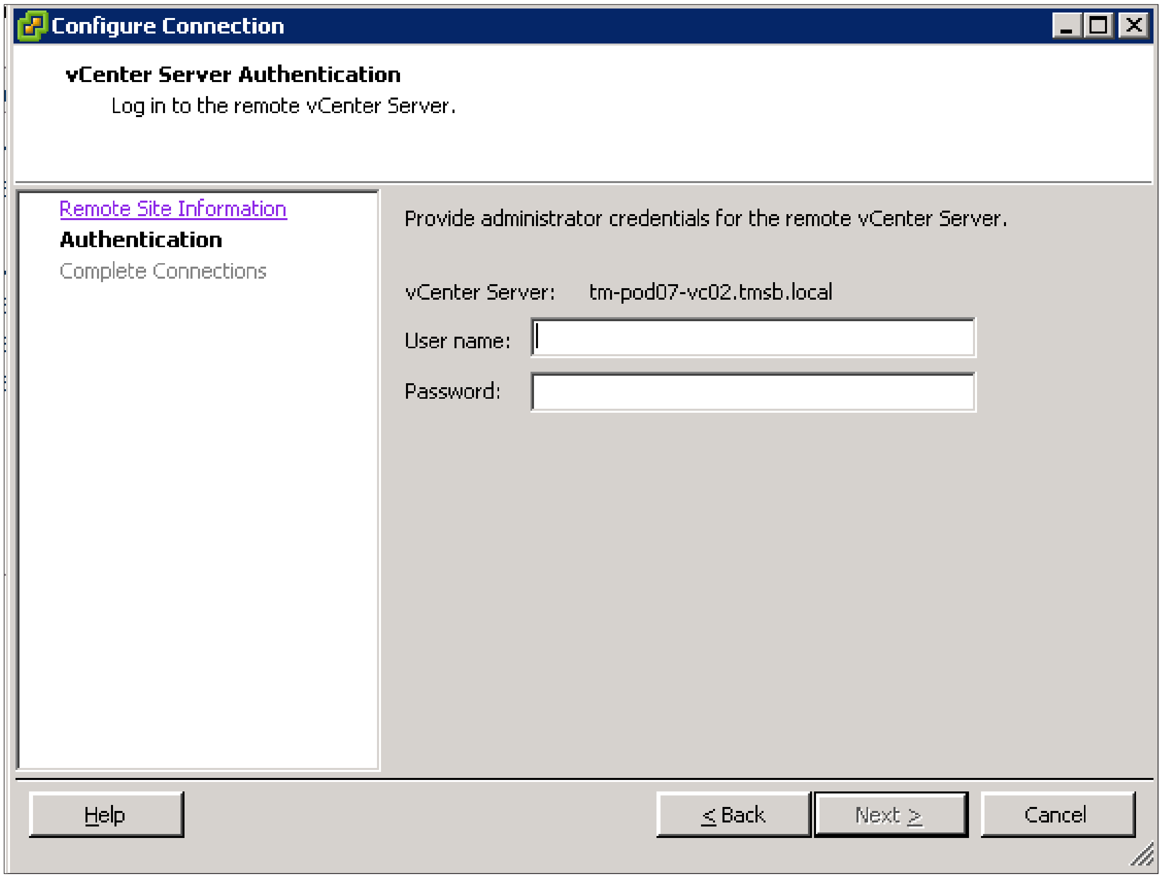

vCenter Server Authentication

Log in to the remote vCenter Server.

Provide administrator credentials for the remote vCenter Server.

vCenter Server: tm-pod07-vc02.tmsb.local

User name:

Password:

Figure 6. Enter vCenter Server Authentication Information

7. The SRM servers will now attempt to pair and establish reciprocity. If any errors occur at this point, you have probably entered either a host name or user name and password incorrectly. Please verify all information and try again. When successful, click **Finish** on the Configure Connection pop-up menu to return to the Site Summary screen.

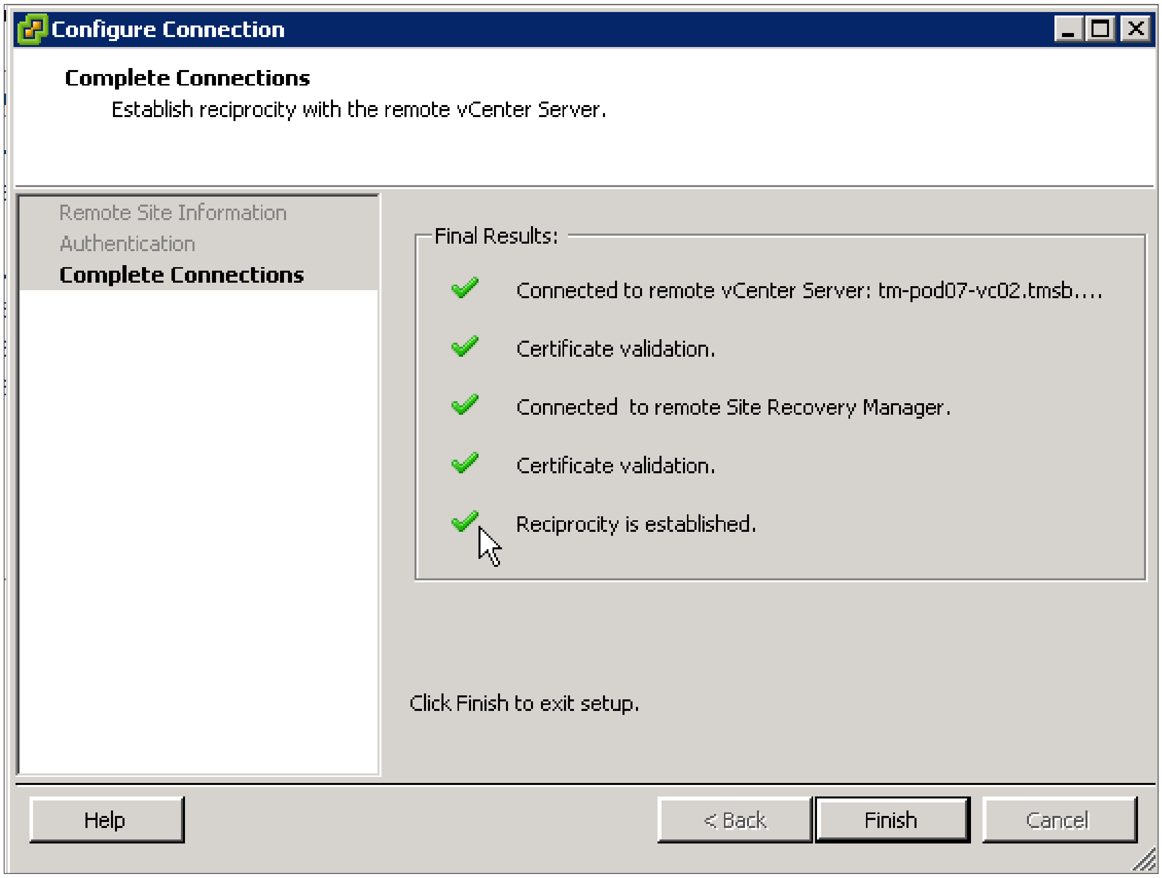

Configure Connection

Complete Connections

Establish reciprocity with the remote vCenter Server.

Final Results:

✓ Connected to remote vCenter Server: tm-pod07-vc02.tmsb....

✓ Certificate validation.

✓ Connected to remote Site Recovery Manager.

✓ Certificate validation.

✓ Reciprocity is established.

Click Finish to exit setup.

Figure 7. Establishing Reciprocity

8. Moments after you return to the Site Summary screen, a pop-up menu will appear prompting for credentials for the remote vCenter server. Enter your credentials for the recovery site vCenter server, and then wait while the paired sites are populated in the SRM Sites screen.

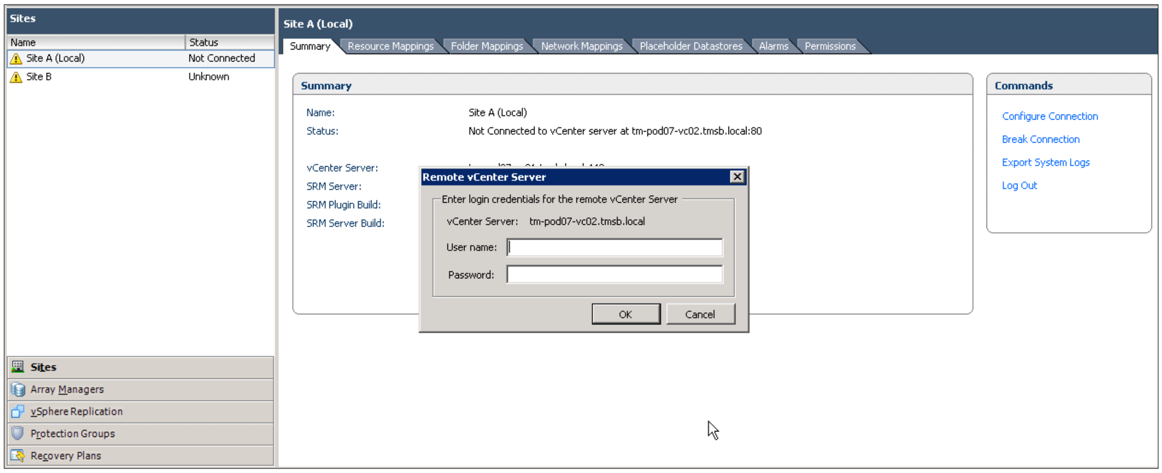

Sites

Name	Status
Site A (Local)	Not Connected
Site B	Unknown

Site A (Local)

Summary

Name: Site A (Local)

Status: Not Connected to vCenter server at tm-pod07-vc02.tmsb.local:80

vCenter Server:

SRM Server:

SRM Plugin Build:

SRM Server Build:

Remote vCenter Server

Enter login credentials for the remote vCenter Server

vCenter Server: tm-pod07-vc02.tmsb.local

User name:

Password:

Figure 8. Remote vCenter Server Authentication

The SRM and vCenter servers at the protected and recovery sites are connected. Connection information is saved in the SRM databases, and persists across logins and host restarts.

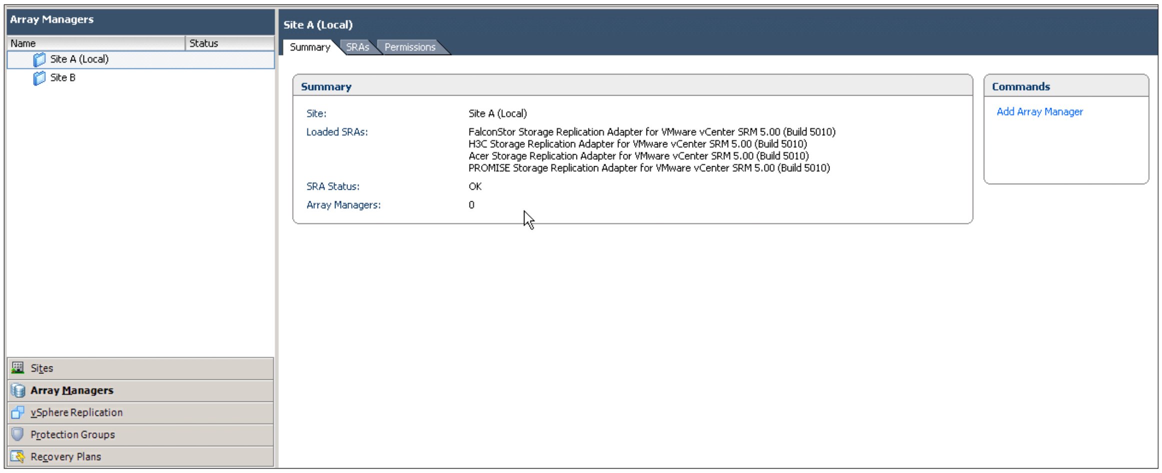

Site: Site A (Local)
Loaded SRAs: FalconStor Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)
H3C Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)
Acer Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)
PROMISE Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)
SRA Status: OK
Array Managers: 0

Figure 9. Paired Protected and Recovery Sites

Step 2: Set up array managers

After you have connected the protected and recovery sites, you must configure their respective array managers so that SRM can discover replicated devices, compute datastore groups, and initiate storage operations.

The array manager configuration wizard leads you through the following steps:

- You provide SRM with connection information and credentials (if needed) for array-management systems at the protected and recovery sites.
- SRM verifies that it can connect to arrays at both sites.
- SRM verifies that it can discover replicated storage devices on these arrays and can identify the datastores that they support.
- SRM computes datastore groups based on virtual machine storage layout and any consistency groups defined by the storage array.

When the configuration process is complete, the wizard presents a list of datastore groups. You typically configure array managers only once, after you have connected the protected and recovery sites. It is not necessary to reconfigure them unless array manager connection information or credentials have changed, or unless you want to use a different set of arrays.

NOTE: The example here uses a particular storage replicated iSCSI datastore. You may have a different storage device type (for example, NFS, iSCSI or FC). In this case, you may see a slight variation of input parameters, depending on the storage device type. The general workflow should still be similar, but individual screenshots may look different in your environment.

Procedure

1. Open the vSphere Client and connect to the vCenter server at the protected site. Log in as a vSphere administrator.
2. Click the **Site Recovery** icon on the vSphere Client Home page.
3. In the main SRM navigation tab on the left frame, click the **Array Managers** line and click the protected site in the top-left frame.
4. Make sure that the SRA you want SRM to use appears in the **SRA** tab in the right frame.

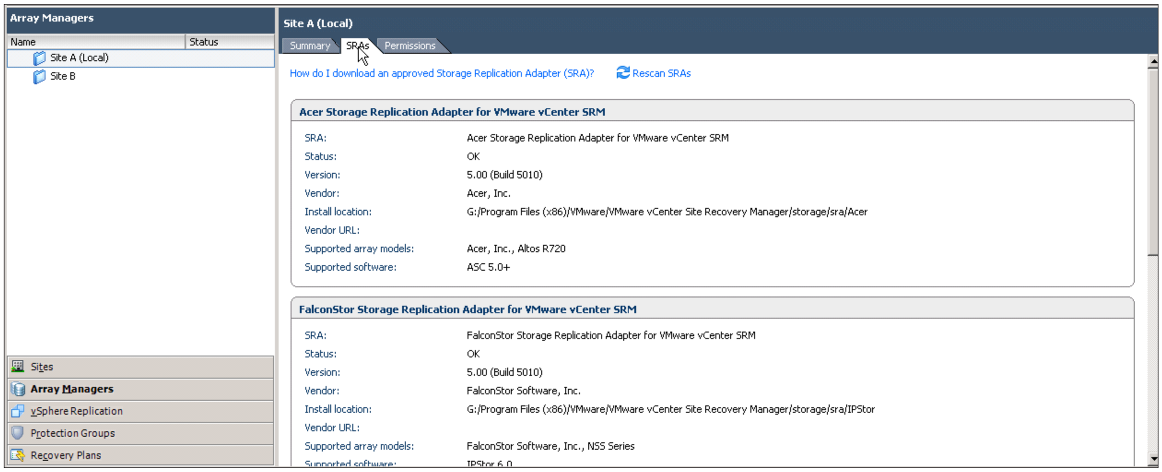

Acer Storage Replication Adapter for VMware vCenter SRM

SRA:	Acer Storage Replication Adapter for VMware vCenter SRM
Status:	OK
Version:	5.00 (Build 5010)
Vendor:	Acer, Inc.
Install location:	G:/Program Files (x86)/VMware/VMware vCenter Site Recovery Manager/storage/sra/Acer
Vendor URL:	
Supported array models:	Acer, Inc., Altos R720
Supported software:	ASC 5.0+

FalconStor Storage Replication Adapter for VMware vCenter SRM

SRA:	FalconStor Storage Replication Adapter for VMware vCenter SRM
Status:	OK
Version:	5.00 (Build 5010)
Vendor:	FalconStor Software, Inc.
Install location:	G:/Program Files (x86)/VMware/VMware vCenter Site Recovery Manager/storage/sra/IPStor
Vendor URL:	
Supported array models:	FalconStor Software, Inc., NSS Series
Supported software:	IPStor 6.0

Figure 10. SRA Tab Showing Load and Installed SRAs

If no SRA is listed, click the **Rescan SRAs** button at the top of the screen. If an SRA is still not listed, then no SRA has been installed on the SRM host. For more information, see the chapter "Install the Storage Replication Adapters" in the *Site Recovery Manager Administration Guide.*

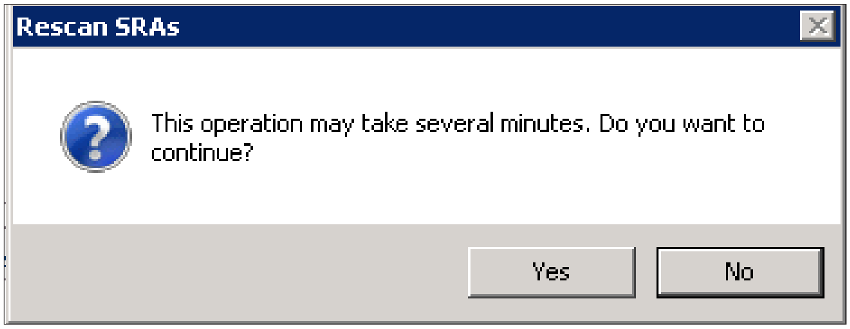

This operation may take several minutes. Do you want to continue?

Figure 11. Rescan SRAs

5. At the main **Array Managers** screen, select the protected site in the top-left screen, and either right-click the site name, or click **Add Array Manager** on the right pane of the summary screen.

6. Enter a specific name for the array manager being added to the site, and select an appropriate SRA from the SRA Type drop-down window. Ensure that you are selecting the correct SRA type for the array manager being added. There might be more than one SRA available for selection.

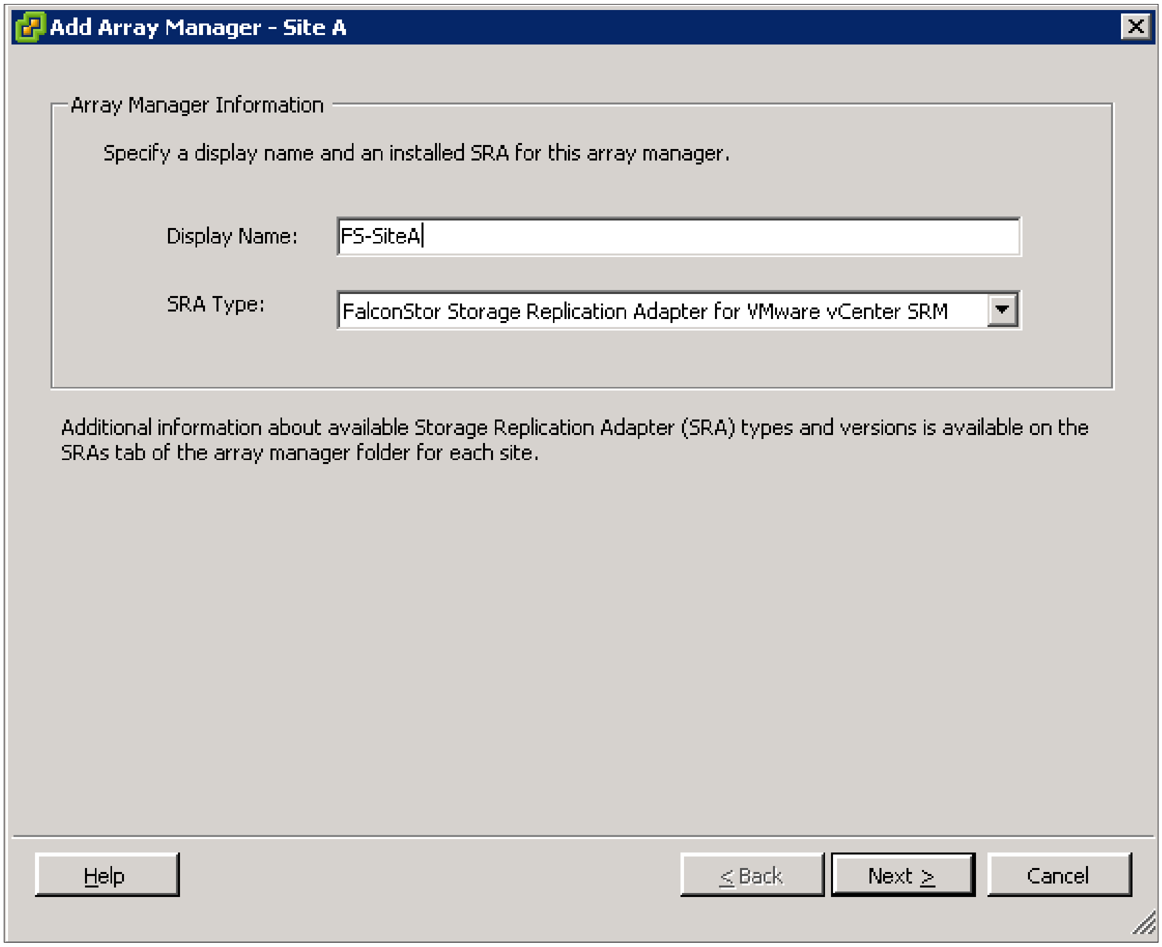

Add Array Manager - Site A

Array Manager Information

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

Display Name: FS-SiteA

SRA Type: FalconStor Storage Replication Adapter for VMware vCenter SRM

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

Figure 12. Adding an Array Name and Choosing an Appropriate SRA

7. Enter configuration and authentication information to connect to the specified array manager. These fields are defined by the SRA. For more information about how to fill them in, see the documentation provided by your SRA vendor.

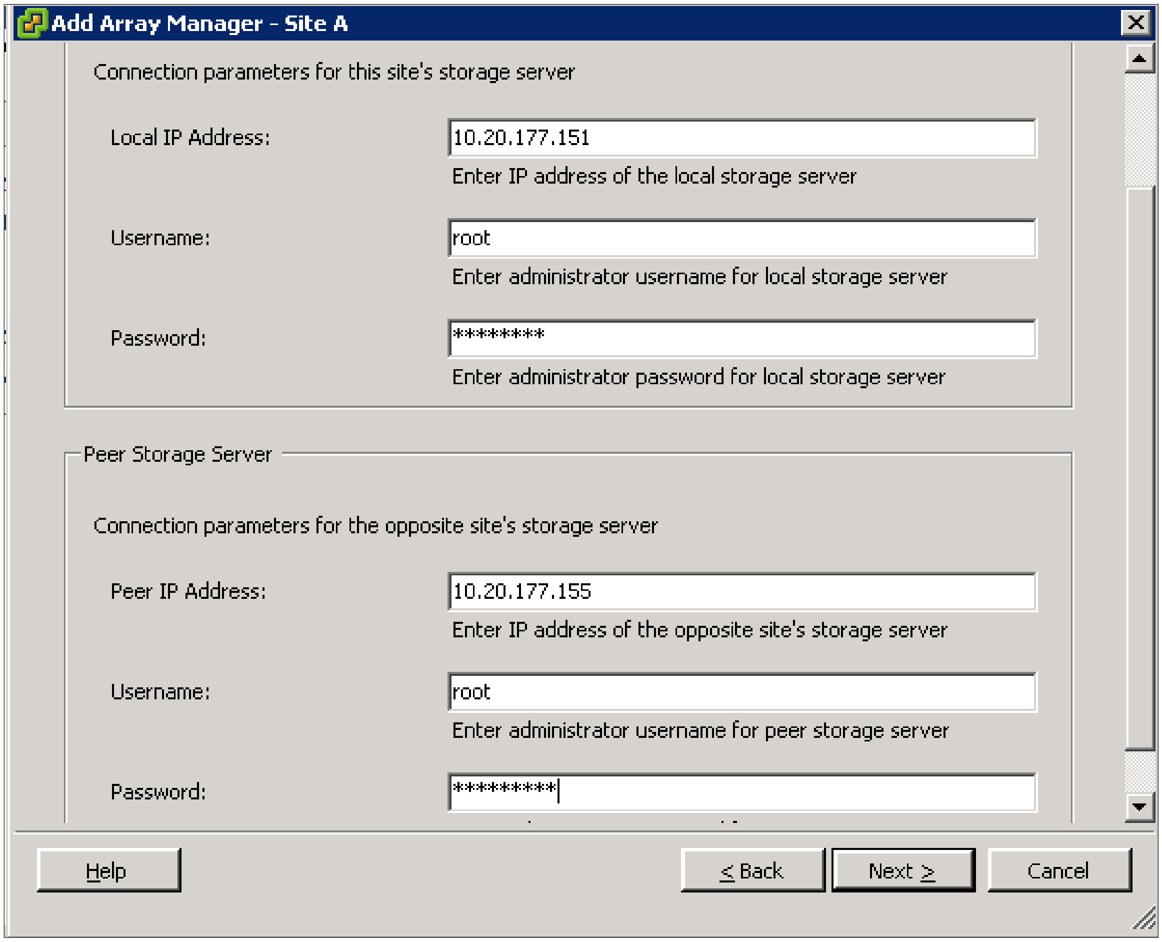

Add Array Manager - Site A

Connection parameters for this site's storage server

Local IP Address: 10.20.177.151
Enter IP address of the local storage server

Username: root
Enter administrator username for local storage server

Password: ********
Enter administrator password for local storage server

Peer Storage Server

Connection parameters for the opposite site's storage server

Peer IP Address: 10.20.177.155
Enter IP address of the opposite site's storage server

Username: root
Enter administrator username for peer storage server

Password: *********

Figure 13. SRA Configuration and Authentication to Array Manager

8. If the information you supplied is correct and SRM can communicate with the array managers through the SRA, click **Finish.**

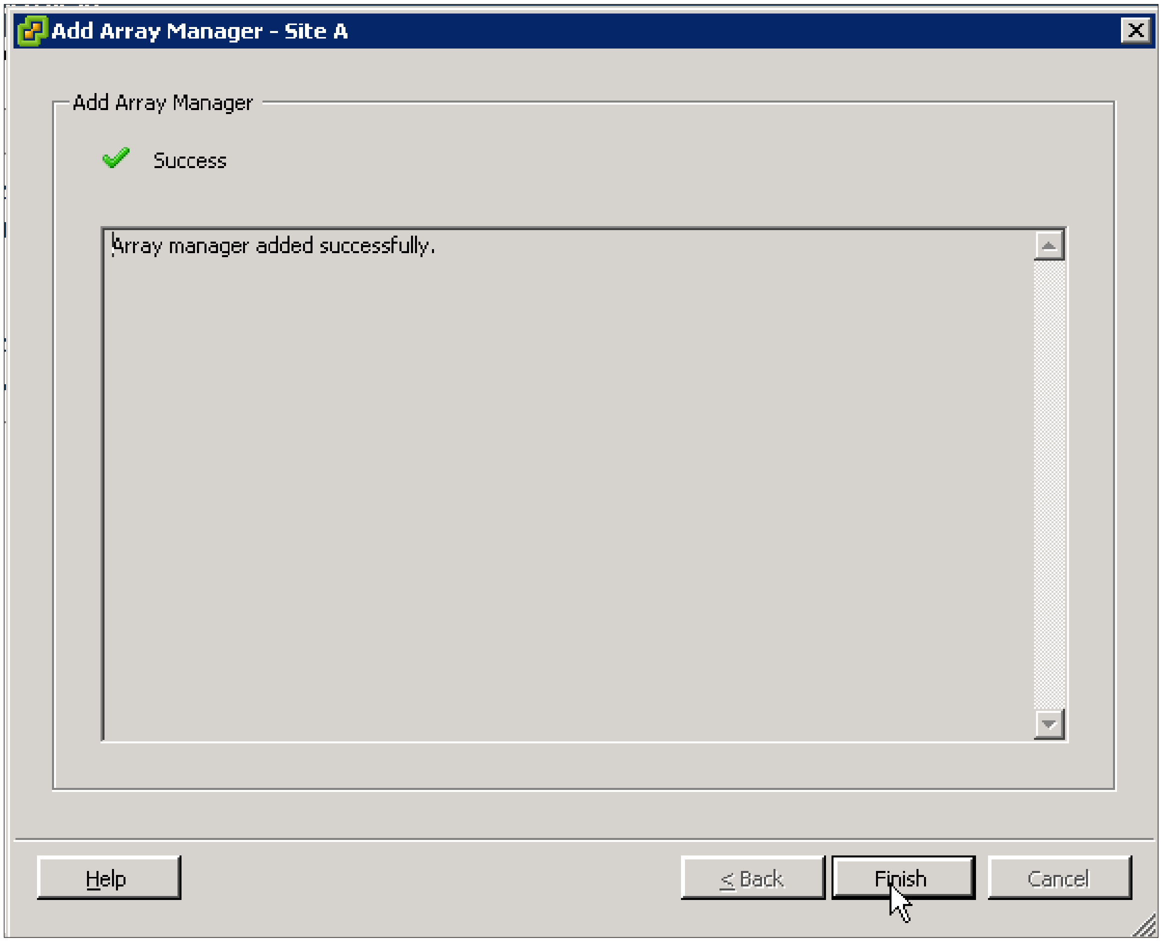

✓ Success

Array manager added successfully.

Figure 14. SRA Configuration Complete

The array manager queries the selected arrays to discover which of their devices are replicated. Detailed information about the selected arrays is available by clicking on the added array that should now appear in the protected site folder.

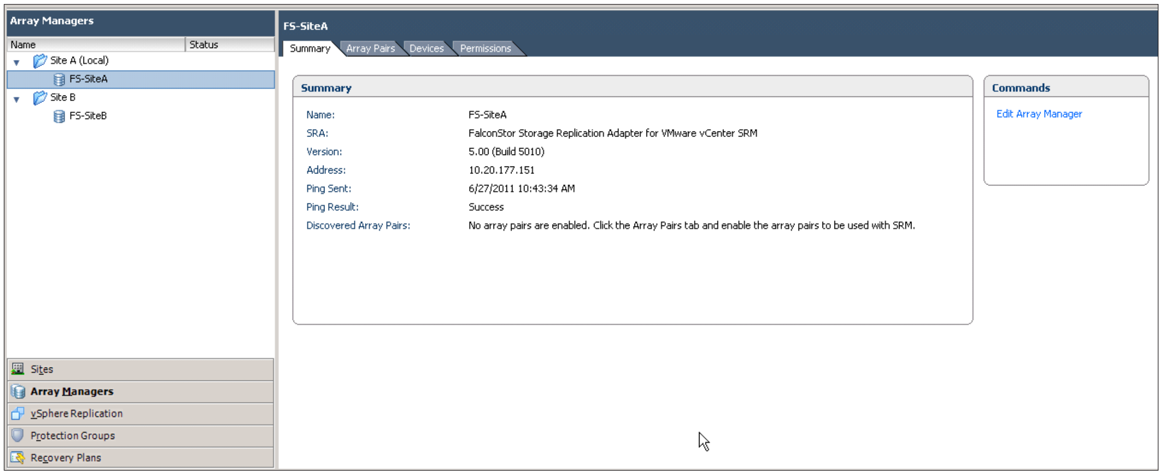

Summary

Name: FS-SiteA
SRA: FalconStor Storage Replication Adapter for VMware vCenter SRM
Version: 5.00 (Build 5010)
Address: 10.20.177.151
Ping Sent: 6/27/2011 10:43:34 AM
Ping Result: Success
Discovered Array Pairs: No array pairs are enabled. Click the Array Pairs tab and enable the array pairs to be used with SRM.

Figure 15. Array-Specific Information for Protected Site

9. Click the name of the recovery site folder to configure array managers at the recovery site.

10. On the **Recovery Site Array Managers** summary page, click **Add Array Manager.**

The procedure for configuring these arrays is identical to the procedure for configuring the arrays at the protected site, described in Step 4 through Step 8. You will enter different information for IP addresses in this step, because you are adding the local array manager for the recovery site.

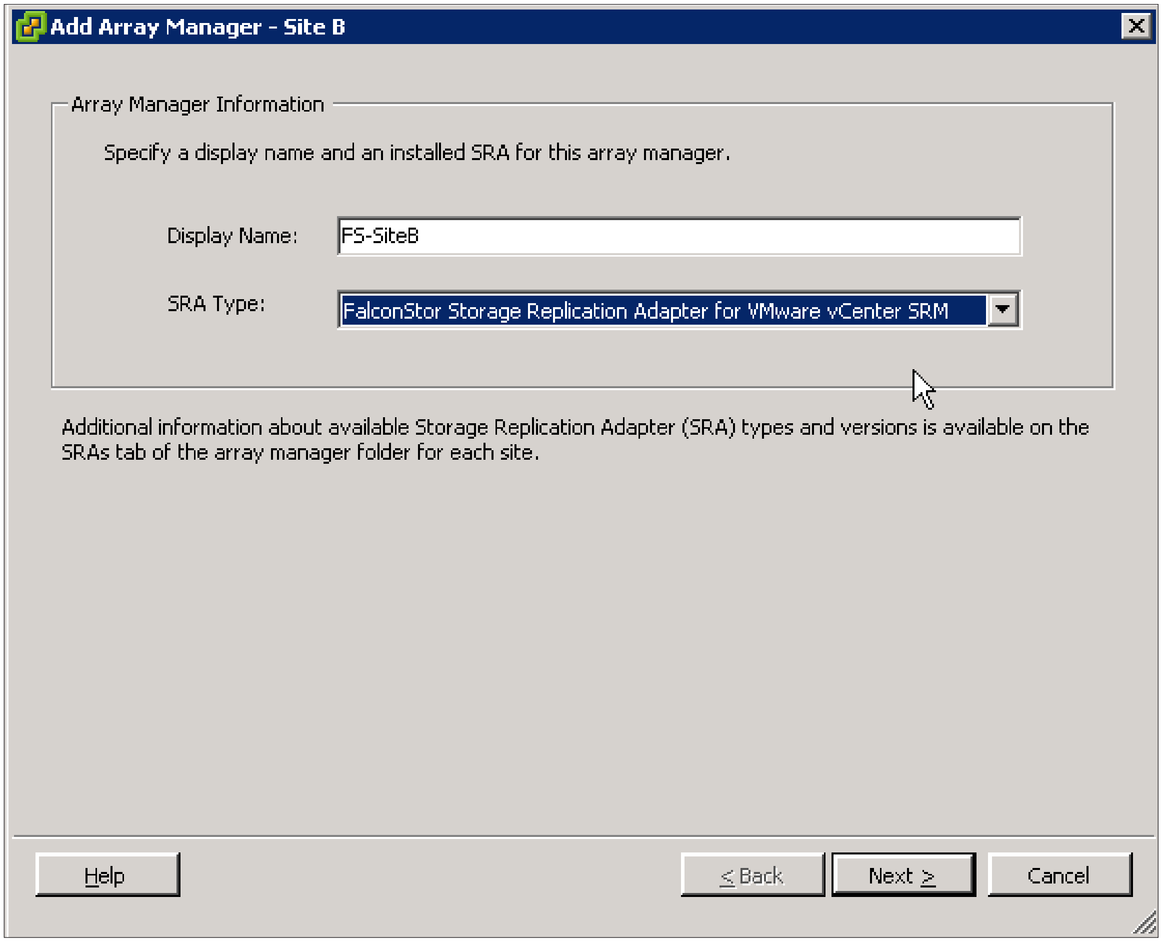

Specify a display name and an installed SRA for this array manager.

Display Name: FS-SiteB

SRA Type: FalconStor Storage Replication Adapter for VMware vCenter SRM

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

Figure 16. Enter Array Manager Information for Recovery Site

11. Click **Finish.**

The **Array Managers** screen should now show two site folders, each with an array configured.

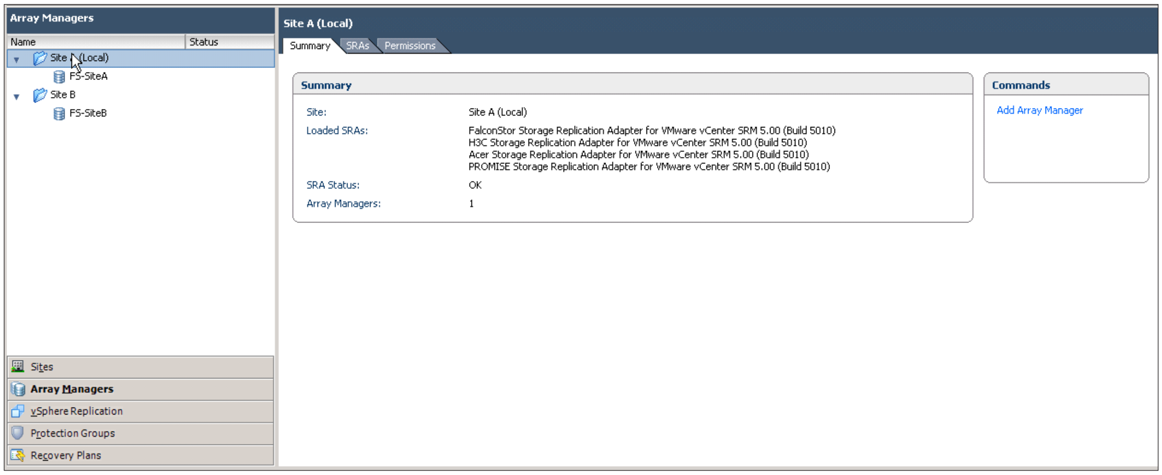

Site: Site A (Local)

Loaded SRAs: FalconStor Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)
H3C Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)
Acer Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)
PROMISE Storage Replication Adapter for VMware vCenter SRM 5.00 (Build 5010)

SRA Status: OK

Array Managers: 1

Figure 17. Arrays Configured for Both Protected and Recovery Sites

12. Click on a configured array on either the protected site or the recovery site.

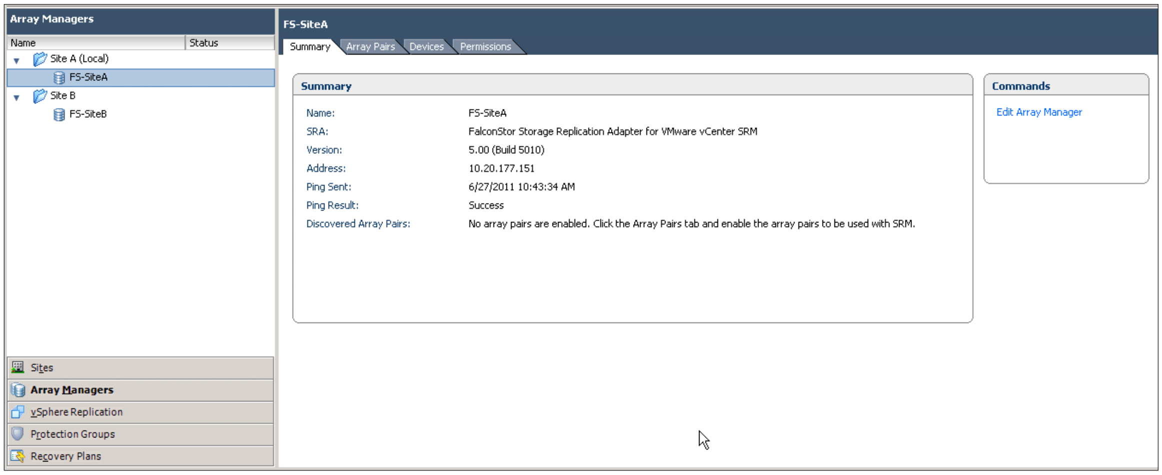

Summary

Name:	FS-SiteA
SRA:	FalconStor Storage Replication Adapter for VMware vCenter SRM
Version:	5.00 (Build 5010)
Address:	10.20.177.151
Ping Sent:	6/27/2011 10:43:34 AM
Ping Result:	Success
Discovered Array Pairs:	No array pairs are enabled. Click the Array Pairs tab and enable the array pairs to be used with SRM.

Figure 18. Selecting an Array

13. Click on the **Array Pairs** tab in the context of this selected array manager. You should see the configured local array and the remote array in the main viewing pane.

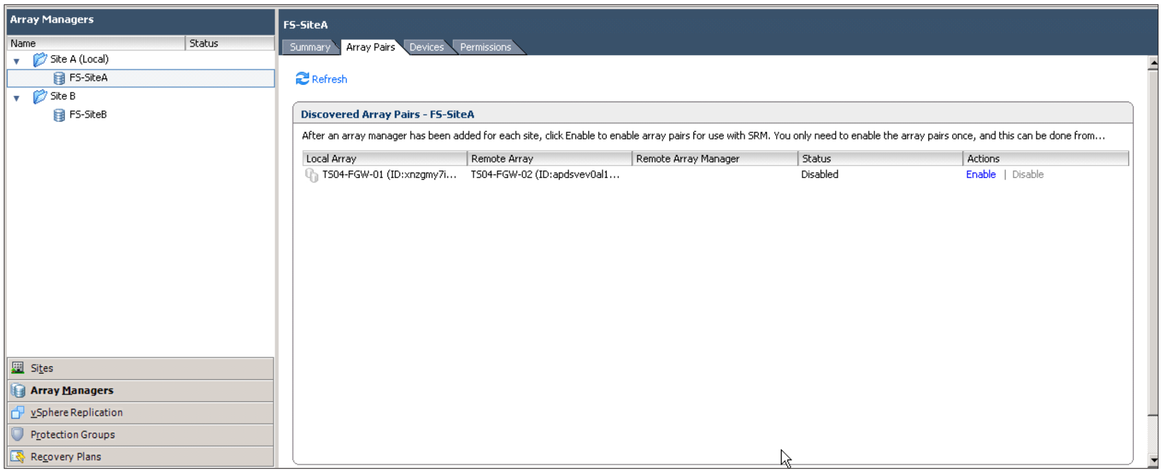

Discovered Array Pairs - FS-SiteA

After an array manager has been added for each site, click Enable to enable array pairs for use with SRM. You only need to enable the array pairs once, and this can be done from...

Local Array	Remote Array	Remote Array Manager	Status	Actions
TS04-FGW-01 (ID:xnzgmy7i...	TS04-FGW-02 (ID:apdsvev0al1...		Disabled	Enable | Disable

Figure 19. Array Pairs

14. Click **Enable** on the right pane under the **Actions** column to enable pairing of the local and remote array managers.

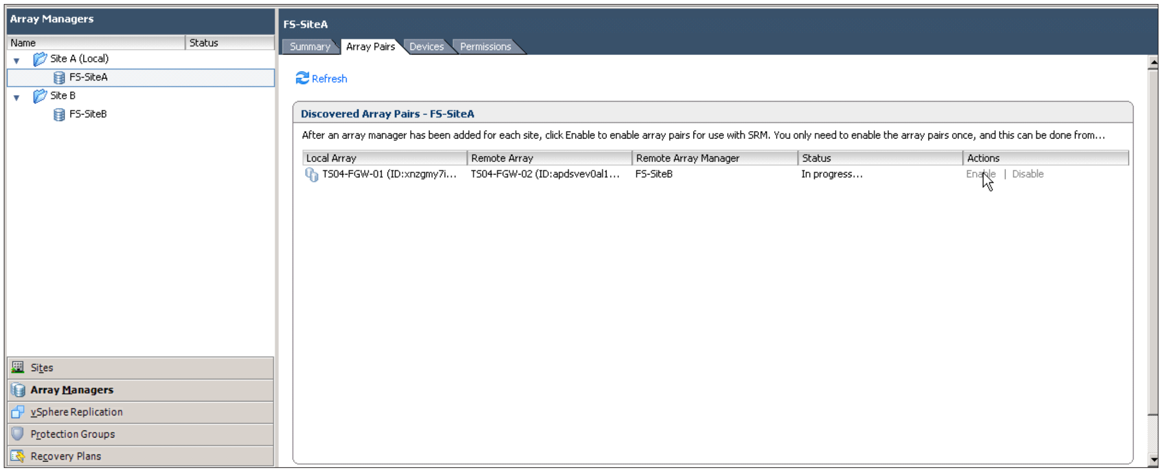

Discovered Array Pairs - FS-SiteA

After an array manager has been added for each site, click Enable to enable array pairs for use with SRM. You only need to enable the array pairs once, and this can be done from...

Local Array	Remote Array	Remote Array Manager	Status	Actions
TS04-FGW-01 (ID:xnzgmy7i...	TS04-FGW-02 (ID:apdsvev0al1...	FS-SiteB	In progress...	Enable | Disable

Figure 20. Enabling Array Pairing

15. Once array pairing is complete, click the **Devices** tab to display the list of replicated datastore groups.

    On the **Devices** tab, you can see which datastores the array manager is replicating, and the current state of replication for those devices. If the list of replicated datastores is not what you expected, you must correct it before continuing.

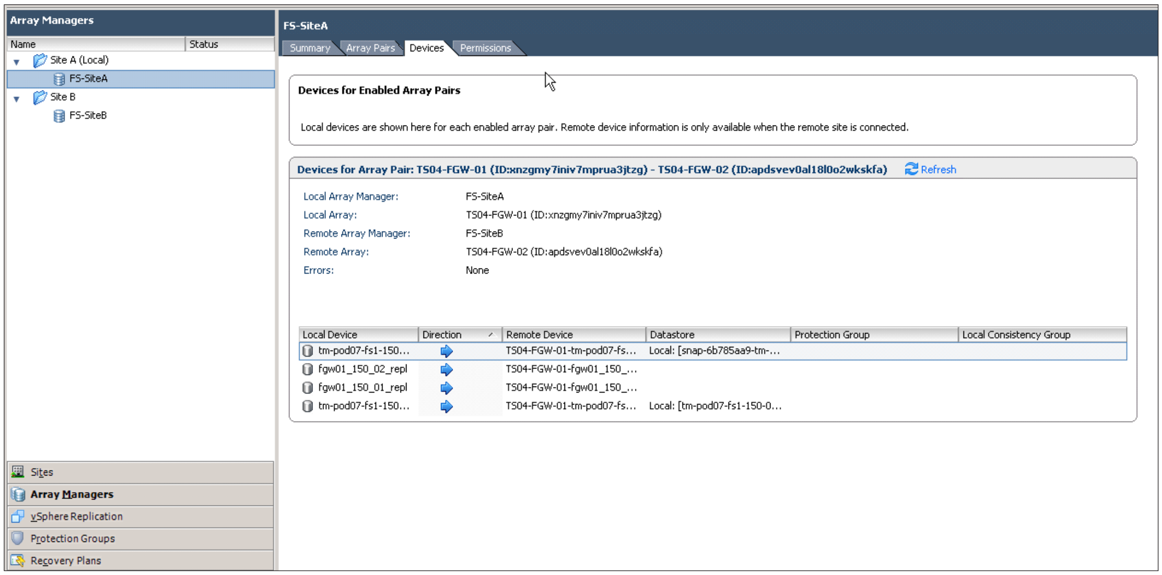

Devices for Enabled Array Pairs

Local devices are shown here for each enabled array pair. Remote device information is only available when the remote site is connected.

Devices for Array Pair: TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg) - TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa)

Local Array Manager: FS-SiteA
Local Array: TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg)
Remote Array Manager: FS-SiteB
Remote Array: TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa)
Errors: None

Local Device	Direction	Remote Device	Datastore	Protection Group	Local Consistency Group
tm-pod07-fs1-150...	➡	TS04-FGW-01-tm-pod07-fs...	Local: [snap-6b785aa9-tm-...		
fgw01_150_02_repl	➡	TS04-FGW-01-fgw01_150_...			
fgw01_150_01_repl	➡	TS04-FGW-01-fgw01_150_...			
tm-pod07-fs1-150...	➡	TS04-FGW-01-tm-pod07-fs...	Local: [tm-pod07-fs1-150-0...		

Figure 21. Devices for Array Pair – View from the Protected Site (Note the Direction of Replication Arrow)

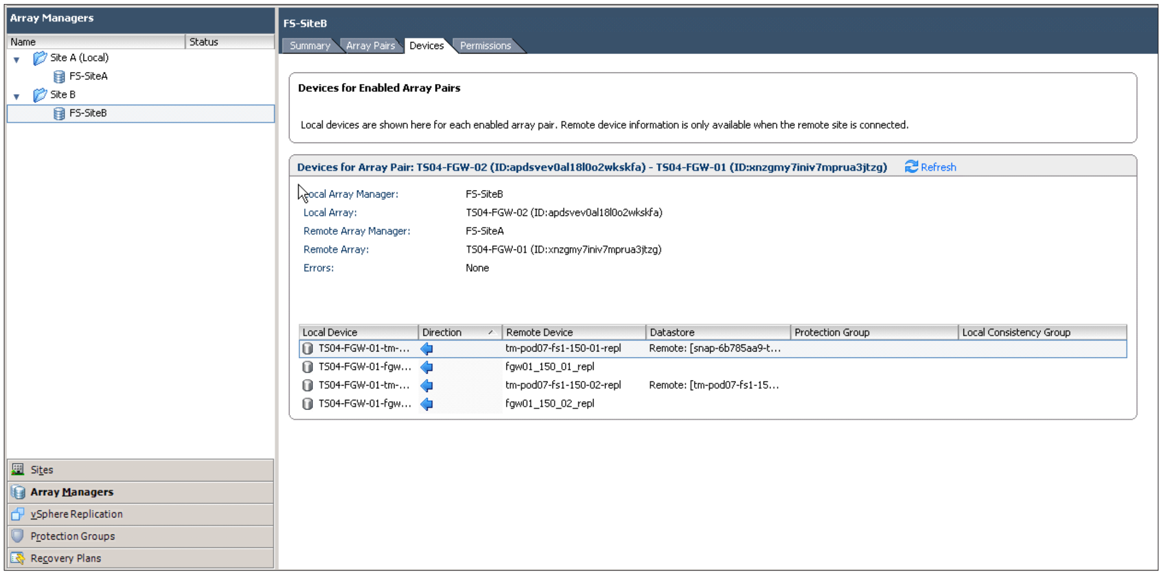

Devices for Enabled Array Pairs

Local devices are shown here for each enabled array pair. Remote device information is only available when the remote site is connected.

Devices for Array Pair: TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa) - TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg)

Local Array Manager:	FS-SiteB
Local Array:	TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa)
Remote Array Manager:	FS-SiteA
Remote Array:	TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg)
Errors:	None

Local Device	Direction	Remote Device	Datastore	Protection Group	Local Consistency Group
TS04-FGW-01-tm-...	←	tm-pod07-fs1-150-01-repl	Remote: [snap-6b785aa9-t...		
TS04-FGW-01-fgw...	←	fgw01_150_01_repl			
TS04-FGW-01-tm-...	←	tm-pod07-fs1-150-02-repl	Remote: [tm-pod07-fs1-15...		
TS04-FGW-01-fgw...	←	fgw01_150_02_repl			

Figure 22. Devices for Array Pair – View from the Recovery Site (Note the Direction of Replication Arrow)

Step 3: Set up inventory mapping

Inventory mappings establish recovery site defaults for the virtual machine folders, networks, and resource pools to which recovered virtual machines are assigned. You create these mappings at the protected site, and they apply to all virtual machines in all protection groups at that site.

Ensure that a placeholder datastore has been created at the recovery site in order to complete this step correctly. A placeholder datastore can contain shadow VMs for members of a protection group. The placeholder datastore should be accessible to all hosts in the recovery cluster. It should not be replicated and can be relatively small.

Inventory mappings are optional, but recommended. They provide a convenient way to specify how resources at the protected site are mapped to resources at the recovery site. These mappings are applied to all members of a protection group when the group is created, and can be reapplied as needed (for example, when new members are added). If you do not create them, you must specify mappings individually for each virtual machine that you add to a protection group. A virtual machine cannot be protected unless it has valid inventory mappings for networks, folders, and resource pools. In addition, you can specify a placeholder datastore at the recovery site that will hold shadow VMs that are used as placeholders for VMs that will be protected.

Do not specify resource mappings for resources that are not used by protected virtual machines.

Procedure

1. Open the vSphere Client and connect to the vCenter server at the protected site. Log in as a vSphere administrator.
2. Click the **Site Recovery** icon on the vSphere Client Home page.
3. In the **Sites** configuration area of the left-hand navigation pane, select the protected site and choose the **Resource Mappings** tab in the main viewing pane.

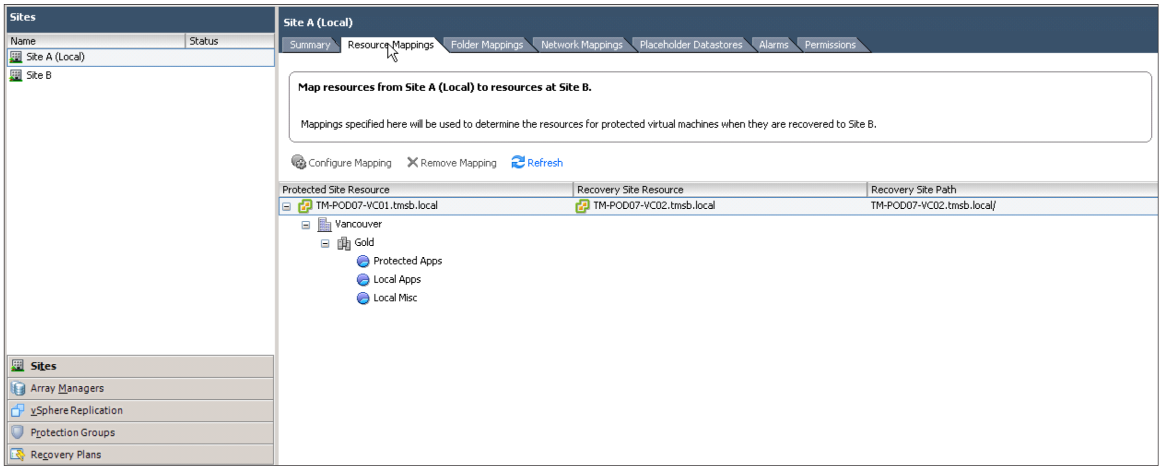

Name	Status
Site A (Local)	
Site B	

Map resources from Site A (Local) to resources at Site B.

Mappings specified here will be used to determine the resources for protected virtual machines when they are recovered to Site B.

Protected Site Resource	Recovery Site Resource	Recovery Site Path
TM-POD07-VC01.tmsb.local	TM-POD07-VC02.tmsb.local	TM-POD07-VC02.tmsb.local/
Vancouver		
Gold		
Protected Apps		
Local Apps		
Local Misc		

Figure 23. Configure Resource Mappings

The **Resource Mappings** page displays a tree of resource pools at the protected site and a corresponding tree of resources at the recovery site.

4. To configure mapping for a resource, click the resource in the **Protected Site Resources** column and click **Configure Mapping.**

5. Expand the top-level folder in the Configure Inventory Mapping window and navigate to the recovery site resource (network, resource pool, or folder) to which you want to map the protected site resource you selected in Step 4. Select the resource and click **OK.**

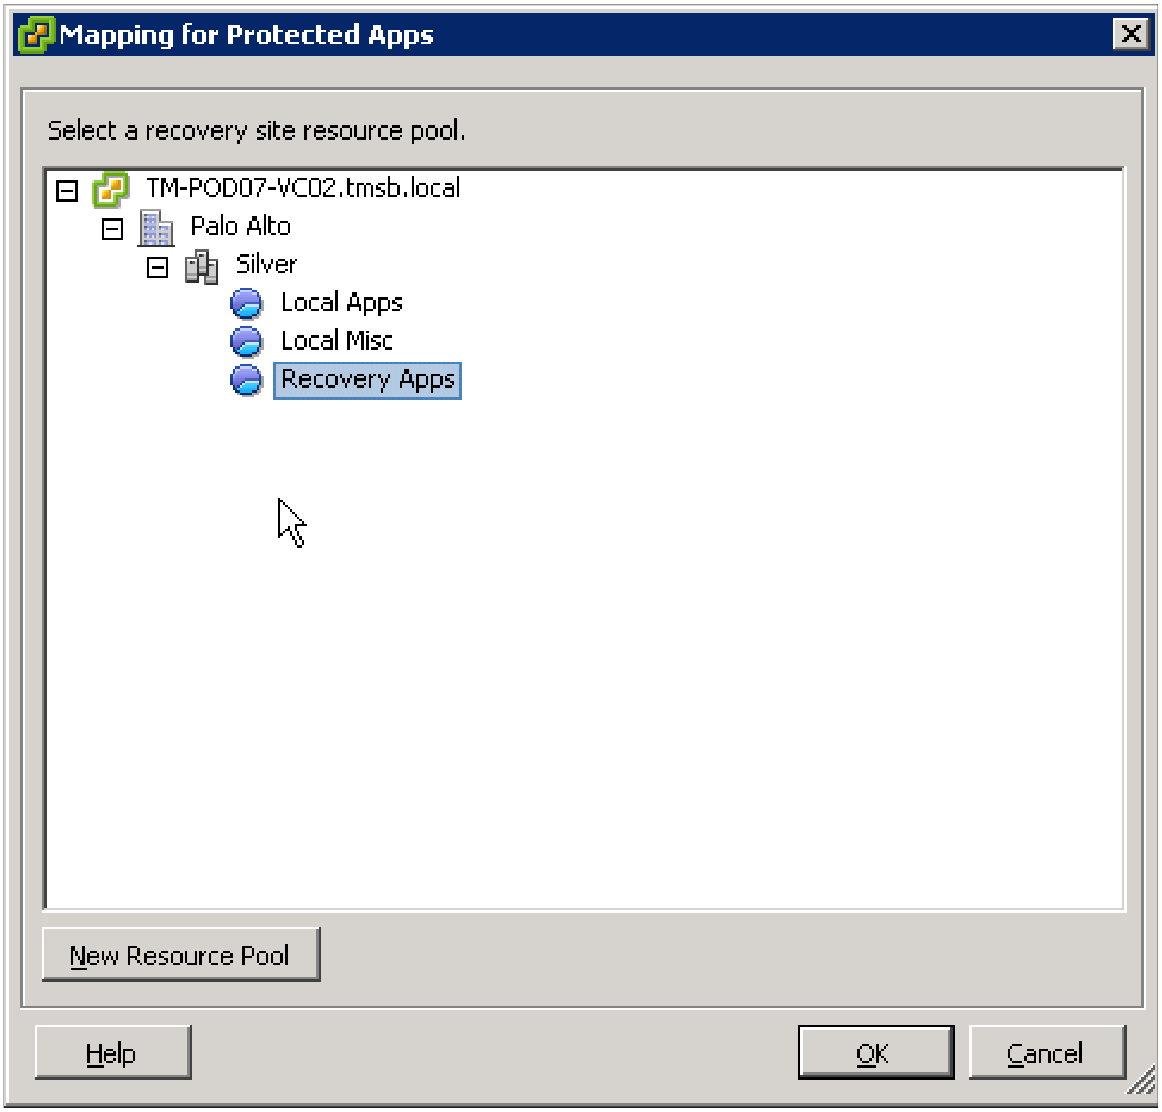

Mapping for Protected Apps

Select a recovery site resource pool.

- TM-POD07-VC02.tmsb.local
  - Palo Alto
    - Silver
      - Local Apps
      - Local Misc
      - Recovery Apps

Figure 24. Choosing a Recovery Site Resource to Map to the "Protected Apps" Resource Pool from the Primary Site

The selected resource is displayed in the **Recovery Site Resources** column, and its path relative to the root of the recovery site vCenter server is displayed in the Recovery Site Path column.

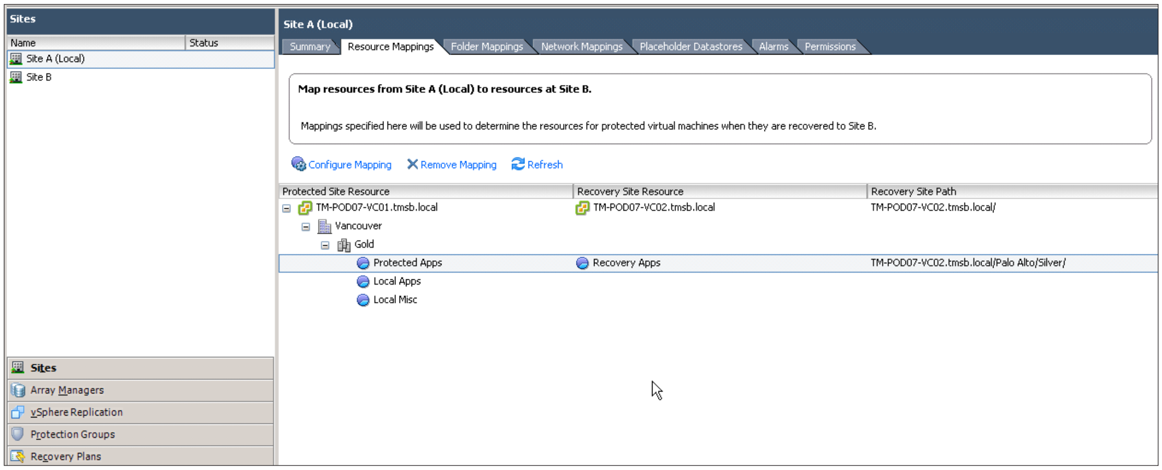

Map resources from Site A (Local) to resources at Site B.

Mappings specified here will be used to determine the resources for protected virtual machines when they are recovered to Site B.

Protected Site Resource	Recovery Site Resource	Recovery Site Path
TM-POD07-VC01.tmsb.local	TM-POD07-VC02.tmsb.local	TM-POD07-VC02.tmsb.local/
Vancouver		
Gold		
Protected Apps	Recovery Apps	TM-POD07-VC02.tmsb.local/Palo Alto/Silver/
Local Apps		
Local Misc		

Figure 25. Inventory Mapping Interface After Mapping a Resource Pool

Ensure that you continue to configure all inventory mapping tabs as needed – resources, folders, networks, and a placeholder datastore mapping to hold shadow VMs at the recovery site.

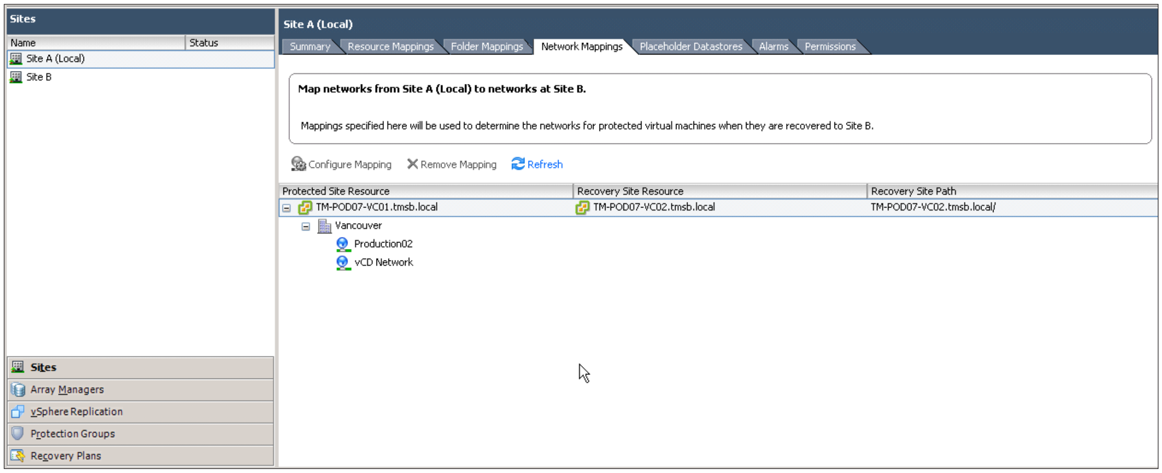

Map networks from Site A (Local) to networks at Site B.

Mappings specified here will be used to determine the networks for protected virtual machines when they are recovered to Site B.

Protected Site Resource	Recovery Site Resource	Recovery Site Path
TM-POD07-VC01.tmsb.local	TM-POD07-VC02.tmsb.local	TM-POD07-VC02.tmsb.local/
Vancouver		
Production02		
vCD Network		

Figure 26. Network Mapping

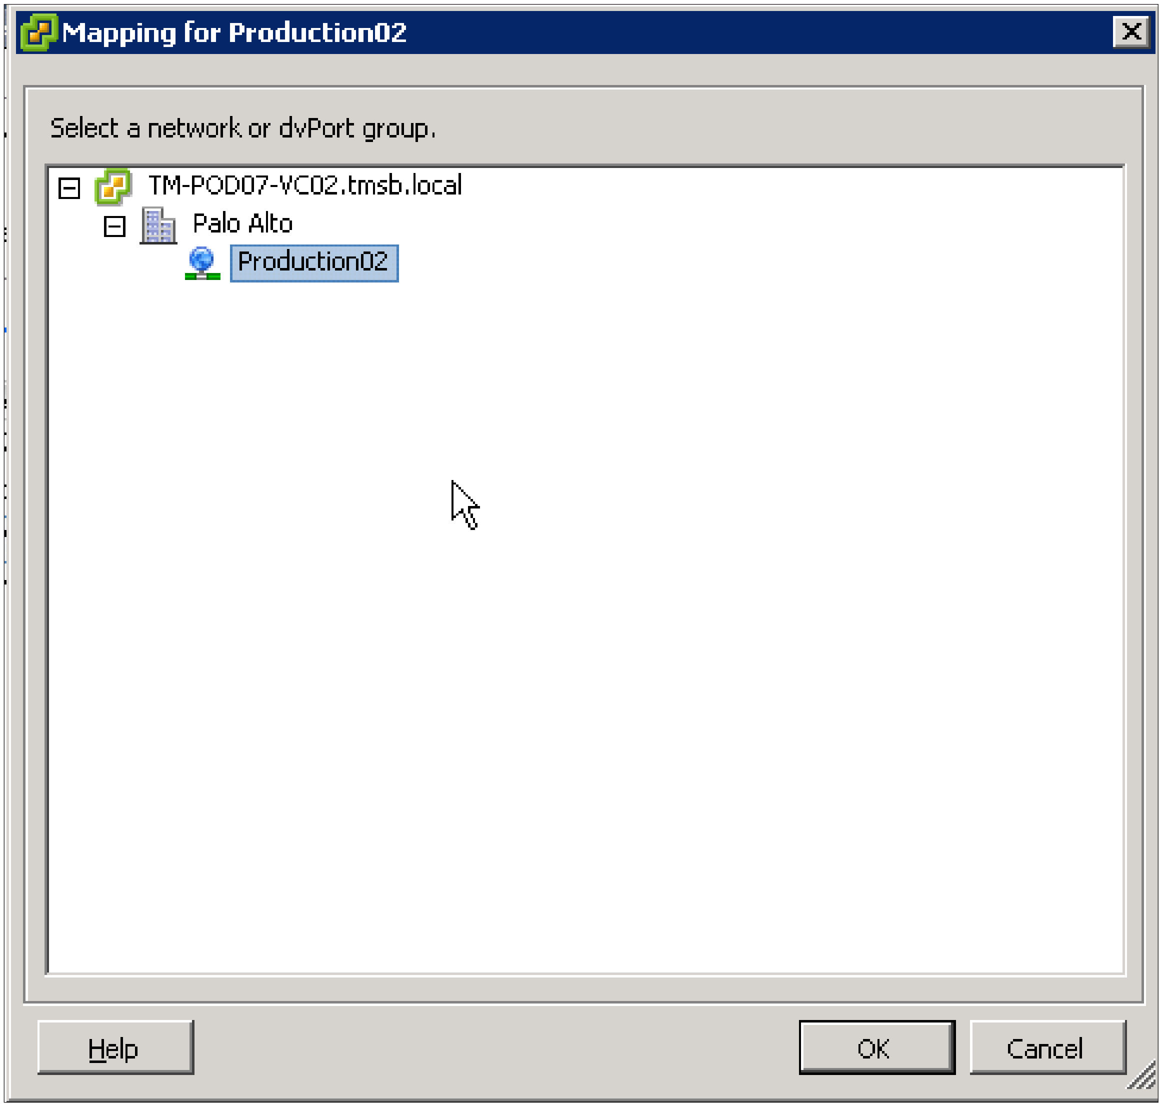

Mapping for Production02

Select a network or dvPort group.

- TM-POD07-VC02.tmsb.local
  - Palo Alto
    - Production02

Figure 27. Selecting a Recovery-Site Network to Map

6. To undo an inventory mapping, select the row to be unconfigured and click **Remove Mapping.**

Step 4: Set up protection group

SRM organizes virtual machines into protection groups based on the datastore group that they use. All virtual machines in a protection group store their files within the same datastore group, and all failover together.

Procedure

1. Open the vSphere Client and connect to the vCenter server at the protected site. Log in as a vSphere administrator.
2. Click the **Site Recovery** icon on the vSphere Client Home page.
3. Select the **Protection Groups** line in the navigation pane on the left side of the screen.

4. Select All Protection Groups in the left pane, and choose the **Summary** tab in the right pane.

5. In the Commands box on the far right, click **Create Protection Group.**

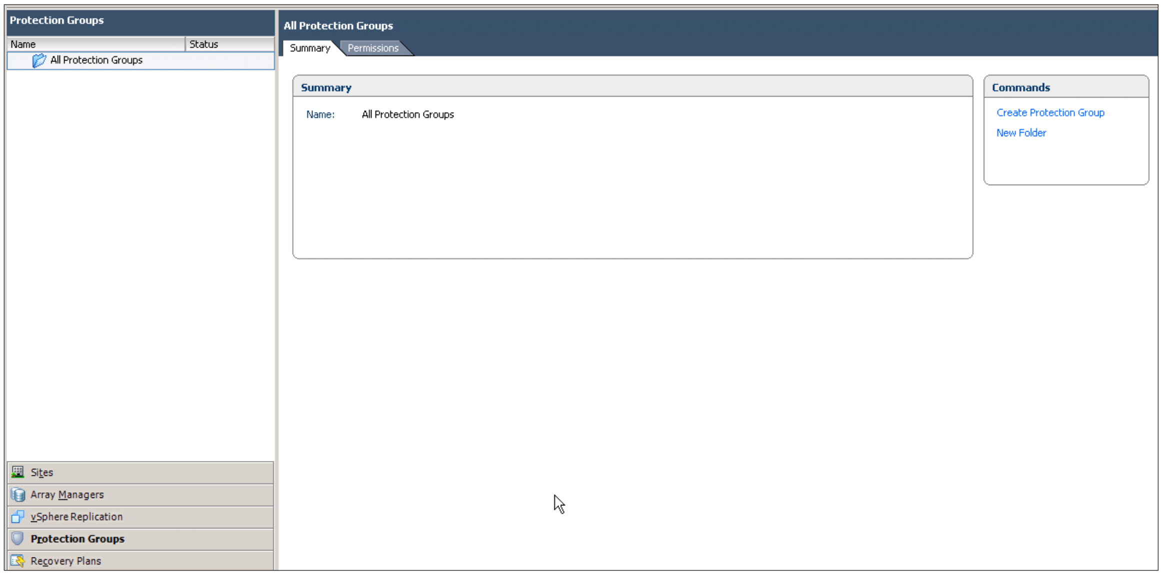

Figure 28. Protection Groups Main Screen

6. On the **Create Protection Group** pop-up screen, select the protected site (Site A) and choose **Array based replication** for the Protection Group Type. The Array Pair section should populate with the datastore group that is replicated as configured earlier in the array pairing step. If nothing appears in the Array Pair section, please return to the Array Managers configuration and rectify any problems. If everything looks correct, click **Next.**

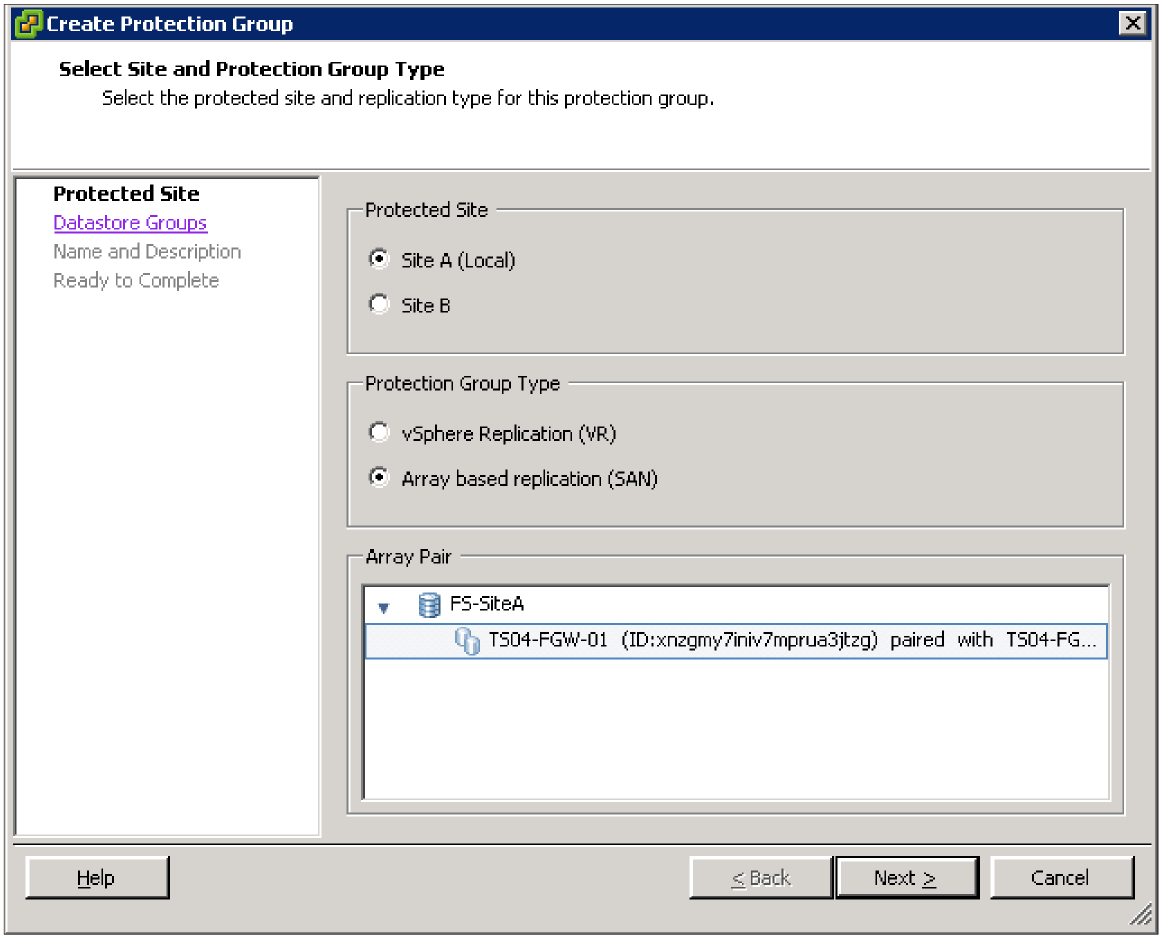

Select Site and Protection Group Type

Select the protected site and replication type for this protection group.

Protected Site

● Site A (Local)

○ Site B

Protection Group Type

○ vSphere Replication (VR)

● Array based replication (SAN)

Array Pair

FS-SiteA

TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg) paired with TS04-FG...

Figure 29. Creating a Protection Group

7. On the **Select One or More Datastore Groups** page of the Create Protection Group wizard, select one or more datastore groups from the list, and then click **Next.**

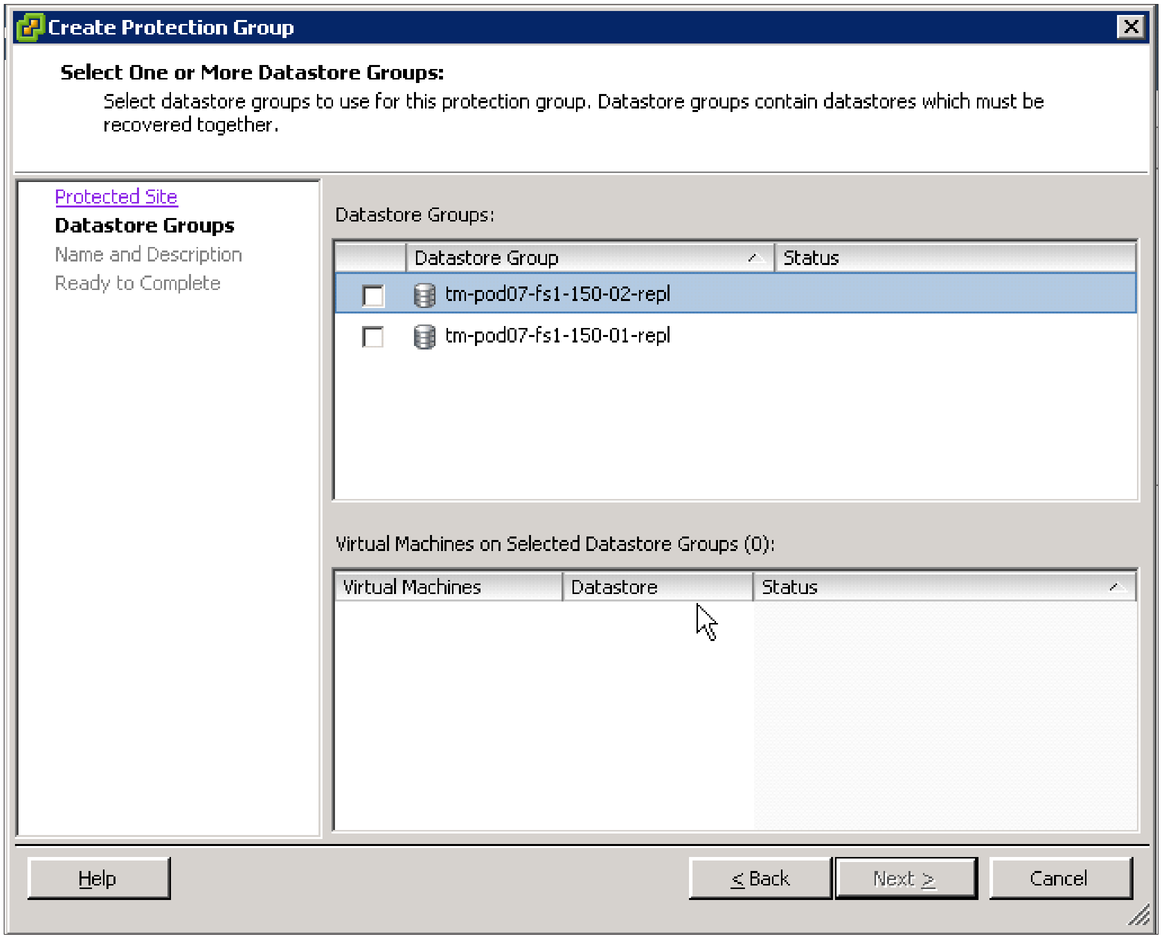

Select One or More Datastore Groups:

Select datastore groups to use for this protection group. Datastore groups contain datastores which must be recovered together.

Datastore Groups:

	Datastore Group	Status
□	tm-pod07-fs1-150-02-repl	
□	tm-pod07-fs1-150-01-repl	

Virtual Machines on Selected Datastore Groups (0):

Virtual Machines	Datastore	Status

Figure 30. Select a Datastore Group for the Protection Group Being Created

The datastore groups listed on this page are the ones that were discovered as replicated datastores when you configured the array managers. Each datastore in the list is replicated to the recovery site, and supports at least one virtual machine at the protected site. When you select a datastore group, the virtual machines that it supports are listed in the **VMs on the selected datastore group** field, and are automatically included in the protection group. You may select more than one datastore group to be part of the protection group. All VMs in all datastores selected will now be handled as one logical protection group with regards to recovery plans. In other words, if more than one datastore group is chosen, the VMs in this protection group will all be failed over together during execution of a recovery plan.

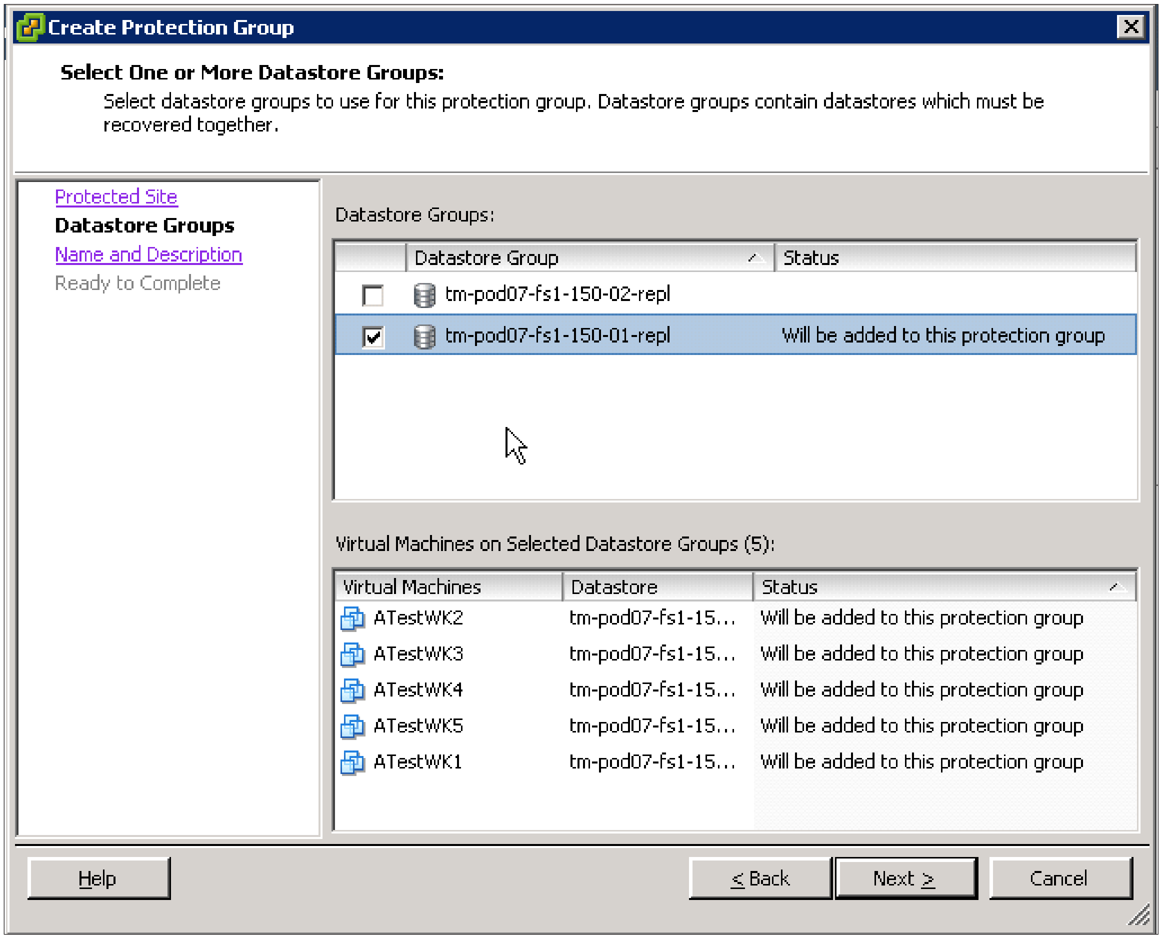

Select One or More Datastore Groups:

Select datastore groups to use for this protection group. Datastore groups contain datastores which must be recovered together.

Datastore Groups:

	Datastore Group	Status
□	tm-pod07-fs1-150-02-repl	
☑	tm-pod07-fs1-150-01-repl	Will be added to this protection group

Virtual Machines on Selected Datastore Groups (5):

Virtual Machines	Datastore	Status
ATestWK2	tm-pod07-fs1-15...	Will be added to this protection group
ATestWK3	tm-pod07-fs1-15...	Will be added to this protection group
ATestWK4	tm-pod07-fs1-15...	Will be added to this protection group
ATestWK5	tm-pod07-fs1-15...	Will be added to this protection group
ATestWK1	tm-pod07-fs1-15...	Will be added to this protection group

Figure 31. Selecting and Showing the VMs Contained in a Datastore Group Chosen for a Protection Group

8. Enter a name and description for the protection group and click **Next.**

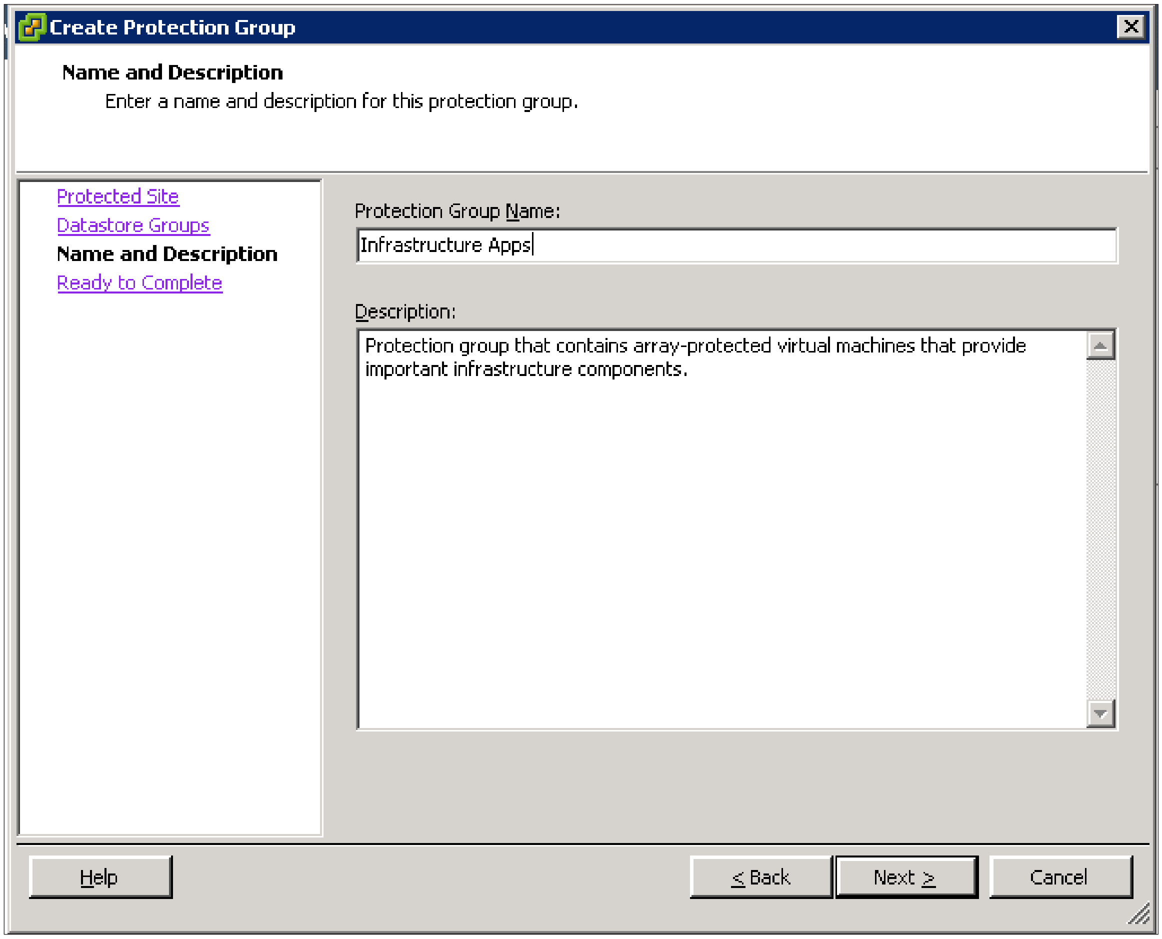

Name and Description

Enter a name and description for this protection group.

Protection Group Name: Infrastructure Apps

Description: Protection group that contains array-protected virtual machines that provide important infrastructure components.

Figure 32. Name and Description of Protection Group

9. Click **Finish** to create the protection group.

SRM creates a protection group that includes all of the virtual machines on the datastore you selected in Step 7. Placeholders are created and inventory mappings are applied for each member of the group. If any group member cannot be mapped to a folder, network, and resource pool on the recovery site, it is listed with a status of **Mapping Missing,** and no placeholder can be created for it.

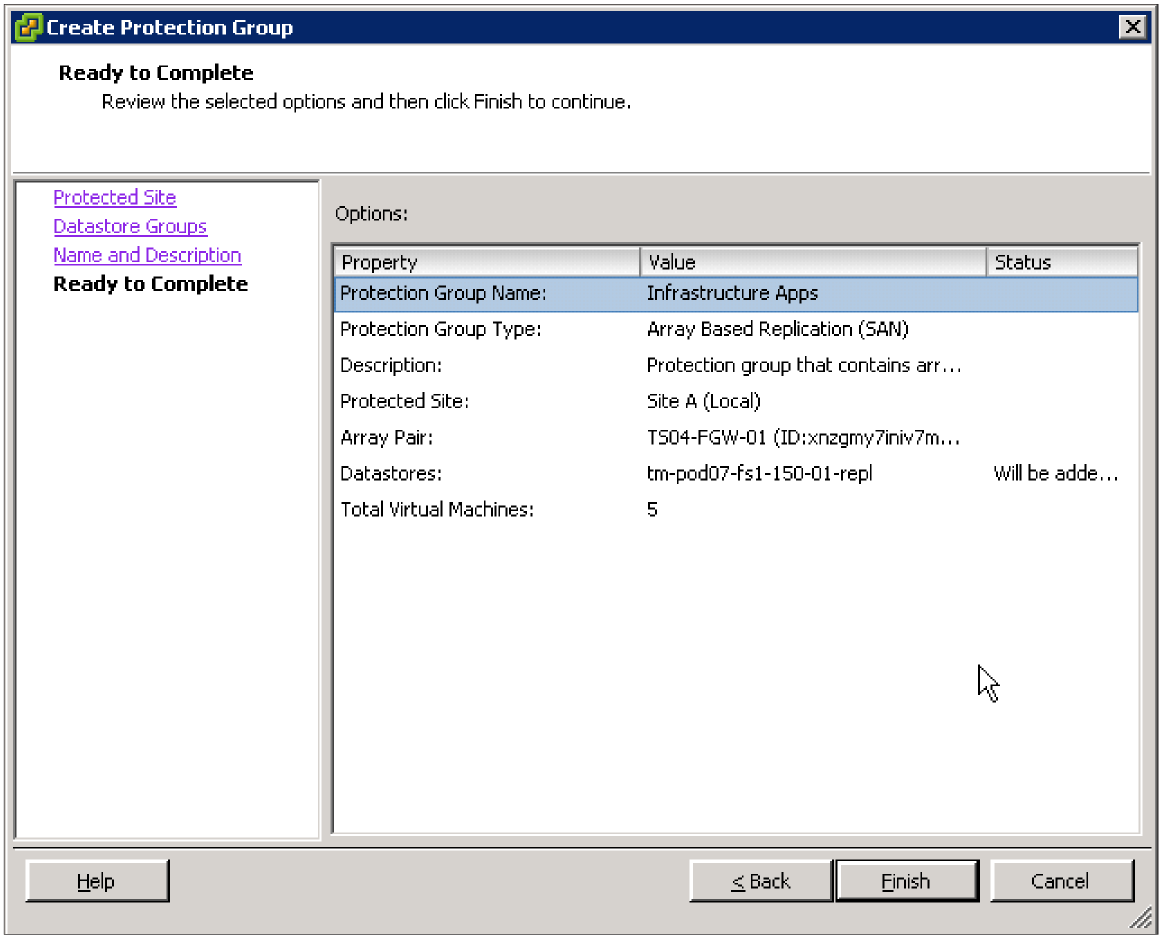

Ready to Complete

Review the selected options and then click Finish to continue.

Options:

Property	Value	Status
Protection Group Name:	Infrastructure Apps	
Protection Group Type:	Array Based Replication (SAN)	
Description:	Protection group that contains arr...	
Protected Site:	Site A (Local)	
Array Pair:	TS04-FGW-01 (ID:xnzgmy7iniv7m...	
Datastores:	tm-pod07-fs1-150-01-repl	Will be adde...
Total Virtual Machines:	5	

Figure 33. Summary of the Protection Group Prior to Completion

10. After the protection group is created (this may take a few moments while protection is configured for the VMs), you may review the status of the protection group by clicking on the name of the protection group in the navigation pane on the left.

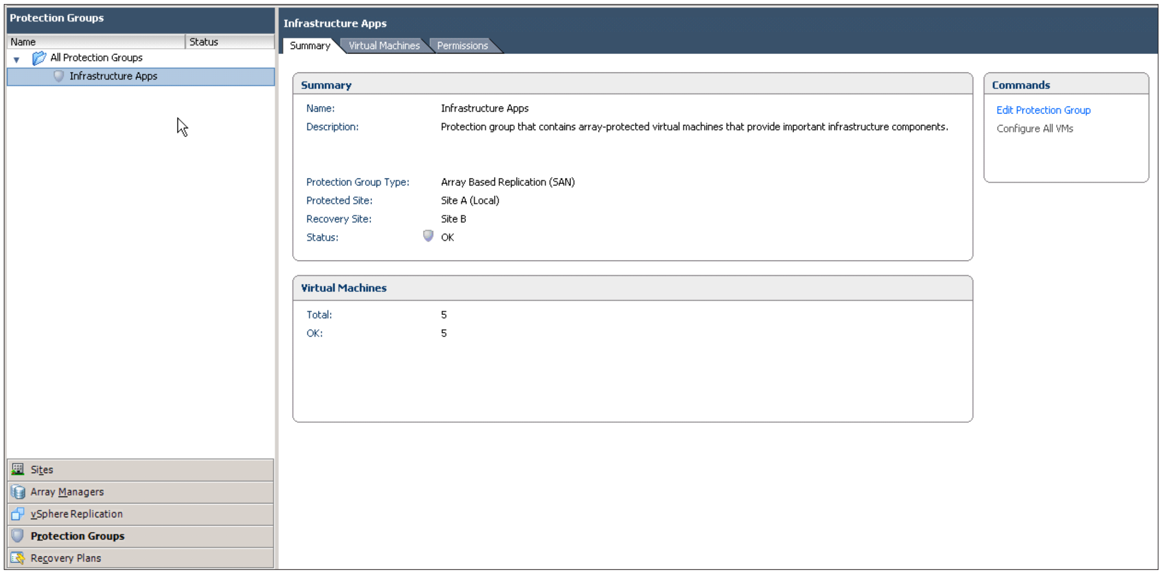

Summary

Name: Infrastructure Apps

Description: Protection group that contains array-protected virtual machines that provide important infrastructure components.

Protection Group Type: Array Based Replication (SAN)

Protected Site: Site A (Local)

Recovery Site: Site B

Status: OK

Virtual Machines

Total: 5

OK: 5

Figure 34. Overview of the Protection Group

You might also examine individual VMs within the protection group to look for the status of protection, missing mappings, or to manually configure the protection for one or all VMs in the protection group by clicking on the name of the protection group in the navigation pane, then selecting the **Virtual Machines** tab on the main pane on the right.

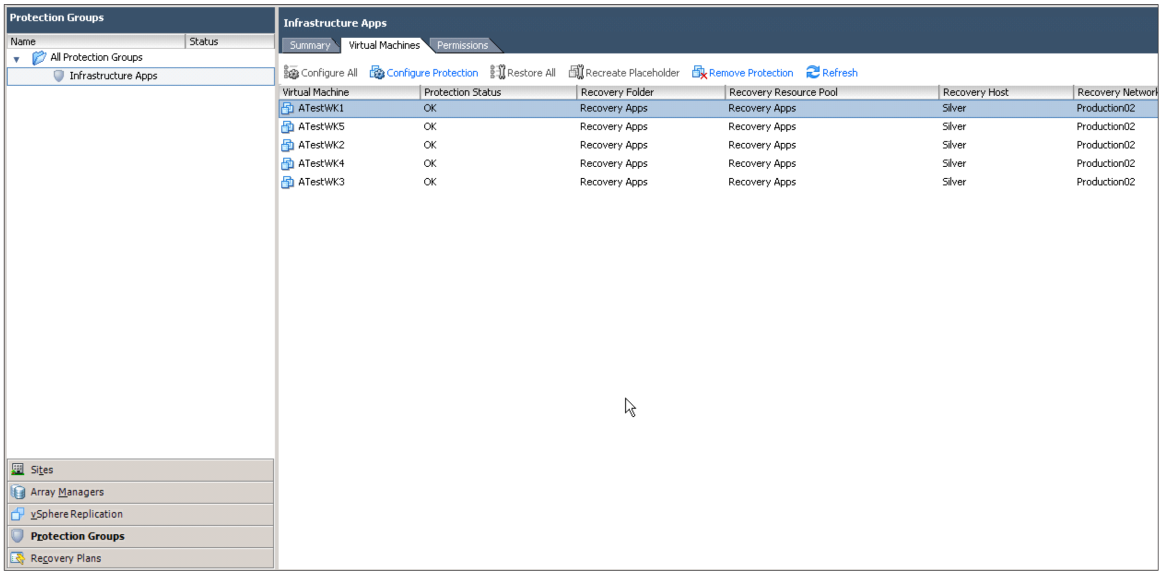

Virtual Machine	Protection Status	Recovery Folder	Recovery Resource Pool	Recovery Host	Recovery Network
ATestWK1	OK	Recovery Apps	Recovery Apps	Silver	Production02
ATestWK5	OK	Recovery Apps	Recovery Apps	Silver	Production02
ATestWK2	OK	Recovery Apps	Recovery Apps	Silver	Production02
ATestWK4	OK	Recovery Apps	Recovery Apps	Silver	Production02
ATestWK3	OK	Recovery Apps	Recovery Apps	Silver	Production02

Figure 35. VM Status within a Protection Group

Step 5: Set up recovery plan

A recovery plan controls the way in which virtual machines in a protection group are recovered. It is stored in the SRM database at the recovery site, and executed by the SRM Server at the recovery site.

A simple recovery plan assigns all virtual machines in a protection group to two networks on the recovery site – a recovery network, and a test network. The recovery network is used in an actual recovery. The test network is a special network that is used only for testing the recovery plan, and does not typically allow the recovered virtual machines to communicate on your corporate network or the Internet. SRM can create a test network that exists only on one ESX server for you, or you can create one yourself. SRM supports a recovery network that spans across the ESX servers at the recovery site (in other words, the vNetwork Distributed Switch, or vDS). In case your recovery plan calls for the need of the vDS, you can create the vDS yourself for testing and failover recovery purposes.

A simple recovery plan includes a number of prescribed steps that use default values to control how protection group members are migrated to the protected site. You can customize a recovery plan to change default values, add steps to the plan itself and to the recovery of individual virtual machines, suspend noncritical virtual machines at the recovery site to make resources available for recovered machines, and so on.

Procedure

1. Open the vSphere Client and connect to the vCenter server at the recovery site. Log in as a vSphere administrator.
2. Click the **Site Recovery** icon on the vSphere Client Home page.
3. In the navigation pane on the left side of the SRM window, select the **Recovery Plans** item at the bottom, select **All Recovery Plans** in the top-left pane, click the **Summary** tab in the main navigation window, and click **Create Recovery Plan** in the **Commands** box on the right side of the screen.

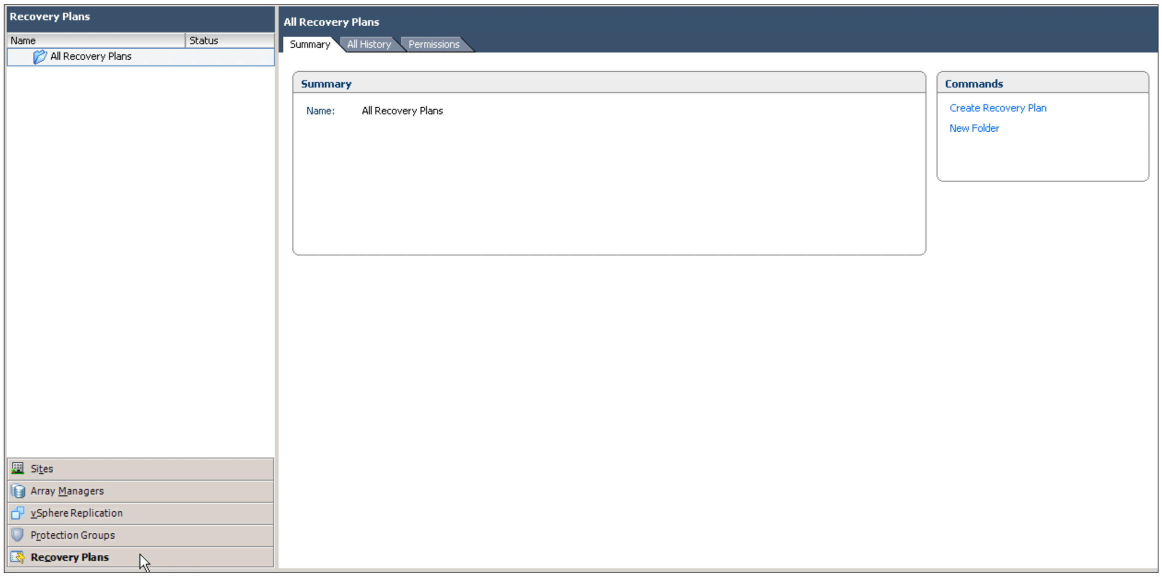

Name: All Recovery Plans

Figure 36. Create a Recovery Plan

4. In the **Create Recovery Plan** wizard, choose a recovery site. Ensure that you are choosing the recovery site in this step (Site B), because this is asking to which site VMs will be failed over, not for the source of virtual machines.

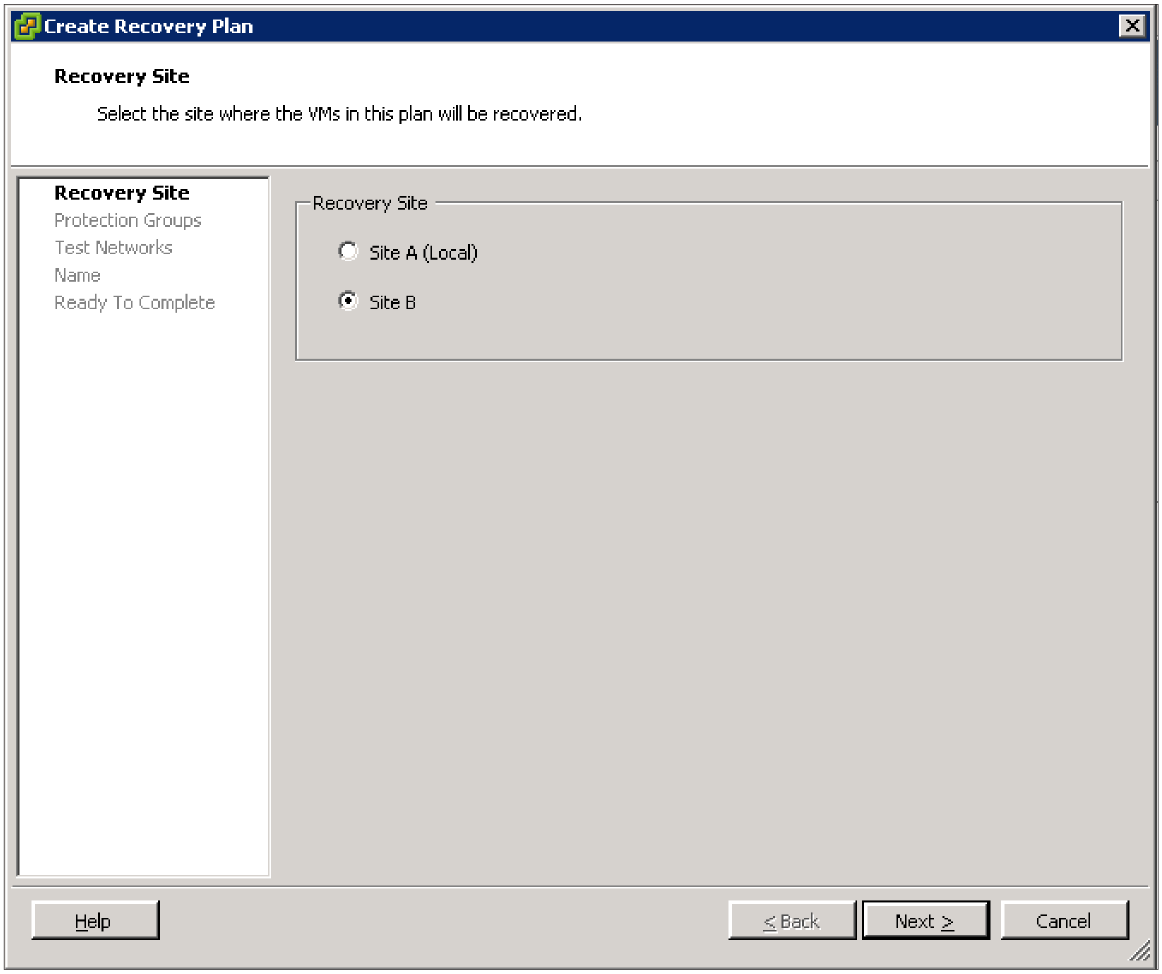

Create Recovery Plan

Recovery Site

Select the site where the VMs in this plan will be recovered.

Recovery Site

○ Site A (Local)

● Site B

Figure 37. Choosing a Recovery Site

5. On the **Select Protection Groups** page of the Create Recovery Plan wizard, select one or more protection groups for the plan to recover, and then click **Next.**

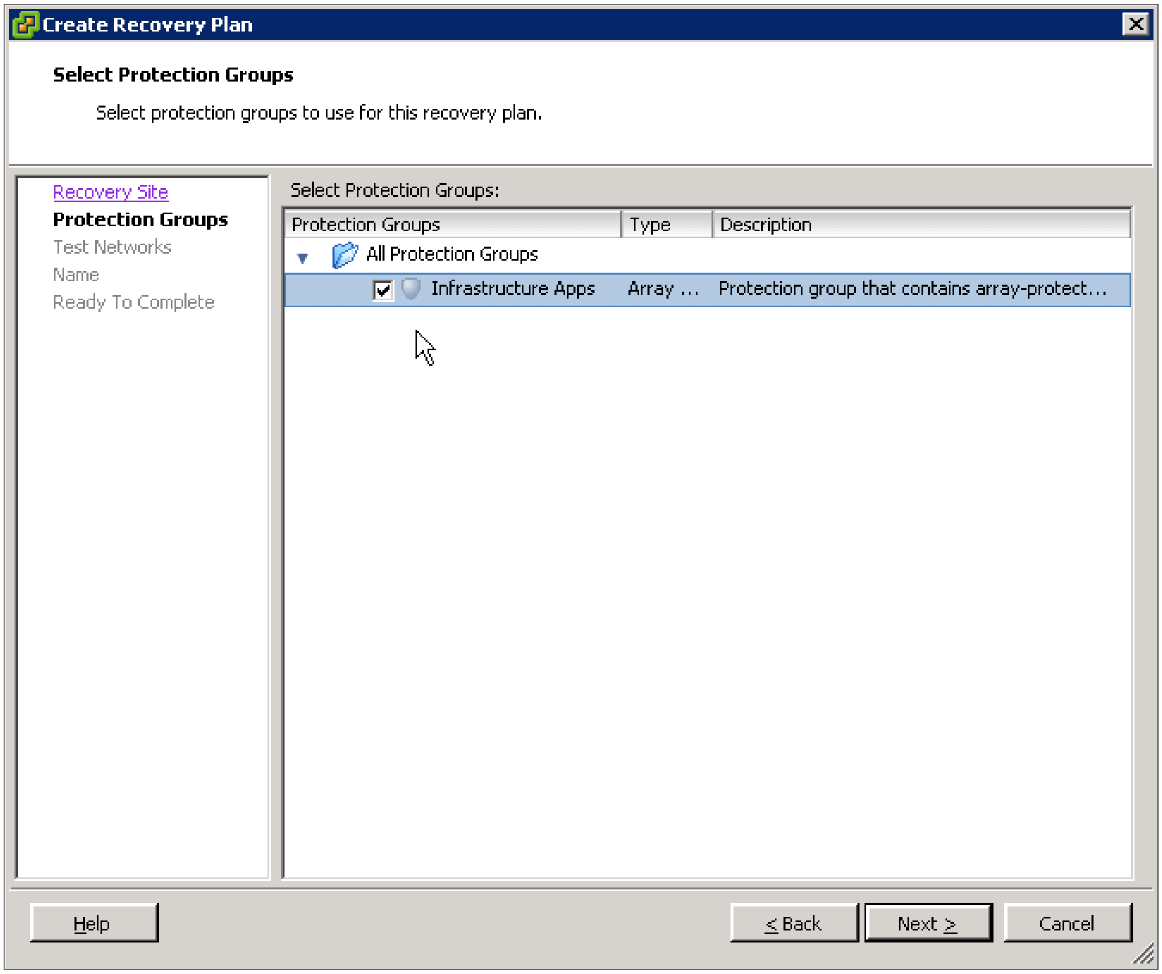

Create Recovery Plan

Select Protection Groups

Select protection groups to use for this recovery plan.

Select Protection Groups:

Protection Groups	Type	Description
All Protection Groups		
☑ Infrastructure Apps	Array …	Protection group that contains array-protect…

Figure 38. Select Protection Group for the Recovery Plan

6. On the **Test Networks** page of the Create Recovery Plan wizard, select a recovery site network to which recovered virtual machines will connect during recovery plan tests, and then click **Next.**

   By default, the test network is specified as **Auto,** which will automatically create an isolated test network on each ESX server that is part of the test. If you would prefer to specify an existing recovery site network as the test network (for example, a vDS port group that spans across your recovery ESX servers, or an isolated VLAN), click **Auto** and select the network from the drop-down control.

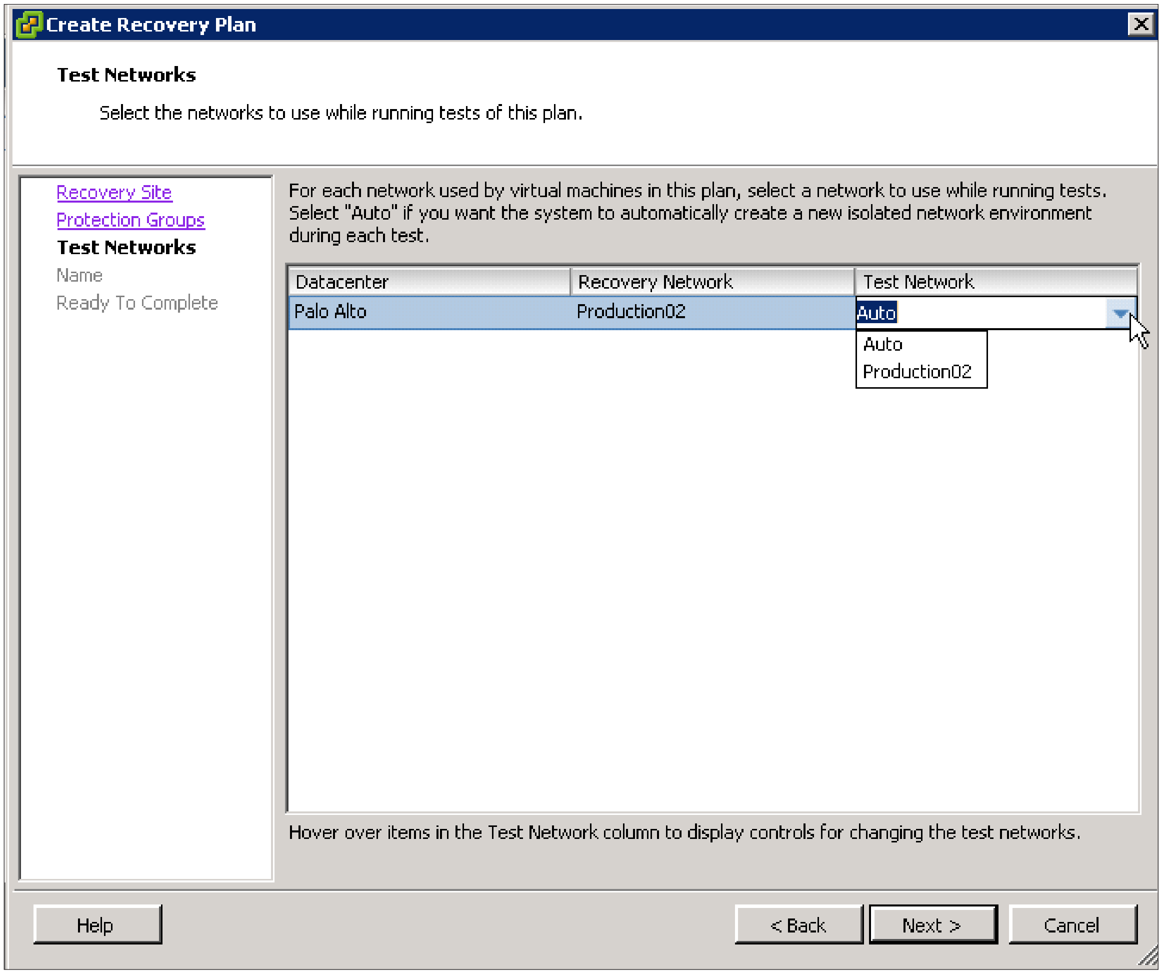

Create Recovery Plan

Test Networks

Select the networks to use while running tests of this plan.

For each network used by virtual machines in this plan, select a network to use while running tests. Select "Auto" if you want the system to automatically create a new isolated network environment during each test.

Datacenter	Recovery Network	Test Network
Palo Alto	Production02	Auto

Hover over items in the Test Network column to display controls for changing the test networks.

Figure 39. Configure Test Network for Recovery Plan

7. Enter a name and description for the recovery plan and click **Next.**
8. Read the summary of the recovery plan in the Ready to Complete screen, and click **Finish** to create the recovery plan.

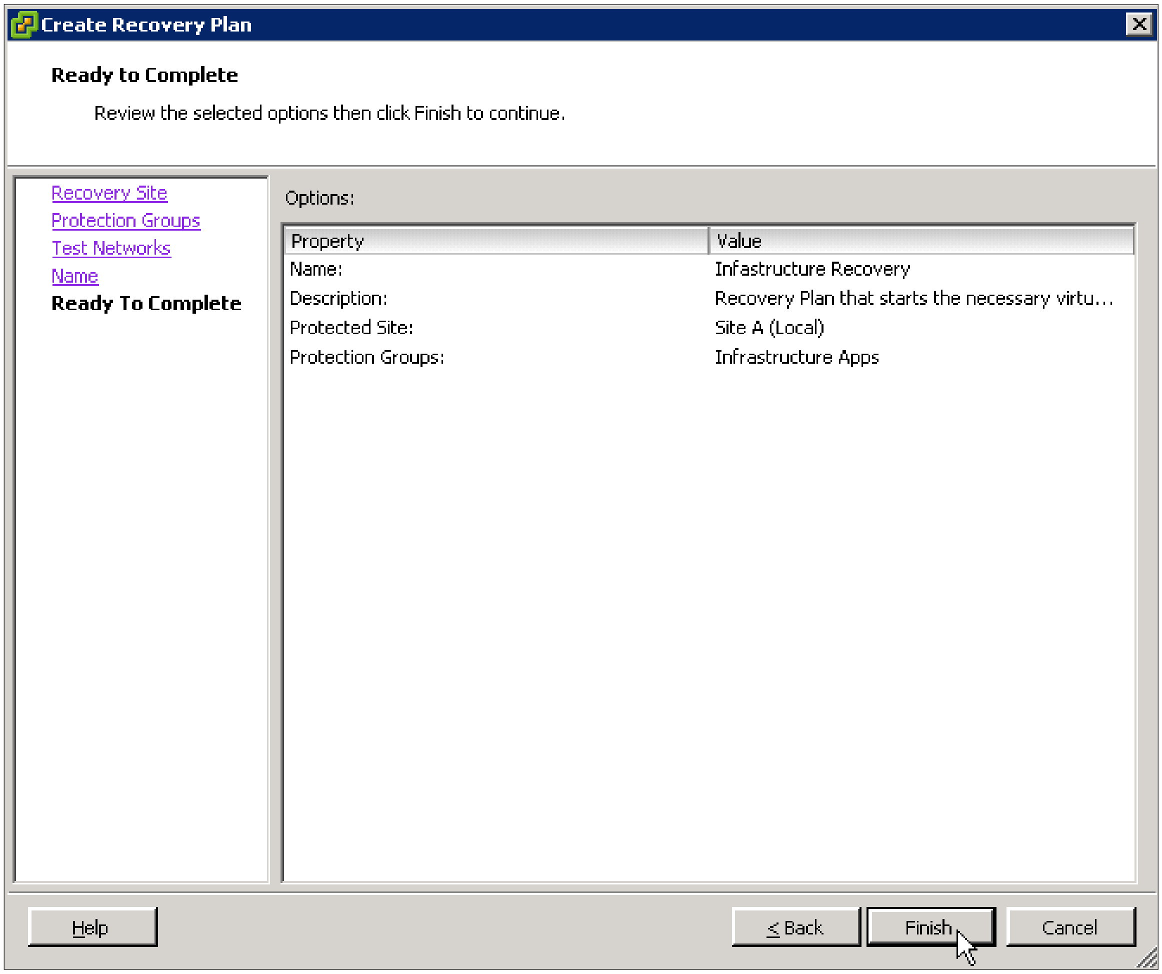

Ready to Complete

Review the selected options then click Finish to continue.

Options:

Property	Value
Name:	Infastructure Recovery
Description:	Recovery Plan that starts the necessary virtu...
Protected Site:	Site A (Local)
Protection Groups:	Infrastructure Apps

Figure 40. Finish Creation of the Recovery Plan

Step 6: Customize IP properties

There are a number of ways to access IP customization for virtual machines within a recovery plan. The easiest way is to select the recovery plan in the navigation field on the left side of the screen, and in the main window, select the **Virtual Machines** tab. Choose the VM to customize, and click the **Configure Recovery** button.

1. You might also browse to the virtual machine to be customized from within the **Recovery Steps** tab of the recovery plan, right-click the VM, and choose **Configure.**

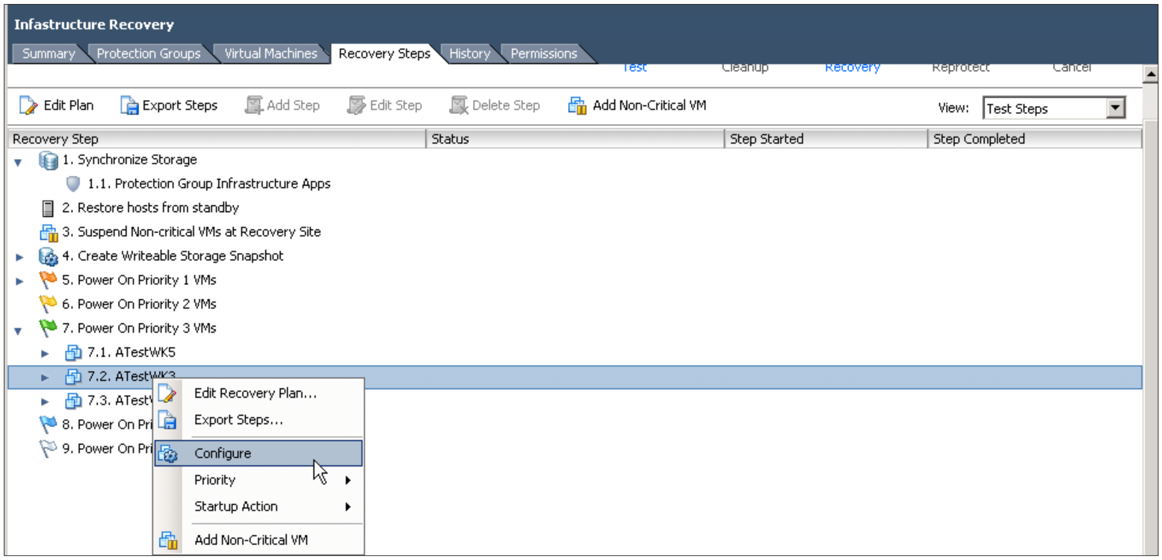

Recovery Step	Status	Step Started	Step Completed
1. Synchronize Storage			
1.1. Protection Group Infrastructure Apps			
2. Restore hosts from standby			
3. Suspend Non-critical VMs at Recovery Site			
4. Create Writeable Storage Snapshot			
5. Power On Priority 1 VMs			
6. Power On Priority 2 VMs			
7. Power On Priority 3 VMs			
7.1. ATestWK5			
7.2. ATestWK3			
7.3. ATest			
8. Power On Pri			
9. Power On Pri			

Figure 41. Right-Click the VM within a Recovery Plan to Access Its Configuration Properties

Regardless of the method chosen to access configuration, the same screen will pop up, showing recovery plan configuration information for this virtual machine.

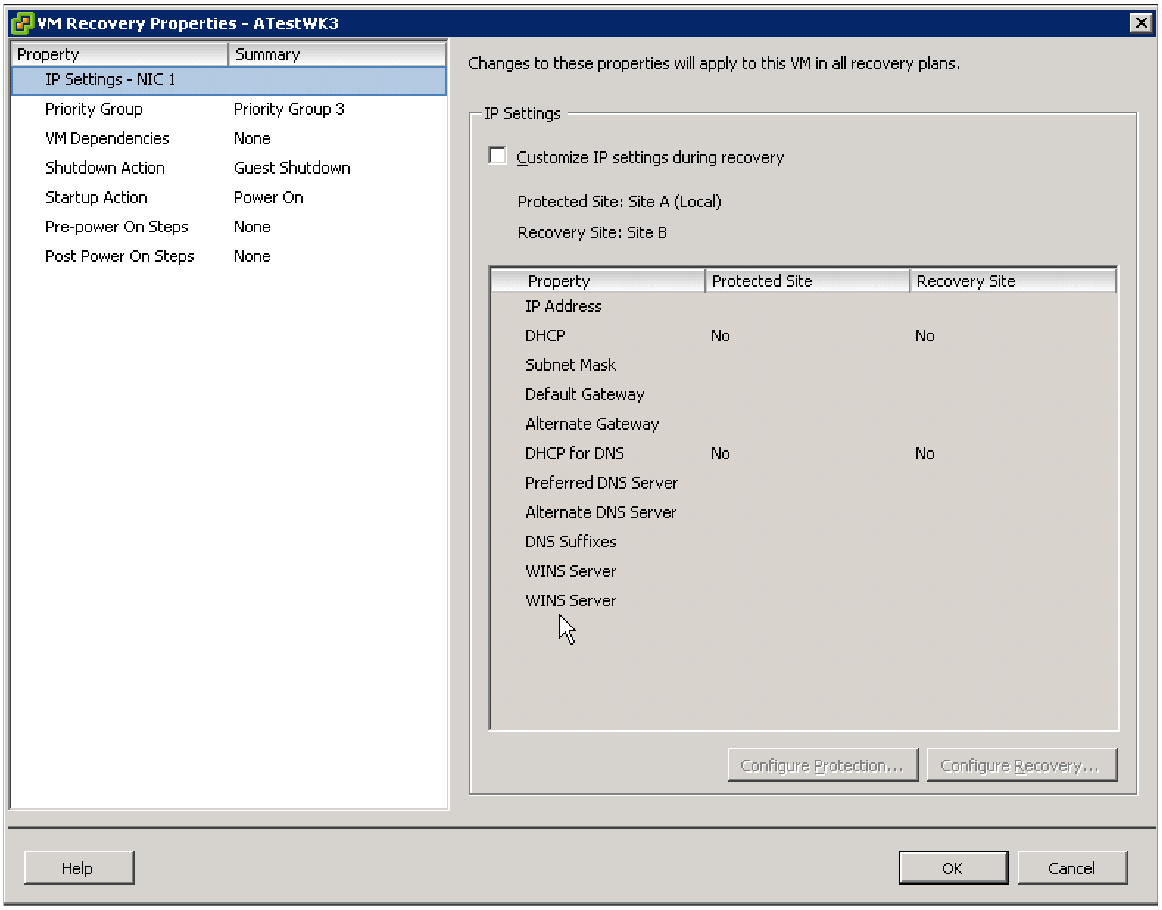

VM Recovery Properties - ATestWK3

Property	Summary
IP Settings - NIC 1	
Priority Group	Priority Group 3
VM Dependencies	None
Shutdown Action	Guest Shutdown
Startup Action	Power On
Pre-power On Steps	None
Post Power On Steps	None

Changes to these properties will apply to this VM in all recovery plans.

IP Settings

□ Customize IP settings during recovery

Protected Site: Site A (Local)

Recovery Site: Site B

Property	Protected Site	Recovery Site
IP Address		
DHCP	No	No
Subnet Mask		
Default Gateway		
Alternate Gateway		
DHCP for DNS	No	No
Preferred DNS Server		
Alternate DNS Server		
DNS Suffixes		
WINS Server		
WINS Server		

Figure 42. Properties for Customization of the VM

SRM also provides a **batch IP customization tool,** dr-ip-customizer.exe. Using dr-ip-customizer enables a very rapid bulk import and changes IP information of many or all virtual machines. Refer to the *Site Recovery Manager Administration Guide* for information regarding dr-ip-customizer.

2. Select the **Customize IP settings during recovery** check to enable customization of network information, then click **Configure Recovery** to customize which IP addresses will be injected to the VM during execution of a recovery plan.

3. Enter all networking information for the virtual machine at the recovery site, including an update of the **DNS** tab and **WINS** tab, if required.

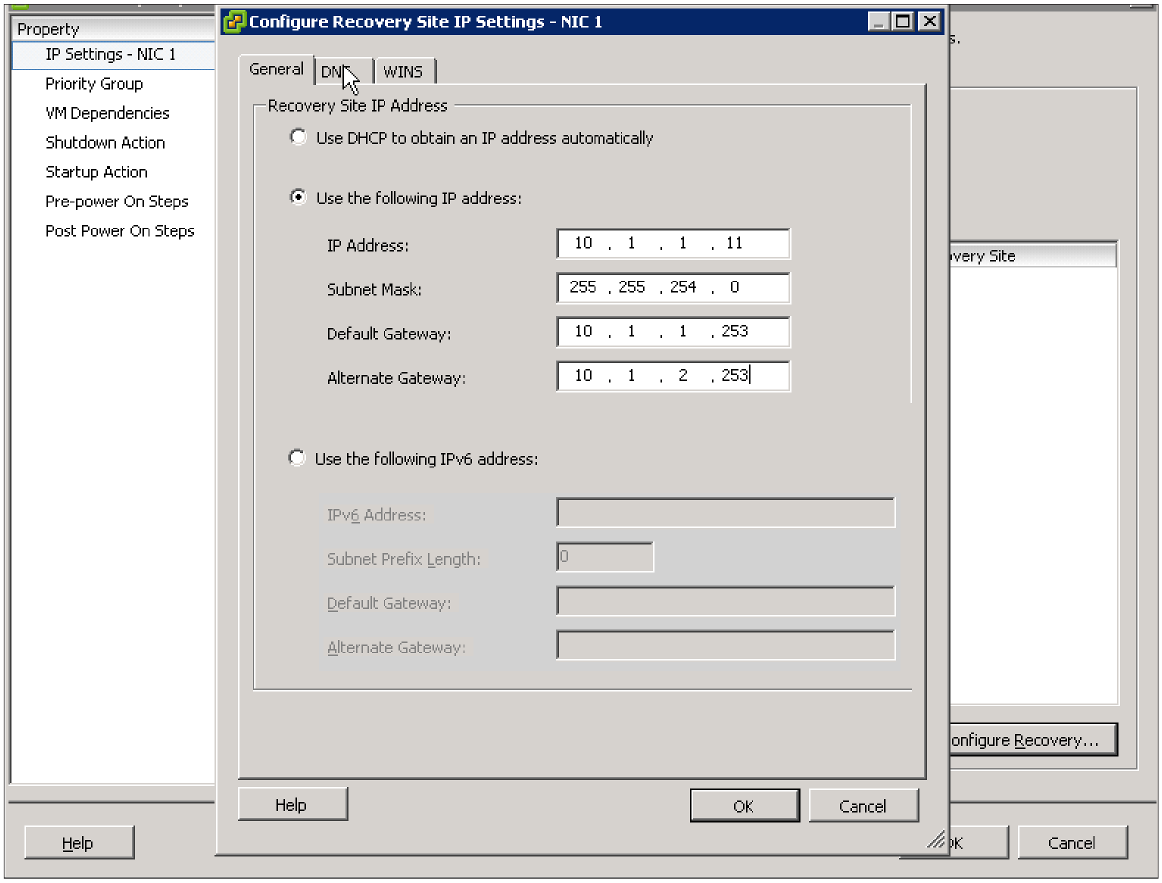

Recovery Site IP Address

○ Use DHCP to obtain an IP address automatically

● Use the following IP address:

IP Address: 10 . 1 . 1 . 11

Subnet Mask: 255 . 255 . 254 . 0

Default Gateway: 10 . 1 . 1 . 253

Alternate Gateway: 10 . 1 . 2 . 253

○ Use the following IPv6 address:

IPv6 Address:

Subnet Prefix Length: 0

Default Gateway:

Alternate Gateway:

Figure 43. Configuring Recovery Site IP Information

4. Click **OK** once the IP information is updated. The networking information is updated for the recovery site. Click **OK** once you are satisfied that the network customization is correct.

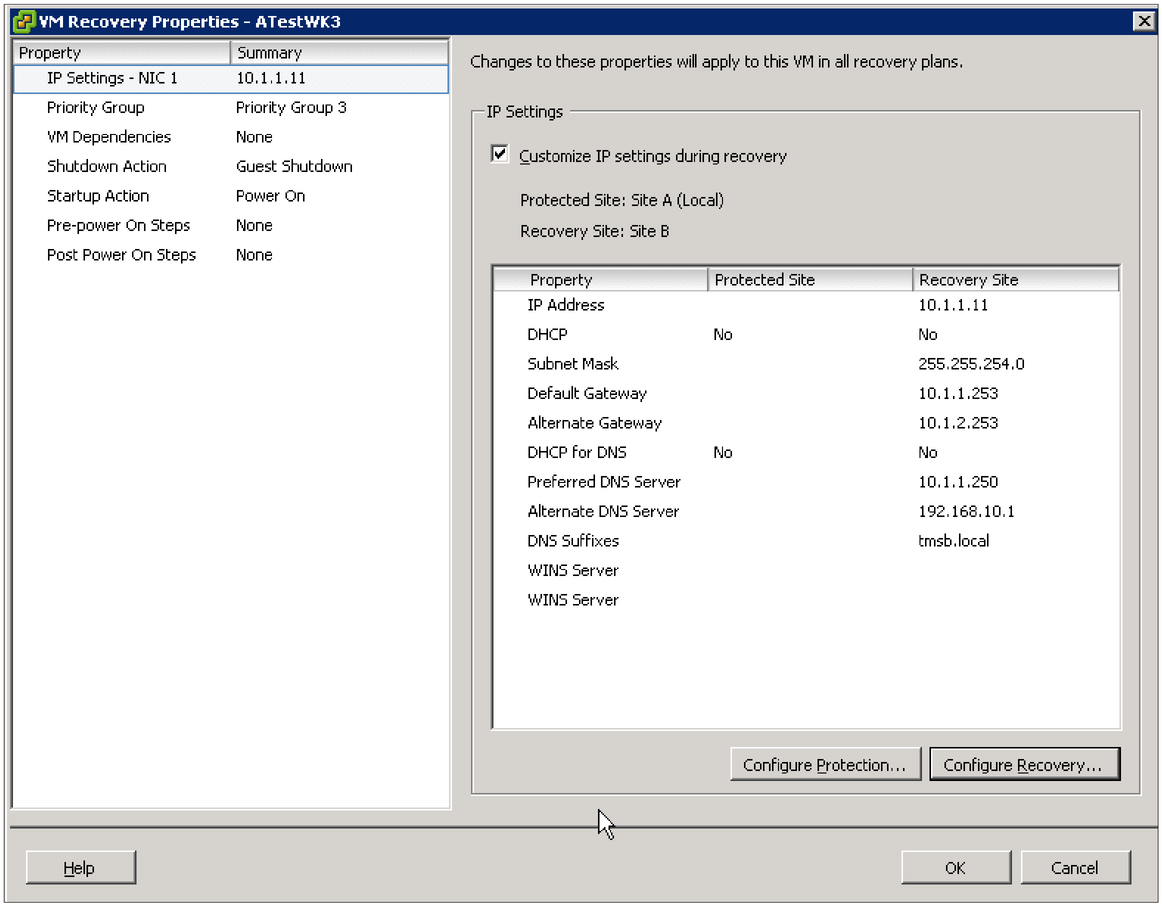

Property	Summary
IP Settings - NIC 1	10.1.1.11
Priority Group	Priority Group 3
VM Dependencies	None
Shutdown Action	Guest Shutdown
Startup Action	Power On
Pre-power On Steps	None
Post Power On Steps	None

Changes to these properties will apply to this VM in all recovery plans.

IP Settings

☑ Customize IP settings during recovery

Protected Site: Site A (Local)

Recovery Site: Site B

Property	Protected Site	Recovery Site
IP Address		10.1.1.11
DHCP	No	No
Subnet Mask		255.255.254.0
Default Gateway		10.1.1.253
Alternate Gateway		10.1.2.253
DHCP for DNS	No	No
Preferred DNS Server		10.1.1.250
Alternate DNS Server		192.168.10.1
DNS Suffixes		tmsb.local
WINS Server		
WINS Server		

Figure 44. Network Recovery Properties

Custom network information can be configured for both the protected and recovery sites. This may be of use if reprotection and automated failback will be used (both of which are only applicable to array-based replication). Feel free to enter IP information for both sites. On a large scale, the command-line interface tool dr-ip-customizer will be of much greater importance for populating mass changes to IP information in VM configurations, but is out of the scope of this evaluation.

Step 7: Configure priority groups and dependencies

SRM includes the ability to set priorities for virtual machines within a recovery plan, as well as the ability of dependencies to set policies around startup sequences for VMs or sets of VMs. Priority groups will dictate which VMs in a recovery plan will start at which stage of the recovery plan. VMs in priority group 1 will all start in parallel (unless restricted by a dependency), and only after all VMs in priority group 1 have started, will VMs in priority group 2 start, and so forth, until priority group 5 is complete. Dependencies may also be set for VMs within a priority group that will enable the administrator to dictate (as a property of a VM itself) which VMs must be running prior to attempting to start the VM with the dependency. This enables multitier applications to have a controlled start sequence.

1. To set virtual machines as part of a specific priority group, click the VM within a recovery plan and click **Configure Recovery,** as though selecting custom IP addressing. Once the properties pop-up menu is open, click on **Priority Group.**

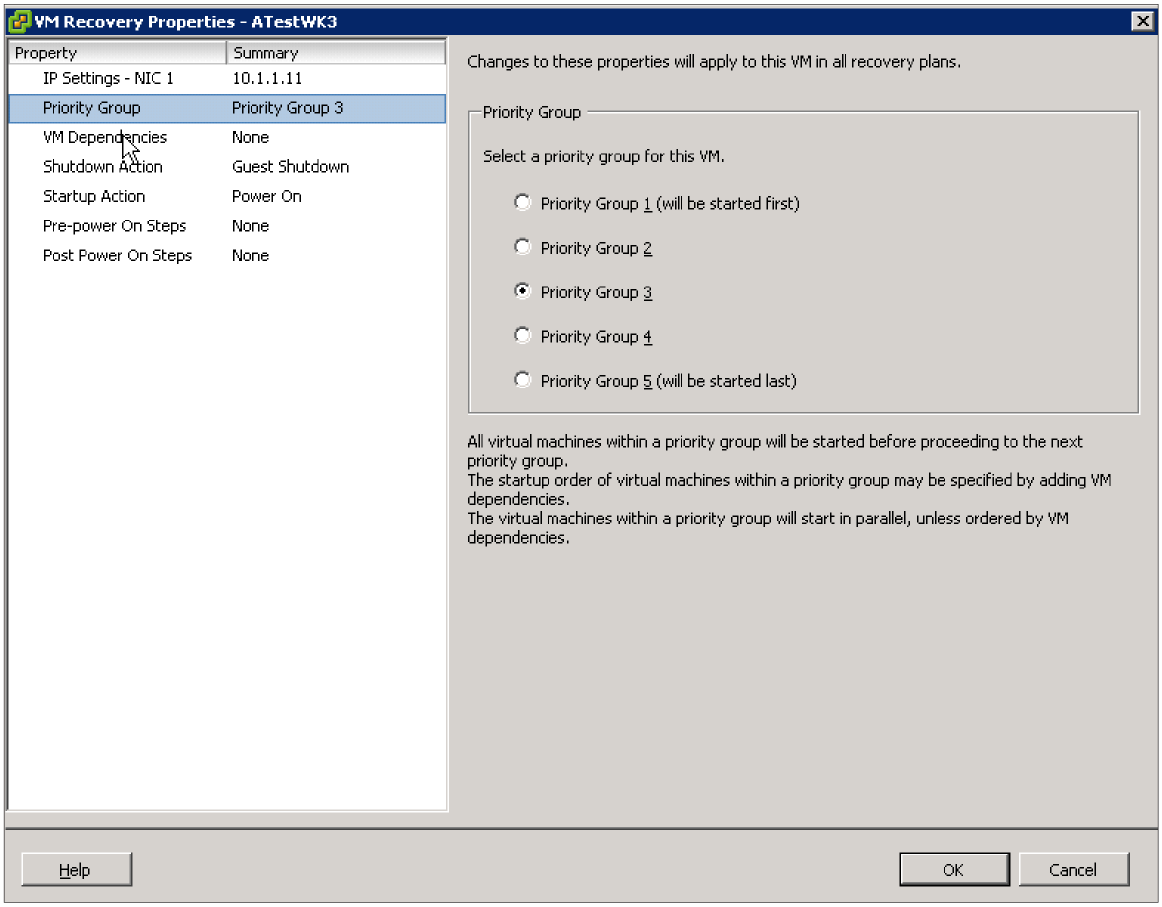

VM Recovery Properties - ATestWK3

Property	Summary
IP Settings - NIC 1	10.1.1.11
Priority Group	Priority Group 3
VM Dependencies	None
Shutdown Action	Guest Shutdown
Startup Action	Power On
Pre-power On Steps	None
Post Power On Steps	None

Changes to these properties will apply to this VM in all recovery plans.

Priority Group

Select a priority group for this VM.

○ Priority Group 1 (will be started first)

○ Priority Group 2

● Priority Group 3

○ Priority Group 4

○ Priority Group 5 (will be started last)

All virtual machines within a priority group will be started before proceeding to the next priority group.
The startup order of virtual machines within a priority group may be specified by adding VM dependencies.
The virtual machines within a priority group will start in parallel, unless ordered by VM dependencies.

Figure 45. Setting a VM's Priority Group

2. Select a priority group for the virtual machine's start sequence. Note the new priority group for the virtual machine. The priority group for a VM might also be configured by right-clicking the VM within the recovery plan and setting it directly.

3. To set a dependency for a virtual machine, click on the **VM Dependencies** line within the **VM Recovery Properties** window. Next, click **Add..** to add a new virtual machine that must be running before this VM will attempt to power on as part of the recovery process.

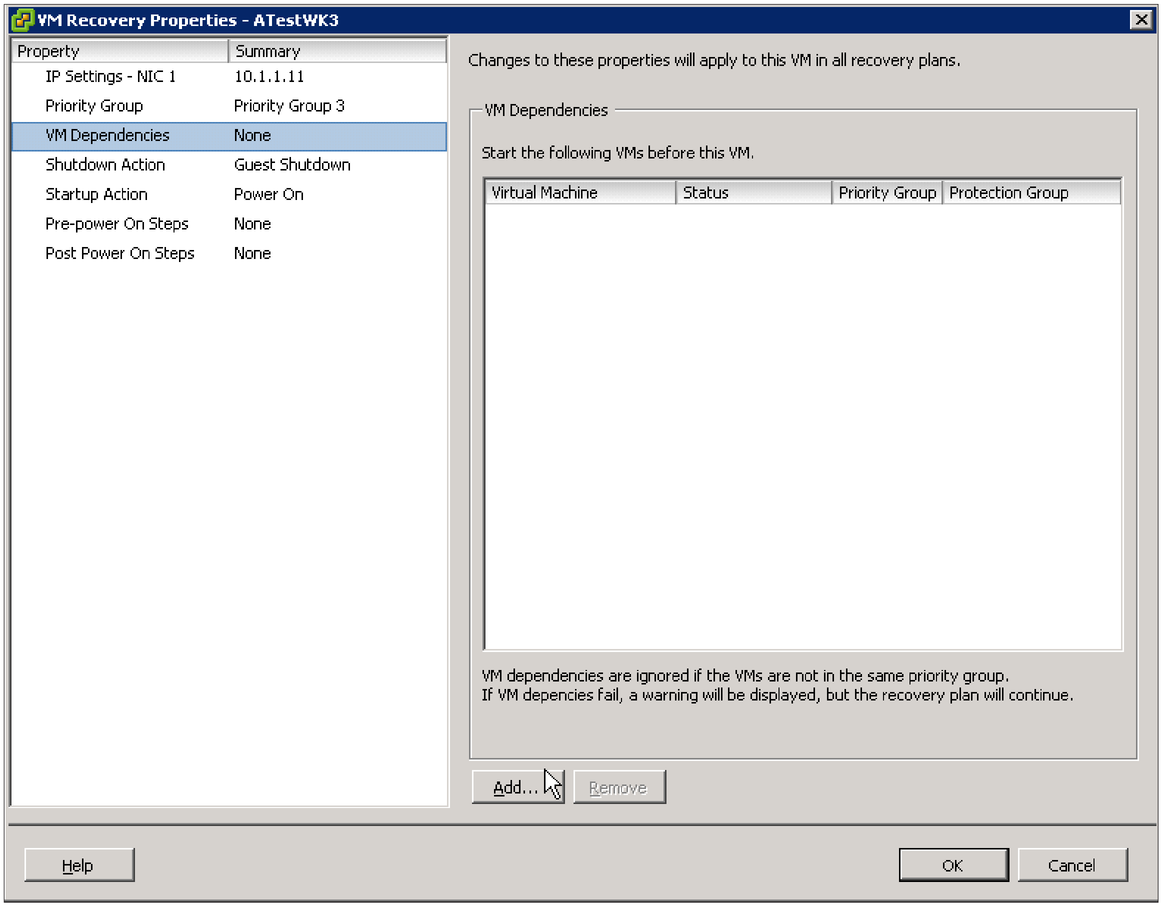

Property	Summary
IP Settings - NIC 1	10.1.1.11
Priority Group	Priority Group 3
VM Dependencies	None
Shutdown Action	Guest Shutdown
Startup Action	Power On
Pre-power On Steps	None
Post Power On Steps	None

Changes to these properties will apply to this VM in all recovery plans.

VM Dependencies

Start the following VMs before this VM.

Virtual Machine	Status	Priority Group	Protection Group

VM dependencies are ignored if the VMs are not in the same priority group.
If VM depencies fail, a warning will be displayed, but the recovery plan will continue.

Figure 46. VM Dependency Settings

4. Select a VM to add as a dependency.

NOTE: The virtual machine being added as a dependency must be a member of the same priority group as the VM whose properties you are editing, or the dependency will be ignored.

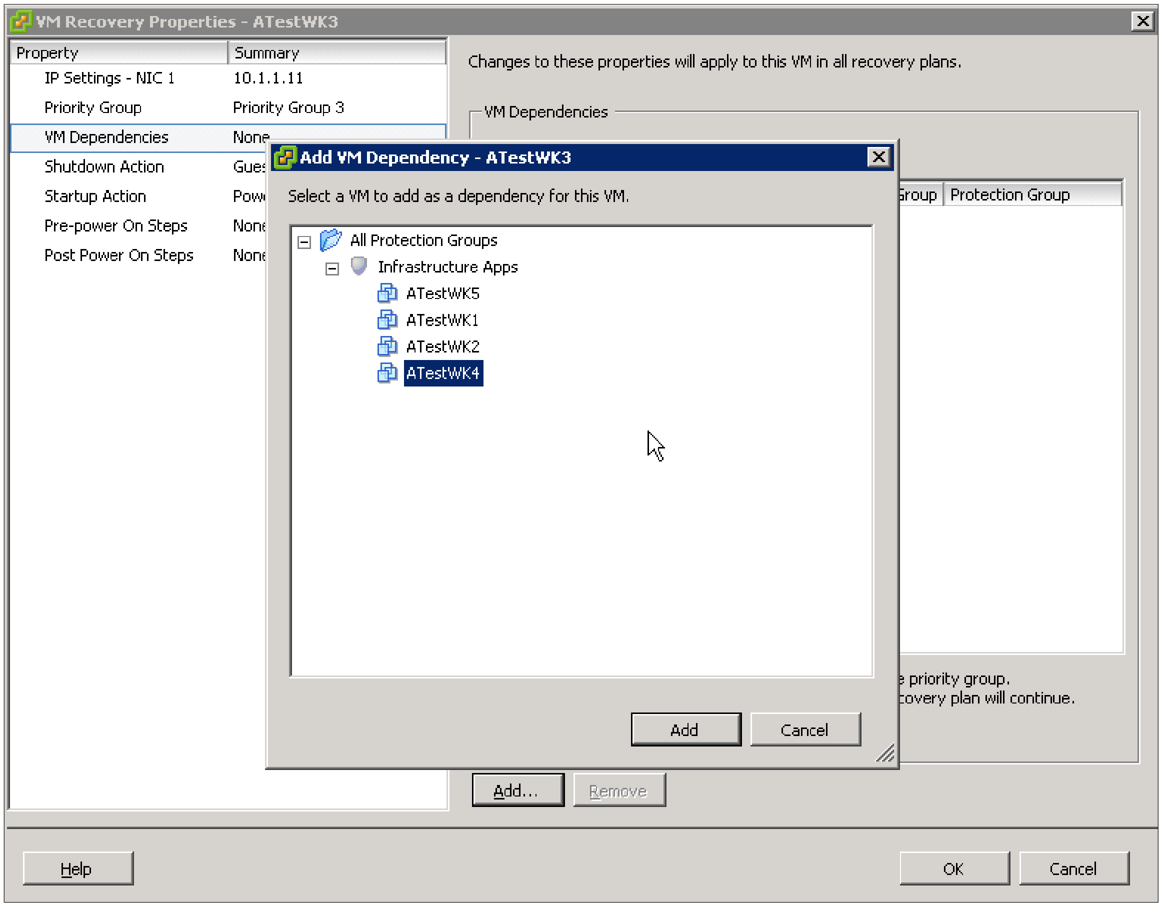

Property	Summary
IP Settings - NIC 1	10.1.1.11
Priority Group	Priority Group 3
VM Dependencies	None
Shutdown Action	Gues
Startup Action	Powe
Pre-power On Steps	None
Post Power On Steps	None

Changes to these properties will apply to this VM in all recovery plans.

Select a VM to add as a dependency for this VM.

e priority group.
covery plan will continue.

Figure 47. Adding a VM as a Dependency for the Current VM

5. After selecting an appropriate VM to act as a parent for dependency, review the dependencies, then click **OK** to continue.

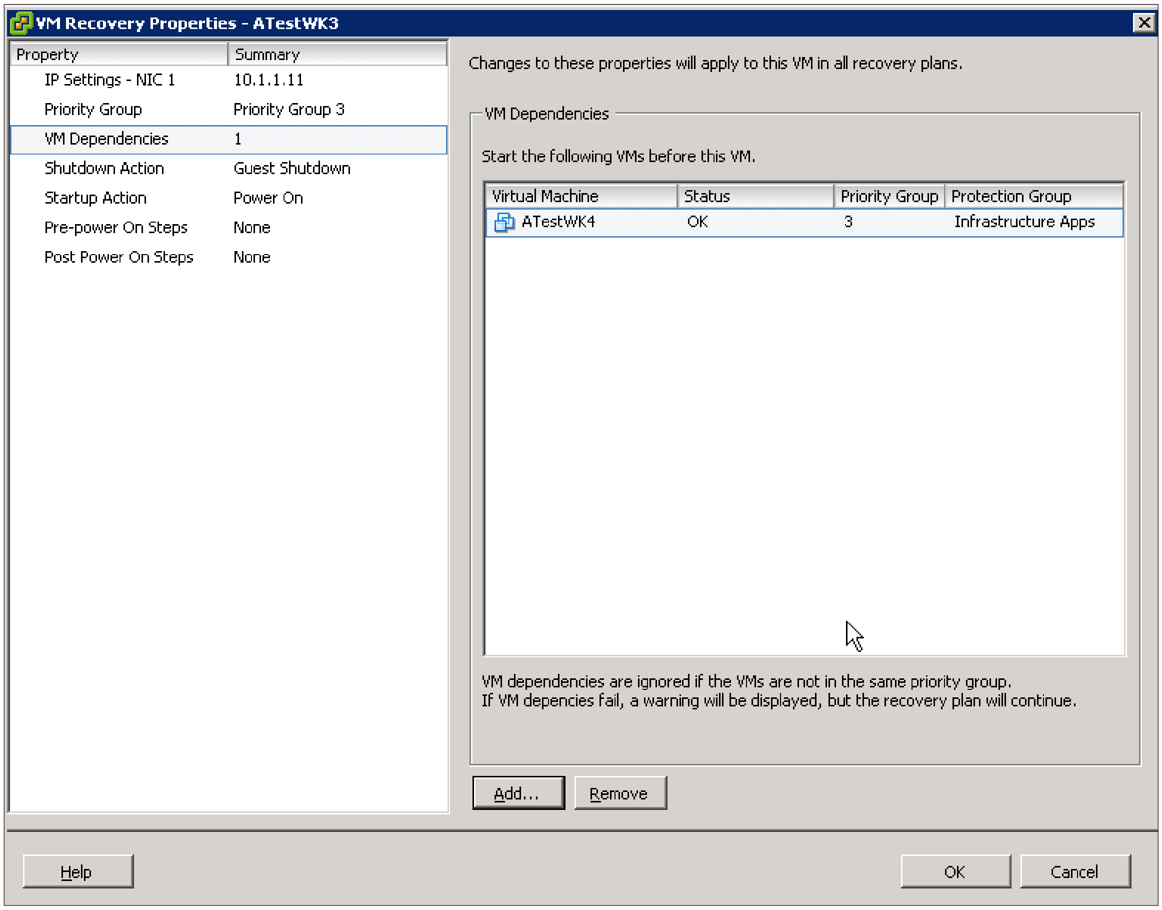

Property	Summary
IP Settings - NIC 1	10.1.1.11
Priority Group	Priority Group 3
VM Dependencies	1
Shutdown Action	Guest Shutdown
Startup Action	Power On
Pre-power On Steps	None
Post Power On Steps	None

Changes to these properties will apply to this VM in all recovery plans.

VM Dependencies

Start the following VMs before this VM.

Virtual Machine	Status	Priority Group	Protection Group
ATestWK4	OK	3	Infrastructure Apps

VM dependencies are ignored if the VMs are not in the same priority group.
If VM depencies fail, a warning will be displayed, but the recovery plan will continue.

Figure 48. VM with a Dependency That Has Been Set

Step 8: Run test recovery

SRM enables you to **test** a recovery plan by simulating a failover of virtual machines from the protected site to the recovery site. The benefit of using SRM to run a failover simulation against a recovery plan is that it allows you to confirm that the recovery plan has been set up correctly for the protected virtual machines. You will be able to confirm that the protected virtual machines start up in the correct order, taking into account the various application service dependencies for the protected virtual machines in your environment.

When you select the option to **test** a recovery plan via SRM, the simulated failover is executed in an isolated environment. This environment includes network and storage infrastructure at the recovery site that is isolated from the protected site (production environment), which ensures that the protected virtual machines at the protected site are not subject to any kind of service interruption during the testing of the recovery plan. SRM will also create a test report that can be used to demonstrate your level of preparedness to the business or individual business units whose services are being protected by SRM, as well as to the auditors and compliance officers, if required.

The simulated failover is completed by resetting the environment so that it is ready for the next event. This could be another simulated failover, or an actual failover for a scheduled BC/DR test, or in response to an event that resulted in the business declaring a disaster.

Procedure

1. Open the vSphere Client and connect to the vCenter server at the recovery site. Log in as a user who has permission to test a recovery plan.
2. Click the **Site Recovery** icon on the vSphere Client Home page.
3. In the **Recovery Plan** section of the navigation screen on the left, select the recovery plan that you want to test. On the main viewing pane, select the **Recovery Steps** tab, and ensure that the **View** drop-down menu is set to show **Test Steps.**
4. Ensure that the test steps indicate the correct sequence and any priority groups that have been set.

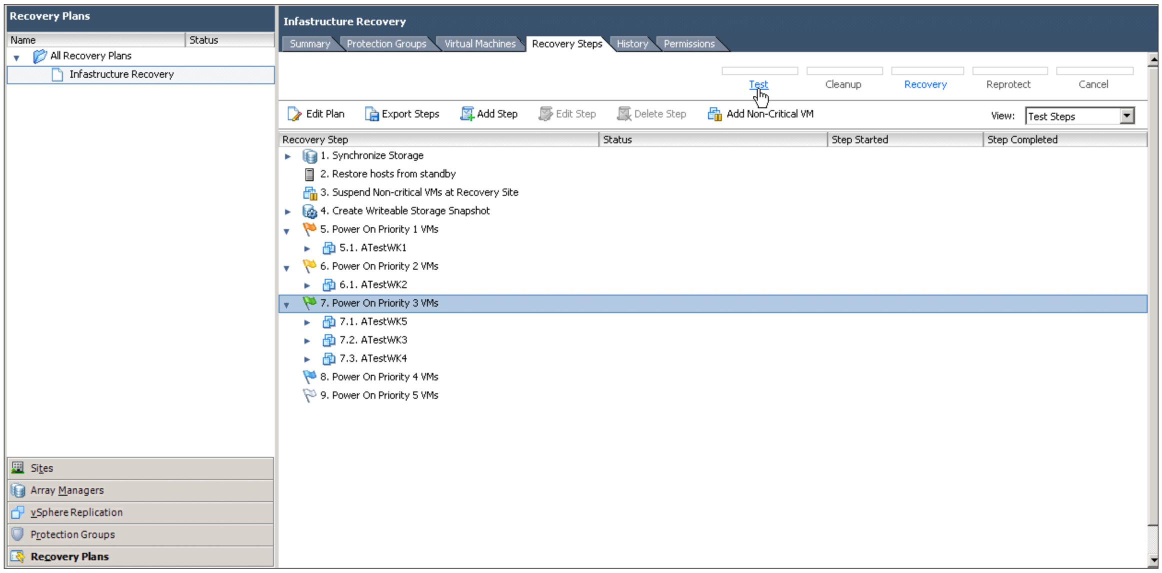

View: Test Steps

Recovery Step	Status	Step Started	Step Completed
1. Synchronize Storage			
2. Restore hosts from standby			
3. Suspend Non-critical VMs at Recovery Site			
4. Create Writeable Storage Snapshot			
5. Power On Priority 1 VMs			
5.1. ATestWK1			
6. Power On Priority 2 VMs			
6.1. ATestWK2			
7. Power On Priority 3 VMs			
7.1. ATestWK5			
7.2. ATestWK3			
7.3. ATestWK4			
8. Power On Priority 4 VMs			
9. Power On Priority 5 VMs			

Figure 49. Preparing to Test a Recovery Plan

5. In the Commands area of the Summary window, click the text labeled **Test.** The **Test** pop-up wizard will open, and prompt you to choose whether you wish to replicate recent changes to the recovery site or not. You may optionally check this box depending on how active the systems are that you are testing. When ready, click **Next** to proceed with the test.

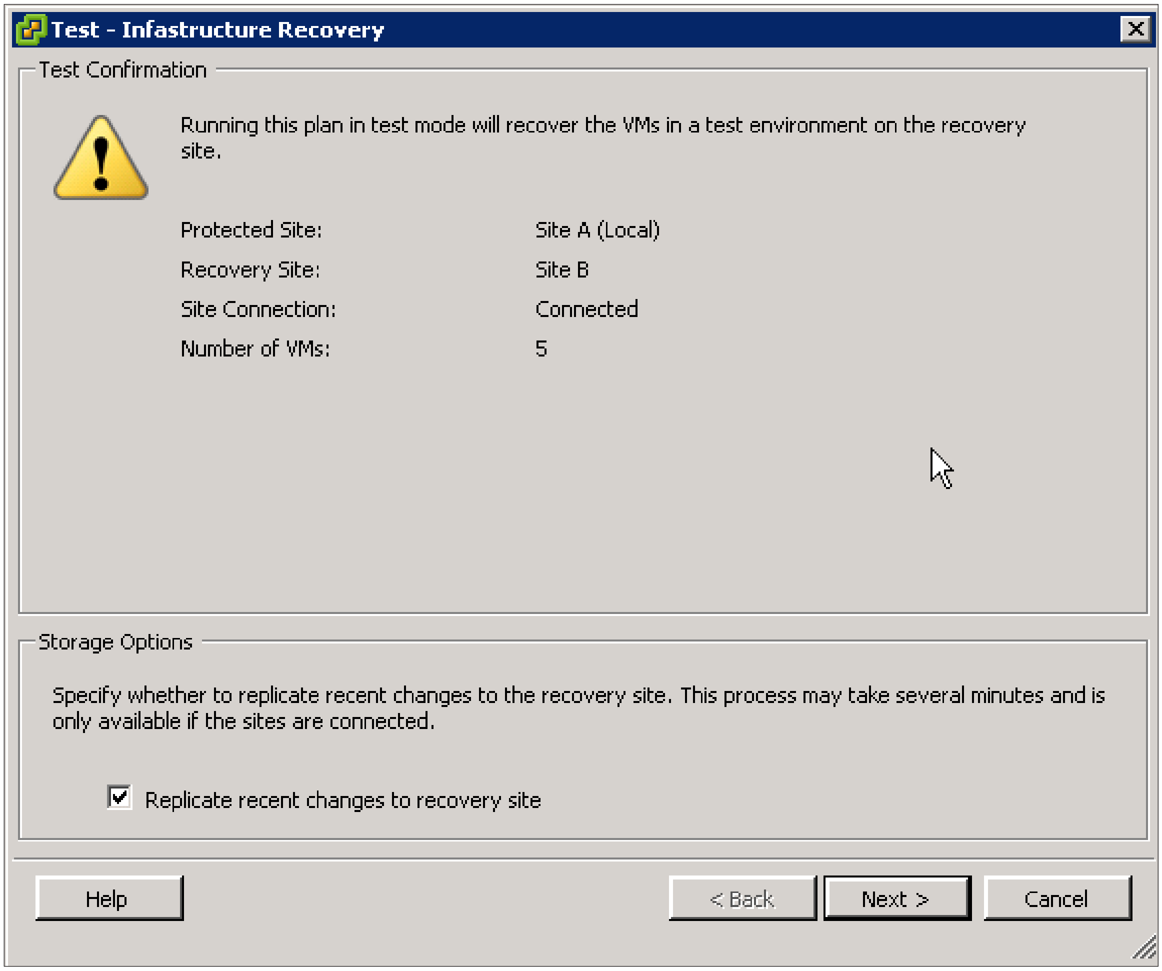

Test Confirmation

Running this plan in test mode will recover the VMs in a test environment on the recovery site.

Protected Site:	Site A (Local)
Recovery Site:	Site B
Site Connection:	Connected
Number of VMs:	5

Storage Options

Specify whether to replicate recent changes to the recovery site. This process may take several minutes and is only available if the sites are connected.

☑ Replicate recent changes to recovery site

Figure 50. Recovery Plan Test Pop-up Screen

5. Review the options selected for the recovery plan test, and click **Start** to initiate the test of the recovery plan failover.

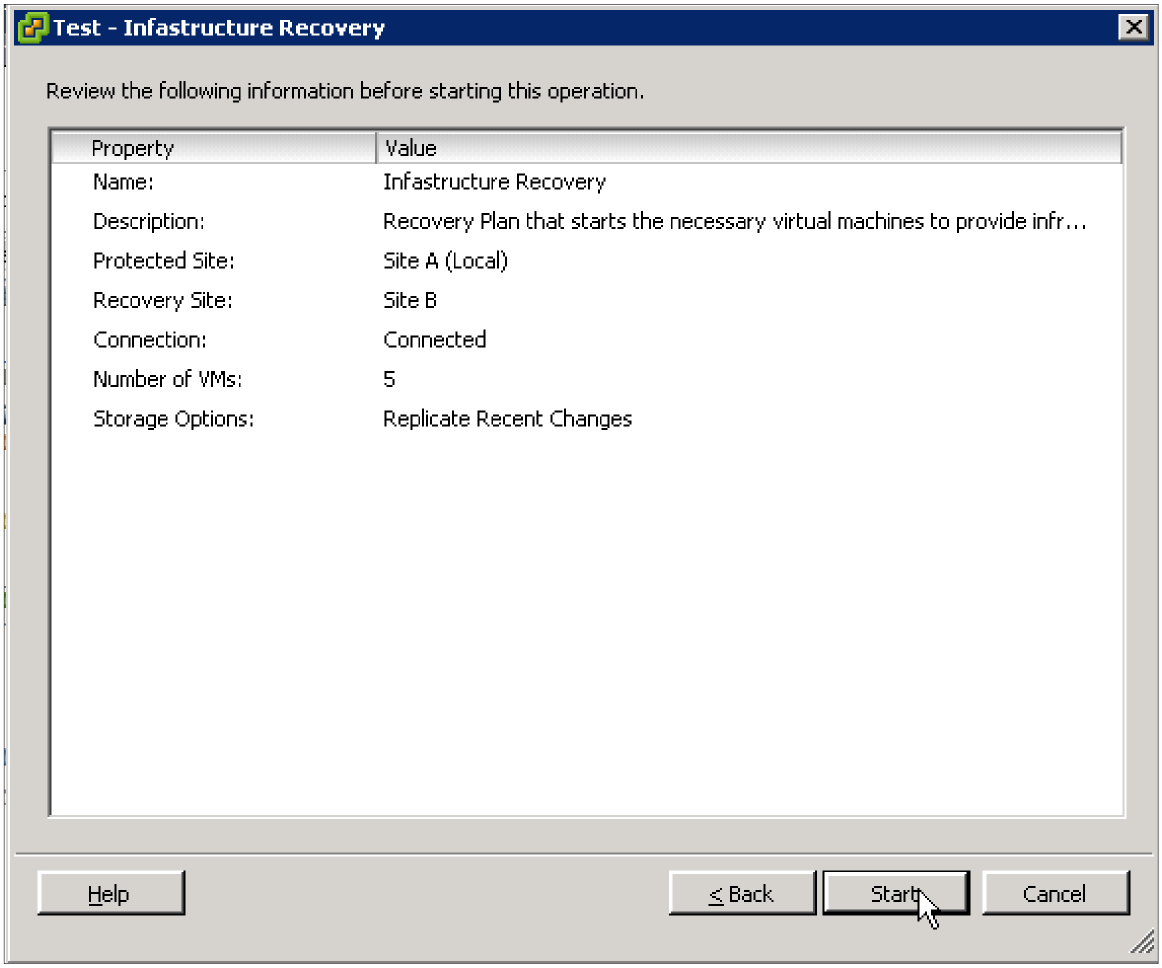

Review the following information before starting this operation.

Property	Value
Name:	Infastructure Recovery
Description:	Recovery Plan that starts the necessary virtual machines to provide infr...
Protected Site:	Site A (Local)
Recovery Site:	Site B
Connection:	Connected
Number of VMs:	5
Storage Options:	Replicate Recent Changes

Figure 51. Click Start to Initiate Test Failover

While the simulated failover test is running, the status of each step that makes up the recovery plan can be monitored by going to the **Recovery Steps** tab in the vSphere Client. This will inform you what steps are currently running as well as what steps were completed.

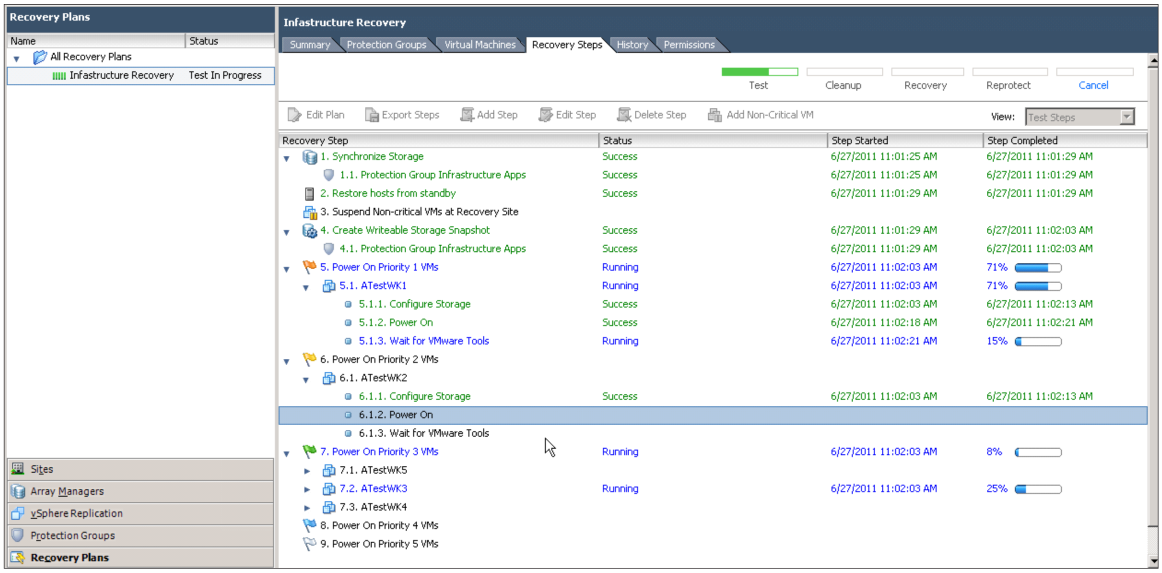

Name	Status
All Recovery Plans	
Infrastructure Recovery	Test In Progress

Recovery Step	Status	Step Started	Step Completed
1. Synchronize Storage	Success	6/27/2011 11:01:25 AM	6/27/2011 11:01:29 AM
1.1. Protection Group Infrastructure Apps	Success	6/27/2011 11:01:25 AM	6/27/2011 11:01:29 AM
2. Restore hosts from standby	Success	6/27/2011 11:01:29 AM	6/27/2011 11:01:29 AM
3. Suspend Non-critical VMs at Recovery Site			
4. Create Writeable Storage Snapshot	Success	6/27/2011 11:01:29 AM	6/27/2011 11:02:03 AM
4.1. Protection Group Infrastructure Apps	Success	6/27/2011 11:01:29 AM	6/27/2011 11:02:03 AM
5. Power On Priority 1 VMs	Running	6/27/2011 11:02:03 AM	71%
5.1. ATestWK1	Running	6/27/2011 11:02:03 AM	71%
5.1.1. Configure Storage	Success	6/27/2011 11:02:03 AM	6/27/2011 11:02:13 AM
5.1.2. Power On	Success	6/27/2011 11:02:18 AM	6/27/2011 11:02:21 AM
5.1.3. Wait for VMware Tools	Running	6/27/2011 11:02:21 AM	15%
6. Power On Priority 2 VMs			
6.1. ATestWK2			
6.1.1. Configure Storage	Success	6/27/2011 11:02:03 AM	6/27/2011 11:02:13 AM
6.1.2. Power On			
6.1.3. Wait for VMware Tools			
7. Power On Priority 3 VMs	Running	6/27/2011 11:02:03 AM	8%
7.1. ATestWK5			
7.2. ATestWK3	Running	6/27/2011 11:02:03 AM	25%
7.3. ATestWK4			
8. Power On Priority 4 VMs			
9. Power On Priority 5 VMs			

Figure 52. Test Steps for Recovery

6. When the test recovery has finished powering on all of the protected virtual machines, it displays a message and requires confirmation before it can continue. Click **Continue** when you are ready for SRM to clean up and finish the test.

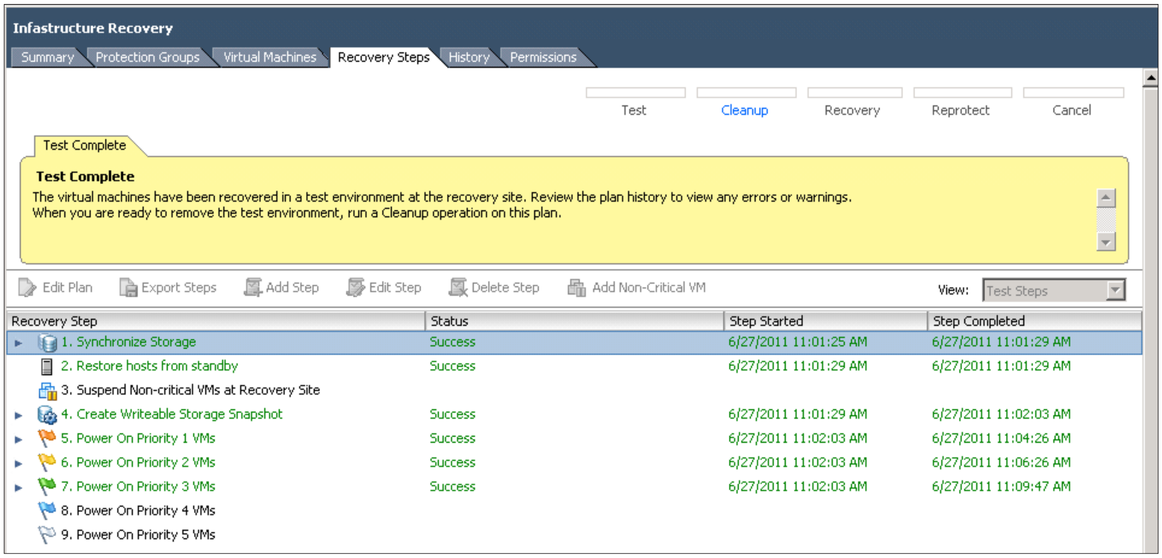

Test Complete

The virtual machines have been recovered in a test environment at the recovery site. Review the plan history to view any errors or warnings.
When you are ready to remove the test environment, run a Cleanup operation on this plan.

Recovery Step	Status	Step Started	Step Completed
1. Synchronize Storage	Success	6/27/2011 11:01:25 AM	6/27/2011 11:01:29 AM
2. Restore hosts from standby	Success	6/27/2011 11:01:29 AM	6/27/2011 11:01:29 AM
3. Suspend Non-critical VMs at Recovery Site			
4. Create Writeable Storage Snapshot	Success	6/27/2011 11:01:29 AM	6/27/2011 11:02:03 AM
5. Power On Priority 1 VMs	Success	6/27/2011 11:02:03 AM	6/27/2011 11:04:26 AM
6. Power On Priority 2 VMs	Success	6/27/2011 11:02:03 AM	6/27/2011 11:06:26 AM
7. Power On Priority 3 VMs	Success	6/27/2011 11:02:03 AM	6/27/2011 11:09:47 AM
8. Power On Priority 4 VMs			
9. Power On Priority 5 VMs			

Figure 53. Recovery Test Complete, Ready to Clean Up

To run an automated cleanup, click the **Cleanup** text button in the action field, review the cleanup actions, and click **Next.** Review the cleanup summary, and click **Start** to begin the cleanup process.

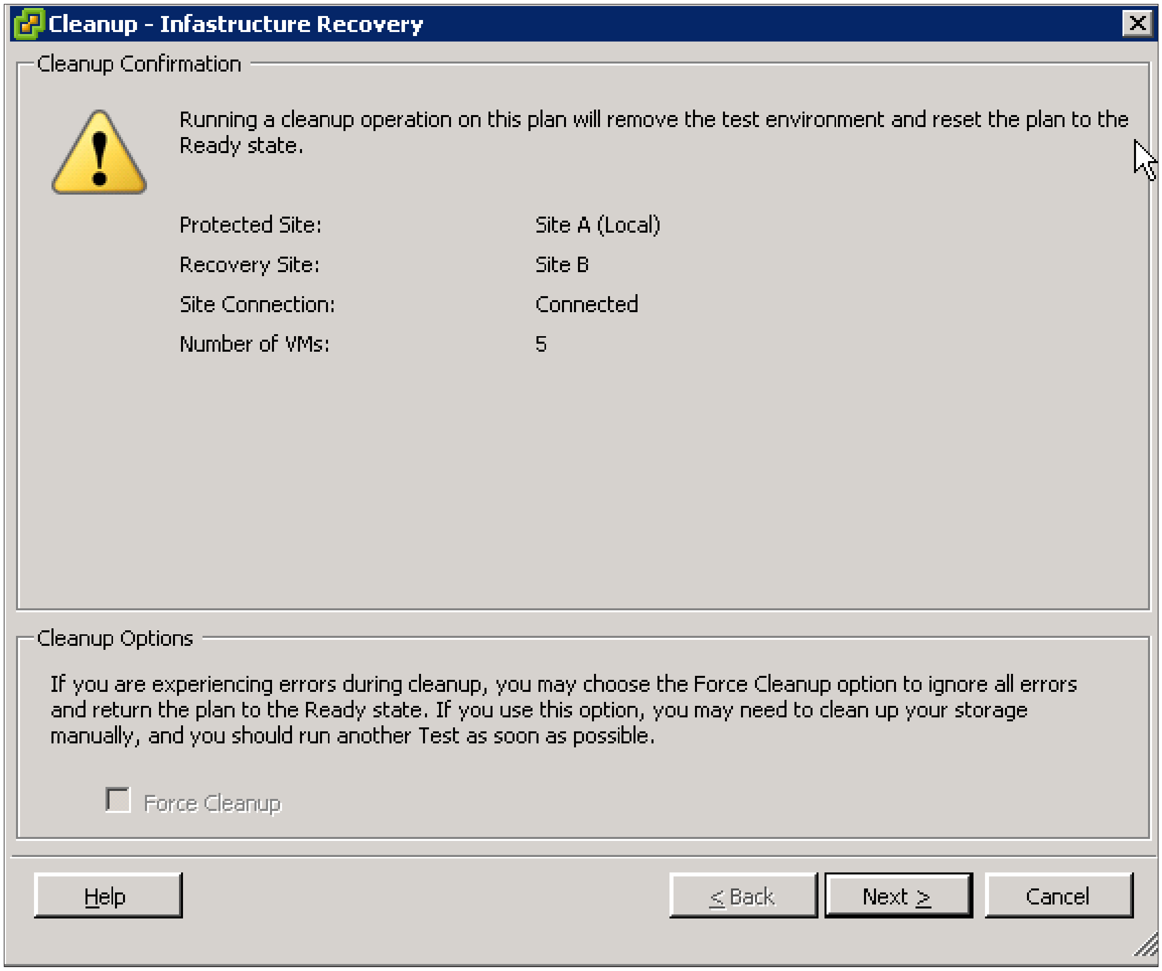

Cleanup - Infastructure Recovery

Cleanup Confirmation

Running a cleanup operation on this plan will remove the test environment and reset the plan to the Ready state.

Protected Site:	Site A (Local)
Recovery Site:	Site B
Site Connection:	Connected
Number of VMs:	5

Cleanup Options

If you are experiencing errors during cleanup, you may choose the Force Cleanup option to ignore all errors and return the plan to the Ready state. If you use this option, you may need to clean up your storage manually, and you should run another Test as soon as possible.

☐ Force Cleanup

Figure 54. Ready to Execute Automated Cleanup

During cleanup, SRM powers down and unregisters the test virtual machines at the recovery site, and then registers the placeholders back.

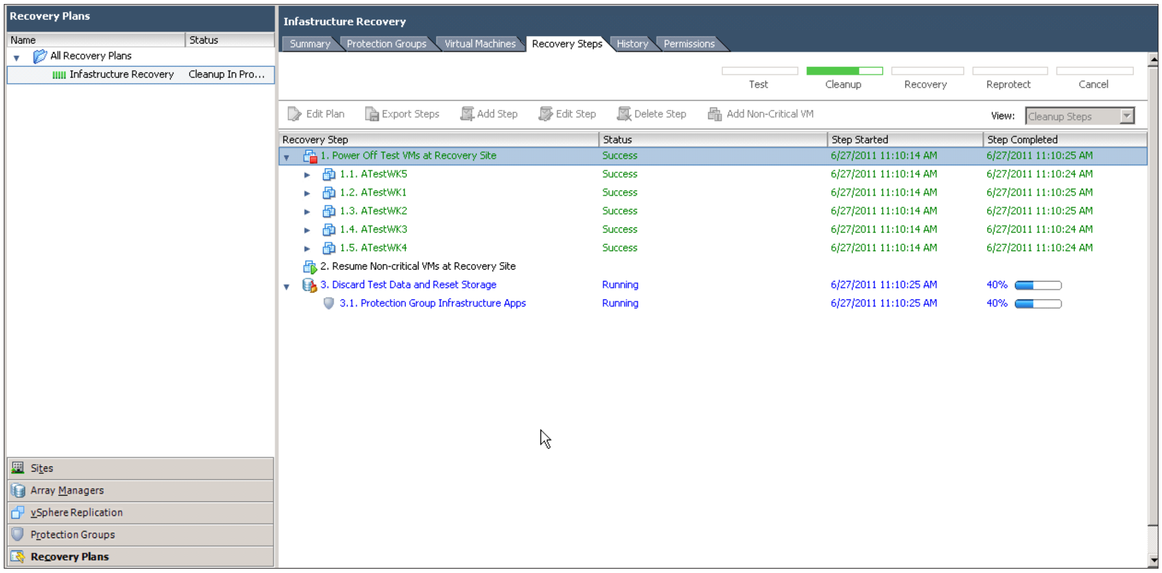

Recovery Plans

Name	Status
All Recovery Plans	
Infastructure Recovery	Cleanup In Pro…

Infastructure Recovery

View: Cleanup Steps

Recovery Step	Status	Step Started	Step Completed
1. Power Off Test VMs at Recovery Site	Success	6/27/2011 11:10:14 AM	6/27/2011 11:10:25 AM
1.1. ATestWK5	Success	6/27/2011 11:10:14 AM	6/27/2011 11:10:24 AM
1.2. ATestWK1	Success	6/27/2011 11:10:14 AM	6/27/2011 11:10:25 AM
1.3. ATestWK2	Success	6/27/2011 11:10:14 AM	6/27/2011 11:10:25 AM
1.4. ATestWK3	Success	6/27/2011 11:10:14 AM	6/27/2011 11:10:24 AM
1.5. ATestWK4	Success	6/27/2011 11:10:14 AM	6/27/2011 11:10:24 AM
2. Resume Non-critical VMs at Recovery Site			
3. Discard Test Data and Reset Storage	Running	6/27/2011 11:10:25 AM	40%
3.1. Protection Group Infrastructure Apps	Running	6/27/2011 11:10:25 AM	40%

Figure 55. Automated Cleanup

7. SRM provides an audit trail via a report that is generated automatically at the end of each SRM test or SRM recovery. The reports are accessible via the **History** tab at the top of the **Recovery Plans** menu.

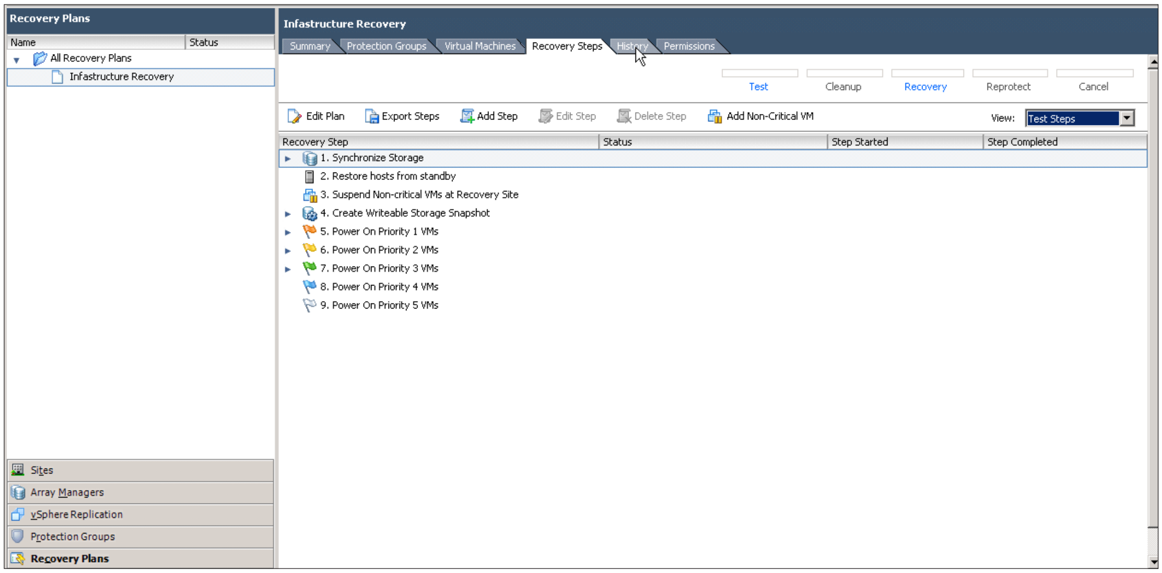

Recovery Plans

Name	Status
All Recovery Plans	
Infastructure Recovery	

Infastructure Recovery

View: Test Steps

Recovery Step	Status	Step Started	Step Completed
1. Synchronize Storage			
2. Restore hosts from standby			
3. Suspend Non-critical VMs at Recovery Site			
4. Create Writeable Storage Snapshot			
5. Power On Priority 1 VMs			
6. Power On Priority 2 VMs			
7. Power On Priority 3 VMs			
8. Power On Priority 4 VMs			
9. Power On Priority 5 VMs			

Figure 56. History Tab

Historical reports can be viewed by clicking the **View** link in the **Actions** column.

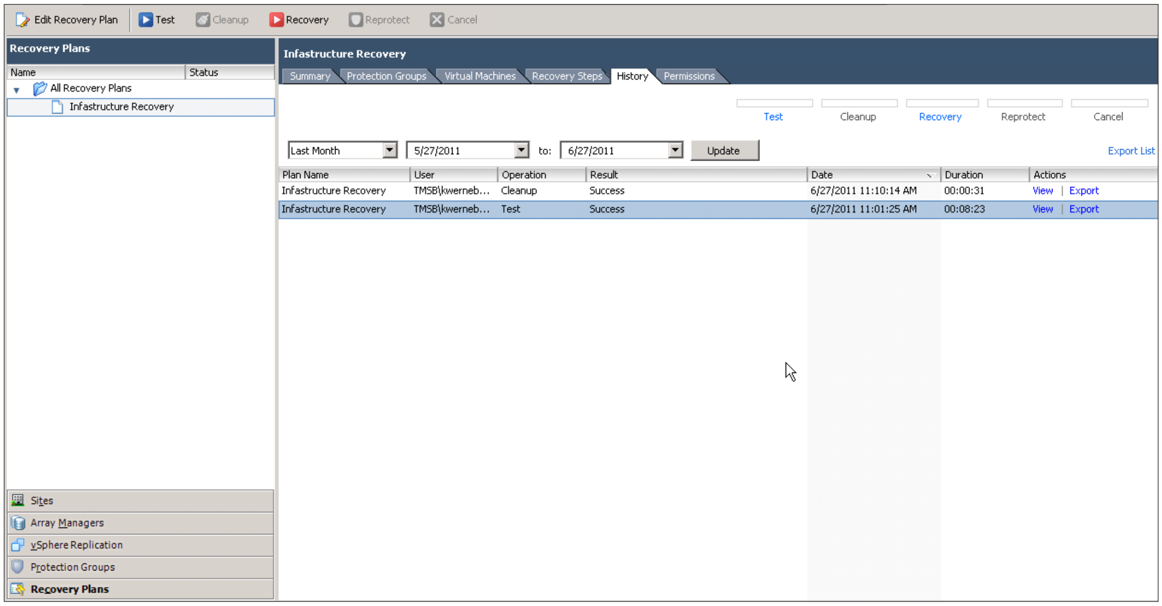

Last Month | 5/27/2011 | to: 6/27/2011 | Update

Plan Name	User	Operation	Result	Date	Duration	Actions
Infastructure Recovery	TMSB\kwerneb…	Cleanup	Success	6/27/2011 11:10:14 AM	00:00:31	View | Export
Infastructure Recovery	TMSB\kwerneb…	Test	Success	6/27/2011 11:01:25 AM	00:08:23	View | Export

Figure 57. History Tab and Actions Column, to View Reports

Clicking **View** on a history report will result in a browser window opening. It contains a log of the steps executed during the test, with the total time it took to execute the recovery plan and the time it took to execute each step in the recovery plan.

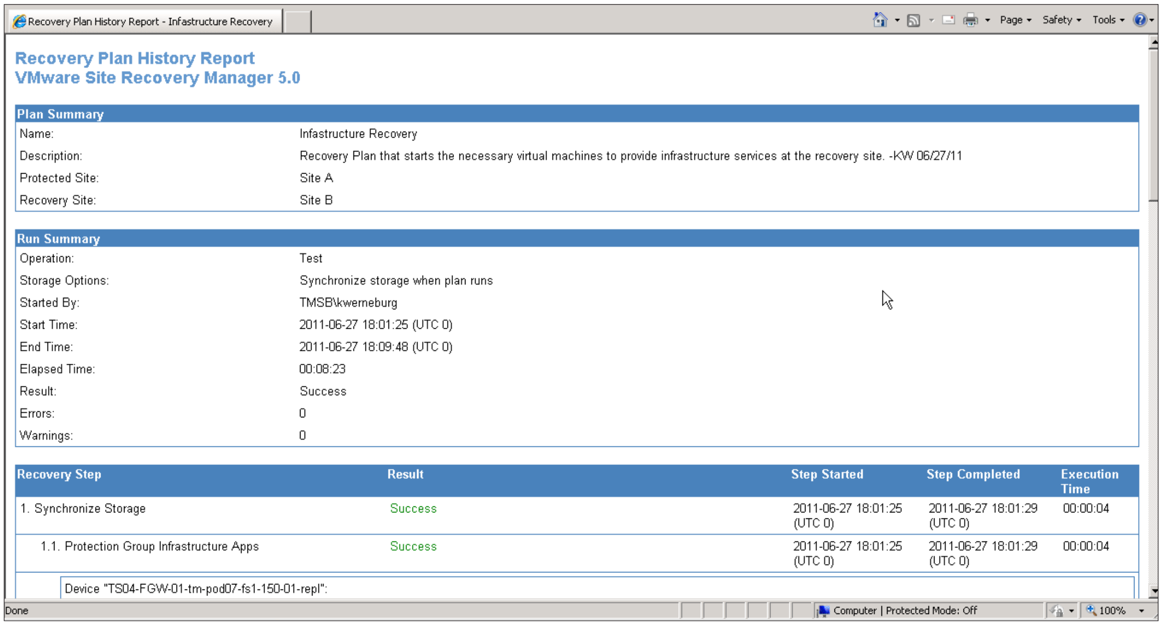
Recovery Plan History Report - Infastructure Recovery
Page
Safety
Tools
Done
Computer | Protected Mode: Off
100%

Recovery Plan History Report
VMware Site Recovery Manager 5.0

Plan Summary	
Name:	Infastructure Recovery
Description:	Recovery Plan that starts the necessary virtual machines to provide infrastructure services at the recovery site. -KW 06/27/11
Protected Site:	Site A
Recovery Site:	Site B

Run Summary	
Operation:	Test
Storage Options:	Synchronize storage when plan runs
Started By:	TMSB\kwerneburg
Start Time:	2011-06-27 18:01:25 (UTC 0)
End Time:	2011-06-27 18:09:48 (UTC 0)
Elapsed Time:	00:08:23
Result:	Success
Errors:	0
Warnings:	0

Recovery Step	Result	Step Started	Step Completed	Execution Time
1. Synchronize Storage	Success	2011-06-27 18:01:25 (UTC 0)	2011-06-27 18:01:29 (UTC 0)	00:00:04
1.1. Protection Group Infrastructure Apps	Success	2011-06-27 18:01:25 (UTC 0)	2011-06-27 18:01:29 (UTC 0)	00:00:04
Device "TS04-FGW-01-tm-pod07-fs1-150-01-repl":				

Figure 58. History Report Displayed in a Browser

Exercise 2. Deploying vSphere Replication

vSphere Replication is a replication engine that is part of SRM 5.0 and requires ESXi 5.0 and later, giving an alternative means of protecting and replicating virtual machines between sites. It is entirely managed within the SRM interface after initial deployment and configuration, and integrates with storage array–based replication to provide full coverage of the virtual environment.

The assumption is that there are multiple databases for vSphere Replication already configured for use, one at each site. In this evaluation guide, we will be using Microsoft SQL Server as a database, and using native SQL authentication for access.

Workflow covered will be as follows:

1. Deploy vSphere Replication Management Servers (VRMS).
2. Configure VRMS.
3. Pair VRMS.
4. Deploy vSphere Replication Server (VRS).
5. Register VRS.
6. Configure virtual machines for protection with vSphere Replication.
7. Create a protection group using vSphere Replication–protected virtual machines.

Step 1: Deploy vSphere Replication Management Server

The VRM servers act as a management framework for vSphere Replication, and therefore it is required that a VRM server be deployed and configured at both protected and recovery sites.

To deploy a VRMS:

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.
3. Click the **Site Recovery** on the vSphere Client Home page under Solutions and Applications.
4. Choose the menu item on the left navigation pane entitled **vSphere Replication,** and select the protected site (Site A). The **Summary View** panel should show no VRM servers, or status. Click on **Deploy VRM Server** in the actions list on the right.

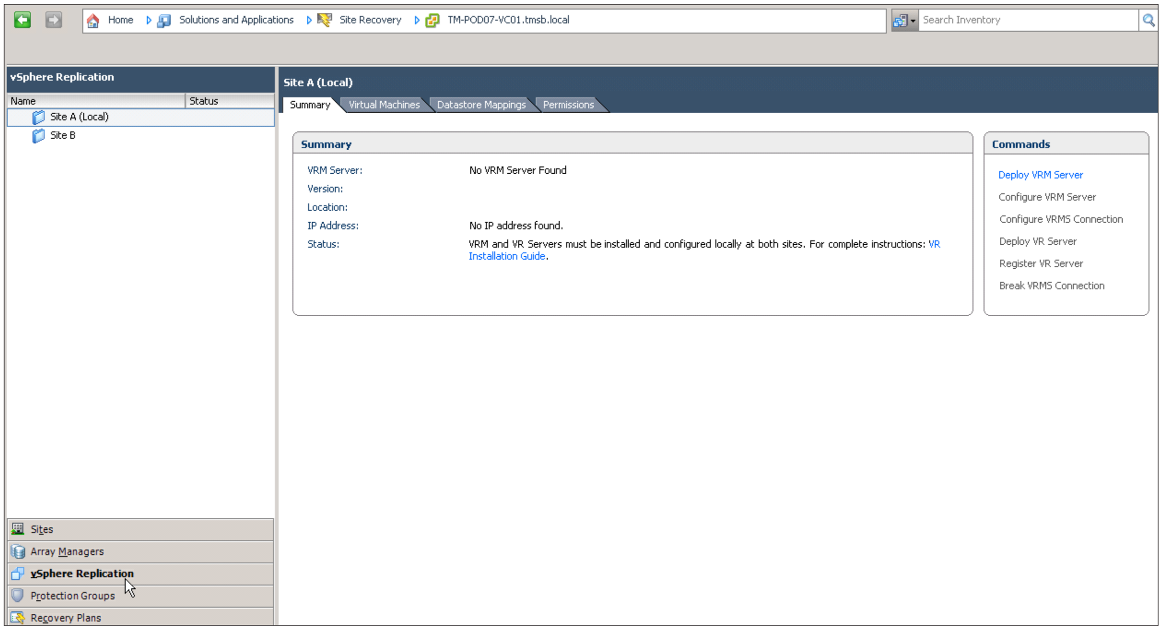

VRM Server: No VRM Server Found

Version:

Location:

IP Address: No IP address found.

Status: VRM and VR Servers must be installed and configured locally at both sites. For complete instructions: VR Installation Guide.

Figure 59. Summary Screen for vSphere Replication

5. Using the OVF wizard to deploy the appliance, click **Next** through the options until you can select a name and location in which to deploy the VRMS appliance. Ensure that at this point you are deploying to the primary protected site (Site A).

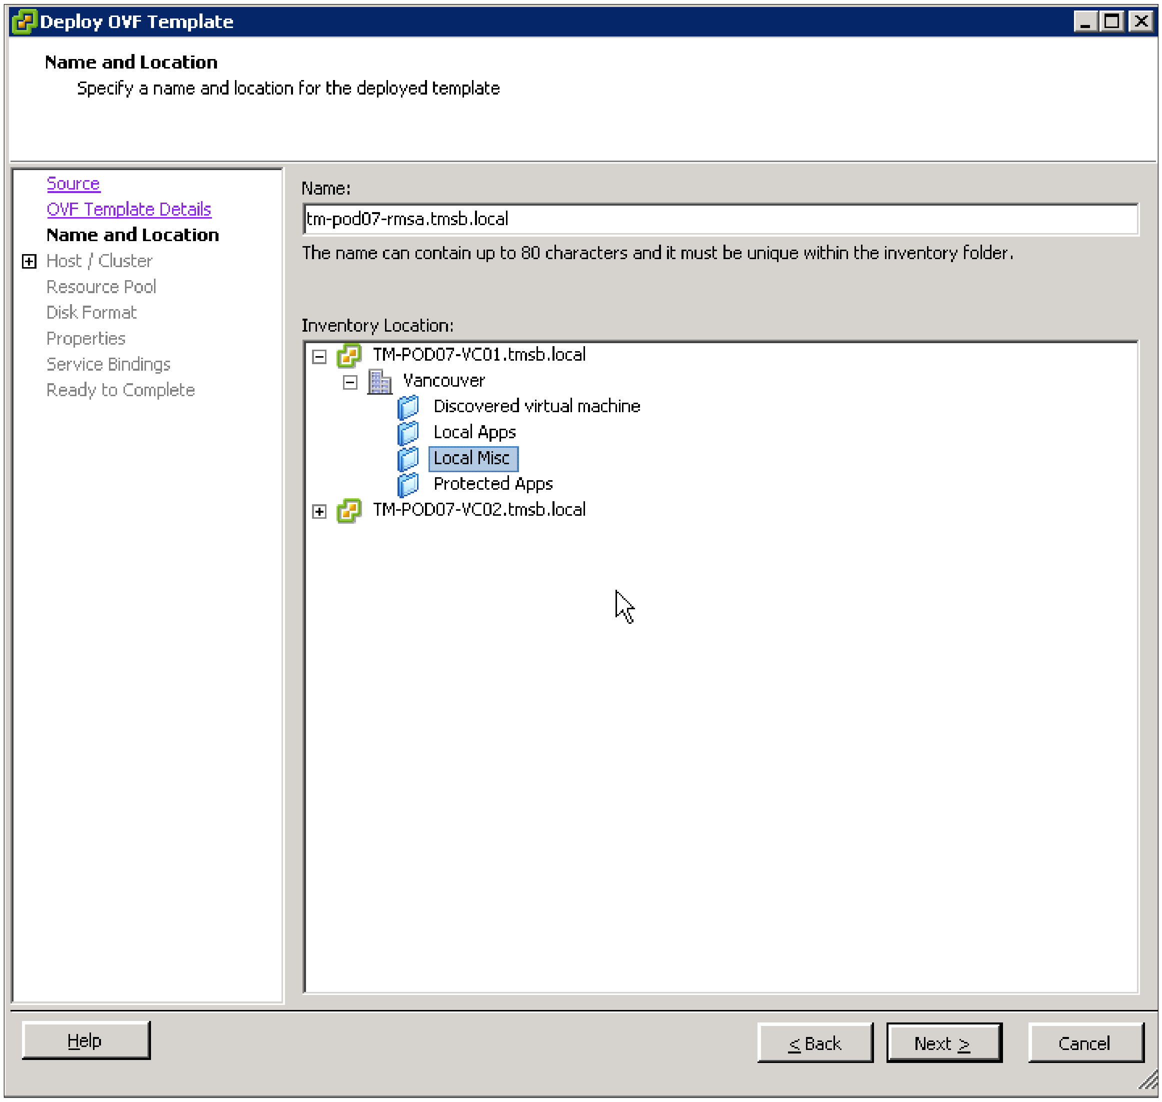

Name and Location

Specify a name and location for the deployed template

Name: tm-pod07-rmsa.tmsb.local

The name can contain up to 80 characters and it must be unique within the inventory folder.

Inventory Location: Local Misc

Figure 60. Name and Location for the VMRS at Site A

6. Choose appropriate hosts, clusters, resource pools, and disk locations, for the VRMS appliance, and ensure that all networking information is correctly populated. Take note of the **root password** you enter and ensure that it is recorded.

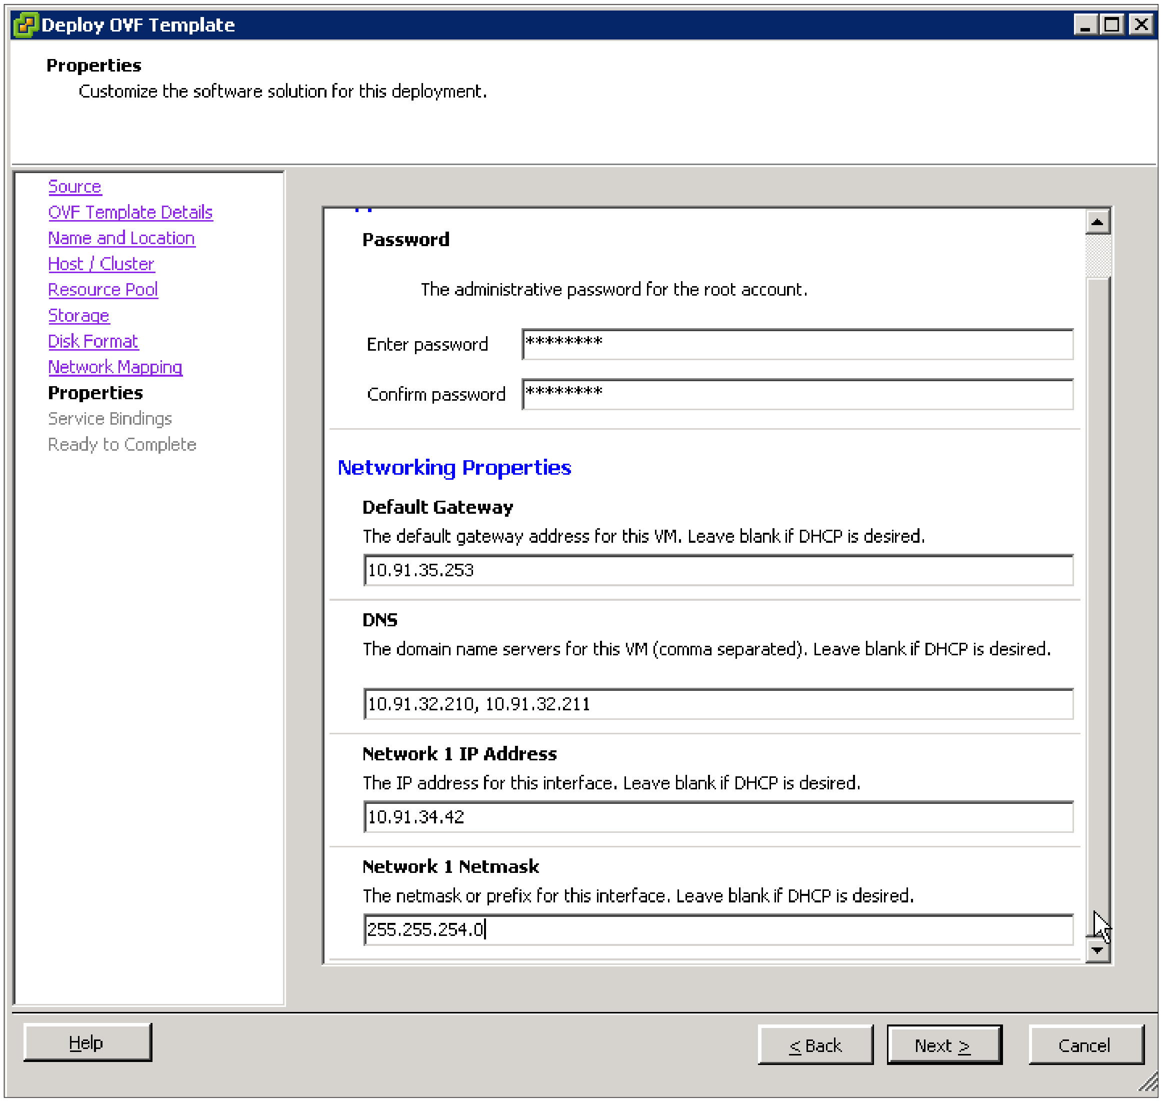

Deploy OVF Template

Properties

Customize the software solution for this deployment.

Password

The administrative password for the root account.

Enter password ********

Confirm password ********

Networking Properties

Default Gateway

The default gateway address for this VM. Leave blank if DHCP is desired.

10.91.35.253

DNS

The domain name servers for this VM (comma separated). Leave blank if DHCP is desired.

10.91.32.210, 10.91.32.211

Network 1 IP Address

The IP address for this interface. Leave blank if DHCP is desired.

10.91.34.42

Network 1 Netmask

The netmask or prefix for this interface. Leave blank if DHCP is desired.

255.255.254.0

Figure 61. Network Details for VRMS

7. Select the correct **vCenter Extension vService** through the drop-down menu. If this is not available or shows an error, it usually indicates that the managed IP address of the vCenter server has not been set. See the knowledge base article http://kb.vmware.com/kb/1008030. Complete managed IP settings in both the protected and recovery site vCenter servers before continuing. If an error has been encountered, you will need to cancel VRMS deployment, fix the settings, and redeploy the VRMS.

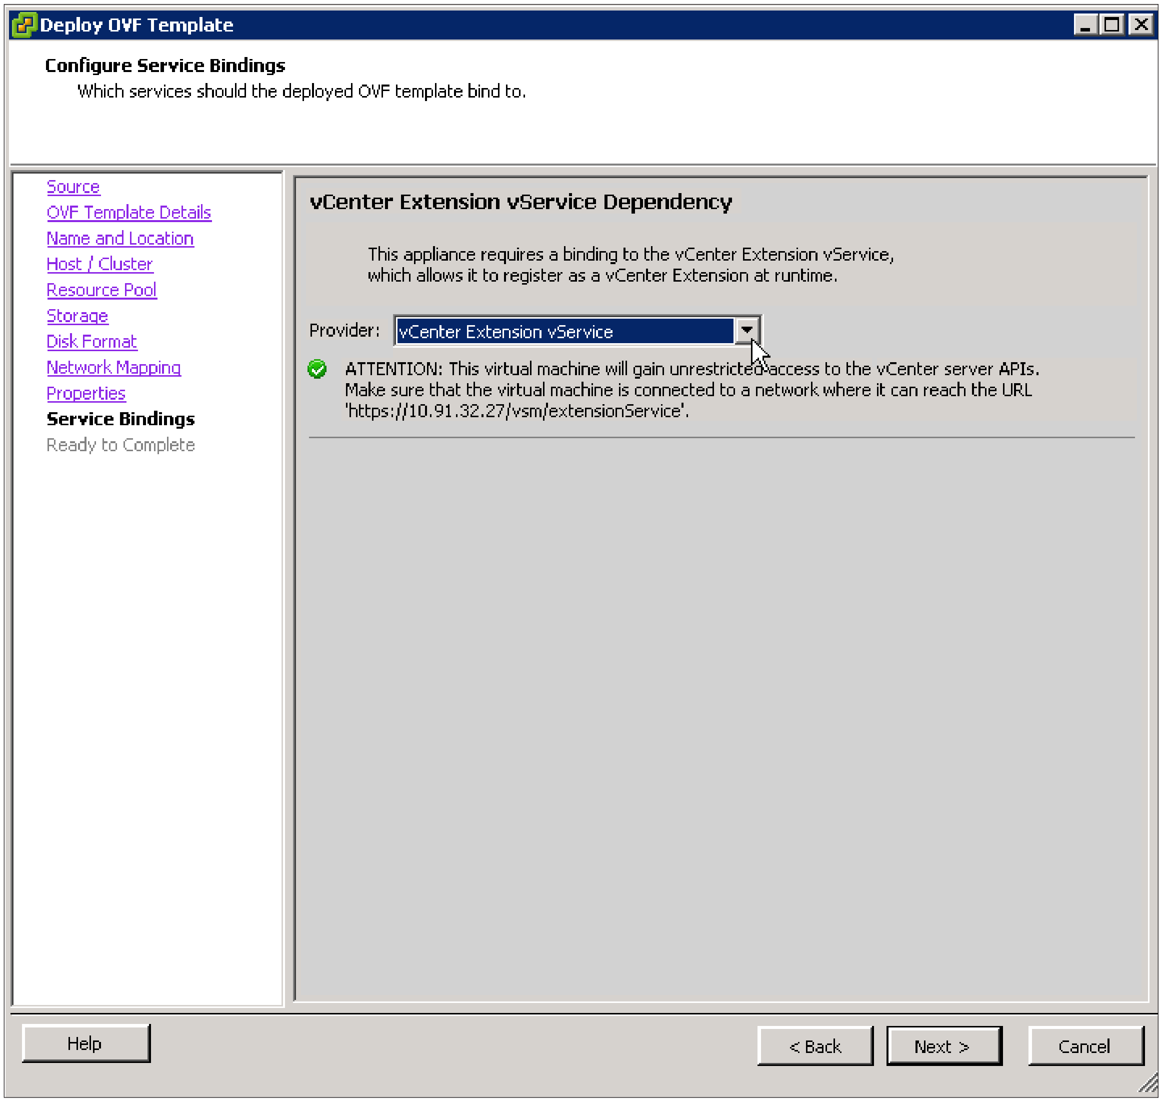

Configure Service Bindings

Which services should the deployed OVF template bind to.

vCenter Extension vService Dependency

This appliance requires a binding to the vCenter Extension vService, which allows it to register as a vCenter Extension at runtime.

Provider: vCenter Extension vService

ATTENTION: This virtual machine will gain unrestricted access to the vCenter server APIs. Make sure that the virtual machine is connected to a network where it can reach the URL 'https://10.91.32.27/vsm/extensionService'.

Figure 62. vCenter Extension Service

8. When you are ready to complete the task, click **Finish** to deploy the VRMS appliance. A pop-up screen will show the state of deployment of the VRMS appliance, and will indicate when the deployment has been completed successfully.

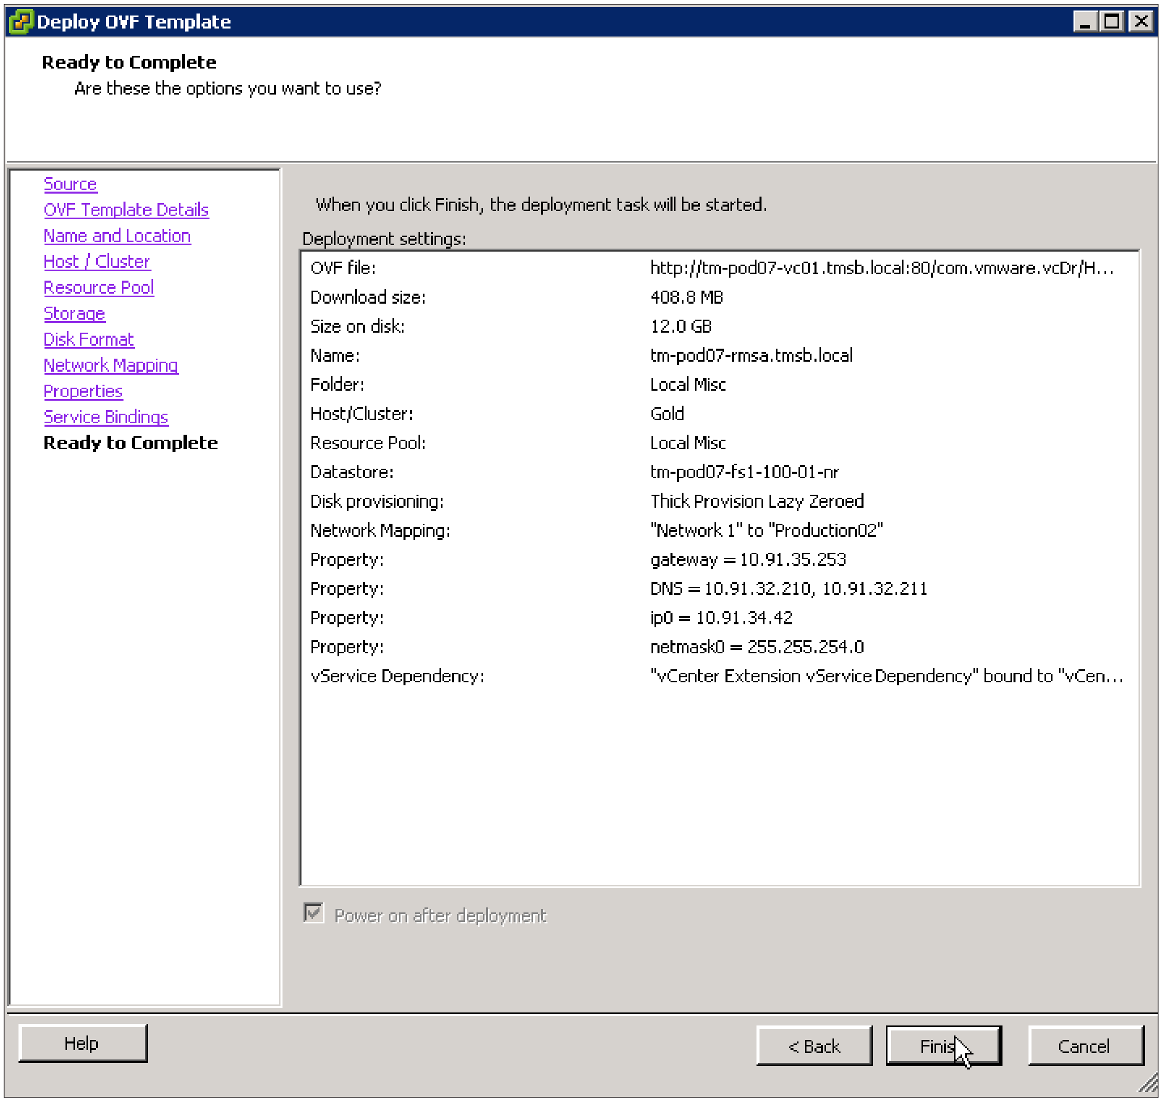

Ready to Complete

Are these the options you want to use?

When you click Finish, the deployment task will be started.

Deployment settings:

OVF file:	http://tm-pod07-vc01.tmsb.local:80/com.vmware.vcDr/H...
Download size:	408.8 MB
Size on disk:	12.0 GB
Name:	tm-pod07-rmsa.tmsb.local
Folder:	Local Misc
Host/Cluster:	Gold
Resource Pool:	Local Misc
Datastore:	tm-pod07-fs1-100-01-nr
Disk provisioning:	Thick Provision Lazy Zeroed
Network Mapping:	"Network 1" to "Production02"
Property:	gateway = 10.91.35.253
Property:	DNS = 10.91.32.210, 10.91.32.211
Property:	ip0 = 10.91.34.42
Property:	netmask0 = 255.255.254.0
vService Dependency:	"vCenter Extension vService Dependency" bound to "vCen...

☑ Power on after deployment

Figure 63. Click Finish When Ready

9. You will now need to deploy a VRMS appliance to the recovery site, following the same process as outlined above in steps 4 through 8. This time, however, please ensure that you are deploying the VRMS to the recovery site (Site B) instead of the protected site.

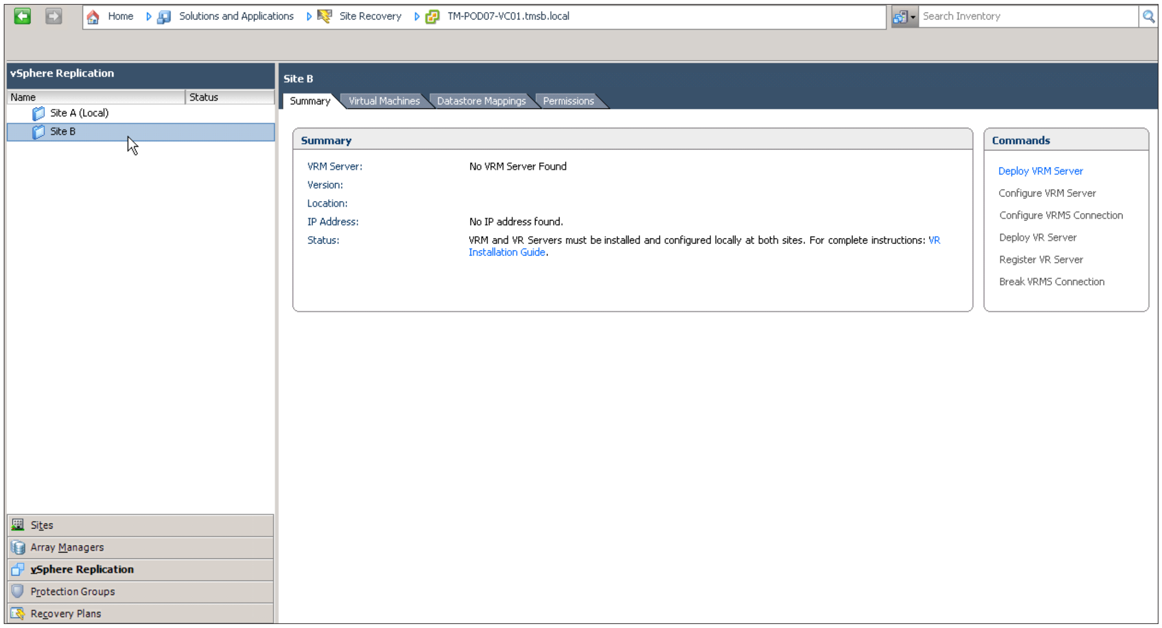

Summary

VRM Server:	No VRM Server Found
Version:	
Location:	
IP Address:	No IP address found.
Status:	VRM and VR Servers must be installed and configured locally at both sites. For complete instructions: VR Installation Guide.

Figure 64. Deploying VRMS to the Recovery Site

10. When you are done, there should be a VRMS deployed to both protected and recovery sites. In the following example, the Site A VRMS is labeled tm-pod07-rmsa and the Site B VRMS is labeled tm-pod07-rmsb.

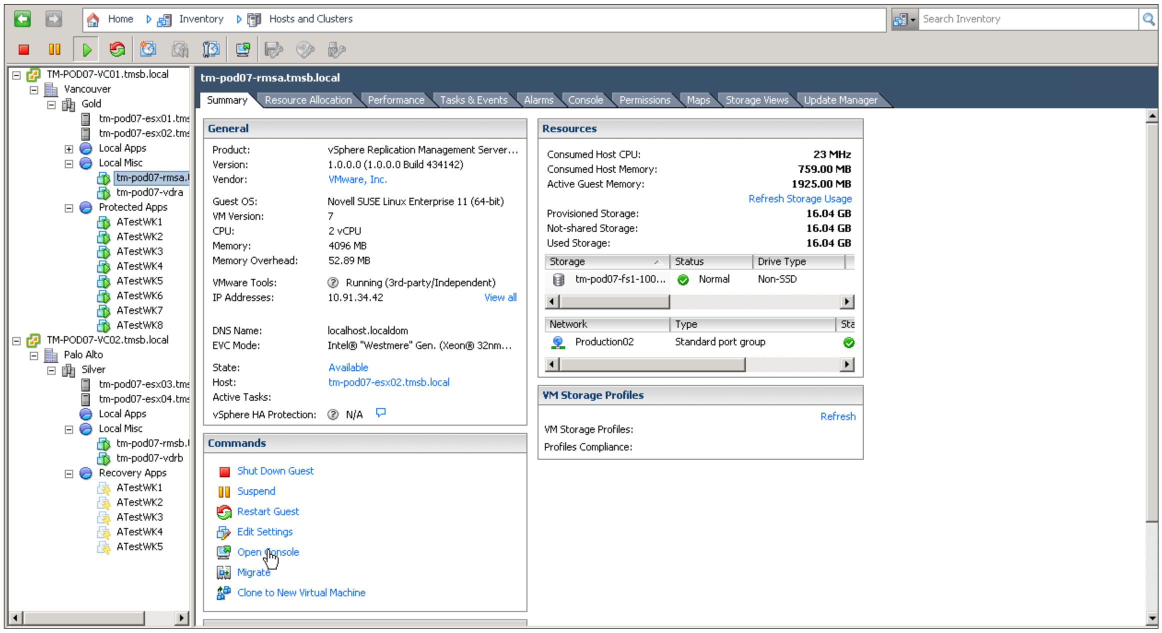

General

Product:	vSphere Replication Management Server…
Version:	1.0.0.0 (1.0.0.0 Build 434142)
Vendor:	VMware, Inc.
Guest OS:	Novell SUSE Linux Enterprise 11 (64-bit)
VM Version:	7
CPU:	2 vCPU
Memory:	4096 MB
Memory Overhead:	52.89 MB
VMware Tools:	Running (3rd-party/Independent)
IP Addresses:	10.91.34.42 View all
DNS Name:	localhost.localdom
EVC Mode:	Intel® "Westmere" Gen. (Xeon® 32nm…
State:	Available
Host:	tm-pod07-esx02.tmsb.local
Active Tasks:	
vSphere HA Protection:	N/A

Resources

Consumed Host CPU:	23 MHz
Consumed Host Memory:	759.00 MB
Active Guest Memory:	1925.00 MB
Provisioned Storage:	16.04 GB
Not-shared Storage:	16.04 GB
Used Storage:	16.04 GB

Storage	Status	Drive Type
tm-pod07-fs1-100…	Normal	Non-SSD

Network	Type	Sta
Production02	Standard port group	✓

VM Storage Profiles

VM Storage Profiles:
Profiles Compliance:

Figure 65. VRMS Appliances Deployed

Step 2: Configuring VRMS appliances

In order to complete the configuration of the management framework, you must complete a few steps. The first step is the configuration of the appliance itself, done through the VM console. The second step is the configuration of the management framework, done through a Web browser interface to the appliances.

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.
3. Click the **Hosts and Clusters** view, and browse to the VRMS appliance deployed to the protected site (Site A). Click **Open Console** to pop the console out of the appliance.

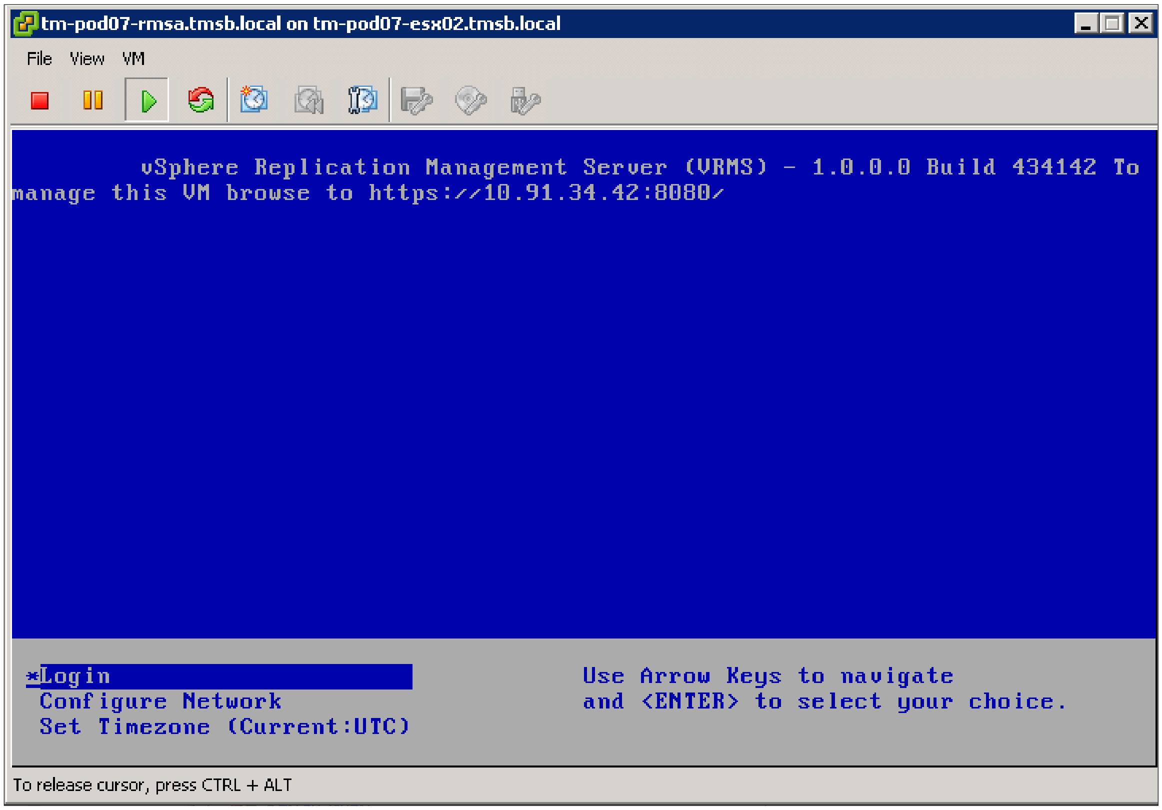
tm-pod07-rmsa.tmsb.local on tm-pod07-esx02.tmsb.local
File
View
VM
To release cursor, press CTRL + ALT

```
        vSphere Replication Management Server (VRMS) - 1.0.0.0 Build 434142 To
manage this VM browse to https://10.91.34.42:8080/

 *Login                                  Use Arrow Keys to navigate
  Configure Network                      and <ENTER> to select your choice.
  Set Timezone (Current:UTC)
```

Figure 66. Configuring VRMS Appliance

4. Using the arrow keys, choose **Configure Network** and press **Enter** or **Return** on your keyboard.

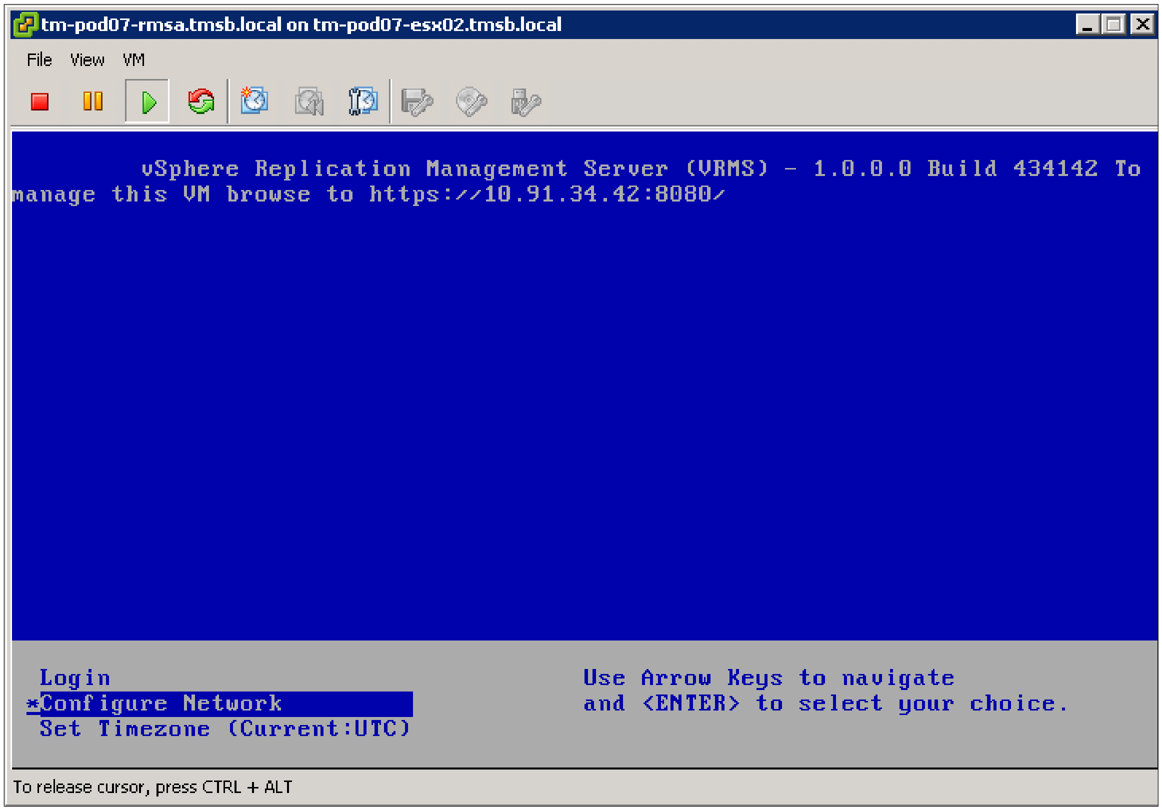
tm-pod07-rmsa.tmsb.local on tm-pod07-esx02.tmsb.local
File View VM

```
         vSphere Replication Management Server (VRMS) - 1.0.0.0 Build 434142 To
manage this VM browse to https://10.91.34.42:8080/

  Login                                  Use Arrow Keys to navigate
 *Configure Network                      and <ENTER> to select your choice.
  Set Timezone (Current:UTC)
```

To release cursor, press CTRL + ALT

Figure 67. Configure the Network

5. At the login screen, type **0** (zero) to review the current network configuration. If the networking information is correct, return to the main menu.

6. Enter **3** (three) to change the host name. Make sure the host name you enter here is noted and recorded, and works correctly with forward, reverse, and fully qualified DNS searches.

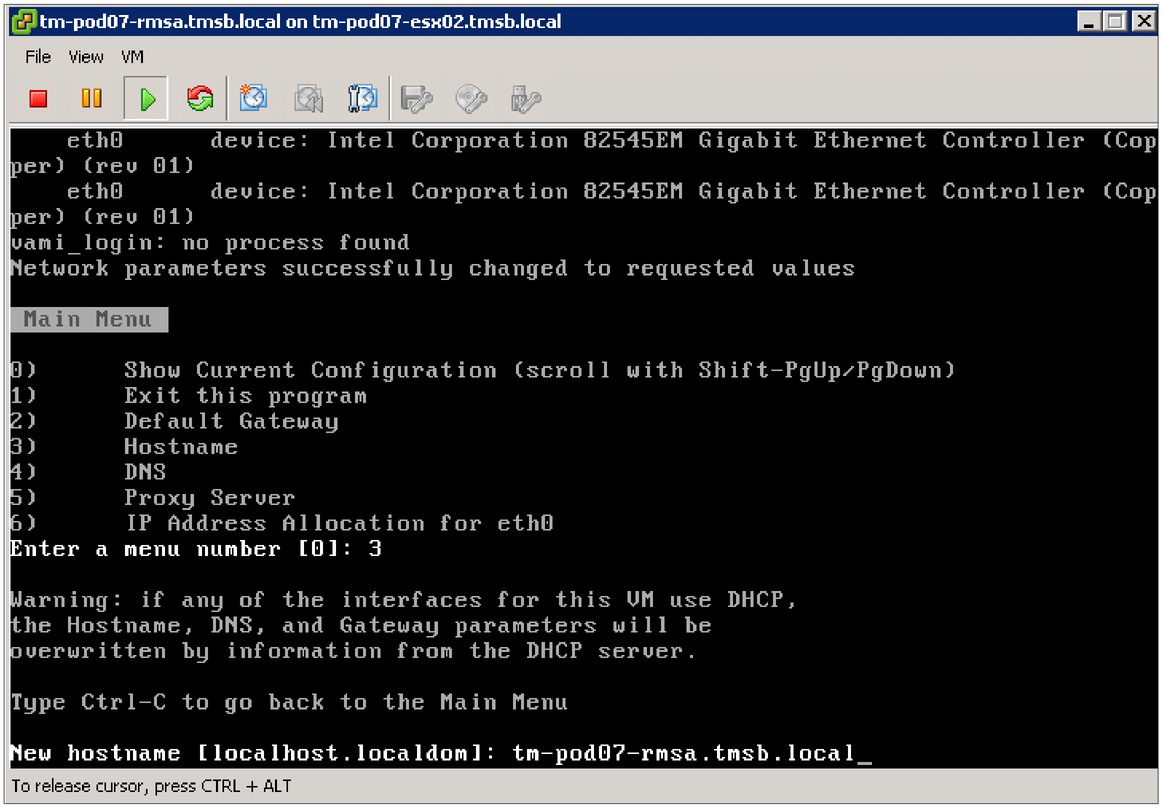

```
    eth0       device: Intel Corporation 82545EM Gigabit Ethernet Controller (Cop
per) (rev 01)
    eth0       device: Intel Corporation 82545EM Gigabit Ethernet Controller (Cop
per) (rev 01)
vami_login: no process found
Network parameters successfully changed to requested values

 Main Menu

0)       Show Current Configuration (scroll with Shift-PgUp/PgDown)
1)       Exit this program
2)       Default Gateway
3)       Hostname
4)       DNS
5)       Proxy Server
6)       IP Address Allocation for eth0
Enter a menu number [0]: 3

Warning: if any of the interfaces for this VM use DHCP,
the Hostname, DNS, and Gateway parameters will be
overwritten by information from the DHCP server.

Type Ctrl-C to go back to the Main Menu

New hostname [localhost.localdom]: tm-pod07-rmsa.tmsb.local_
```

Figure 68. Configuring the Host Name

7. Press **1** (one) to write changes and exit the interface. If necessary, at the main appliance menu, use the arrow keys to select **Set Timezone** to update the time zone of the appliance to your current geography.

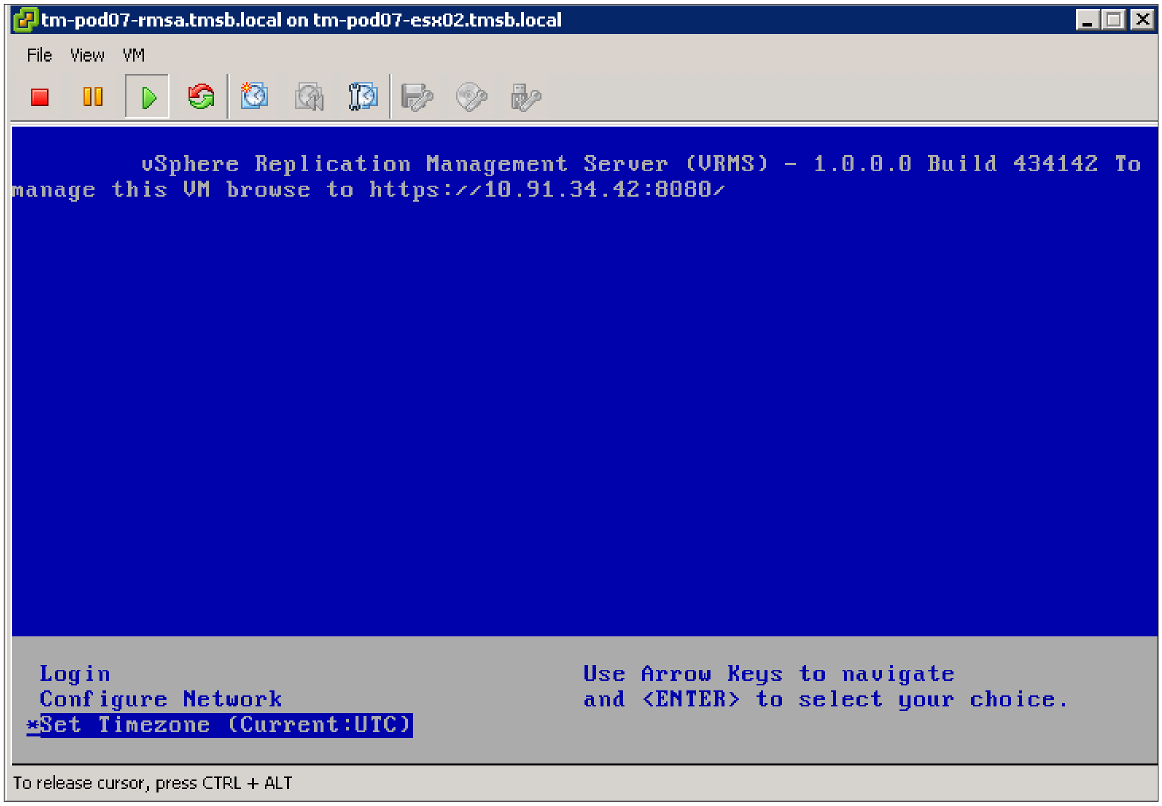

tm-pod07-rmsa.tmsb.local on tm-pod07-esx02.tmsb.local
File View VM

```
         vSphere Replication Management Server (VRMS) - 1.0.0.0 Build 434142 To
manage this VM browse to https://10.91.34.42:8080/

  Login                                        Use Arrow Keys to navigate
  Configure Network                            and <ENTER> to select your choice.
 *Set Timezone (Current:UTC)
```

To release cursor, press CTRL + ALT

Figure 69. Configuring the Timezone

Step 3: Configuring VRMS through Web management interface

After deployment and configuration of a VRMS appliance, configuration of the management itself will allow us to register the VRMS with the vCenter server and connect to the database previously configured for use with vSphere Replication.

To configure the VRMS, follow these steps:

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.
3. Click the **Site Recovery icon** on the vSphere Client Home page under Solutions and Applications.
4. Choose the menu item on the left navigation pane entitled **vSphere Replication,** and select the protected site (Site A). Click on **Configure VRM Server** in the actions list on the right. This will launch a Web browser to the VRMS appliance for configuration.

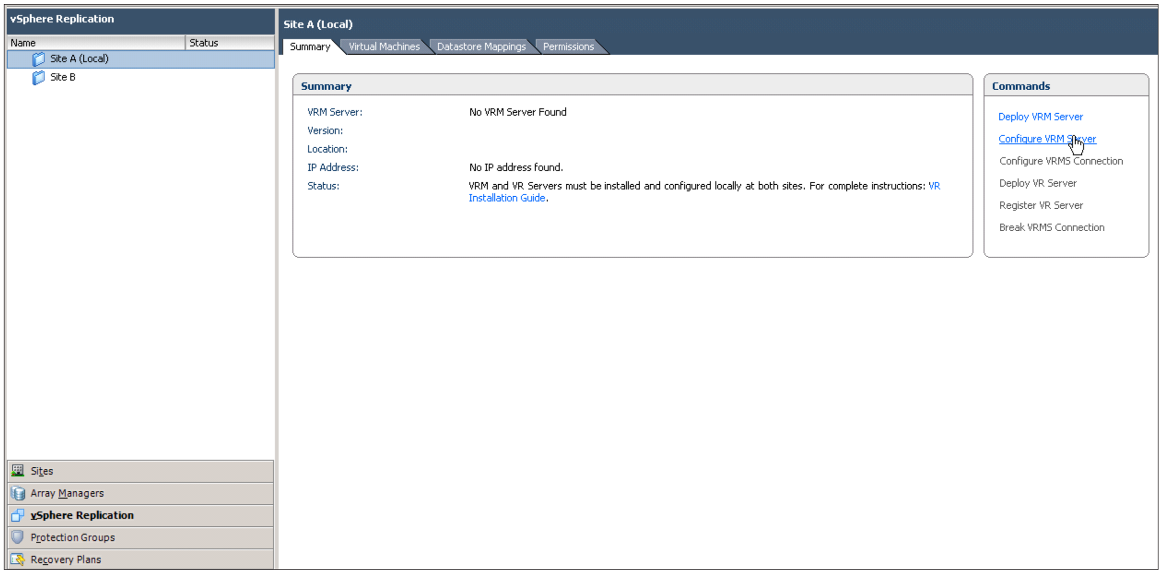

Summary

VRM Server: No VRM Server Found
Version:
Location:
IP Address: No IP address found.
Status: VRM and VR Servers must be installed and configured locally at both sites. For complete instructions: VR Installation Guide.

Figure 70. Configure the VRMS

5. Log in to the appliance using the user name **root** and the password you chose for the appliance in Step 1.

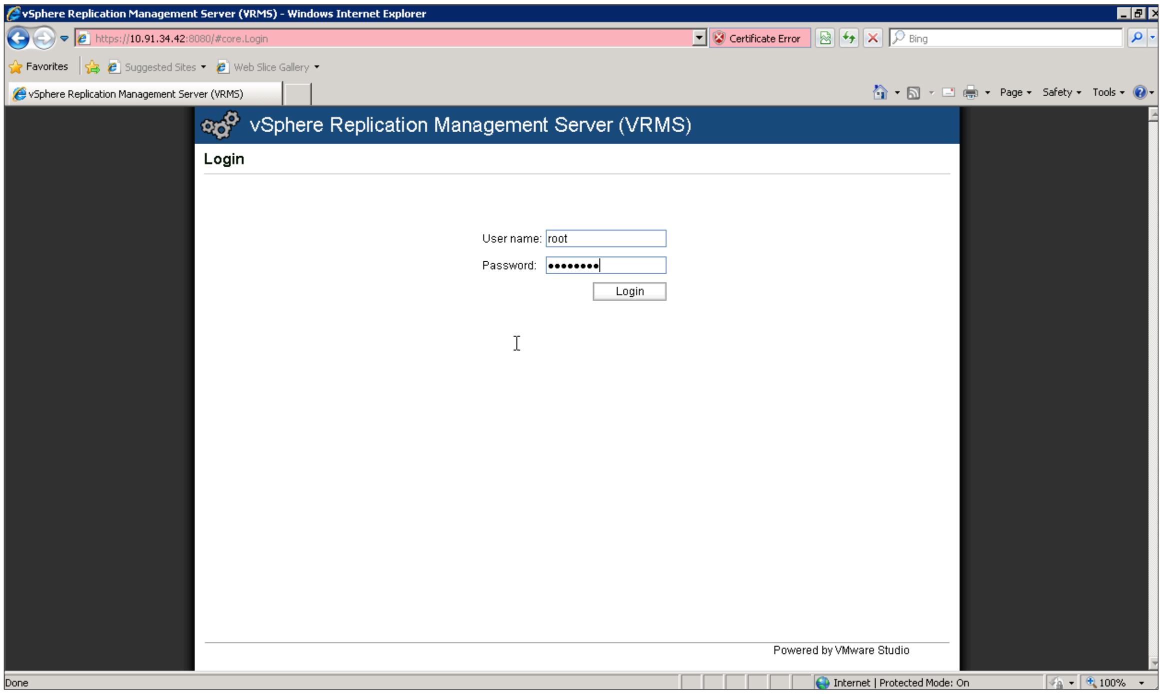
vSphere Replication Management Server (VRMS) - Windows Internet Explorer
https://10.91.34.42:8080/#core.Login
Certificate Error
Bing
Favorites
Suggested Sites
Web Slice Gallery
vSphere Replication Management Server (VRMS)
Page
Safety
Tools
vSphere Replication Management Server (VRMS)
Login
Login
Powered by VMware Studio
Done
Internet | Protected Mode: On
100%

User name: root
Password: ••••••••

Figure 71. Log in to VRMS Web Interface

6. Once logged in, within the **VRM** menu tab, choose the submenu tab labeled **Configuration**.

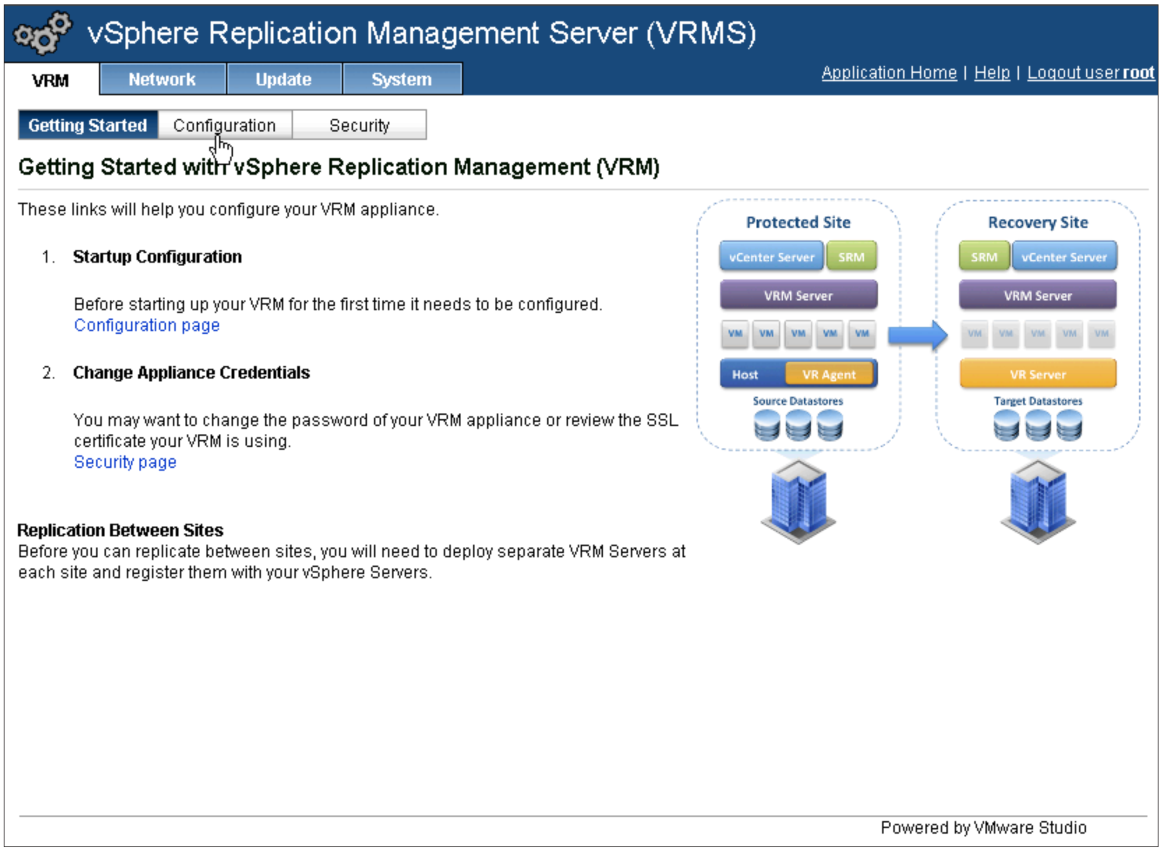

Getting Started with vSphere Replication Management (VRM)

These links will help you configure your VRM appliance.

1. **Startup Configuration**

   Before starting up your VRM for the first time it needs to be configured.
   Configuration page

2. **Change Appliance Credentials**

   You may want to change the password of your VRM appliance or review the SSL certificate your VRM is using.
   Security page

Replication Between Sites
Before you can replicate between sites, you will need to deploy separate VRM Servers at each site and register them with your vSphere Servers.

Figure 72. VRM Interface

7. In this location, you will need to enter information regarding the database used for vSphere Replication, vCenter information, and credentials to access both. In this evaluation guide, we are assuming the use of Microsoft SQL Server and the use of SQL Native authentication, although your environment may differ in terms of details for configuration. Select **Manual configuration** for Configuration Mode, and enter relevant database information and authentication data. Choose the **IP address** from the drop-down menu for the current VRMS appliance, which you are currently using, and provide a unique **VRMS site name.**

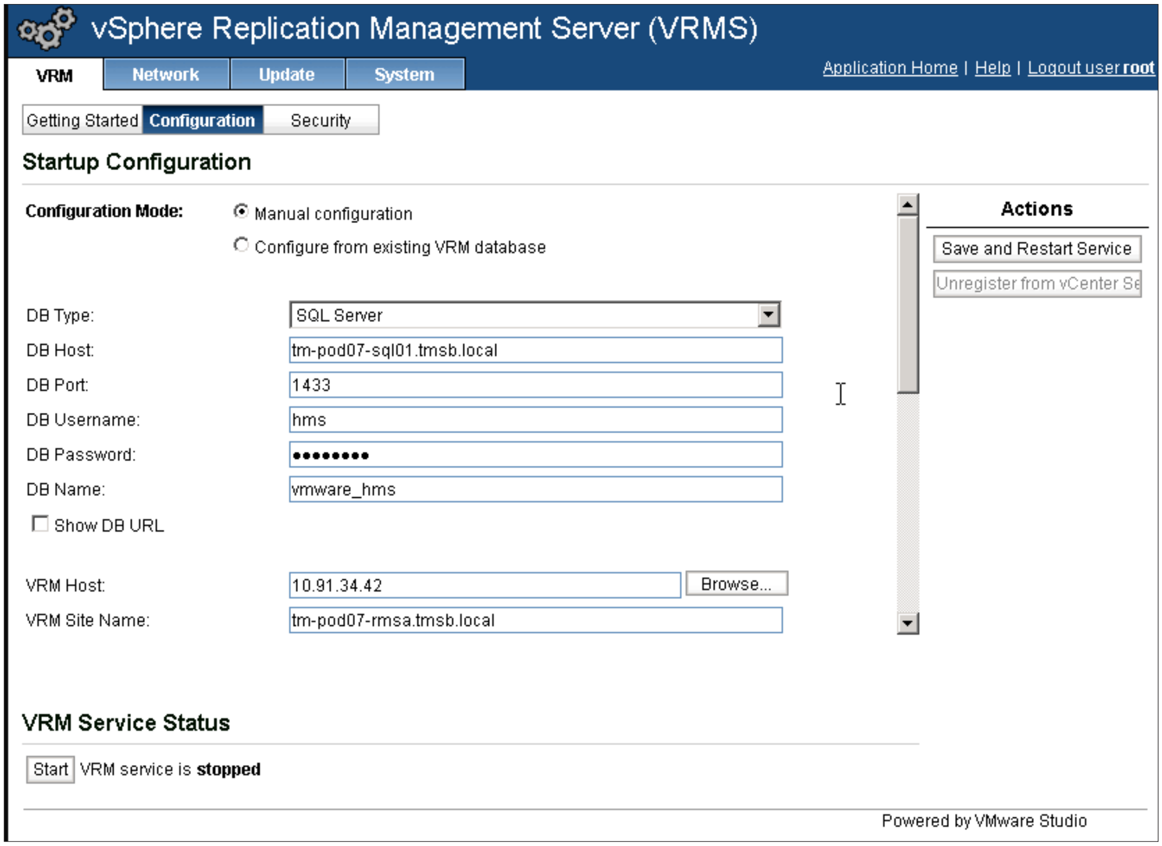

Startup Configuration

Configuration Mode: ● Manual configuration ○ Configure from existing VRM database

DB Type: SQL Server

DB Host: tm-pod07-sql01.tmsb.local

DB Port: 1433

DB Username: hms

DB Password: ••••••••

DB Name: vmware_hms

☐ Show DB URL

VRM Host: 10.91.34.42

VRM Site Name: tm-pod07-rmsa.tmsb.local

VRM Service Status

Start VRM service is **stopped**

Figure 73. VRMS Configuration – First Screen

8. After configuring this information, scroll down and notice the vCenter configuration information that is necessary. It is very important that the **vCenter Server Address** field is correctly populated. If you have used IP addresses for all site pairing activities, continue to use IP addresses in this location. If your SRM sites were paired with host names or fully qualified domain names, **it is important that you do the same at this location** as you did when pairing the sites. VRMS requires naming consistency throughout the process in order to function correctly. Enter the **vCenter Server Address** for the site you are currently using, which should be the protected site (Site A) vCenter. Click on **Generate and Install** an SSL Certificate after all the information is filled out.

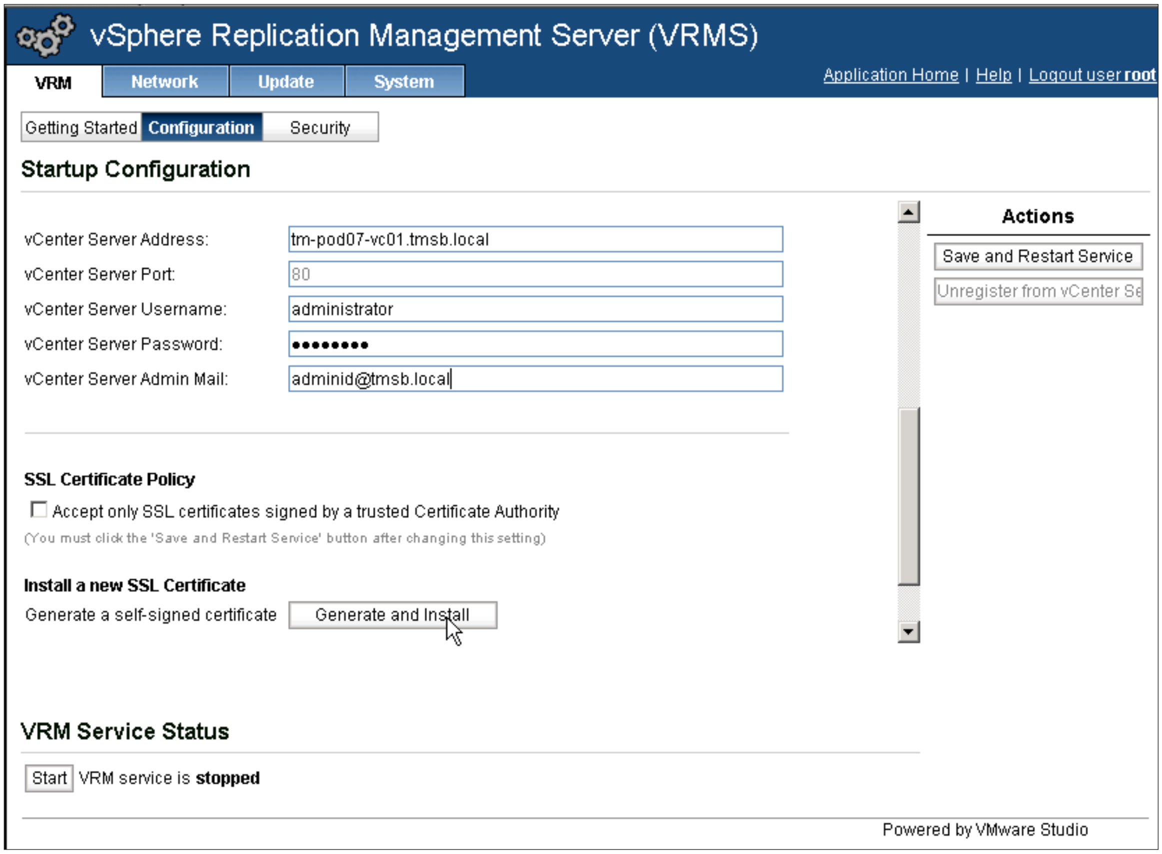

Startup Configuration

vCenter Server Address: tm-pod07-vc01.tmsb.local

vCenter Server Port: 80

vCenter Server Username: administrator

vCenter Server Password: ••••••••

vCenter Server Admin Mail: adminid@tmsb.local

SSL Certificate Policy

□ Accept only SSL certificates signed by a trusted Certificate Authority

(You must click the 'Save and Restart Service' button after changing this setting)

Install a new SSL Certificate

Generate a self-signed certificate [Generate and Install]

VRM Service Status

[Start] VRM service is **stopped**

Figure 74. vCenter Configuration Information in VRMS

9. When all the information is correctly filled in and the SSL Certificate is generated, click on the **Save and Restart Service** button. This will register the VRMS with vCenter and connect to the supplied database to run the initial configuration of vSphere Replication.

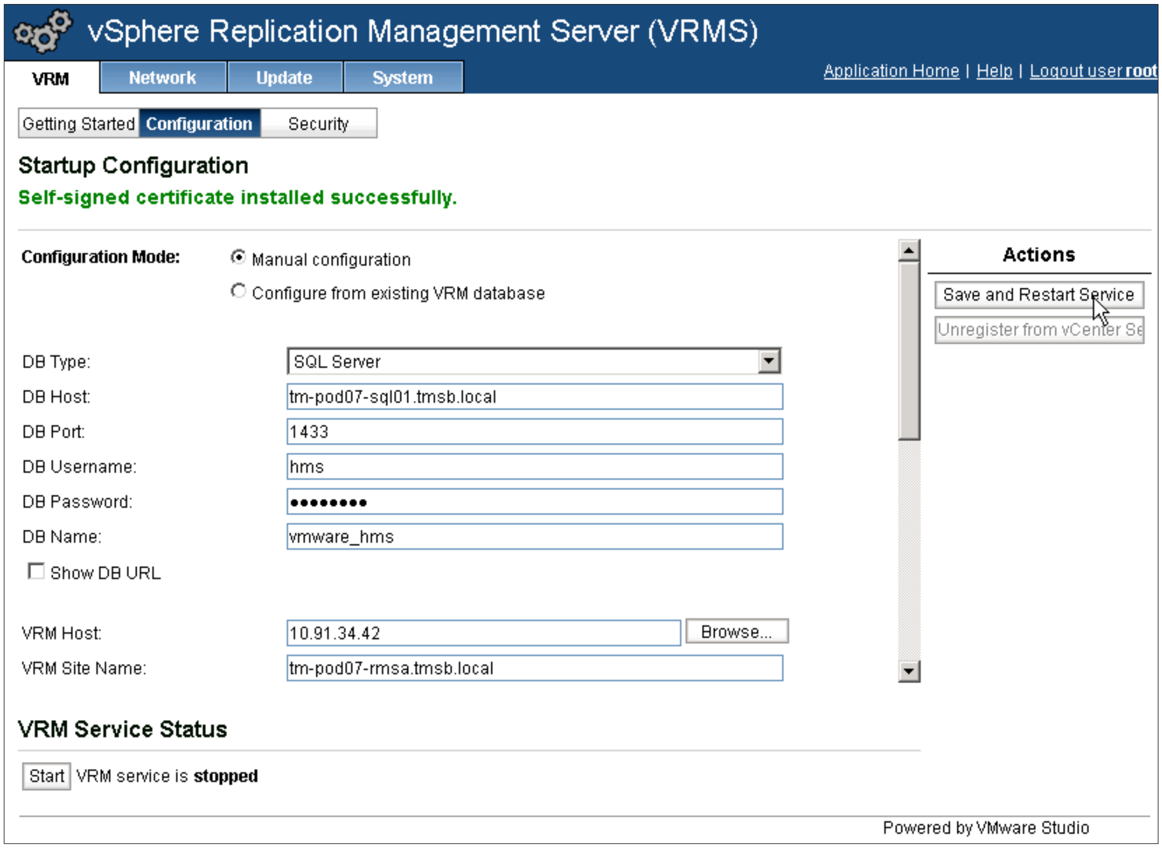

Startup Configuration

Self-signed certificate installed successfully.

Configuration Mode: ● Manual configuration ○ Configure from existing VRM database

DB Type: SQL Server

DB Host: tm-pod07-sql01.tmsb.local

DB Port: 1433

DB Username: hms

DB Password: ••••••••

DB Name: vmware_hms

□ Show DB URL

VRM Host: 10.91.34.42

VRM Site Name: tm-pod07-rmsa.tmsb.local

VRM Service Status

VRM service is **stopped**

Figure 75. Save and Restart Service

10. A pop-up dialogue box will ask you to confirm the vCenter SSL Certificate. Press **Accept** to continue.

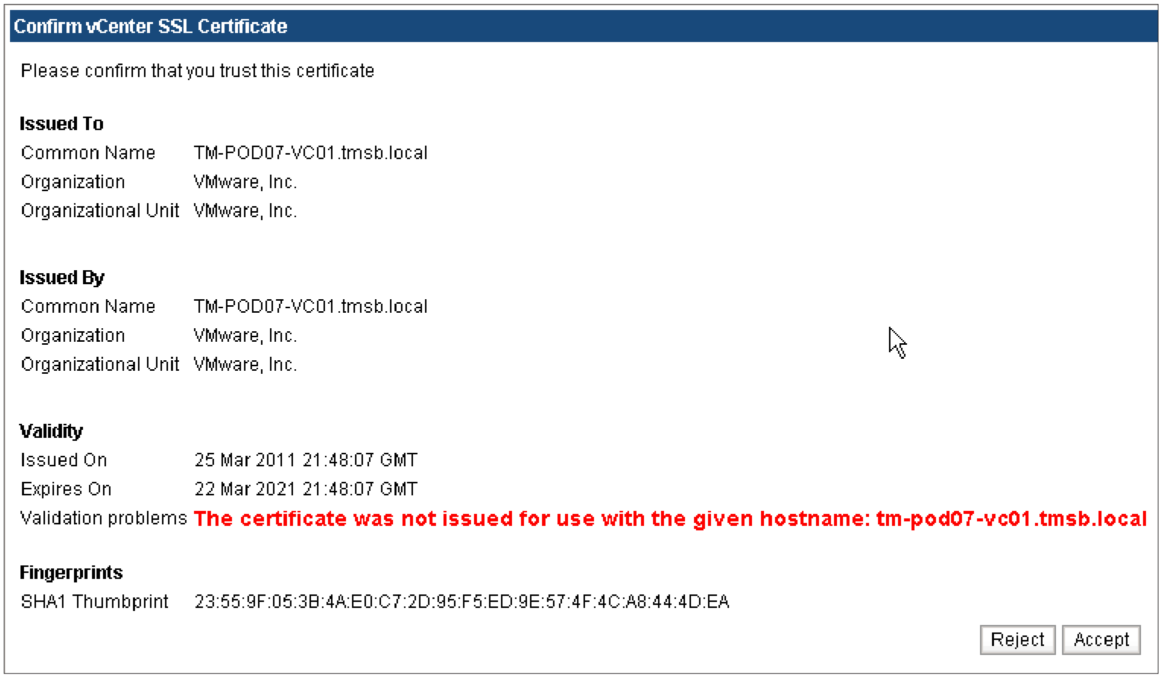

Confirm vCenter SSL Certificate

Please confirm that you trust this certificate

Issued To

Common Name	TM-POD07-VC01.tmsb.local
Organization	VMware, Inc.
Organizational Unit	VMware, Inc.

Issued By

Common Name	TM-POD07-VC01.tmsb.local
Organization	VMware, Inc.
Organizational Unit	VMware, Inc.

Validity

Issued On	25 Mar 2011 21:48:07 GMT
Expires On	22 Mar 2021 21:48:07 GMT
Validation problems	**The certificate was not issued for use with the given hostname: tm-pod07-vc01.tmsb.local**

Fingerprints

SHA1 Thumbprint	23:55:9F:05:3B:4A:E0:C7:2D:95:F5:ED:9E:57:4F:4C:A8:44:4D:EA

Figure 76. vCenter SSL Certificate

11. If VRMS has been successful communicating with vCenter and the database, it will return to the Configuration screen with a green message labeled **Successfully saved the startup configuration.** This may take a few minutes to return. Wait until a message is generated, whether it is the green success message or an error. If an error message is generated, re-examine both the database and vCenter information carefully and try again.

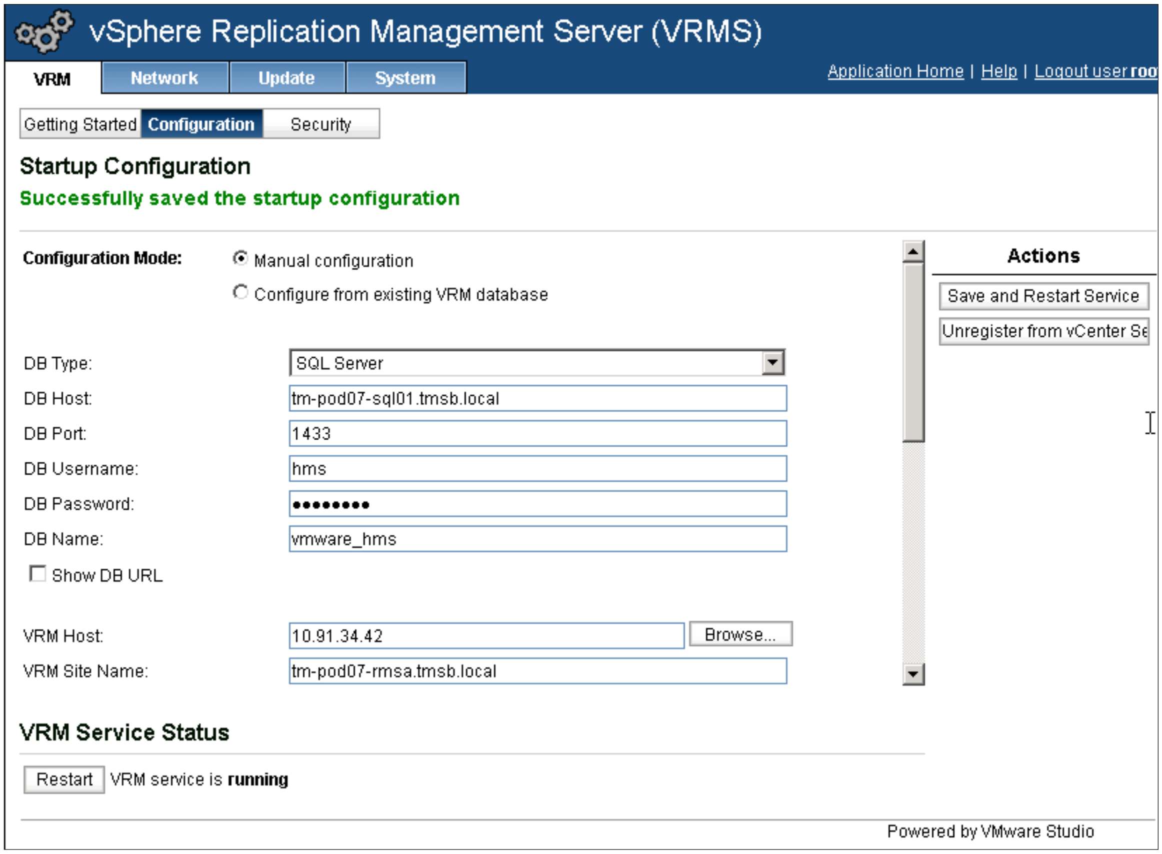

vSphere Replication Management Server (VRMS)
VRM
Network
Update
System
Application Home | Help | Logout user roo
Getting Started
Configuration
Security

Startup Configuration

Successfully saved the startup configuration

Configuration Mode: ● Manual configuration ○ Configure from existing VRM database

DB Type: SQL Server
DB Host: tm-pod07-sql01.tmsb.local
DB Port: 1433
DB Username: hms
DB Password: ●●●●●●●●●
DB Name: vmware_hms
□ Show DB URL

VRM Host: 10.91.34.42
VRM Site Name: tm-pod07-rmsa.tmsb.local

Browse...
Actions
Save and Restart Service
Unregister from vCenter Se

VRM Service Status

Restart VRM service is **running**

Figure 77. Successful Configuration

12. You can now log out of the VRMS Web interface and return to the vSphere Client. You may expect to see a number of certificate security warnings that indicate secure connections are now being attempted between both vCenters and the VRMS appliances. This will occur once per session when opening the SRM interface. Choose to both **Install** the certificate and **Ignore** errors.

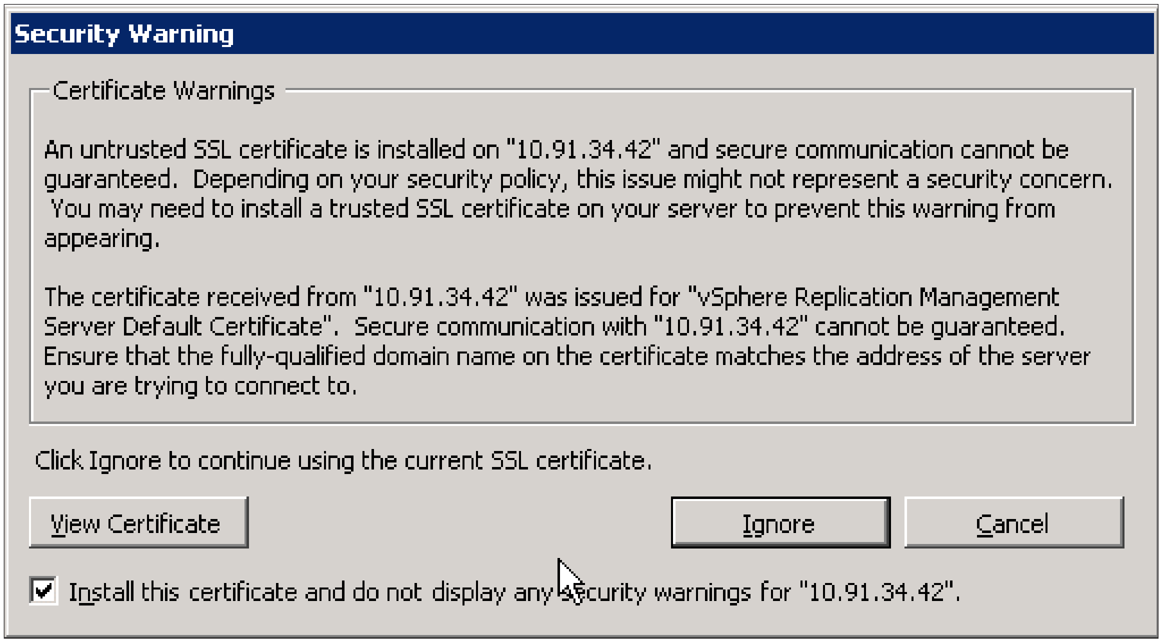

Security Warning

Certificate Warnings

An untrusted SSL certificate is installed on "10.91.34.42" and secure communication cannot be guaranteed. Depending on your security policy, this issue might not represent a security concern. You may need to install a trusted SSL certificate on your server to prevent this warning from appearing.

The certificate received from "10.91.34.42" was issued for "vSphere Replication Management Server Default Certificate". Secure communication with "10.91.34.42" cannot be guaranteed. Ensure that the fully-qualified domain name on the certificate matches the address of the server you are trying to connect to.

Click Ignore to continue using the current SSL certificate.

☑ Install this certificate and do not display any security warnings for "10.91.34.42".

Figure 78. SSL Certificate Warnings

Step 4: Configure VRMS appliance and VRMS Web management for recovery site

1. Return to the vSphere Client in the **SRM** solution page, and repeat the configuration of the VRMS appliance at the **recovery site** and configuration of the Web management interface at the **recovery site.**

Remember, when configuring the VRMS at the recovery site, you will need to enter information specific to that site. Enter a unique VRMS database, the recovery site vCenter, and so on. Follow the preceding process outlined in both **Step 2: Configuring VRMS appliances** and **Step 3: Configuring VRMS through Web management interface,** while ensuring that you use correct site-specific information.

Step 5: Configuring VRMS pairing connection

After both VRMS appliances are responding, they must be connected to one another in order to create the framework for replication. This step will configure the connection between VRMS servers.

To configure the VRMS connection, follow these steps:

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.
3. Click the **Site Recovery icon** on the vSphere Client Home page under Solutions and Applications.
4. Choose the menu item on the left navigation pane entitled **vSphere Replication,** and select the protected site (Site A). Click **Configure VRMS Connection** in the actions list on the right.

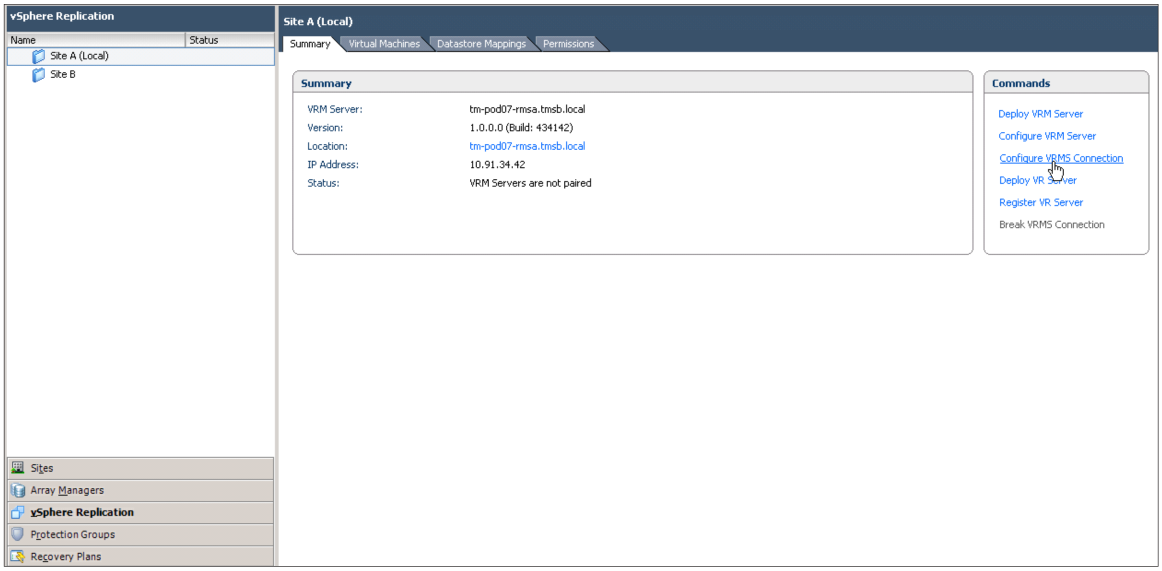

VRM Server: tm-pod07-rmsa.tmsb.local
Version: 1.0.0.0 (Build: 434142)
Location: tm-pod07-rmsa.tmsb.local
IP Address: 10.91.34.42
Status: VRM Servers are not paired

Figure 79. Configure VMRS Connection

5. SRM will query if you want to configure the VRMS connection. Click **Yes** to continue.

6. SRM might issue a server certificate error. This is strictly because you did not use signed certificates earlier and is completely normal. Click **OK** to continue.

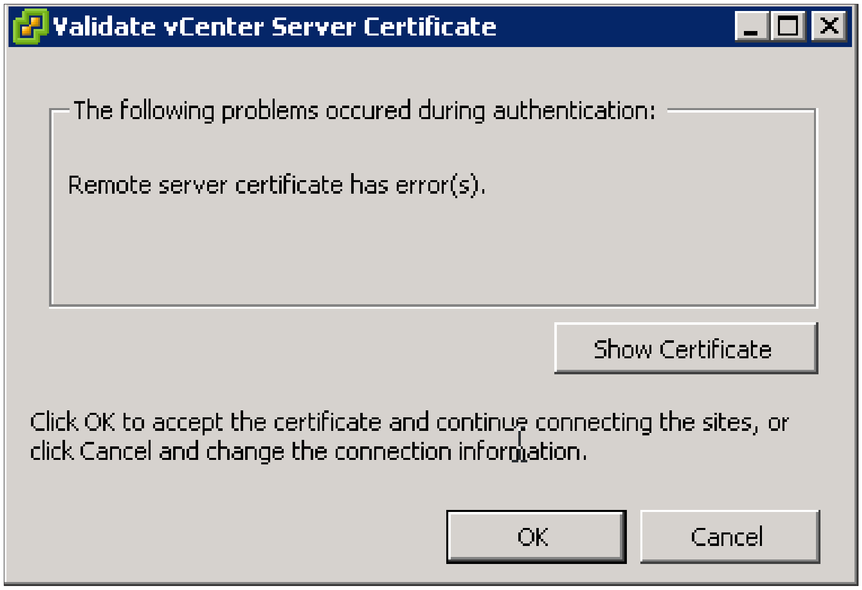

The following problems occured during authentication:

Remote server certificate has error(s).

Click OK to accept the certificate and continue connecting the sites, or click Cancel and change the connection information.

Figure 80. VMRS Certificate Errors

7. SRM will prompt you for login credentials for the remote vCenter server. Provide your credentials, and click **OK** to continue. There might be another server certificate error as in Figure 80. For this error as well, press **OK** to continue.

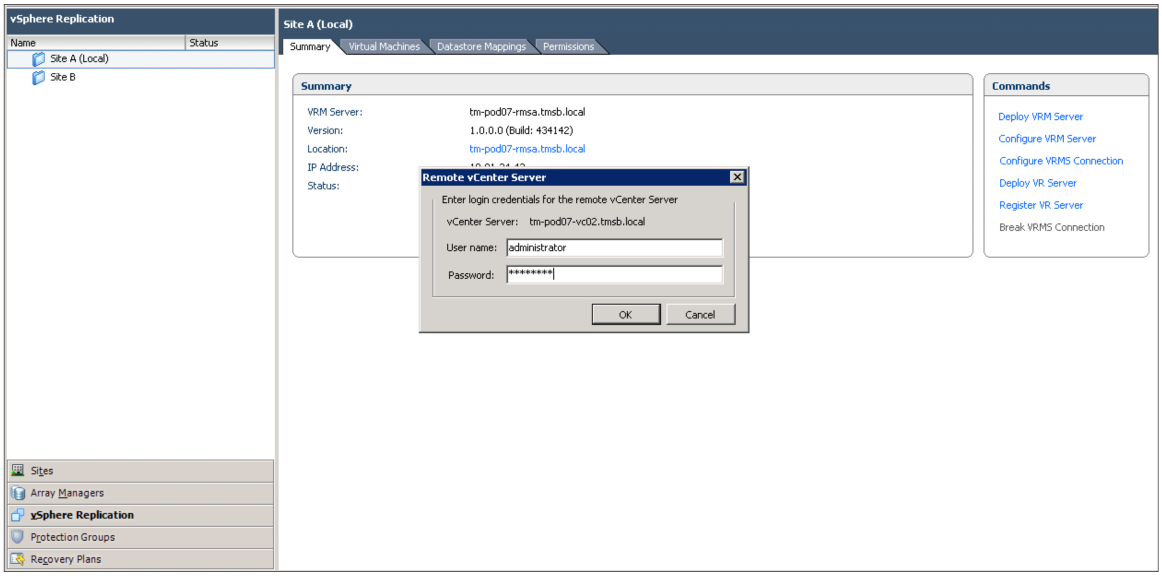

Summary

VRM Server: tm-pod07-rmsa.tmsb.local
Version: 1.0.0.0 (Build: 434142)
Location: tm-pod07-rmsa.tmsb.local
IP Address:
Status:

Remote vCenter Server

Enter login credentials for the remote vCenter Server

vCenter Server: tm-pod07-vc02.tmsb.local
User name: administrator
Password: ********

Figure 81. Credential Authorization Prompt

8. After a momentary delay while configuring communication between sites, a success message that configuration pairing succeeded should appear. Press **OK** to continue.

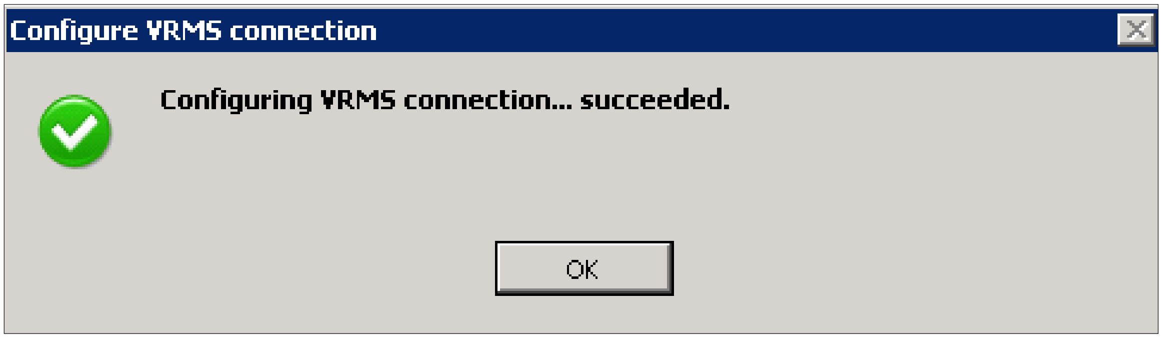

Configure VRMS connection

Configuring VRMS connection... succeeded.

Figure 82. Successful Configuration

9. Both sites within the SRM vSphere Replication **Summary** screen should now show information about the VRM server, such as Location and IP Address. Importantly, **Status** should display **Connected** on both sites.

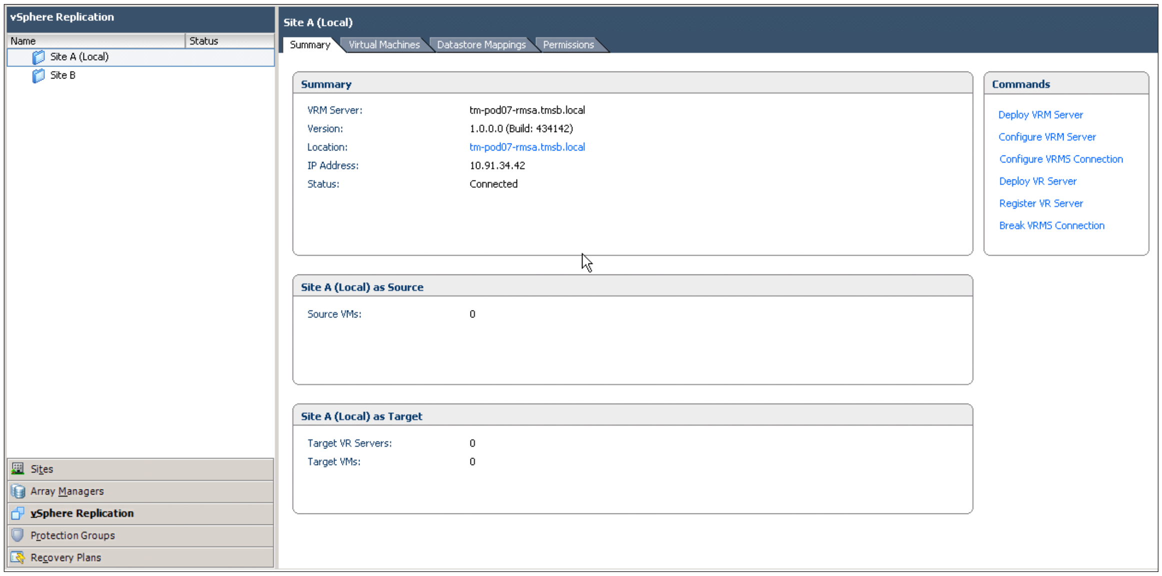

Summary

VRM Server: tm-pod07-rmsa.tmsb.local
Version: 1.0.0.0 (Build: 434142)
Location: tm-pod07-rmsa.tmsb.local
IP Address: 10.91.34.42
Status: Connected

Site A (Local) as Source

Source VMs: 0

Site A (Local) as Target

Target VR Servers: 0
Target VMs: 0

Figure 83. Connected Status – Site A

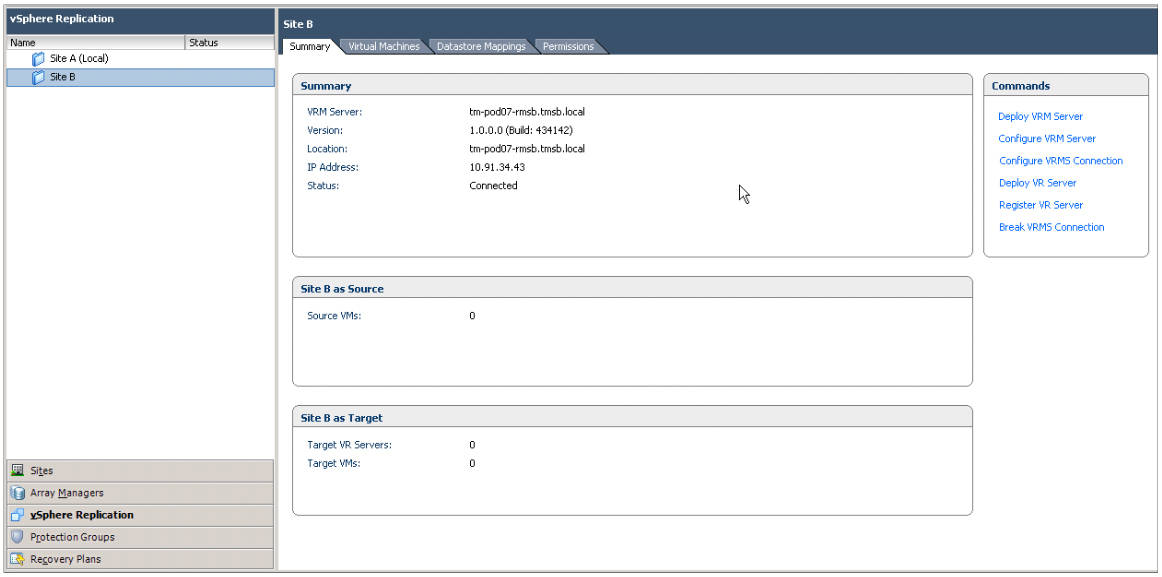

Summary

VRM Server: tm-pod07-rmsb.tmsb.local
Version: 1.0.0.0 (Build: 434142)
Location: tm-pod07-rmsb.tmsb.local
IP Address: 10.91.34.43
Status: Connected

Site B as Source

Source VMs: 0

Site B as Target

Target VR Servers: 0
Target VMs: 0

Figure 84. Connected Status – Site B

Step 6: Deploying a vSphere Replication Server (VRS)

The VRS acts as a recipient of changed blocks captured by vSphere Replication. The VRMS directs the VR agents on the ESXi hosts at the protected site (Site A) to pass changed blocks to the VRS that resides at the recovery site (Site B). Therefore, there is a requirement that at least one VRS must be deployed at the recovery site (Site B). If bidirectional protection using VR is required, there must be a VRS at the protected site (Site A) as well. For this evaluation, we will only deploy and register a single VRS at the recovery site.

To deploy a VRS, follow these steps:

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.
3. Click the **Site Recovery icon** on the vSphere Client Home page under Solutions and Applications.
4. Choose the menu item on the left navigation pane entitled **vSphere Replication,** and select the protected site (Site A). Click **Deploy VR Server** in the actions list on the right. This will launch a deployment wizard.

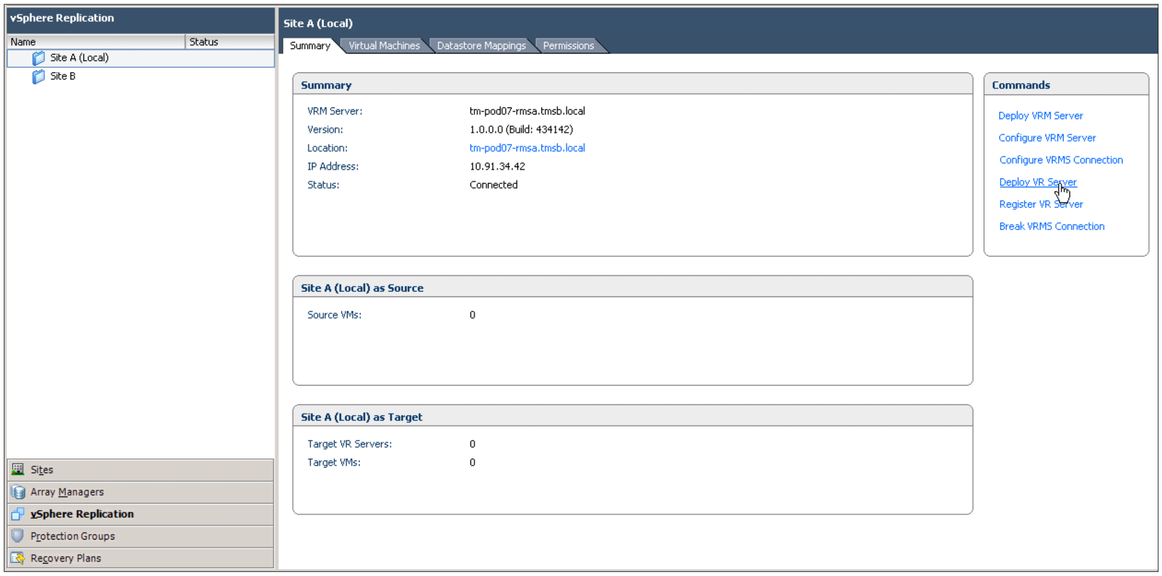

Summary

VRM Server: tm-pod07-rmsa.tmsb.local
Version: 1.0.0.0 (Build: 434142)
Location: tm-pod07-rmsa.tmsb.local
IP Address: 10.91.34.42
Status: Connected

Site A (Local) as Source

Source VMs: 0

Site A (Local) as Target

Target VR Servers: 0
Target VMs: 0

Figure 85. Deploy VR Server

5. Follow the prompts to deploy the VR Server (VRS) at the **recovery site.**

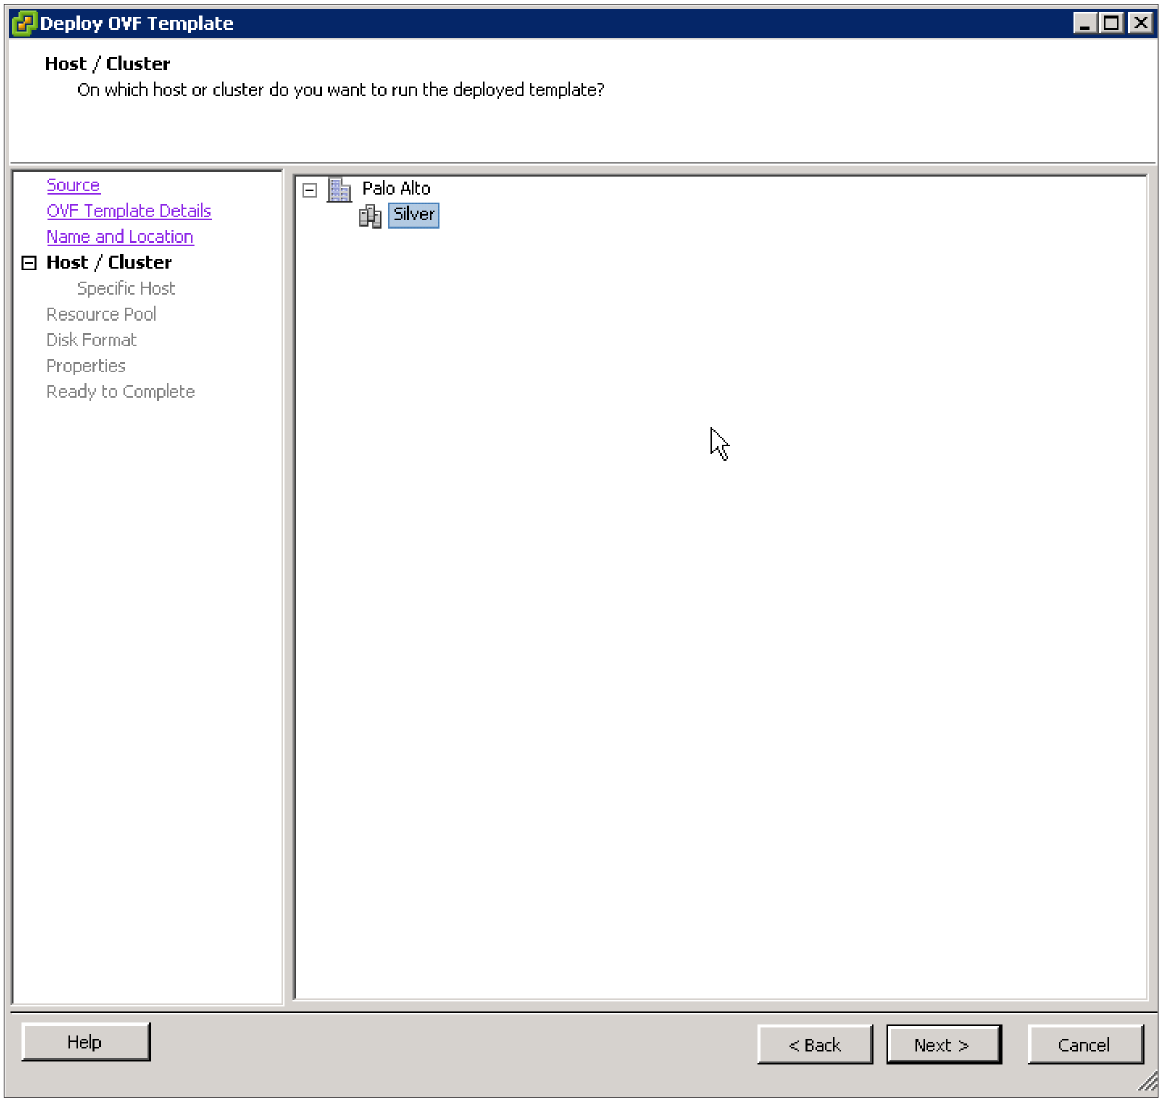

Deploy OVF Template

Host / Cluster

On which host or cluster do you want to run the deployed template?

Figure 86. Deploying VRS at the Recovery Site (Site B)

6. Name the VRS appropriately and optionally deploy it into a specific folder and resource pool at the **recovery site.**

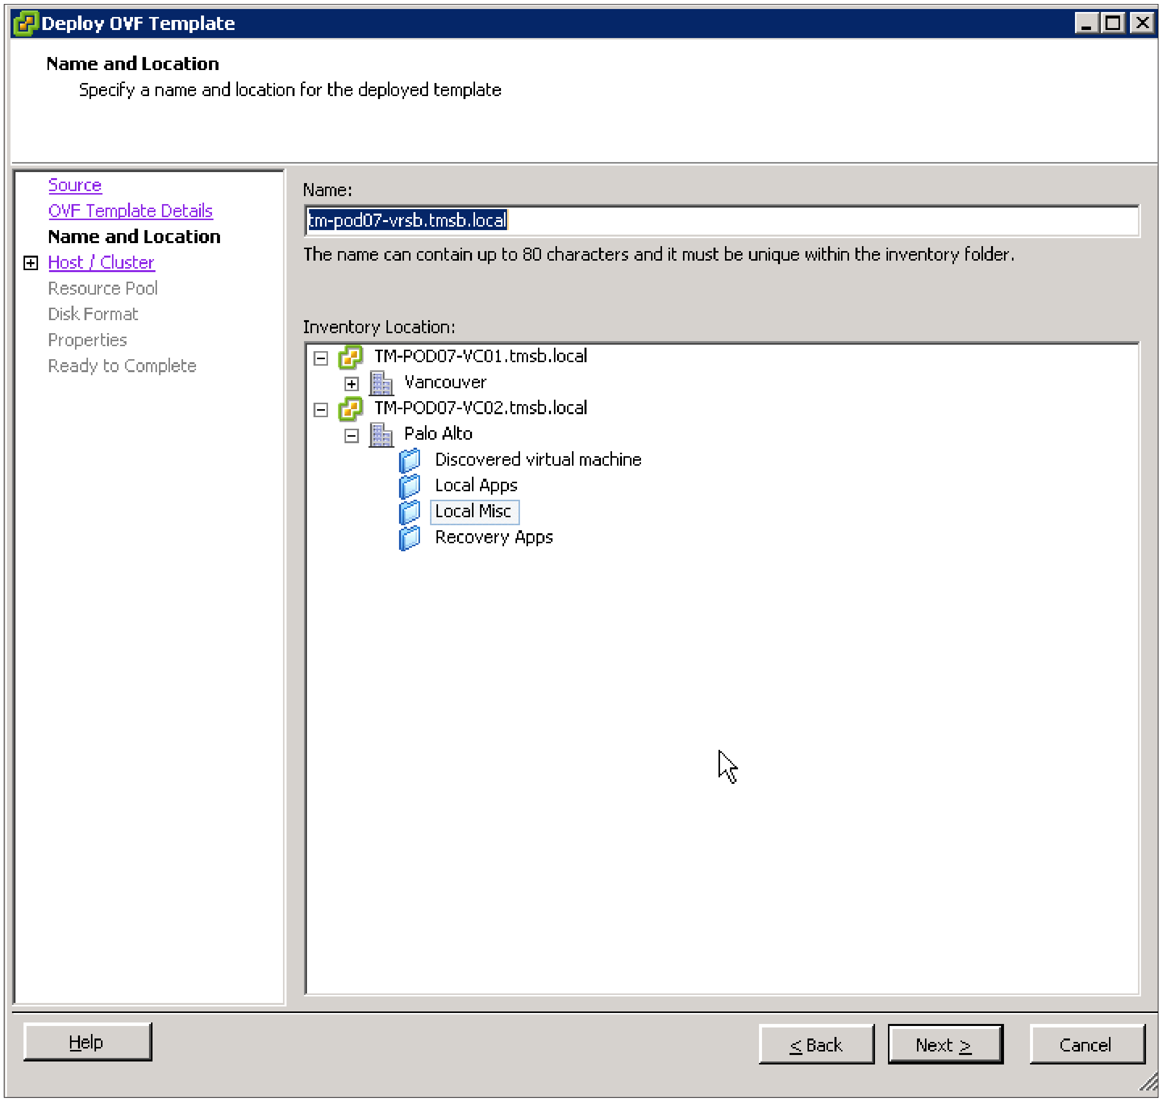

Deploy OVF Template

Name and Location

Specify a name and location for the deployed template

Name: tm-pod07-vrsb.tmsb.local

The name can contain up to 80 characters and it must be unique within the inventory folder.

Inventory Location:

Figure 87. Name and Place for the VRS at the Recovery Site (Site B)

7. Ensure that the name and IP addresses given to the VRS appliance work correctly with forward, reverse, short and FQDN DNS lookups.

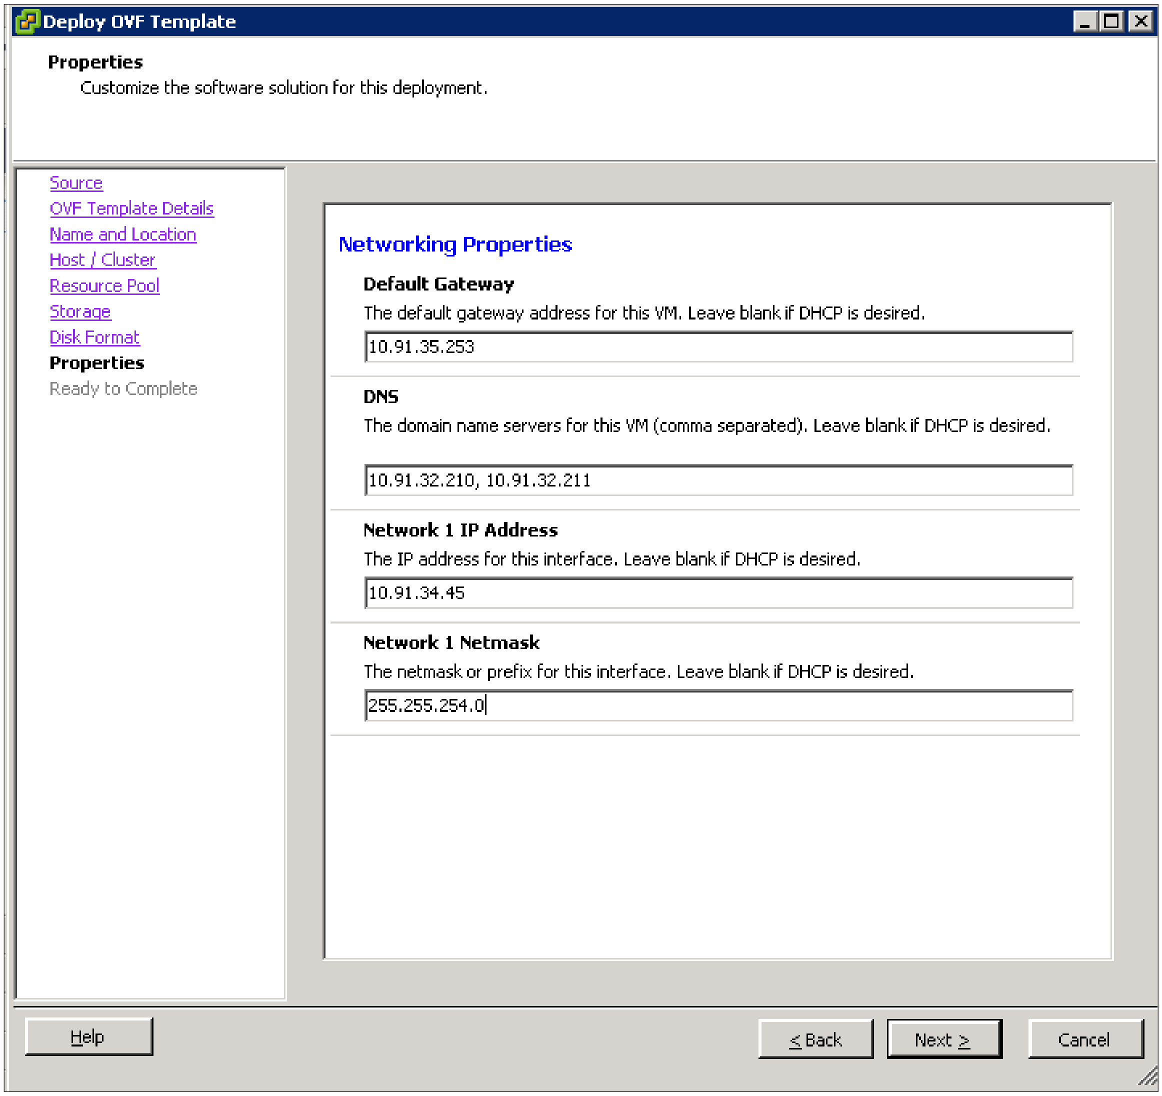

Properties

Customize the software solution for this deployment.

Networking Properties

Default Gateway

The default gateway address for this VM. Leave blank if DHCP is desired.

10.91.35.253

DNS

The domain name servers for this VM (comma separated). Leave blank if DHCP is desired.

10.91.32.210, 10.91.32.211

Network 1 IP Address

The IP address for this interface. Leave blank if DHCP is desired.

10.91.34.45

Network 1 Netmask

The netmask or prefix for this interface. Leave blank if DHCP is desired.

255.255.254.0

Figure 88. Networking Information for VR Server

8. When all the information is complete and correct, click **Finish** to deploy the VR Server.

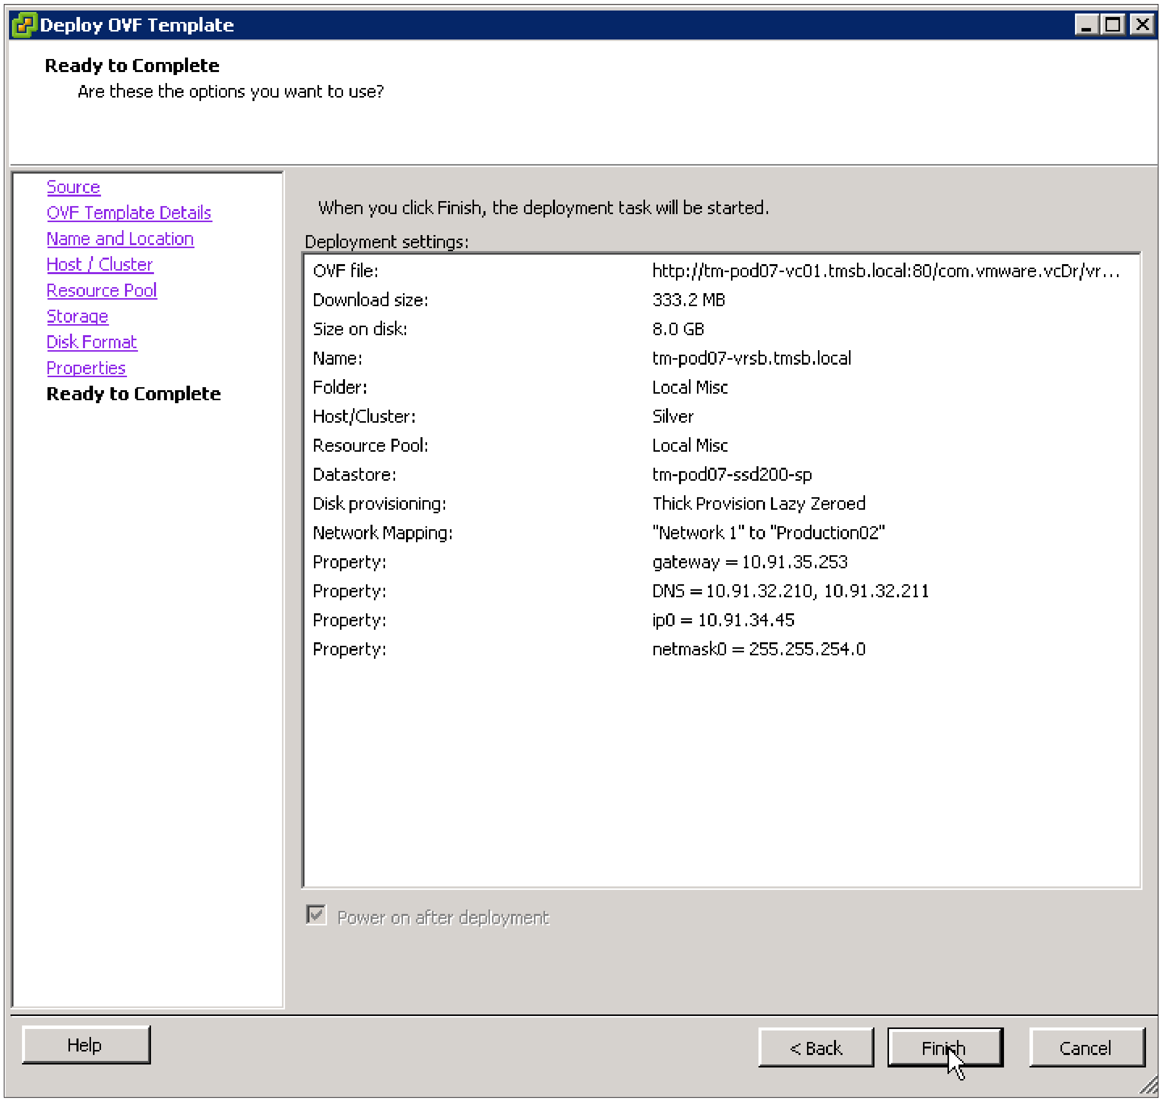

Ready to Complete

Are these the options you want to use?

When you click Finish, the deployment task will be started.

Deployment settings:

OVF file:	http://tm-pod07-vc01.tmsb.local:80/com.vmware.vcDr/vr...
Download size:	333.2 MB
Size on disk:	8.0 GB
Name:	tm-pod07-vrsb.tmsb.local
Folder:	Local Misc
Host/Cluster:	Silver
Resource Pool:	Local Misc
Datastore:	tm-pod07-ssd200-sp
Disk provisioning:	Thick Provision Lazy Zeroed
Network Mapping:	"Network 1" to "Production02"
Property:	gateway = 10.91.35.253
Property:	DNS = 10.91.32.210, 10.91.32.211
Property:	ip0 = 10.91.34.45
Property:	netmask0 = 255.255.254.0

☑ Power on after deployment

Figure 89. Deploy the VR Server After Reviewing All Information

Step 7: Register a VR Server

To complete the process of deployment and configuration of vSphere Replication, you **must register the VRS** to the VRMS framework to list it as a valid destination for changed blocks.

To register the VRS, follow these steps:

1. Open the vSphere Client and connect to the vCenter server at the protected site.

2. Log in as a vSphere administrator.

3. Click the **Site Recovery icon** on the vSphere Client Home page under Solutions and Applications.

4. Choose the menu item on the left navigation pane entitled **vSphere Replication,** and select the **recovery site (Site B).** Click **Register VR Server** in the actions list on the right.

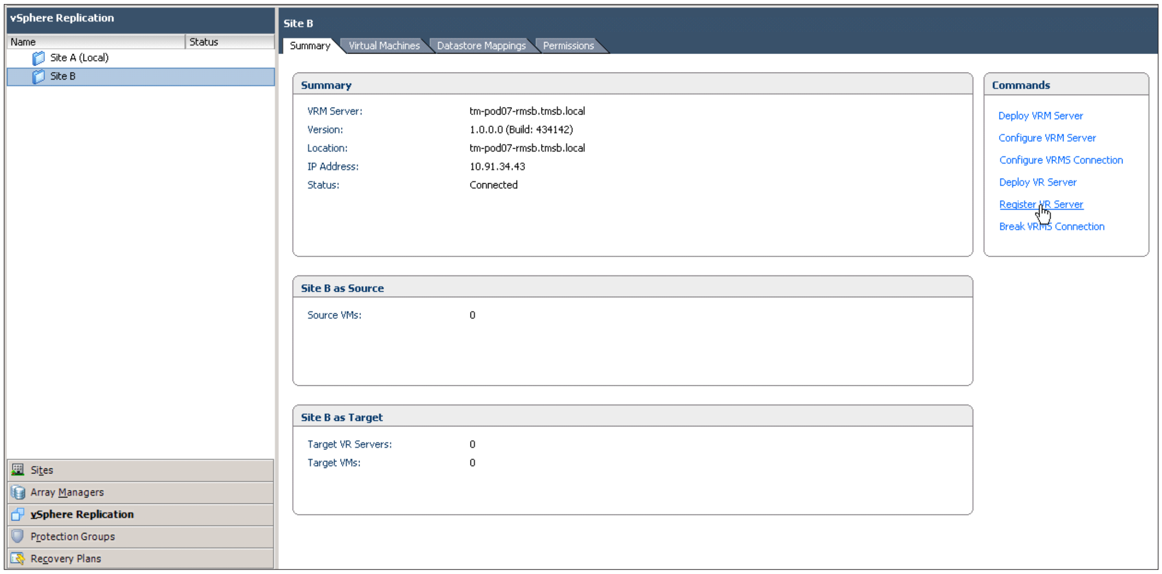

Summary

VRM Server: tm-pod07-rmsb.tmsb.local
Version: 1.0.0.0 (Build: 434142)
Location: tm-pod07-rmsb.tmsb.local
IP Address: 10.91.34.43
Status: Connected

Site B as Source

Source VMs: 0

Site B as Target

Target VR Servers: 0
Target VMs: 0

Figure 90. Register VR Server on the Recovery Site

5. The **Register VR Server pop-up screen** will show a tree of VMs available at the recovery site to register. Find the VR Server that was deployed in Step 6, select it and press **OK** to register the VRS.

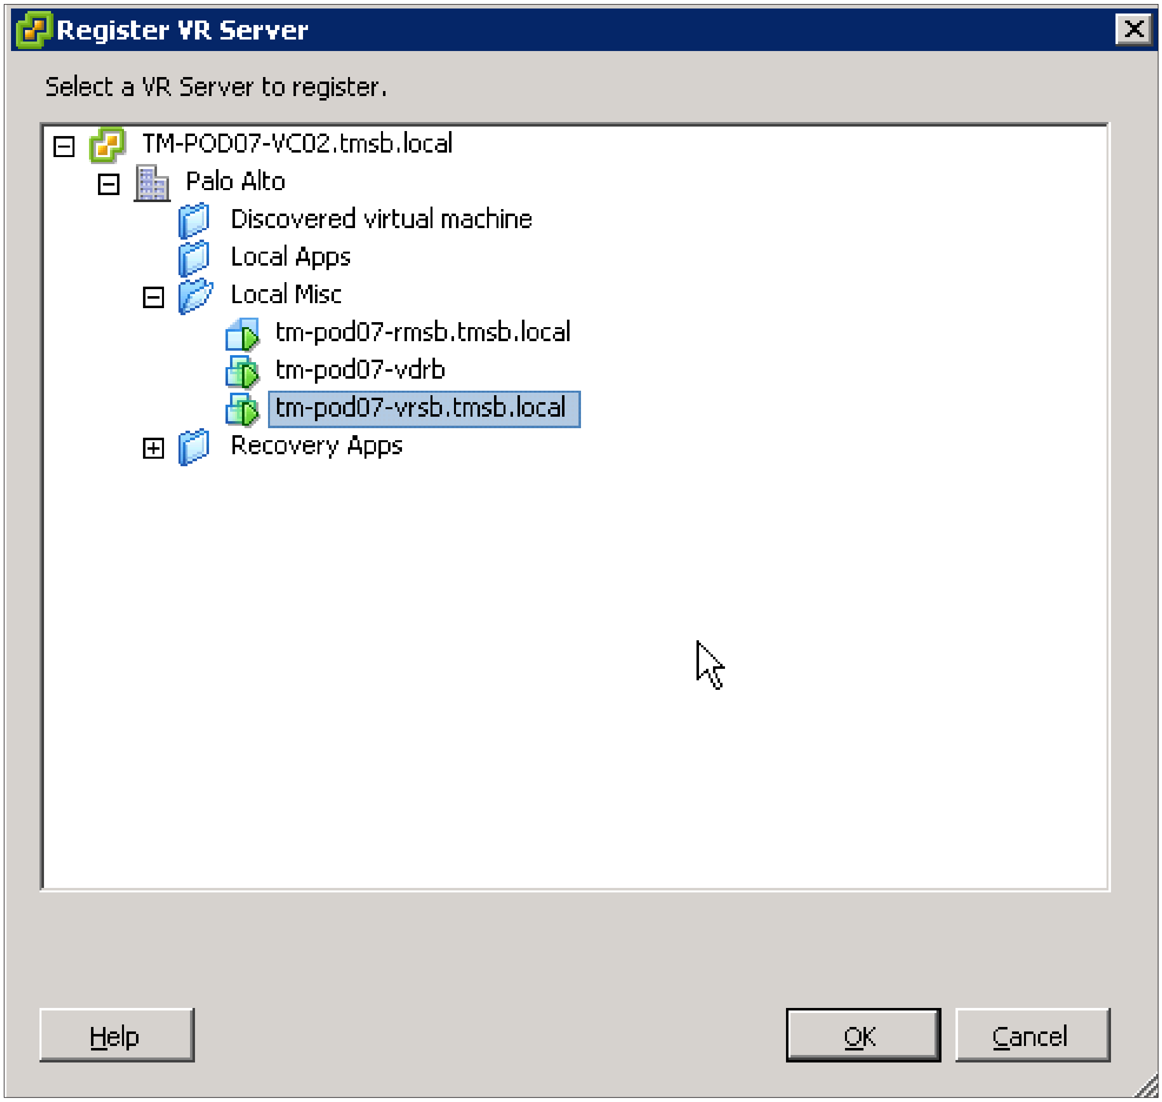

Register VR Server

Select a VR Server to register.

- TM-POD07-VC02.tmsb.local
  - Palo Alto
    - Discovered virtual machine
    - Local Apps
    - Local Misc
      - tm-pod07-rmsb.tmsb.local
      - tm-pod07-vdrb
      - tm-pod07-vrsb.tmsb.local
    - Recovery Apps

Figure 91. Register VR Server Pop-up Screen

6. When a prompt appears verifying that you wish to register the selected VM as a VRS, ensure that you have chosen the correct VM, and press **Yes** to continue. You may see server certificate warning messages. If so, press **OK** to continue.

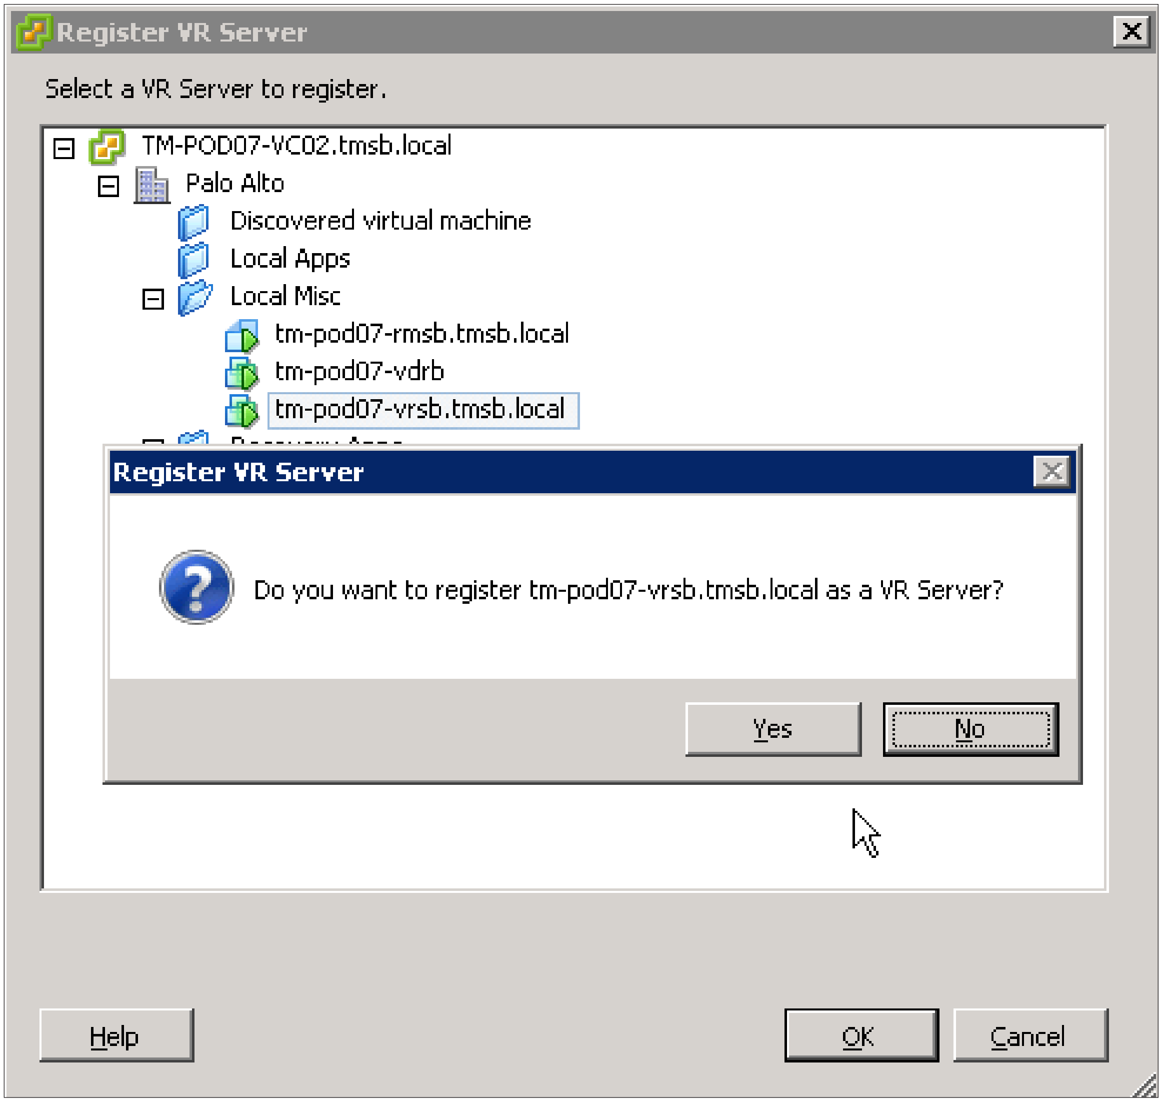

Register VR Server

Select a VR Server to register.

Register VR Server

Do you want to register tm-pod07-vrsb.tmsb.local as a VR Server?

Figure 92. Ensure That the Correct VM Is Selected

7. When the VRS is successfully registered, the following screen will notify you of success.

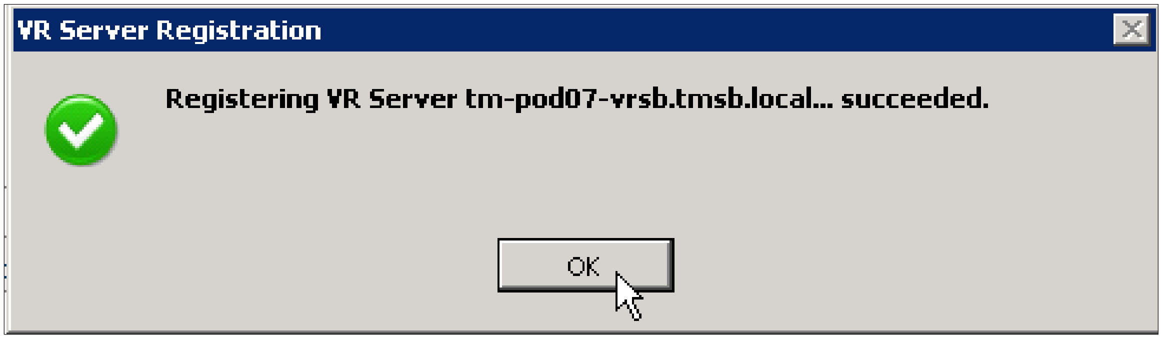

VR Server Registration

Registering VR Server tm-pod07-vrsb.tmsb.local... succeeded.

Figure 93. Successful Registration of the VRS

At this point, the newly registered VRS will appear in the **vSphere Replication navigation window** within the folder representing the recovery site (Site B). You might click the VRS to see where information will be populated once vSphere Replication is configured for virtual machines. The status should show **Connected** with zero virtual machines listed.

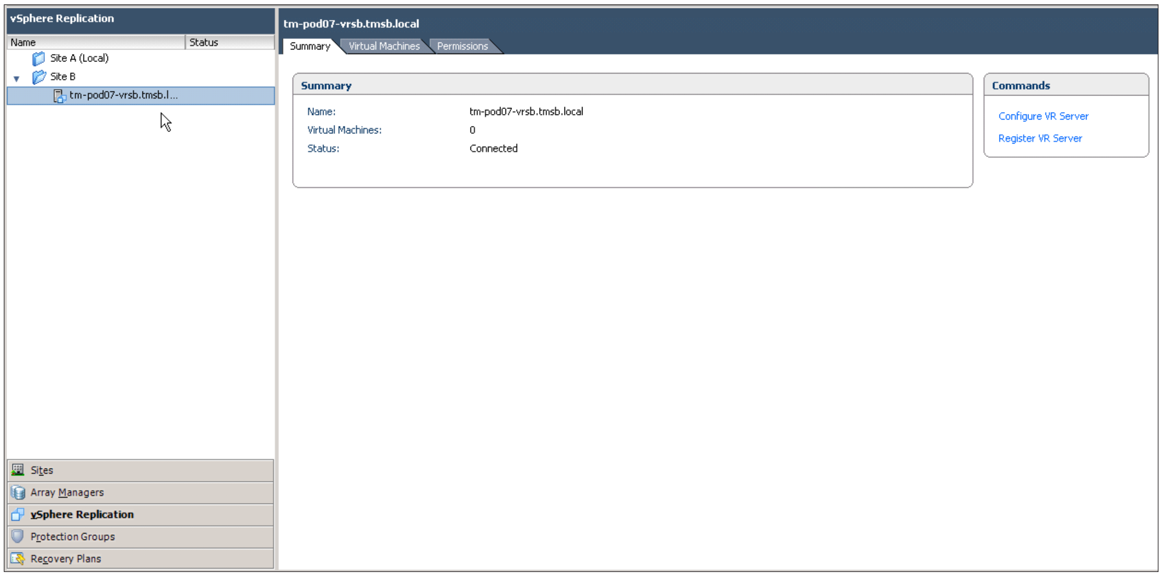

Name: tm-pod07-vrsb.tmsb.local
Virtual Machines: 0
Status: Connected

Figure 94. The VRS Registered to Site B

This concludes the installation and configuration of vSphere Replication components. Replication and protection is now configured as a property of the virtual machines themselves.

Step 8: Configuring protection for a vSphere Replication–protected VM

Once the vSphere Replication framework has been deployed and configured, virtual machines can now be set up for replication, and added to a protection group and recovery plan.

In order to configure protection for a VM using vSphere Replication, follow these steps:

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.
3. Click the **Hosts and Clusters** icon on the vSphere Client Home page under Inventory.

4. Expand the list of virtual machines at the protected site (Site A) and select a virtual machine to be protected with vSphere Replication. Choose a VM that has been previously identified as one not residing on replicated storage and one that is not currently part of a protection group. Right-click the virtual machine and select **vSphere Replication** at the bottom of the pop-up screen.

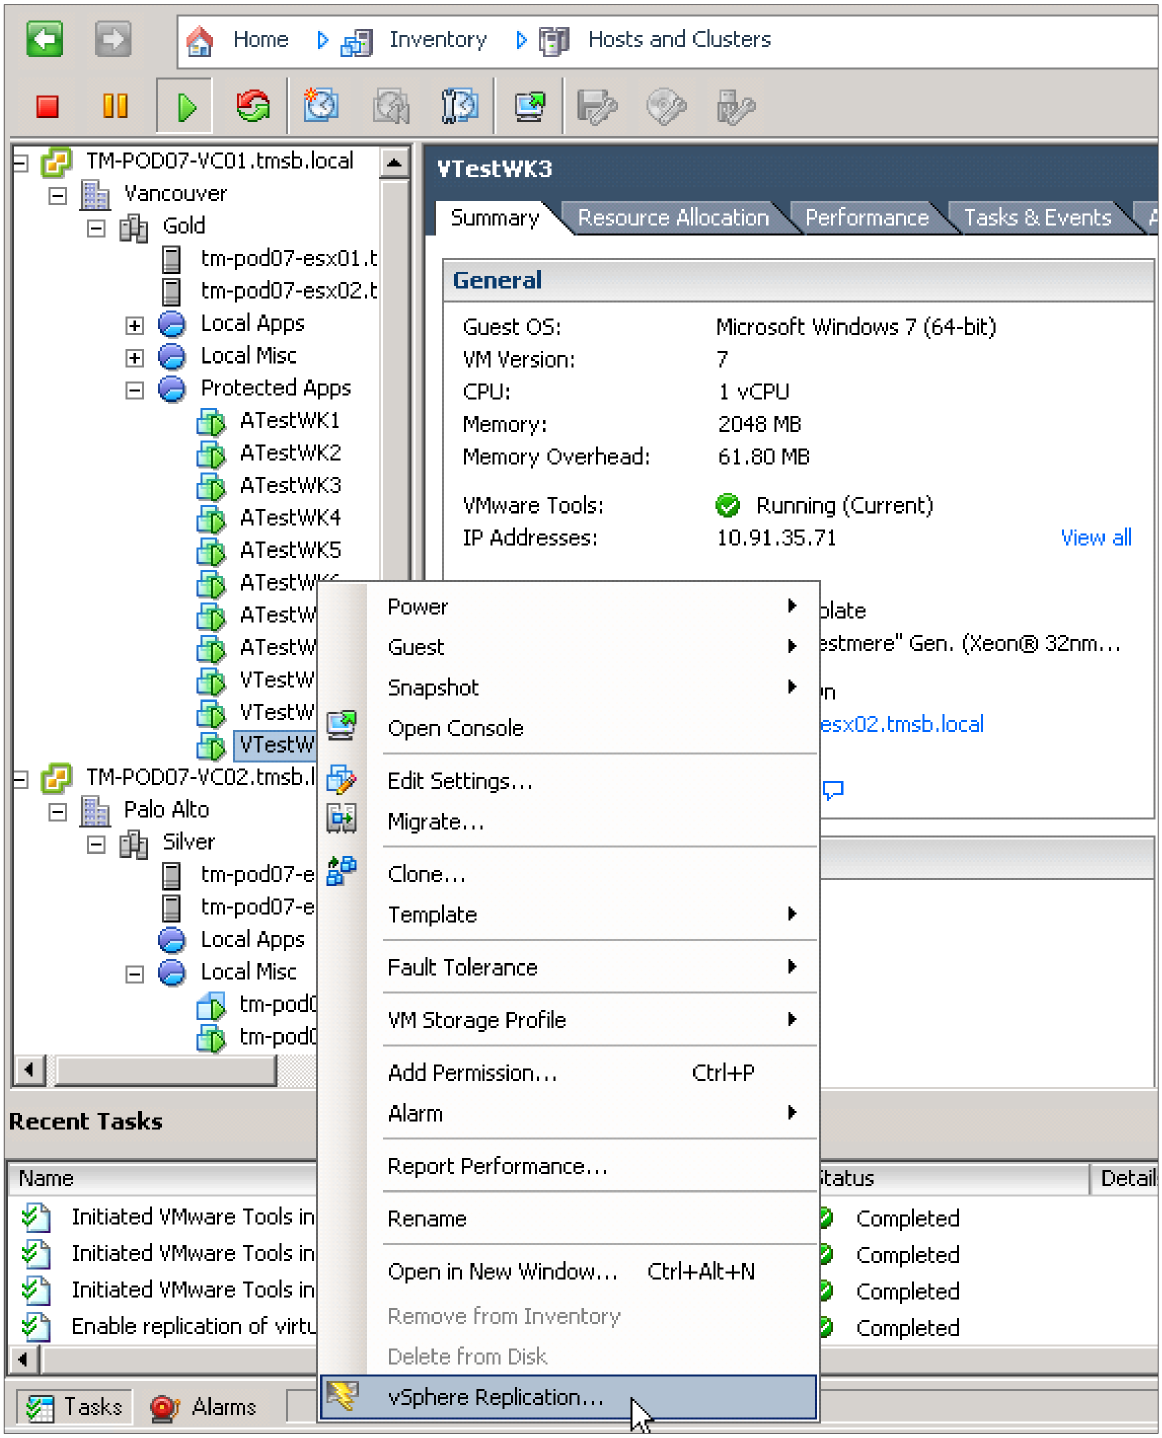

General

Guest OS: Microsoft Windows 7 (64-bit)
VM Version: 7
CPU: 1 vCPU
Memory: 2048 MB
Memory Overhead: 61.80 MB

VMware Tools: Running (Current)
IP Addresses: 10.91.35.71 View all

plate
estmere" Gen. (Xeon® 32nm...
n
esx02.tmsb.local

Recent Tasks

Name	tatus	Detail
Initiated VMware Tools in	Completed	
Initiated VMware Tools in	Completed	
Initiated VMware Tools in	Completed	
Enable replication of virtu	Completed	

Figure 95. Selecting vSphere Replication for an Individual Virtual Machine

The vSphere Replication configuration item is available in many locations as a context-specific option when right-clicking a VM.

5. Click on **vSphere Replication** to bring up the **Configure Replication** menu for the selected virtual machine. You can set the recovery point objective (RPO) for the VM either by using a slider to range between 15 minutes and 24 hours, or by selecting the RPO by choosing drop-down items in the RPO window. Initially, choose a four-hour RPO for this VM.

You might also choose to use VSS quiescing for Microsoft Windows virtual machines to assist with application, OS, or file system quiescing. This feature will not be explored in this evaluation guide.

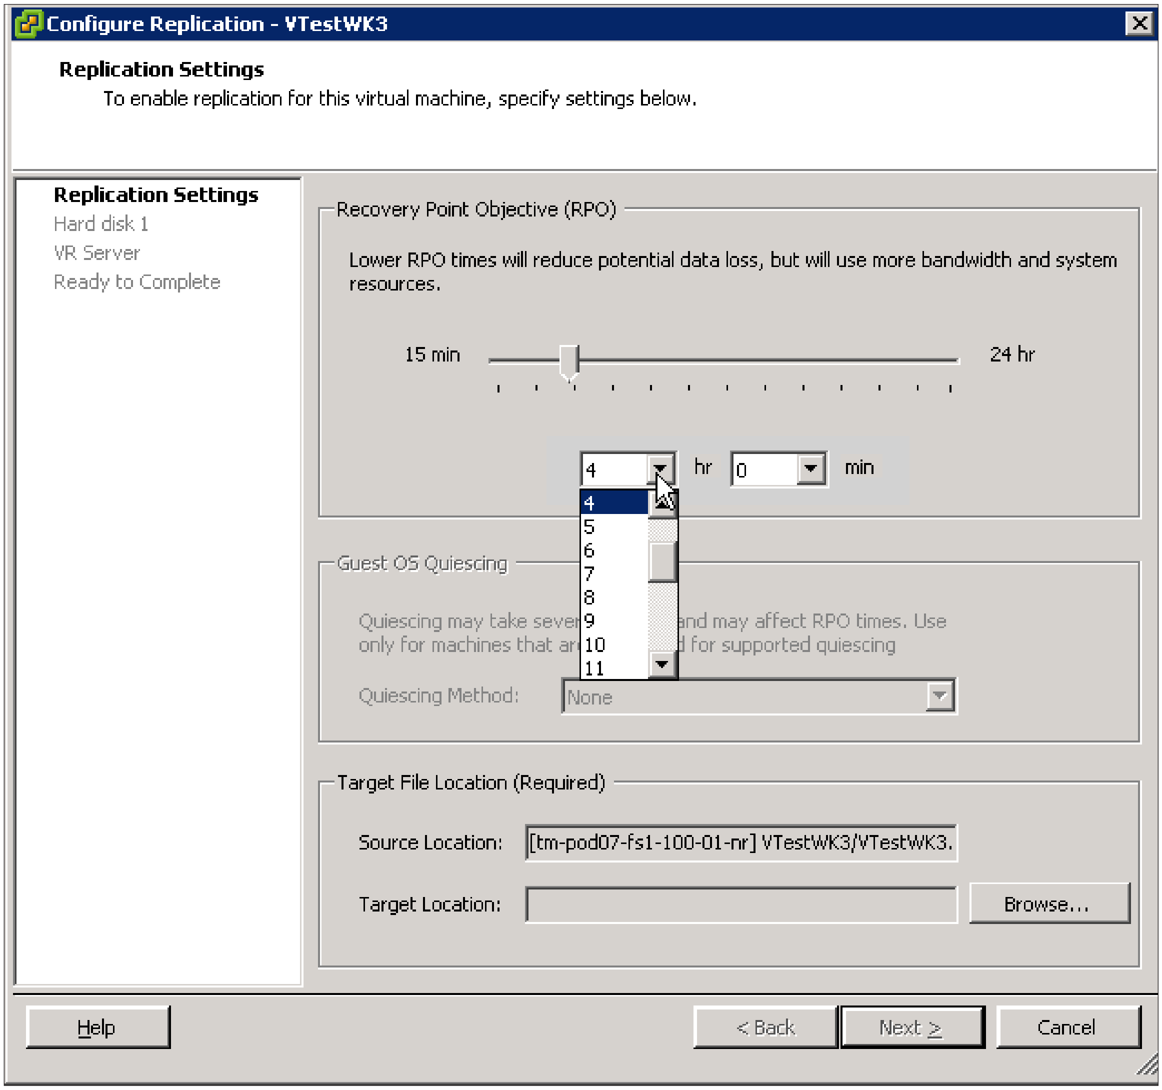

Configure Replication - VTestWK3

Replication Settings

To enable replication for this virtual machine, specify settings below.

Recovery Point Objective (RPO)

Lower RPO times will reduce potential data loss, but will use more bandwidth and system resources.

15 min — 24 hr

4 hr 0 min

Guest OS Quiescing

Quiescing may take sever… and may affect RPO times. Use only for machines that ar… for supported quiescing

Quiescing Method: None

Target File Location (Required)

Source Location: [tm-pod07-fs1-100-01-nr] VTestWK3/VTestWK3.

Target Location:

Figure 96. Choosing a Four-Hour RPO for the Protected VM

6. In the Target File Location window, click **Browse** to select a destination datastore at the **recovery site.** A pop-up window for the **Target VM Location** will open. Expand the available datastores at the **recovery site** and select a **nonreplicated** destination for the VM. This is the datastore to which the VM will be replicated.

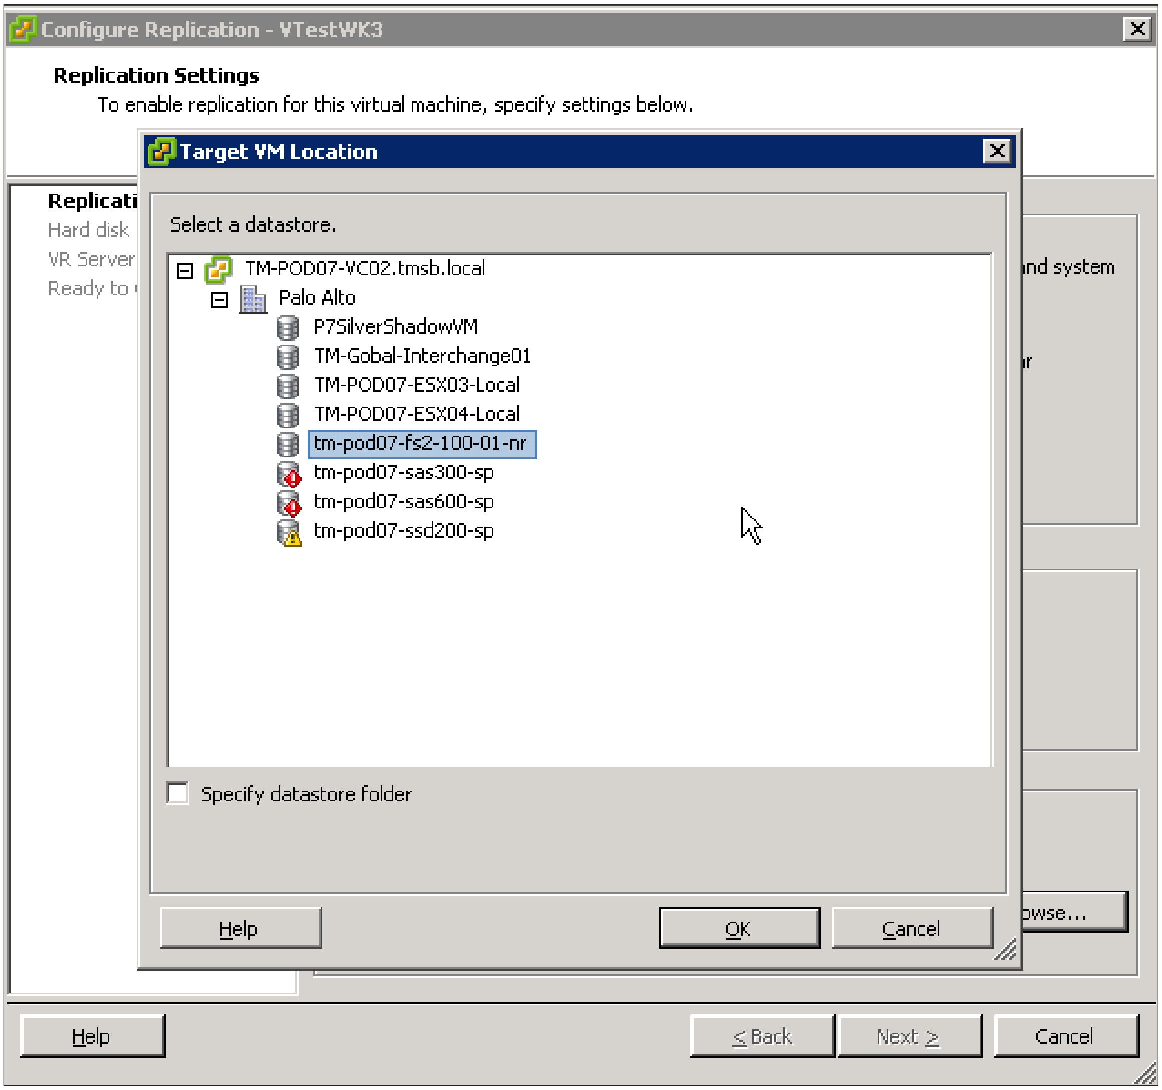

Configure Replication - VTestWK3

Replication Settings

To enable replication for this virtual machine, specify settings below.

Target VM Location

Select a datastore.

- TM-POD07-VC02.tmsb.local
  - Palo Alto
    - P7SilverShadowVM
    - TM-Gobal-Interchange01
    - TM-POD07-ESX03-Local
    - TM-POD07-ESX04-Local
    - tm-pod07-fs2-100-01-nr
    - tm-pod07-sas300-sp
    - tm-pod07-sas600-sp
    - tm-pod07-ssd200-sp

☐ Specify datastore folder

Figure 97. Choosing a Target VM Location at the Recovery Site (Site B)

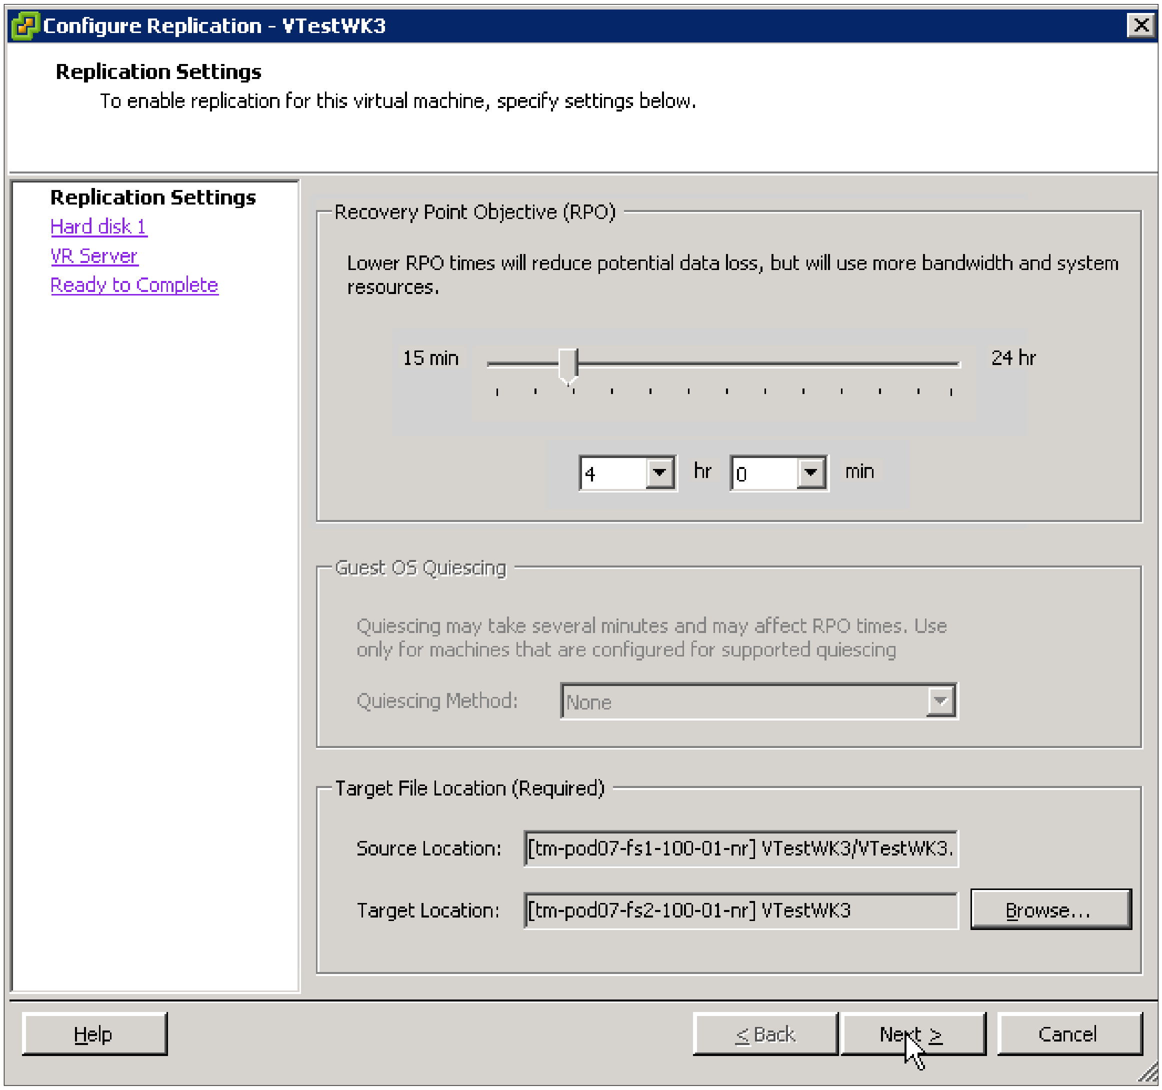

Configure Replication - VTestWK3

Replication Settings

To enable replication for this virtual machine, specify settings below.

Recovery Point Objective (RPO)

Lower RPO times will reduce potential data loss, but will use more bandwidth and system resources.

4 hr 0 min

Guest OS Quiescing

Quiescing may take several minutes and may affect RPO times. Use only for machines that are configured for supported quiescing

Quiescing Method: None

Target File Location (Required)

Source Location: [tm-pod07-fs1-100-01-nr] VTestWK3/VTestWK3.

Target Location: [tm-pod07-fs2-100-01-nr] VTestWK3

Figure 98. A Selected Datastore Has Been Chosen for the VM

7. Click **Next.** Note in the **Hard Disk Options** there is the option to enable or disable replication. With this mechanism, you may configure replication for a virtual machine but choose not to turn it on until, for example, a change window allows it. If the **Target Disk File Location** and **Target Disk Type** selections are correct, click **Next.** If there is more than one Virtual Machine Disk (VMDK) associated with this VM, there will be more screens similar to this one to allow you to make changes to the destination location, replication enablement, and disk type for each unique VMDK.

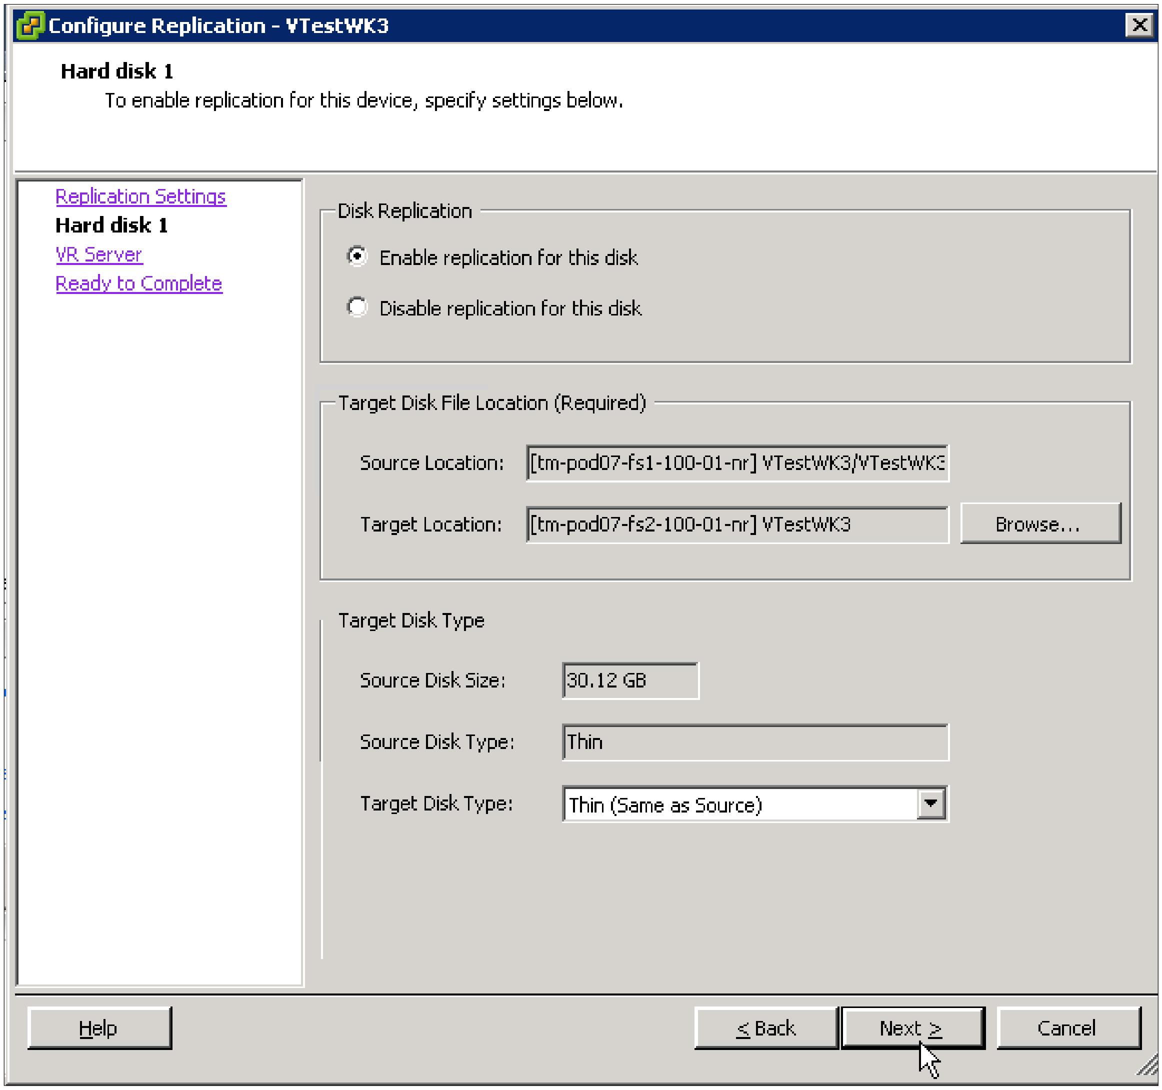

Configure Replication - VTestWK3

Hard disk 1

To enable replication for this device, specify settings below.

Disk Replication

● Enable replication for this disk

○ Disable replication for this disk

Target Disk File Location (Required)

Source Location: [tm-pod07-fs1-100-01-nr] VTestWK3/VTestWK3

Target Location: [tm-pod07-fs2-100-01-nr] VTestWK3

Target Disk Type

Source Disk Size: 30.12 GB

Source Disk Type: Thin

Target Disk Type: Thin (Same as Source)

Figure 99. Hard Disk Configuration for a VMDK in a VM

8. Select the **VR Server** deployed and registered earlier as the target for VR copies.

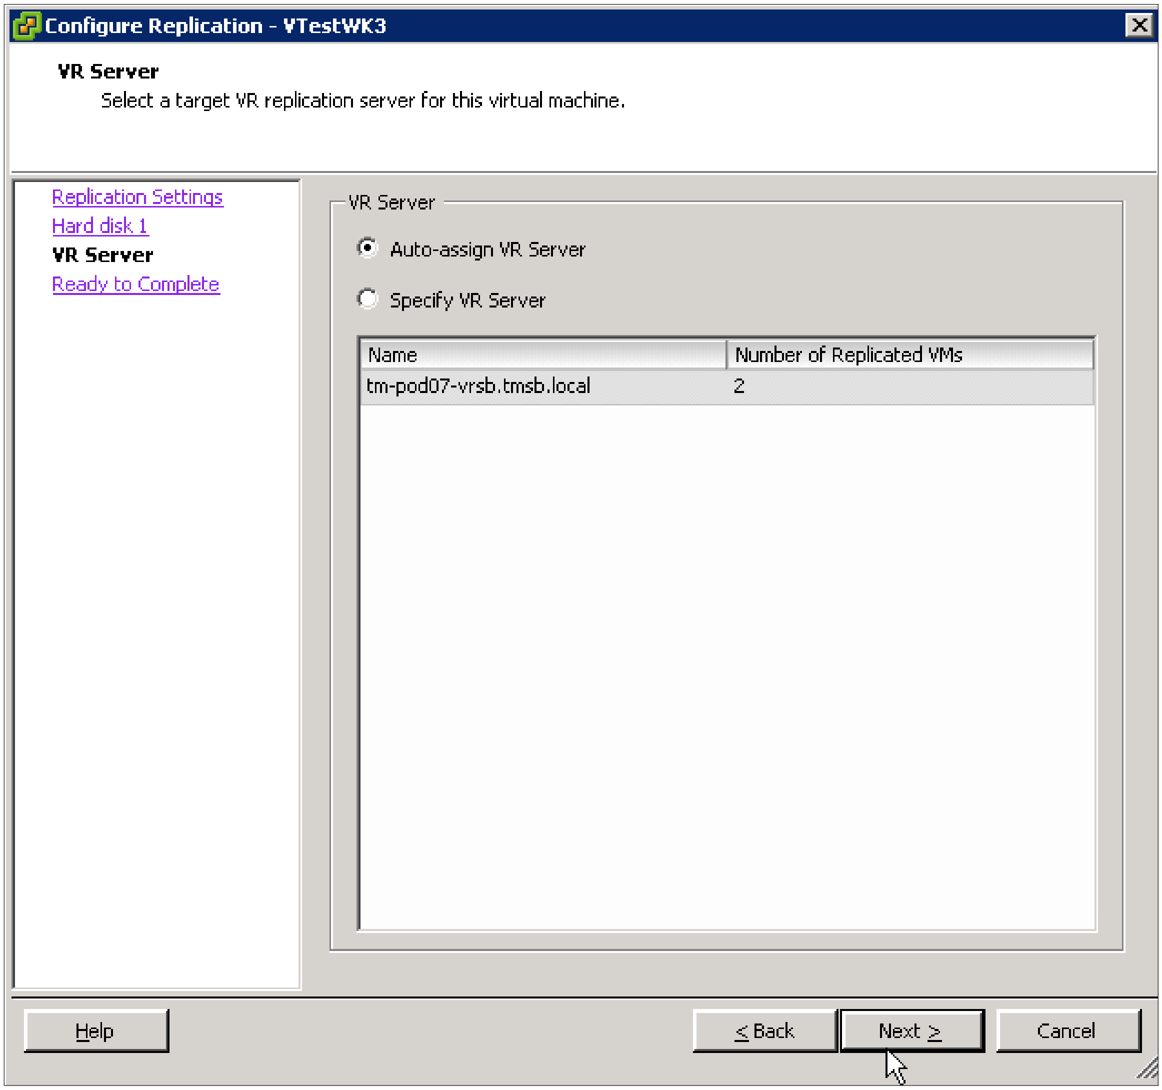

Configure Replication - VTestWK3

VR Server

Select a target VR replication server for this virtual machine.

● Auto-assign VR Server

○ Specify VR Server

Name	Number of Replicated VMs
tm-pod07-vrsb.tmsb.local	2

Figure 100. Selecting the VR Server

9. Review the options selected, ensuring that both the target destination and the target VR Server are at the recovery site (Site B). Click **Finish** to configure replication for this VM.

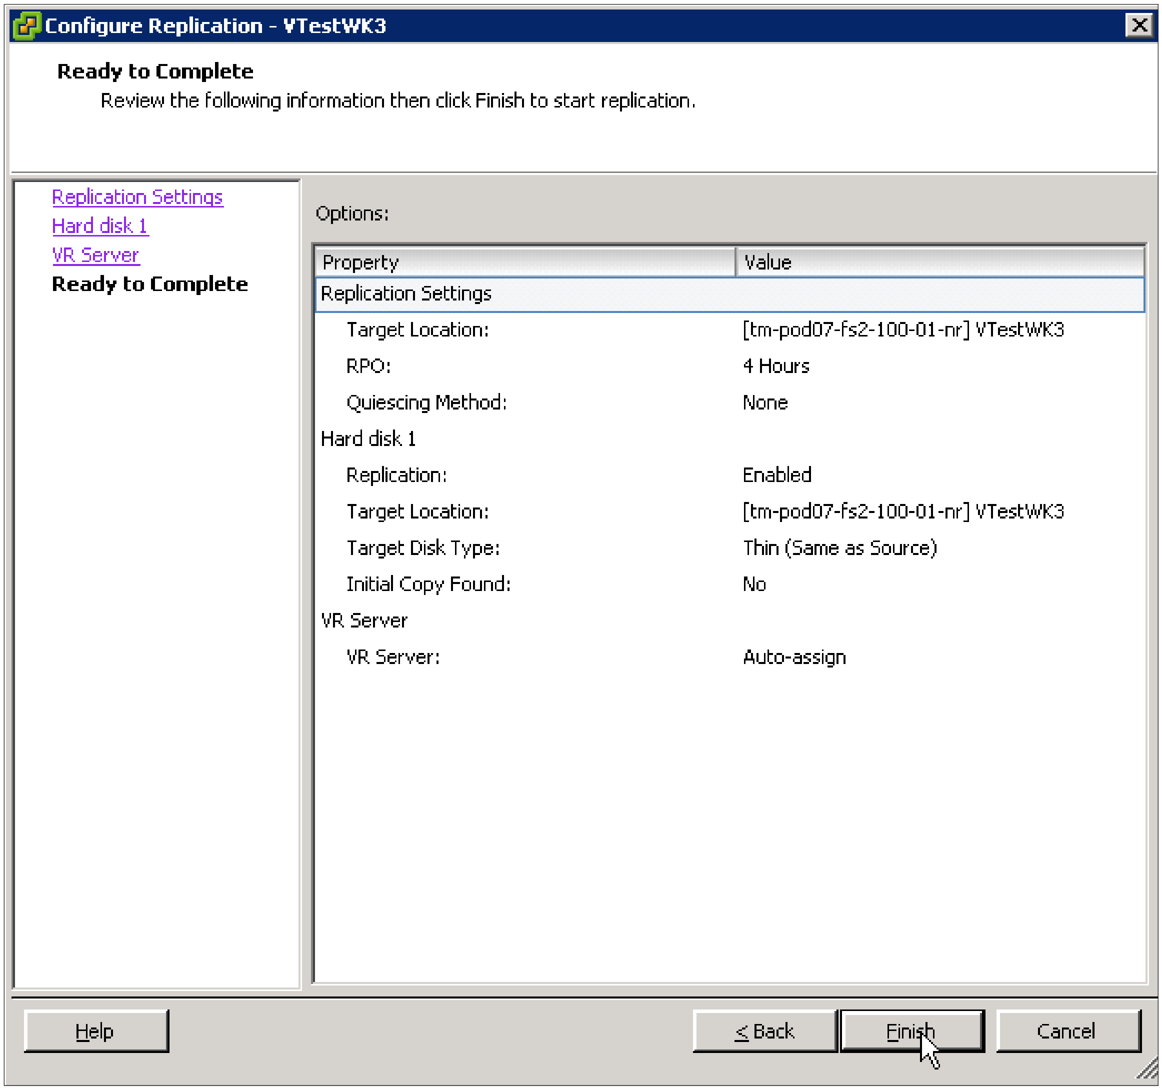

Ready to Complete

Review the following information then click Finish to start replication.

Options:

Property	Value
Replication Settings	
Target Location:	[tm-pod07-fs2-100-01-nr] VTestWK3
RPO:	4 Hours
Quiescing Method:	None
Hard disk 1	
Replication:	Enabled
Target Location:	[tm-pod07-fs2-100-01-nr] VTestWK3
Target Disk Type:	Thin (Same as Source)
Initial Copy Found:	No
VR Server	
VR Server:	Auto-assign

Figure 101. Final Review of VM Replication

Replication will now be configured for the VM and if it is successful the **Configuring Replication** pop-up screen will indicate success.

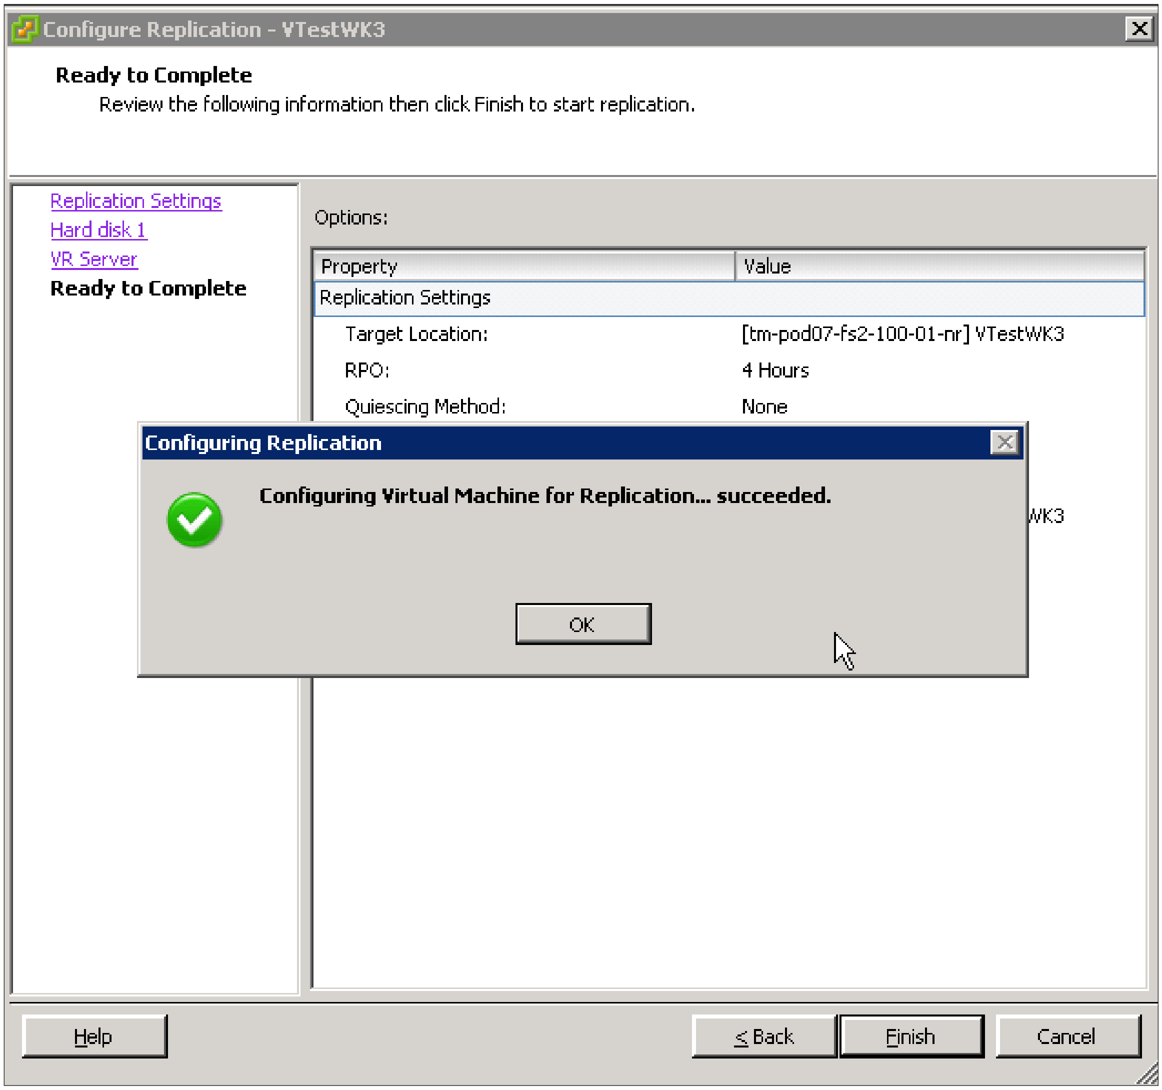

Configure Replication - VTestWK3

Ready to Complete

Review the following information then click Finish to start replication.

Options:

Property	Value
Replication Settings	
Target Location:	[tm-pod07-fs2-100-01-nr] VTestWK3
RPO:	4 Hours
Quiescing Method:	None

Configuring Replication

Configuring Virtual Machine for Replication... succeeded.

Figure 102. Successful Configuration of vSphere Replication for a VM

10. To ensure that replication has begun for the selected VM, return to the vSphere Client home page and click the **Site Recovery** icon. Choose the **vSphere Replication** navigation line on the left panel, expand the folder for **Site B** and click the registered VRS. Click the **Virtual Machines** tab in the main panel and it will show the current status of replication for the VM configured for replication in the previous steps.

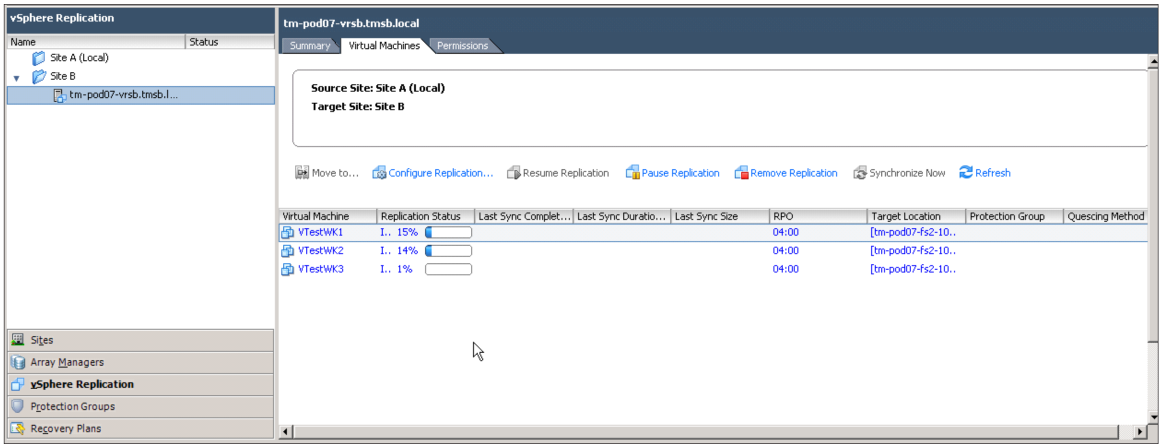

Source Site: Site A (Local)
Target Site: Site B

Virtual Machine	Replication Status	Last Sync Complet...	Last Sync Duratio...	Last Sync Size	RPO	Target Location	Protection Group	Quescing Method
VTestWK1	I.. 15%				04:00	[tm-pod07-fs2-10..		
VTestWK2	I.. 14%				04:00	[tm-pod07-fs2-10..		
VTestWK3	I.. 1%				04:00	[tm-pod07-fs2-10..		

Figure 103. Replication Status for VR-Protected VMs

It is recommended that you do not continue to the next exercise until replication of the protected VM is complete. This may take minutes, or hours, depending on the size of the VM and the network speed between sites. You may continue to create a protection group and recovery plan without errors, but if reconfiguration of vSphere Replication is necessary, it is a good idea to track the replication of the VM before continuing.

You may also repeat this process for more virtual machines, but for the purpose of this evaluation guide, we will assume no more than three VMs have been protected by vSphere Replication.

Step 9: Creating a protection group for VR-protected VMs

After configuration of the protection of a virtual machine is complete, the next step is to create a unique **protection group** for vSphere Replication–protected VMs.

To create a protection group for VR-protected VMs, follow these steps:

1. Return to the Home page of the vSphere Client and select **Site Recovery** from the Solutions and Applications menu.
2. Choose the **Protection Groups** line from the left navigation panel, and select the All Protection Groups item.
3. On the far-right Commands window, click the **Create Protection Group** command.

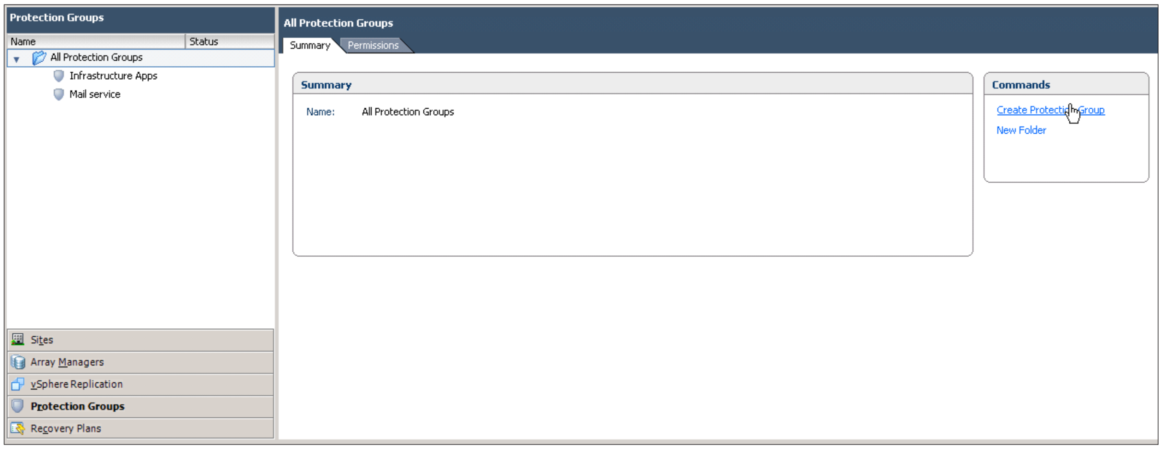

Name: All Protection Groups

Figure 104. Create a Protection Group

4. Select **Site A** for the Protected Site.

5. Select **vSphere Replication** for the Protection Group Type.

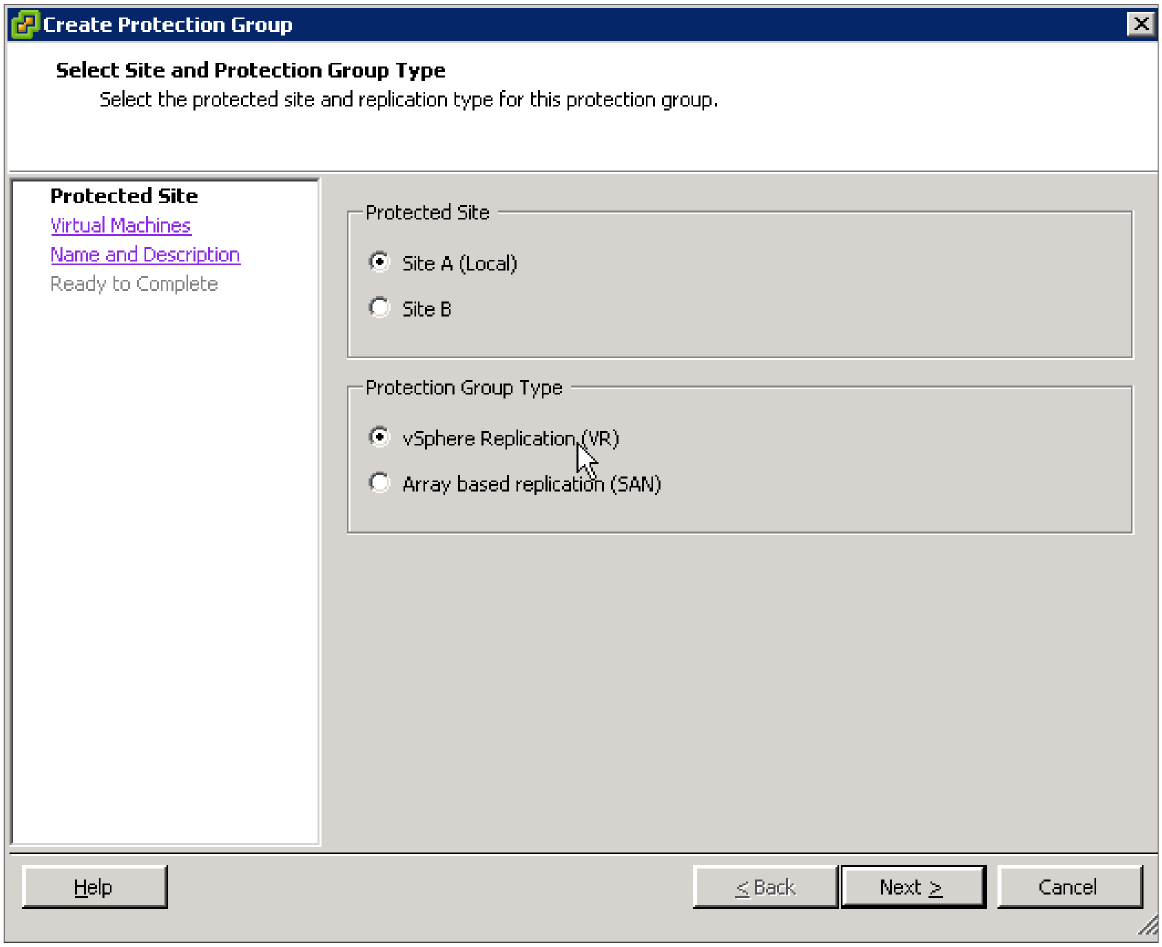

Create Protection Group

Select Site and Protection Group Type

Select the protected site and replication type for this protection group.

Protected Site

● Site A (Local)

○ Site B

Protection Group Type

● vSphere Replication (VR)

○ Array based replication (SAN)

Figure 105. Creating a vSphere Replication Protection Group

6. Under Replicated Virtual Machines, choose one or more VMs that were chosen for vSphere Replication in the previous exercise. If no VMs are visible, then vSphere Replication was not configured correctly. If this is the case, return to the previous exercise and ensure that VMs were configured correctly. If VMs are available, choose as many VR-protected VMs as you wish for this protection group. Also note the status of the VMs being chosen. They may still be completing an initial full synchronization, or they may have completed synchronization if you chose to wait for successful replication in the previous exercise. You may choose to set up multiple different protection groups for VMs, for example, if they serve different business or service requirements and will be part of different recovery plans. Click **Next.**

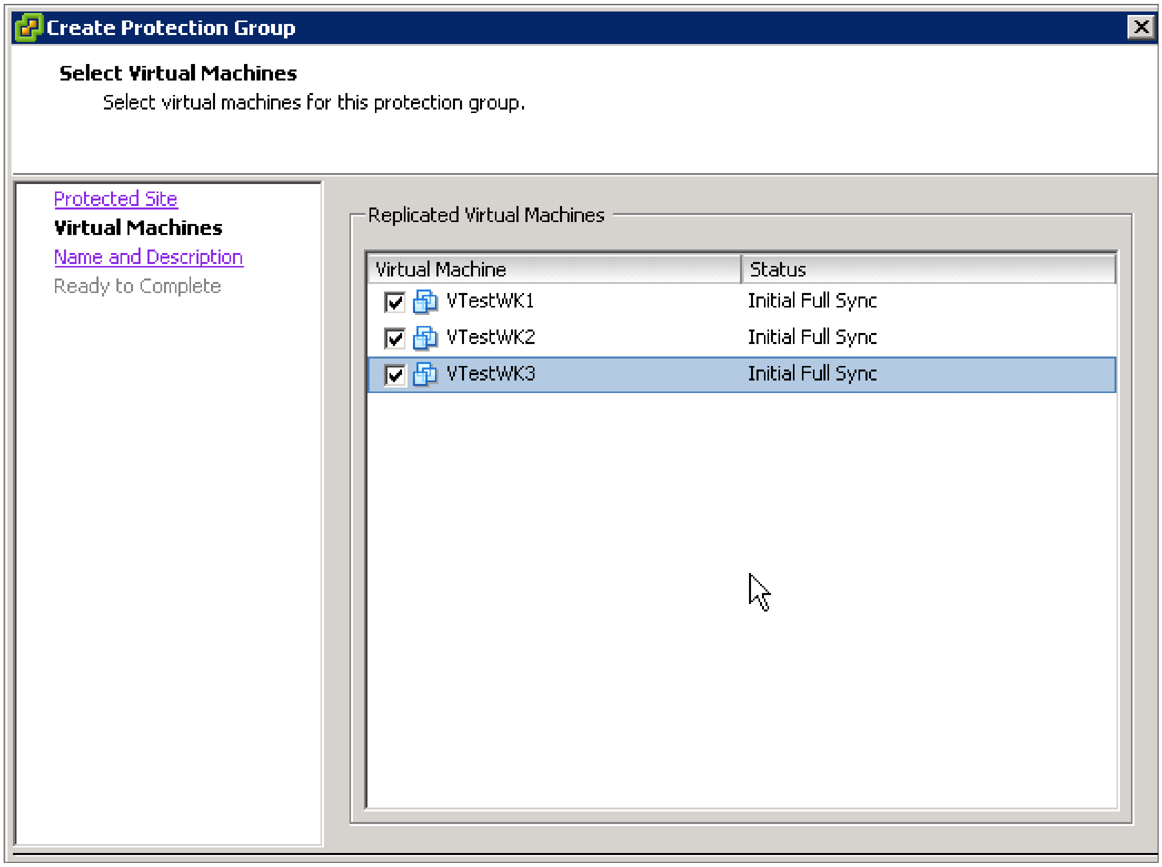

Create Protection Group

Select Virtual Machines

Select virtual machines for this protection group.

Replicated Virtual Machines

Virtual Machine	Status
☑ VTestWK1	Initial Full Sync
☑ VTestWK2	Initial Full Sync
☑ VTestWK3	Initial Full Sync

Figure 106. Selecting VR-Protected VMs for a Protection Group

7. Provide a meaningful name and description for the collection of VMs selected for this protection group.

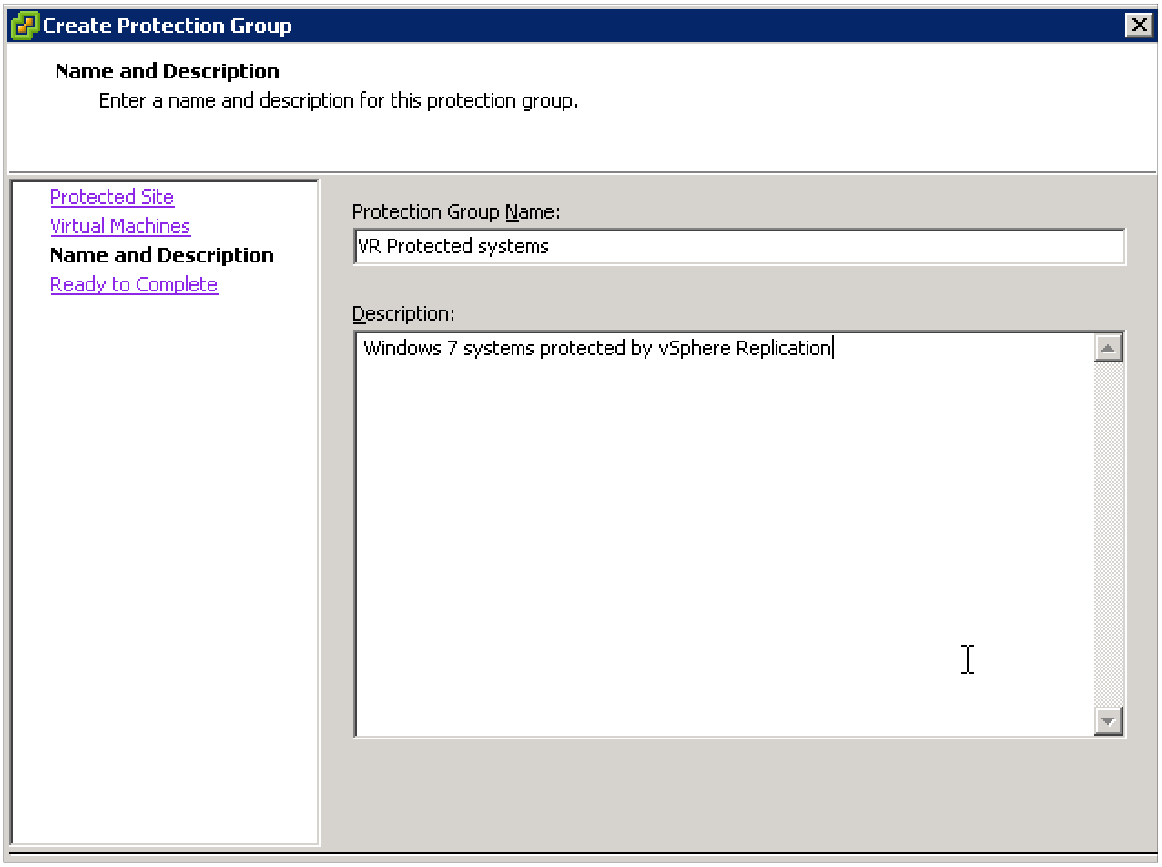

Name and Description

Enter a name and description for this protection group.

Protection Group Name: VR Protected systems

Description: Windows 7 systems protected by vSphere Replication

Figure 107. Naming the Protection Group

8. Review the options and click **Finish** to complete the creation of the vSphere Replication protection group. This will return you to the **Protection Groups** menu in the SRM plug-in of the vSphere Client. You should see a new protection group populated in the left pane. Select this protection group and click the **Virtual Machines** tab in the main screen to see more detail about the VMs in this protection group. Take note of the **Protection Status** before continuing.

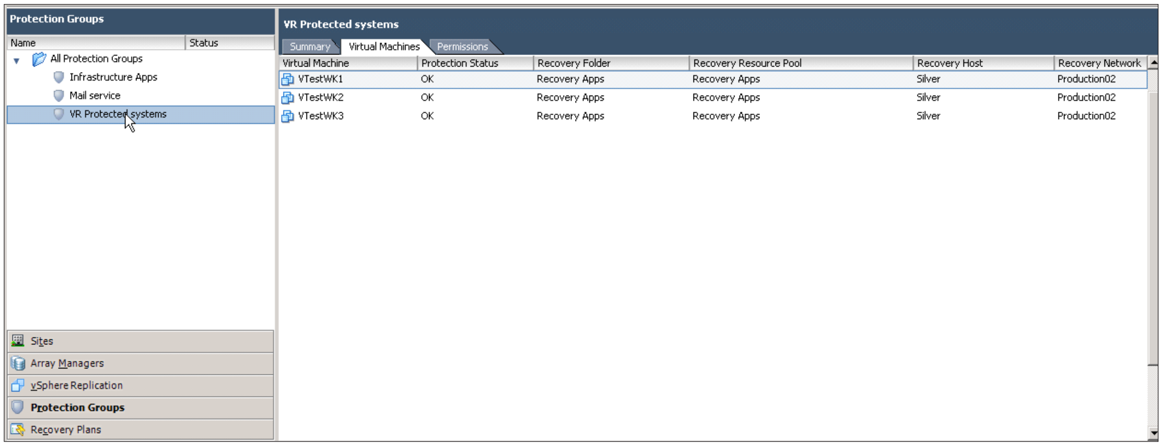

Virtual Machine	Protection Status	Recovery Folder	Recovery Resource Pool	Recovery Host	Recovery Network
VTestWK1	OK	Recovery Apps	Recovery Apps	Silver	Production02
VTestWK2	OK	Recovery Apps	Recovery Apps	Silver	Production02
VTestWK3	OK	Recovery Apps	Recovery Apps	Silver	Production02

Figure 108. VR-Protected Systems in the Newly Created Protection Group

Step 10: Creating a recovery plan for VR-based protection groups

After creating the protection group, you can now create a VR-specific recovery plan. Although you are able to add vSphere Replication–based protection groups to existing recovery plans, or to add these PGs to recovery plans that also use array-based protection groups, this is **not recommended.** Protection and failover of the VMs will work correctly, but reprotection and automated failback of VR-protected VMs within these scenarios will not work. This will lead to errors when running the recovery plan for failback. For the purposes of this evaluation guide, we will create a separate recovery plan for vSphere Replication–based protection groups.

To create a recovery plan, follow the steps listed earlier in this guide, and select the newly created protection group.

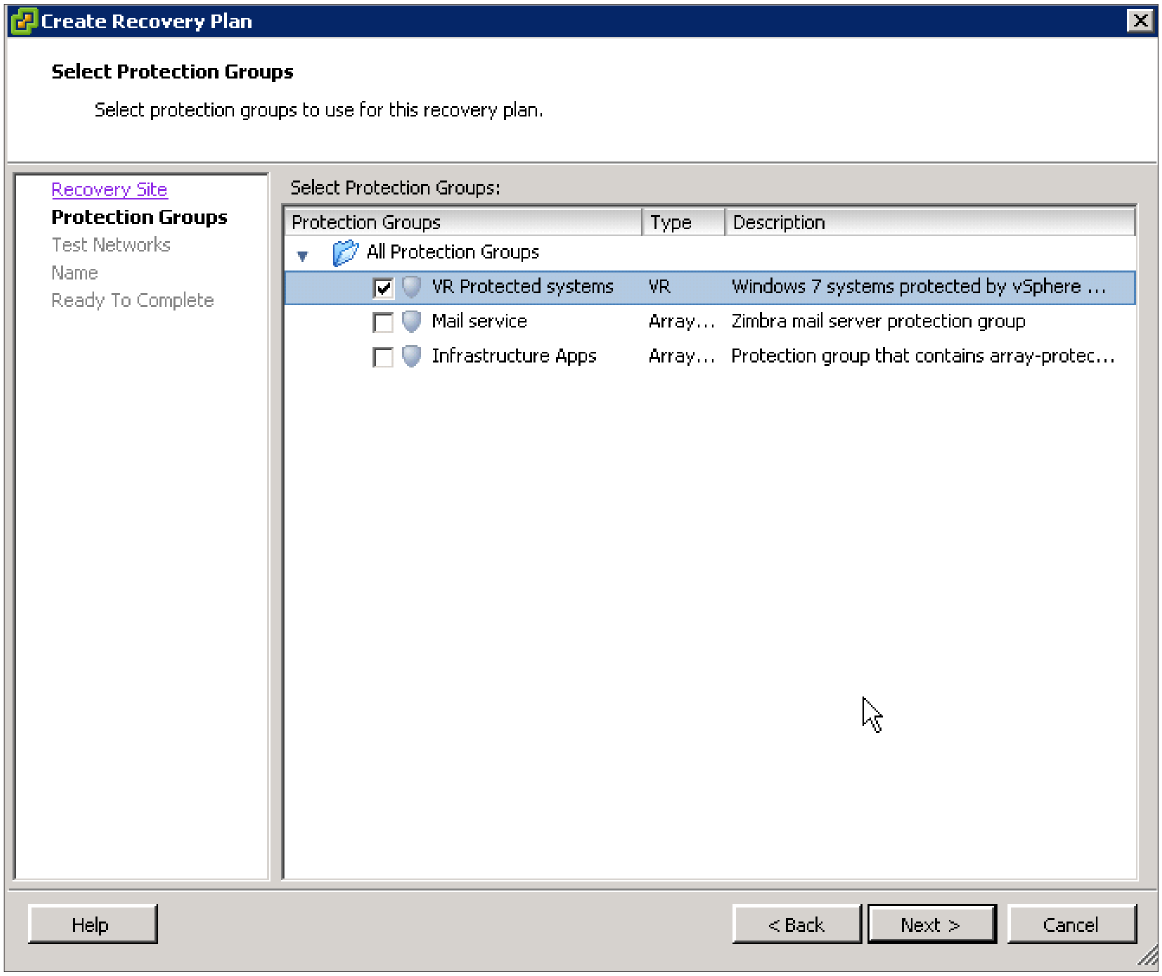

Select Protection Groups

Select protection groups to use for this recovery plan.

Select Protection Groups:

Protection Groups	Type	Description
All Protection Groups		
☑ VR Protected systems	VR	Windows 7 systems protected by vSphere ...
☐ Mail service	Array...	Zimbra mail server protection group
☐ Infrastructure Apps	Array...	Protection group that contains array-protec...

Figure 109. Create a New Recovery Plan, Selecting the Protection Group Created in This Exercise

This concludes the deployment and configuration of vSphere Replication. At this point, you have accomplished the five following tasks: deployed and configured the management framework for VR; deployed and registered the VRS appliance that receives replication; configured VMs for protection and replication with VR; created a protection group; and created a recovery plan for vSphere Replication–protected virtual machines.

Exercise 3. Configuring Site Recovery Manager Alarms

Awareness of the SRM alarms is an important part of understanding how SRM works across the protected and recovery sites. During the SRM product evaluation, it is recommended that, wherever possible and without impact to your production environment, you create failures or conditions in the protected and recovery site that will result in the generation of SRM alarms. The generation of these SRM alarms will serve as validation that SRM is monitoring both the protected and recovery sites correctly.

Each SRM server monitors the CPU utilization, disk space, and memory consumption of the guest on which it is running, and also maintains a heartbeat with its peer SRM server. vCenter events are sent if any of these measures falls outside of the configured bounds.

SRM supports the configuration of event-triggered alarms so that you can associate a notification action with any given SRM alarm event. These alarms are configured via the SRM UI.

SRM supports the following alarm notification actions:

- **Send a notification email** to a specific email address.
- **Send a notification trap** to vCenter trap receivers.
- **Run a script** on the vCenter server.

Refer to the chapter "Customizing Site Recovery Manager" in the *Site Recovery Manager Administration Guide* that details how to set up the preceding alarm actions listed.

Failure of either site generates the following events that can be associated with vCenter alarms:

- Problems with the local site (for example, resource constraints)
- Problems with remote site (for example, inability to ping a remote site that may indicate a disaster)
- Remote site failure, which is reflected in the SRM alarm events and will not automatically trigger a recovery – this must be initiated manually

SRM is configured to raise vCenter events for the following conditions:

- Disk space is low.
- CPU use exceeds limit.
- Memory is low.

As a starting point during the SRM evaluation, it is recommended that you complete the action setup for the following SRM alarm events listed for the protected and recovery sites. You should be able to trigger these events in your environment without impacting your production environment. The goal is that you can see firsthand how SRM responds and notifies you when you are subjected to one of the failure events listed.

Recommended alarms are as follows:

- VM Discovered – A virtual machine has been discovered on replicated storage.
- VM Not Protected – One or more devices that contain virtual machines must be configured for protection within SRM.
- Recovery Plan Prompt Display – A prompt that requires an answer has been displayed during the execution of a recovery plan.
- Remote Site Down – The remote site has stopped responding.
- Recovery Plan Destroyed – A recovery plan has been deleted.
- Recovery Plan Started/Recovery Plan Execute Test Begin – A notification of the start of a recovery plan or test of a recovery plan has been sent.

As you become more familiar with SRM and its associated workflows that allow you to **Test** your recovery plans as well as **Run** your recovery plan, which results in the failover of services from your protected site to your recovery site, it is recommended that you work through the list of **SRM alarm events.** These can be accessed via the **Alarms** tab, as depicted in Figure 110, and can enable the appropriate notification **Actions** for any additional SRM alarm event that you deem important for your environment.

Configure Site Recovery Manager Alarms

SRM Alarms	Configure action for an SRM alarm	Configure action for Remote Site Down alarm 1. Configure alarm action to send out notification email.	10 minutes

Step 1: Configure alarm action to send out notification email

Procedure

1. Open the vSphere Client and connect to the vCenter server at the recovery site. Log in as a vSphere administrator.
2. Click the **Site Recovery** icon on the vSphere Client Home page.
3. In the main window, click the **Alarms** tab to display the list of SRM alarms.

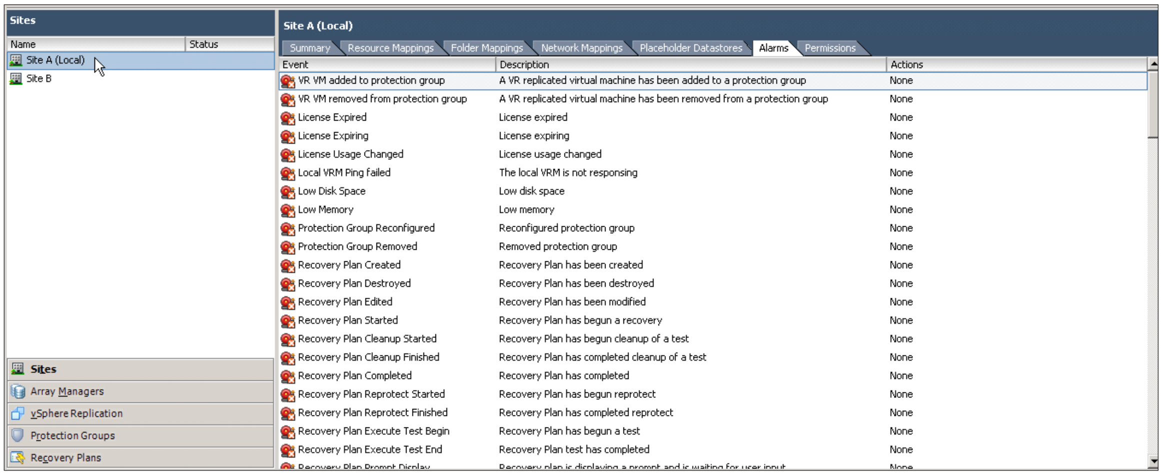

Event	Description	Actions
VR VM added to protection group	A VR replicated virtual machine has been added to a protection group	None
VR VM removed from protection group	A VR replicated virtual machine has been removed from a protection group	None
License Expired	License expired	None
License Expiring	License expiring	None
License Usage Changed	License usage changed	None
Local VRM Ping failed	The local VRM is not responsing	None
Low Disk Space	Low disk space	None
Low Memory	Low memory	None
Protection Group Reconfigured	Reconfigured protection group	None
Protection Group Removed	Removed protection group	None
Recovery Plan Created	Recovery Plan has been created	None
Recovery Plan Destroyed	Recovery Plan has been destroyed	None
Recovery Plan Edited	Recovery Plan has been modified	None
Recovery Plan Started	Recovery Plan has begun a recovery	None
Recovery Plan Cleanup Started	Recovery Plan has begun cleanup of a test	None
Recovery Plan Cleanup Finished	Recovery Plan has completed cleanup of a test	None
Recovery Plan Completed	Recovery Plan has completed	None
Recovery Plan Reprotect Started	Recovery Plan has begun reprotect	None
Recovery Plan Reprotect Finished	Recovery Plan has completed reprotect	None
Recovery Plan Execute Test Begin	Recovery Plan has begun a test	None
Recovery Plan Execute Test End	Recovery Plan test has completed	None
Recovery Plan Prompt Display	Recovery plan is displaying a prompt and is waiting for user input	None

Figure 110. Site Recovery Manager Alarms Tab

4. Right-click **Remote Site Down** and click **Edit Settings.**

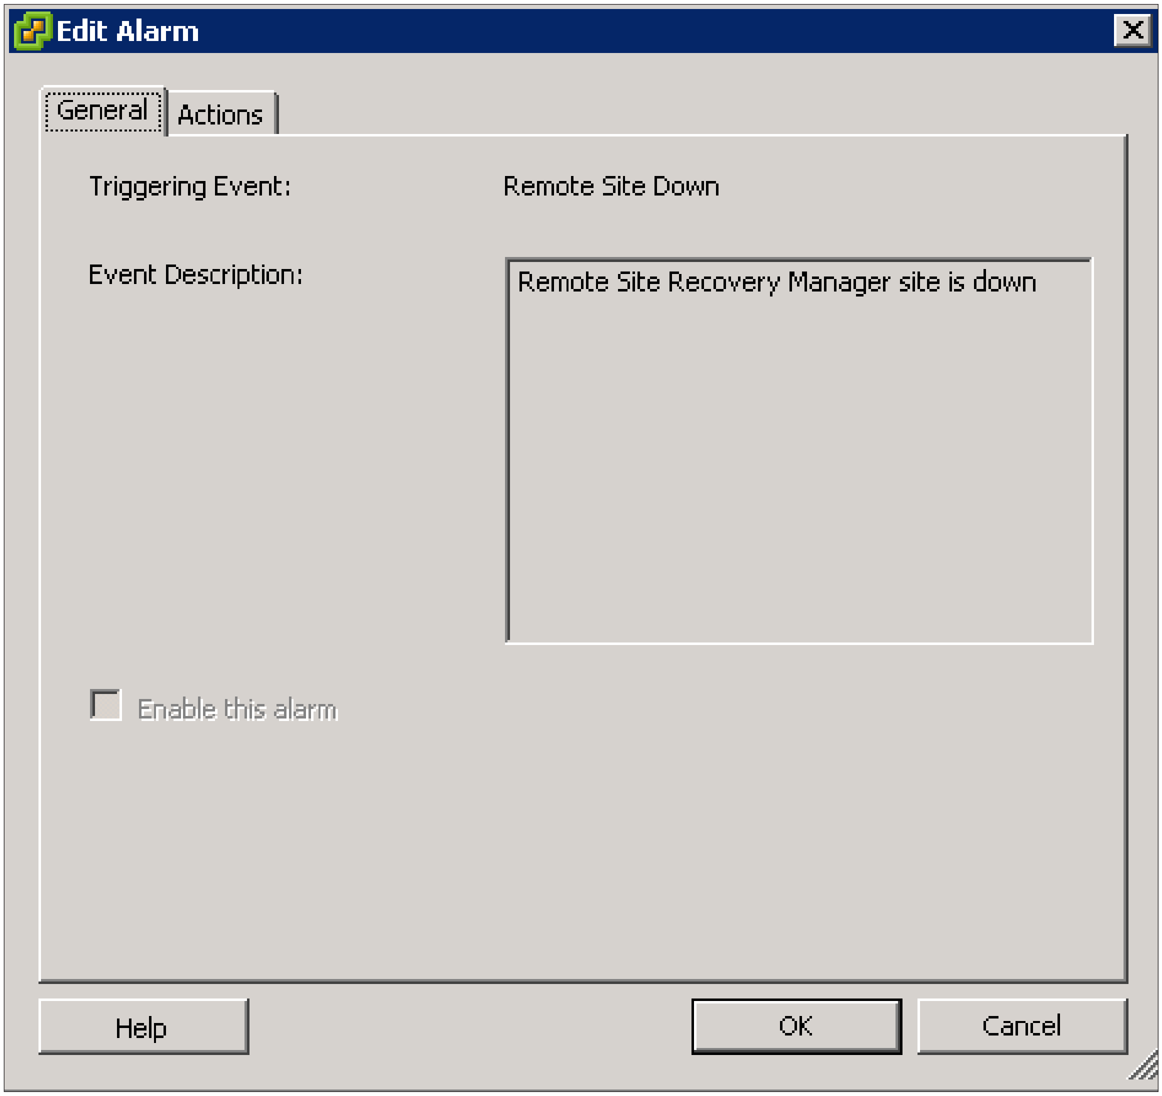

Triggering Event: Remote Site Down

Event Description: Remote Site Recovery Manager site is down

☐ Enable this alarm

Figure 111. Edit Settings for “Remote Site Down” Alarm

5. In the Edit Settings dialog box, click the Actions tab. In the Actions window, click **Add** to add an action.

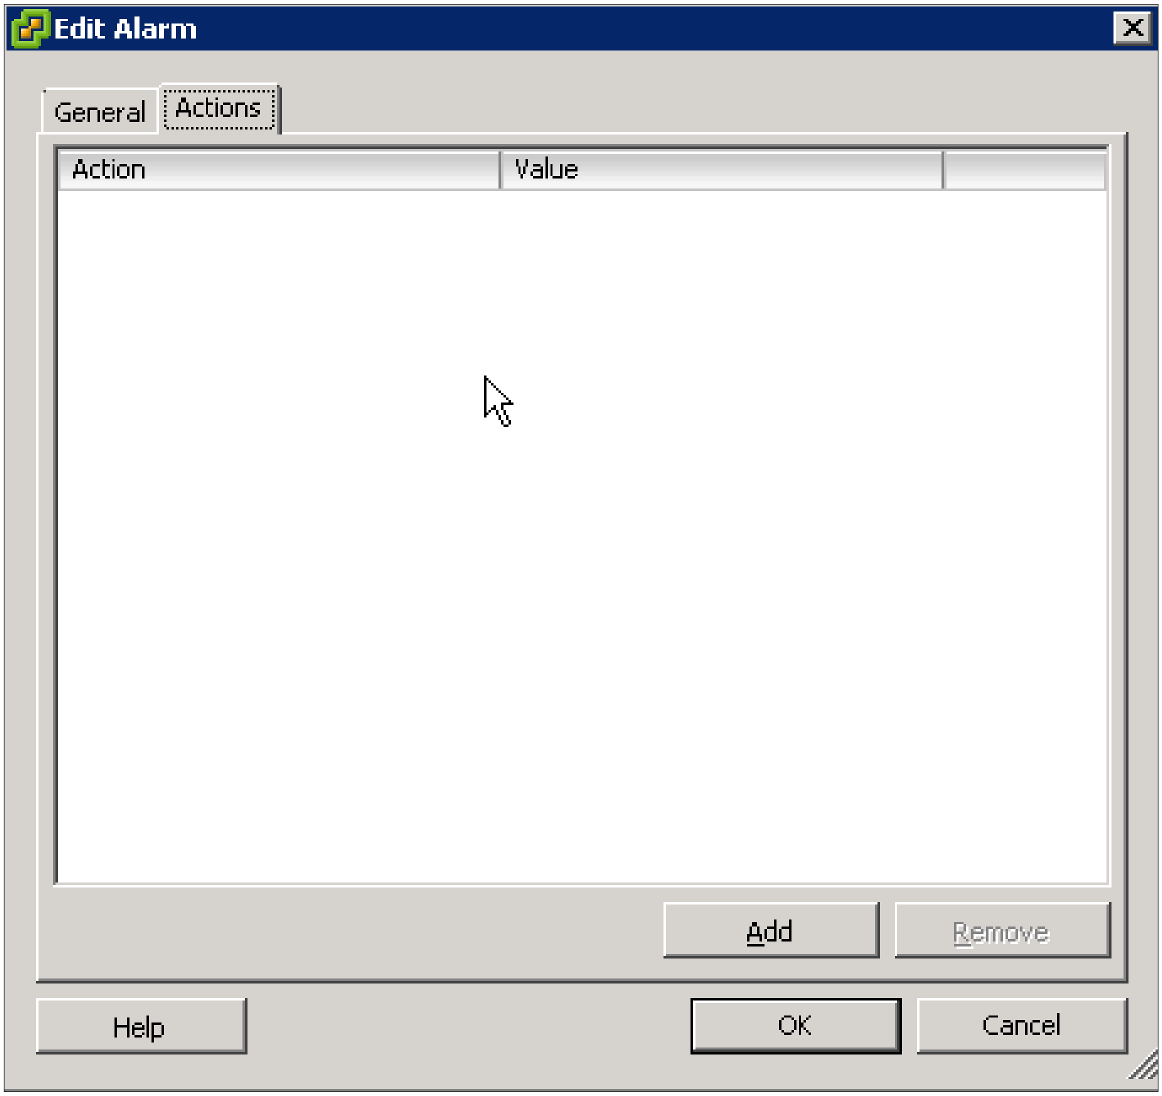

Figure 112. Add Action for "Remote Site Down" Alarm

Use the default action **Send a notification email** and type an email address in the Value column. (To change this action, click it and select a different action from the drop-down box.)

*NOTE: In order for SRM alarm actions to send an SNMP trap or to send an email, the vCenter server must be configured correctly. To configure mail and SNMP settings in vCenter, appropriately configure mail servers and trap destination in **vCenter Server Settings** on the **Home/Administration** screen in your vSphere Client.*

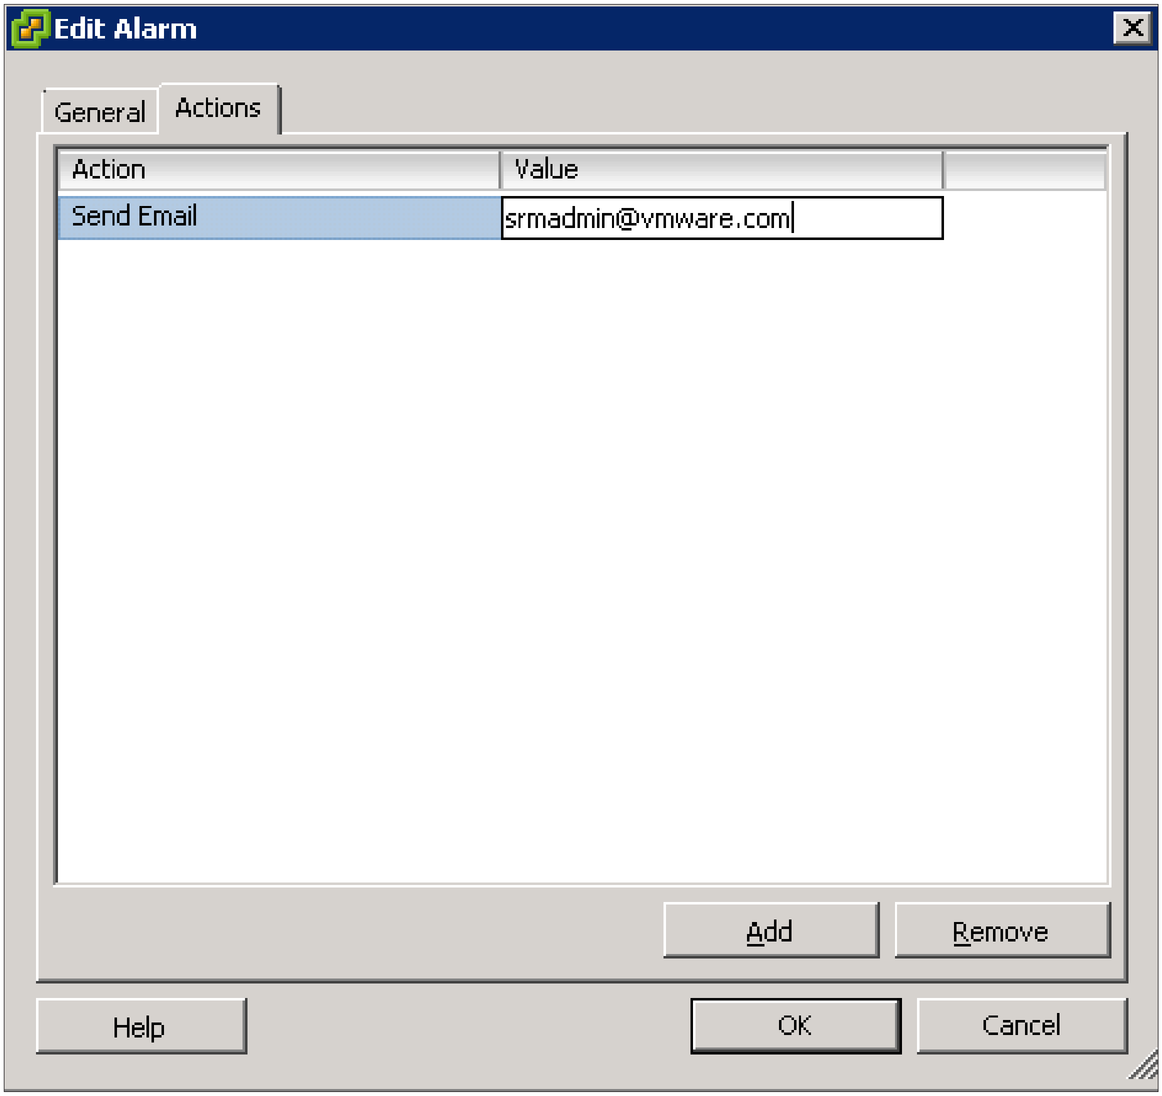

Action	Value
Send Email	srmadmin@vmware.com

Figure 113. Select Action for "Remote Site Down" Alarm

Exercise 4. Running a Recovery Plan

SRM enables you to **Run** a recovery plan that will result in the actual failover of virtual machines from the protected site. Similar to test recovery, failover operations are triggered via a button in the SRM UI on the recovery site. The failover process via SRM is rapid, repeatable, reliable, manageable, and auditable.

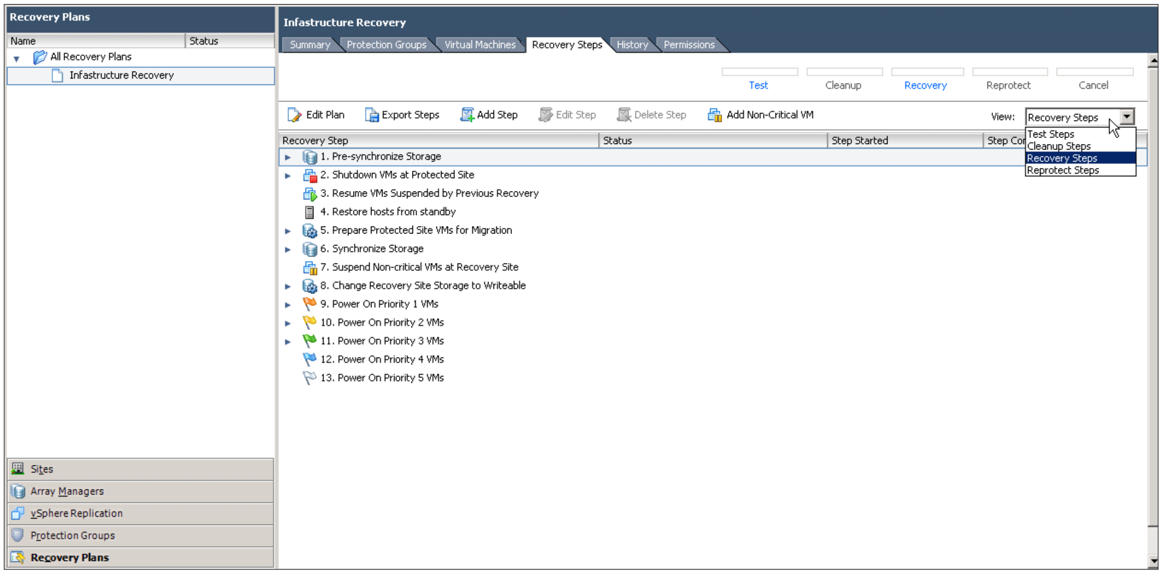

Recovery Step	Status	Step Started	Step Co
1. Pre-synchronize Storage			
2. Shutdown VMs at Protected Site			
3. Resume VMs Suspended by Previous Recovery			
4. Restore hosts from standby			
5. Prepare Protected Site VMs for Migration			
6. Synchronize Storage			
7. Suspend Non-critical VMs at Recovery Site			
8. Change Recovery Site Storage to Writeable			
9. Power On Priority 1 VMs			
10. Power On Priority 2 VMs			
11. Power On Priority 3 VMs			
12. Power On Priority 4 VMs			
13. Power On Priority 5 VMs			

Figure 114. Trigger Failover

This example will show you how to work through an actual failover leveraging the SRM **Run a recovery** plan option.

Step 1: Execute failover

In Figure 114, the Site Recovery Manager UI lists the recovery plan Infrastructure Recovery that was created in Section 1.

1. After selecting **Recovery Steps** in the **View** drop-down box (instead of Test Steps or other options), there are two ways to initiate the actual failover. You can either click the red **Recovery** button with the white arrow on the menu bar at the top of the pane or click the blue **Recovery** text shown as an optional command above the recovery plan steps.

   The **Recovery** dialog box represented by Figure 115 warns you that you are about to run the recovery plan, which will result in changes to the protected virtual machines and the infrastructure of both the protected and recovery site datacenters.

2. Click the **selector check box** to confirm that you understand the implications of running your recovery plan. You might choose to run the recovery plan as either a **planned migration** (which will halt in case of errors) or as a **disaster recovery,** which will not stop if errors are encountered.

3. For the purposes of this guide, select **disaster recovery** and then click **Next** to start the failover of protected virtual machines from the protected site to the recovery site.

   The **Recovery** dialog box also provides a summary of the **Recovery Plan Information.** This includes the recovery plan that is going to be run, the names of the protected and recovery sites, the number of protected virtual machines that will be failed over, and a connectivity status from the recovery site back to the protected site.

4. When satisfied that the information is complete, click **Start** to begin execution of the recovery plan.

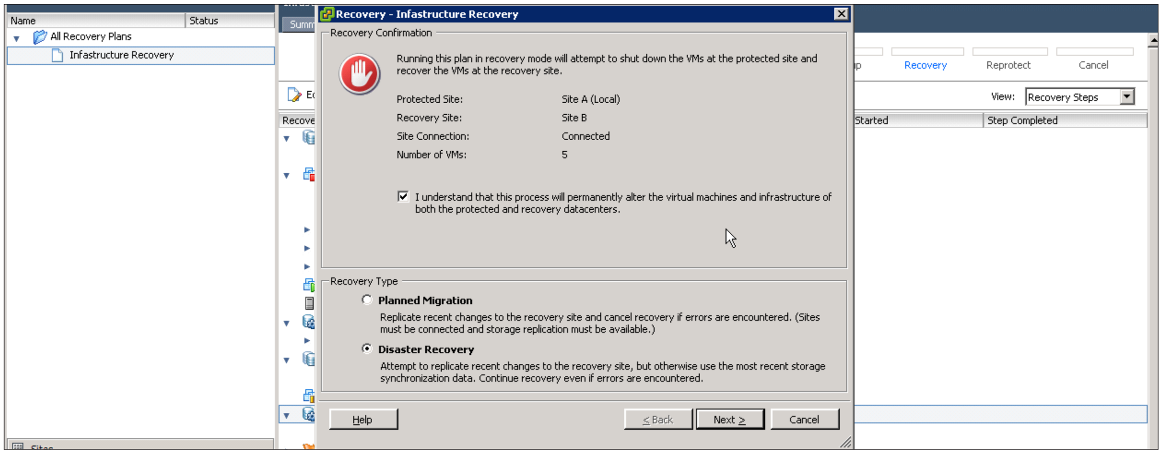

Recovery - Infrastructure Recovery

Recovery Confirmation

Running this plan in recovery mode will attempt to shut down the VMs at the protected site and recover the VMs at the recovery site.

Protected Site:	Site A (Local)
Recovery Site:	Site B
Site Connection:	Connected
Number of VMs:	5

☑ I understand that this process will permanently alter the virtual machines and infrastructure of both the protected and recovery datacenters.

Recovery Type

○ **Planned Migration**
Replicate recent changes to the recovery site and cancel recovery if errors are encountered. (Sites must be connected and storage replication must be available.)

● **Disaster Recovery**
Attempt to replicate recent changes to the recovery site, but otherwise use the most recent storage synchronization data. Continue recovery even if errors are encountered.

Figure 115. Options for Running a Recovery Plan

While the failover is being executed, the status of each step that makes up the recovery plan can be monitored by going to the **Recovery Steps** tab of the SRM UI on the recovery site. The UI informs you which steps are currently **Running** as well as which steps were completed. There are some steps in a recovery plan that will only be executed during a simulated test. **Test only** identifies these steps under the **Mode** column. There are also some steps that will only be executed during an actual failover. These steps are identified by **Recovery only** under the **Mode** column.

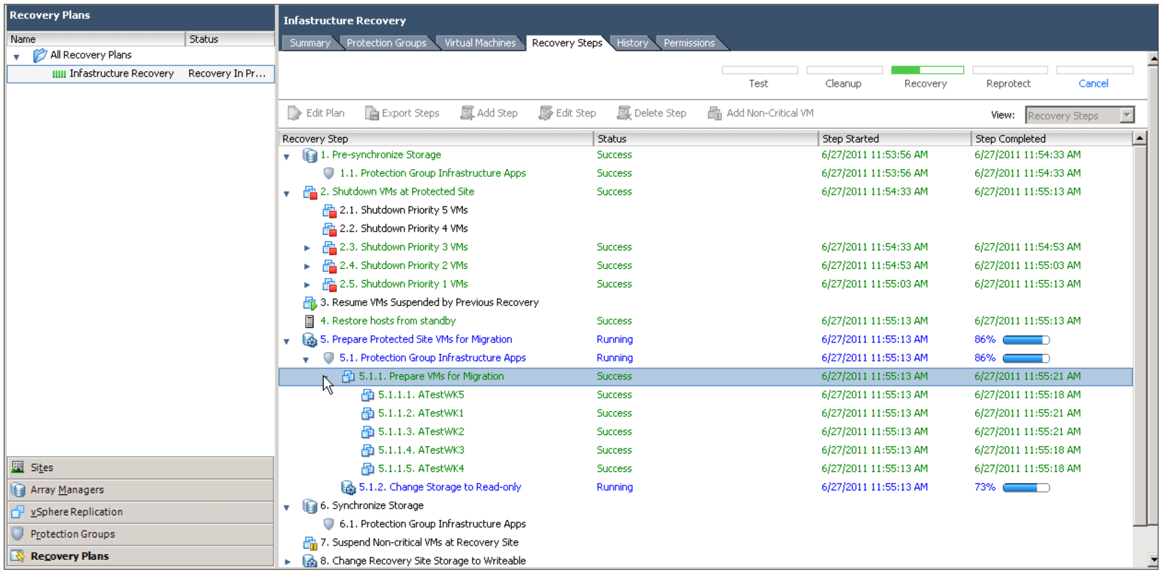

Recovery Step	Status	Step Started	Step Completed
1. Pre-synchronize Storage	Success	6/27/2011 11:53:56 AM	6/27/2011 11:54:33 AM
1.1. Protection Group Infrastructure Apps	Success	6/27/2011 11:53:56 AM	6/27/2011 11:54:33 AM
2. Shutdown VMs at Protected Site	Success	6/27/2011 11:54:33 AM	6/27/2011 11:55:13 AM
2.1. Shutdown Priority 5 VMs			
2.2. Shutdown Priority 4 VMs			
2.3. Shutdown Priority 3 VMs	Success	6/27/2011 11:54:33 AM	6/27/2011 11:54:53 AM
2.4. Shutdown Priority 2 VMs	Success	6/27/2011 11:54:53 AM	6/27/2011 11:55:03 AM
2.5. Shutdown Priority 1 VMs	Success	6/27/2011 11:55:03 AM	6/27/2011 11:55:13 AM
3. Resume VMs Suspended by Previous Recovery			
4. Restore hosts from standby	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:13 AM
5. Prepare Protected Site VMs for Migration	Running	6/27/2011 11:55:13 AM	86%
5.1. Protection Group Infrastructure Apps	Running	6/27/2011 11:55:13 AM	86%
5.1.1. Prepare VMs for Migration	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:21 AM
5.1.1.1. ATestWK5	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:18 AM
5.1.1.2. ATestWK1	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:21 AM
5.1.1.3. ATestWK2	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:21 AM
5.1.1.4. ATestWK3	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:18 AM
5.1.1.5. ATestWK4	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:18 AM
5.1.2. Change Storage to Read-only	Running	6/27/2011 11:55:13 AM	73%
6. Synchronize Storage			
6.1. Protection Group Infrastructure Apps			
7. Suspend Non-critical VMs at Recovery Site			
8. Change Recovery Site Storage to Writeable			

Figure 116. A Running Recovery Plan

While a recovery is running, you can track the status from multiple locations. Figure 116 shows the detailed **Recovery Steps** interface. From the **Summary** screen you can also see the current status of a recovery plan, as well as historical information.

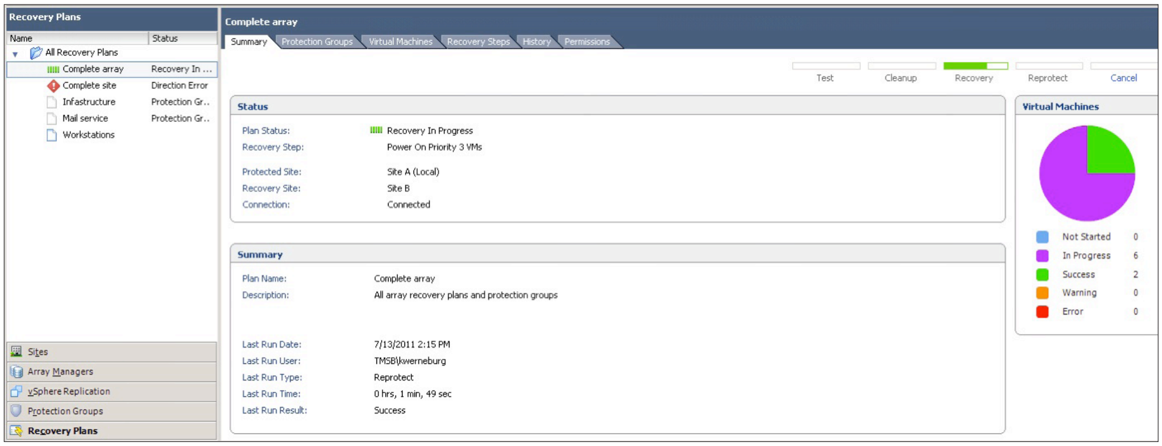

Name	Status
All Recovery Plans	
Complete array	Recovery In ...
Complete site	Direction Error
Infastructure	Protection Gr...
Mail service	Protection Gr...
Workstations	

Status

Plan Status: Recovery In Progress
Recovery Step: Power On Priority 3 VMs
Protected Site: Site A (Local)
Recovery Site: Site B
Connection: Connected

Summary

Plan Name: Complete array
Description: All array recovery plans and protection groups

Last Run Date: 7/13/2011 2:15 PM
Last Run User: TMSB\kwerneburg
Last Run Type: Reprotect
Last Run Time: 0 hrs, 1 min, 49 sec
Last Run Result: Success

🟦	Not Started	0
🟪	In Progress	6
🟩	Success	2
🟧	Warning	0
🟥	Error	0

Figure 117. A Running Recovery Plan Seen from the Summary Screen

Once all the protected virtual machines have been failed over and reported as powered on, you are ready to start validating that all application services restarted cleanly at the recovery site. Once you have completed the validation of the failed over application services at the recovery site, you are now in a position to report the successful failover to the business and enable the respective business users to access the application services, which are now being hosted on the recovery site.

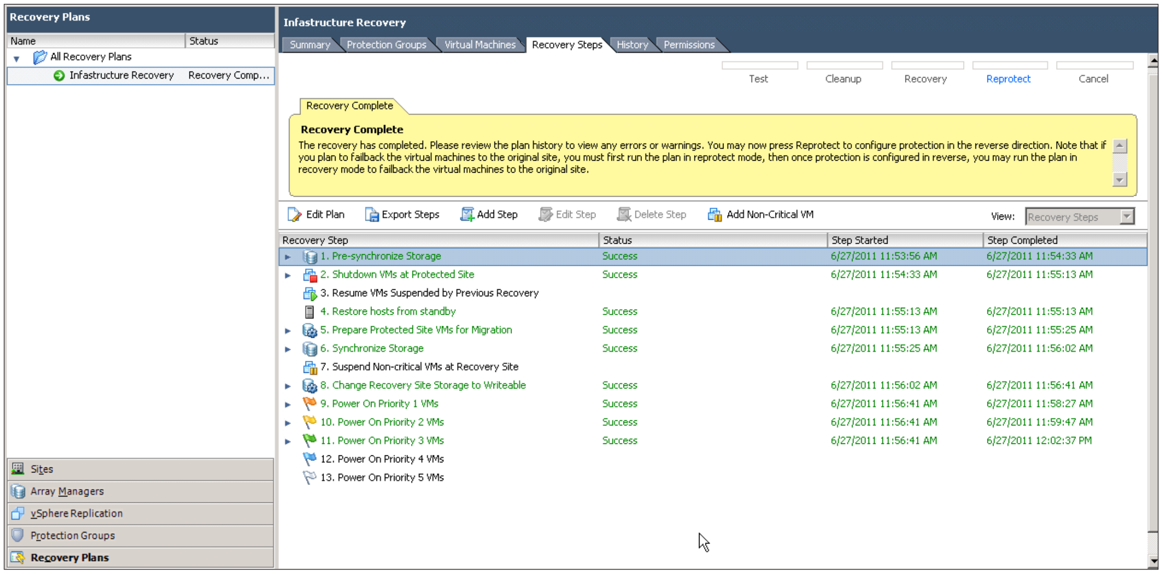

Name	Status
All Recovery Plans	
Infastructure Recovery	Recovery Comp...

Recovery Complete

The recovery has completed. Please review the plan history to view any errors or warnings. You may now press Reprotect to configure protection in the reverse direction. Note that if you plan to failback the virtual machines to the original site, you must first run the plan in reprotect mode, then once protection is configured in reverse, you may run the plan in recovery mode to failback the virtual machines to the original site.

Recovery Step	Status	Step Started	Step Completed
1. Pre-synchronize Storage	Success	6/27/2011 11:53:56 AM	6/27/2011 11:54:33 AM
2. Shutdown VMs at Protected Site	Success	6/27/2011 11:54:33 AM	6/27/2011 11:55:13 AM
3. Resume VMs Suspended by Previous Recovery			
4. Restore hosts from standby	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:13 AM
5. Prepare Protected Site VMs for Migration	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:25 AM
6. Synchronize Storage	Success	6/27/2011 11:55:25 AM	6/27/2011 11:56:02 AM
7. Suspend Non-critical VMs at Recovery Site			
8. Change Recovery Site Storage to Writeable	Success	6/27/2011 11:56:02 AM	6/27/2011 11:56:41 AM
9. Power On Priority 1 VMs	Success	6/27/2011 11:56:41 AM	6/27/2011 11:58:27 AM
10. Power On Priority 2 VMs	Success	6/27/2011 11:56:41 AM	6/27/2011 11:59:47 AM
11. Power On Priority 3 VMs	Success	6/27/2011 11:56:41 AM	6/27/2011 12:02:37 PM
12. Power On Priority 4 VMs			
13. Power On Priority 5 VMs			

Figure 118. Recovery Complete

SRM automatically generates a report for each recovery plan execution. In this instance, the report is for an SRM **Run** operation against the recovery plan that was selected.

5. The report is accessible via the **History** tab and can be viewed by clicking the **View** link under the **Actions** column.

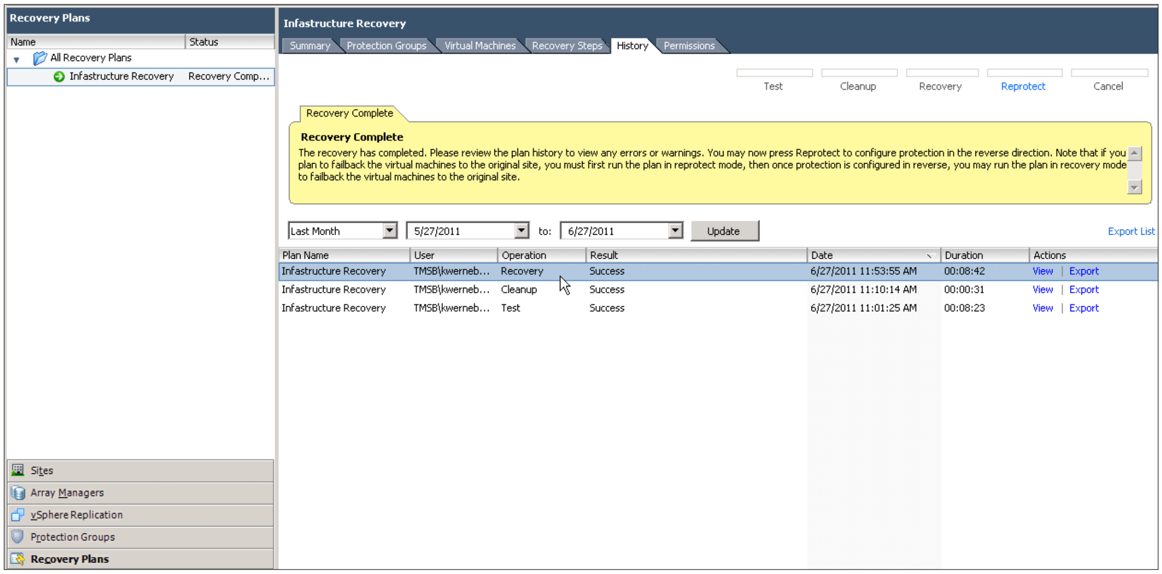

Recovery Complete

The recovery has completed. Please review the plan history to view any errors or warnings. You may now press Reprotect to configure protection in the reverse direction. Note that if you plan to failback the virtual machines to the original site, you must first run the plan in reprotect mode, then once protection is configured in reverse, you may run the plan in recovery mode to failback the virtual machines to the original site.

Plan Name	User	Operation	Result	Date	Duration	Actions
Infastructure Recovery	TMSB\kwerneb…	Recovery	Success	6/27/2011 11:53:55 AM	00:08:42	View | Export
Infastructure Recovery	TMSB\kwerneb…	Cleanup	Success	6/27/2011 11:10:14 AM	00:00:31	View | Export
Infastructure Recovery	TMSB\kwerneb…	Test	Success	6/27/2011 11:01:25 AM	00:08:23	View | Export

Figure 119. History Report for Recovery Plan

The steps to failback services from the recovery site back to the protected site once the disaster event is over are outlined in the next exercise.

The following is a recap of the high-level tasks executed by SRM when performing a failover of virtual machines from the protected site to the recovery site via the **Run a recovery plan** option. SRM automates many of the tasks required at the time of failover. With the push of one button, SRM does the following:

- Powers down the protected virtual machines if there is connectivity between sites and they are online.
- Suspends data replication and Read/Write enables the replica datastores.
- Rescans the ESX servers at the recovery site.
- Registers the replicated protected virtual machines.
- Suspends nonessential virtual machines at the recovery site if specified to free up resources for the protected virtual machines being failed over.
- Completes power-up of replicated protected virtual machines in accordance with the recovery plan.

Exercise 5. Automating Failback

Following a DR event or a planned migration, it may be beneficial or necessary to ensure that the environment is once again protected and replicated back to the initial primary site (Site A). This ensures that the environment is protected against any further unrecoverable service interruptions, and it also enables an automated failback to the primary site. SRM can be configured so that, with the use of a single button, the entire environment that has been recovered can be reprotected again back to the initial site.

Automatic reprotection of the environment is only supported for protection groups that are using array-based replication, because the reprotect process must use a Storage Replication Adapter (SRA) to reverse replication of an array.

Step 1: Reprotect the environment

To automatically reprotect an environment, as follows, you must be looking at the context of a **completed recovery plan:**

1. Open the vSphere Client and connect to the vCenter server at the protected site.
2. Log in as a vSphere administrator.
3. Click the **Site Recovery icon** on the vSphere Client Home page under Solutions and Applications.
4. Navigate to a recovery plan that has completed successfully by clicking **Recovery Plans** in the left pane, and selecting the recovery plan that has completed a successful failover. If you are continuing from the previous exercise, you should already be on this window.
5. Click the blue **Reprotect** button in the top task bar or click the blue **Reprotect** text in the available actions listed above the recovery plan.

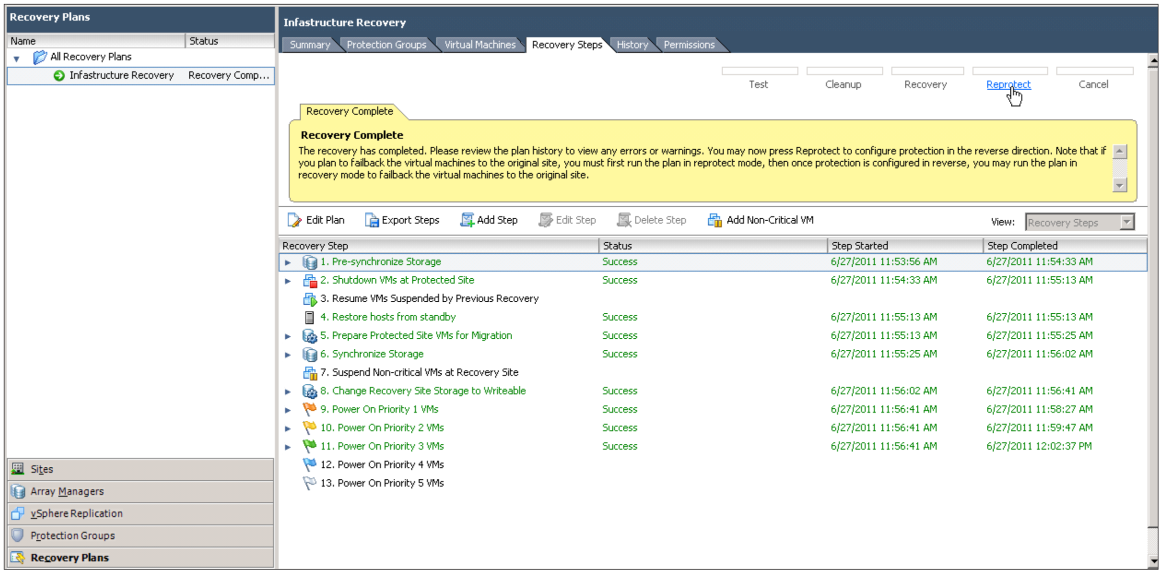

Recovery Complete

The recovery has completed. Please review the plan history to view any errors or warnings. You may now press Reprotect to configure protection in the reverse direction. Note that if you plan to failback the virtual machines to the original site, you must first run the plan in reprotect mode, then once protection is configured in reverse, you may run the plan in recovery mode to failback the virtual machines to the original site.

Recovery Step	Status	Step Started	Step Completed
1. Pre-synchronize Storage	Success	6/27/2011 11:53:56 AM	6/27/2011 11:54:33 AM
2. Shutdown VMs at Protected Site	Success	6/27/2011 11:54:33 AM	6/27/2011 11:55:13 AM
3. Resume VMs Suspended by Previous Recovery			
4. Restore hosts from standby	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:13 AM
5. Prepare Protected Site VMs for Migration	Success	6/27/2011 11:55:13 AM	6/27/2011 11:55:25 AM
6. Synchronize Storage	Success	6/27/2011 11:55:25 AM	6/27/2011 11:56:02 AM
7. Suspend Non-critical VMs at Recovery Site			
8. Change Recovery Site Storage to Writeable	Success	6/27/2011 11:56:02 AM	6/27/2011 11:56:41 AM
9. Power On Priority 1 VMs	Success	6/27/2011 11:56:41 AM	6/27/2011 11:58:27 AM
10. Power On Priority 2 VMs	Success	6/27/2011 11:56:41 AM	6/27/2011 11:59:47 AM
11. Power On Priority 3 VMs	Success	6/27/2011 11:56:41 AM	6/27/2011 12:02:37 PM
12. Power On Priority 4 VMs			
13. Power On Priority 5 VMs			

Figure 120. Automated Reprotect

6. Click the acknowledgement check box, indicating that you understand the operation cannot be undone, and click **Next.**
7. Review the summary information regarding the reprotect action and, if satisfied that the options are correct, click **Start** to initiate the reprotect of the environment.

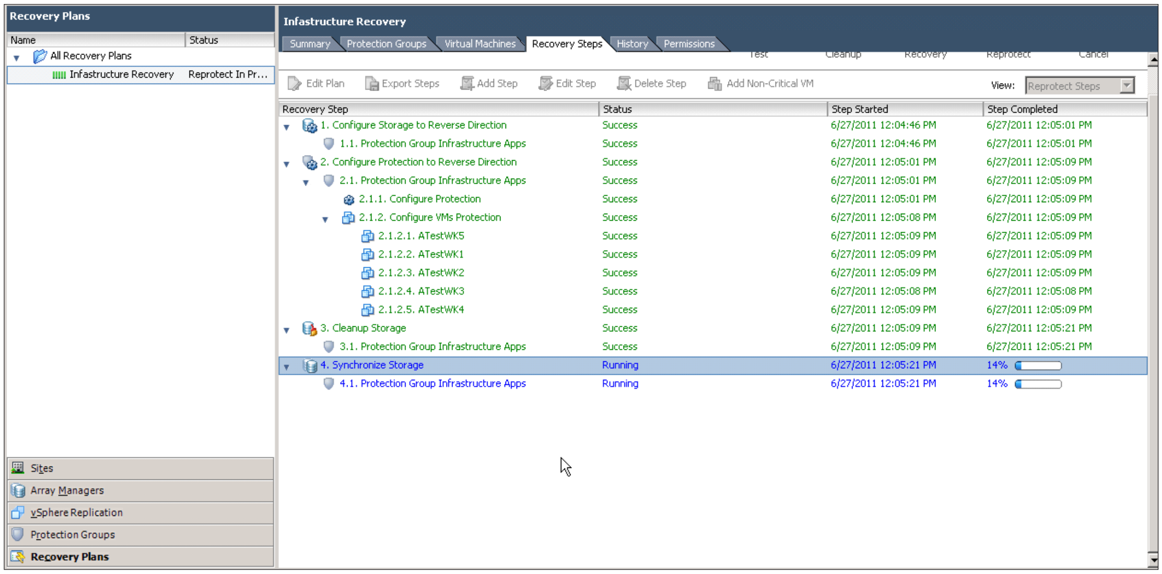

Recovery Step	Status	Step Started	Step Completed
1. Configure Storage to Reverse Direction	Success	6/27/2011 12:04:46 PM	6/27/2011 12:05:01 PM
1.1. Protection Group Infrastructure Apps	Success	6/27/2011 12:04:46 PM	6/27/2011 12:05:01 PM
2. Configure Protection to Reverse Direction	Success	6/27/2011 12:05:01 PM	6/27/2011 12:05:09 PM
2.1. Protection Group Infrastructure Apps	Success	6/27/2011 12:05:01 PM	6/27/2011 12:05:09 PM
2.1.1. Configure Protection	Success	6/27/2011 12:05:01 PM	6/27/2011 12:05:09 PM
2.1.2. Configure VMs Protection	Success	6/27/2011 12:05:08 PM	6/27/2011 12:05:09 PM
2.1.2.1. ATestWK5	Success	6/27/2011 12:05:09 PM	6/27/2011 12:05:09 PM
2.1.2.2. ATestWK1	Success	6/27/2011 12:05:09 PM	6/27/2011 12:05:09 PM
2.1.2.3. ATestWK2	Success	6/27/2011 12:05:09 PM	6/27/2011 12:05:09 PM
2.1.2.4. ATestWK3	Success	6/27/2011 12:05:08 PM	6/27/2011 12:05:08 PM
2.1.2.5. ATestWK4	Success	6/27/2011 12:05:09 PM	6/27/2011 12:05:09 PM
3. Cleanup Storage	Success	6/27/2011 12:05:09 PM	6/27/2011 12:05:21 PM
3.1. Protection Group Infrastructure Apps	Success	6/27/2011 12:05:09 PM	6/27/2011 12:05:21 PM
4. Synchronize Storage	Running	6/27/2011 12:05:21 PM	14%
4.1. Protection Group Infrastructure Apps	Running	6/27/2011 12:05:21 PM	14%

Figure 121. Reprotect in Progress

The reprotect action will use the SRAs to communicate with the arrays that are associated with the protection groups used by the recovery plan to first ensure that replication can be established in the "reverse" direction – from Site B to Site A. Once this is established, it will direct the array to replicate the protection group, and ensure that protection for all VMs in the recovery plan can be configured. If protection can be established, data is now synchronized between Site B and Site A to ensure that the environment is now protected again and ready for an automated failback to the initial primary site.

Reprotecting the environment will ensure that the now active readable/writable VMs at Site B are replicated, and it will also create **shadow VMs** for the replicas of these systems that are now held at Site A. If you look through your VM inventory at both sites, you will see the unique lightning bolt icons representing the placeholder shadow VMs at Site A and the active VMs themselves at Site B. If the location of these objects looks incorrect, you may have to revisit the inventory mappings used early in the guide to ensure that VMs and shadow VMs are being positioned in the correct location of inventory.

Step 2: Failback to the original site

Failback in SRM is the process of reprotecting the environment, and executing the same recovery plan that was used for initial failover to ensure that the same steps as used in a failover are run, but in the opposite direction.

Once virtual machines have been successfully recovered by SRM, the next step will at some point in time be a failback, to return the environment to its primary site of operations, or to distribute workloads between sites.

The failback scenario covered as part of this evaluation will involve failing back to a site that is still in a good state after the DR event (in other words, the same equipment and configuration that was failed over from has remained). If you suffer a total site loss of the site you failed over from, then additional steps must obviously be followed before you can failback, as you must do to recreate the environment at the lost site before commencing any failback. If the equipment is completely replaced, a reprotect and failback will not be an option, because the array pairs will have changed and the protection groups must be recreated. For the evaluation guide, we will assume that the same gear is in place.

A summary of the workflow is as follows:

1. The failover recovery plan from Site A (protected site) to Site B (recovery site) is run.
2. Virtual machines in the recovery plan Infrastructure Recovery are successfully failed over to Site B.

3. Virtual machines in the recovery plan are now powered on and running successfully at Site B.
4. **Reprotect** has been run and the environment is now protected once more back to Site A, and the direction of replication reversed successfully.
5. A **test** of the recovery plan that was used to failover is run.
6. The recovery plan is executed as a full **Recovery.**

It is as simple to execute a failback as it is to run the initial recovery. Presuming the reprotect worked correctly, you can proceed to run the same recovery plan once again, ensuring that you are using **test** mode, to determine if an automated failback will run correctly.

After running the test, clean up the environment and run the recovery plan. In this situation, you may wish to run in **planned migration** mode, because a failback usually indicates a controlled environment that is not as constrained by RTO and is more focused on data consistency and predictability.

When running the recovery plan for failback, take note of the **Protected Site** and **Recovery Site** in the summary of information before clicking **Start.** It should reflect the appropriate sites, with **protected** reflecting your **Site B** and **recovery** reflecting your **Site A.**

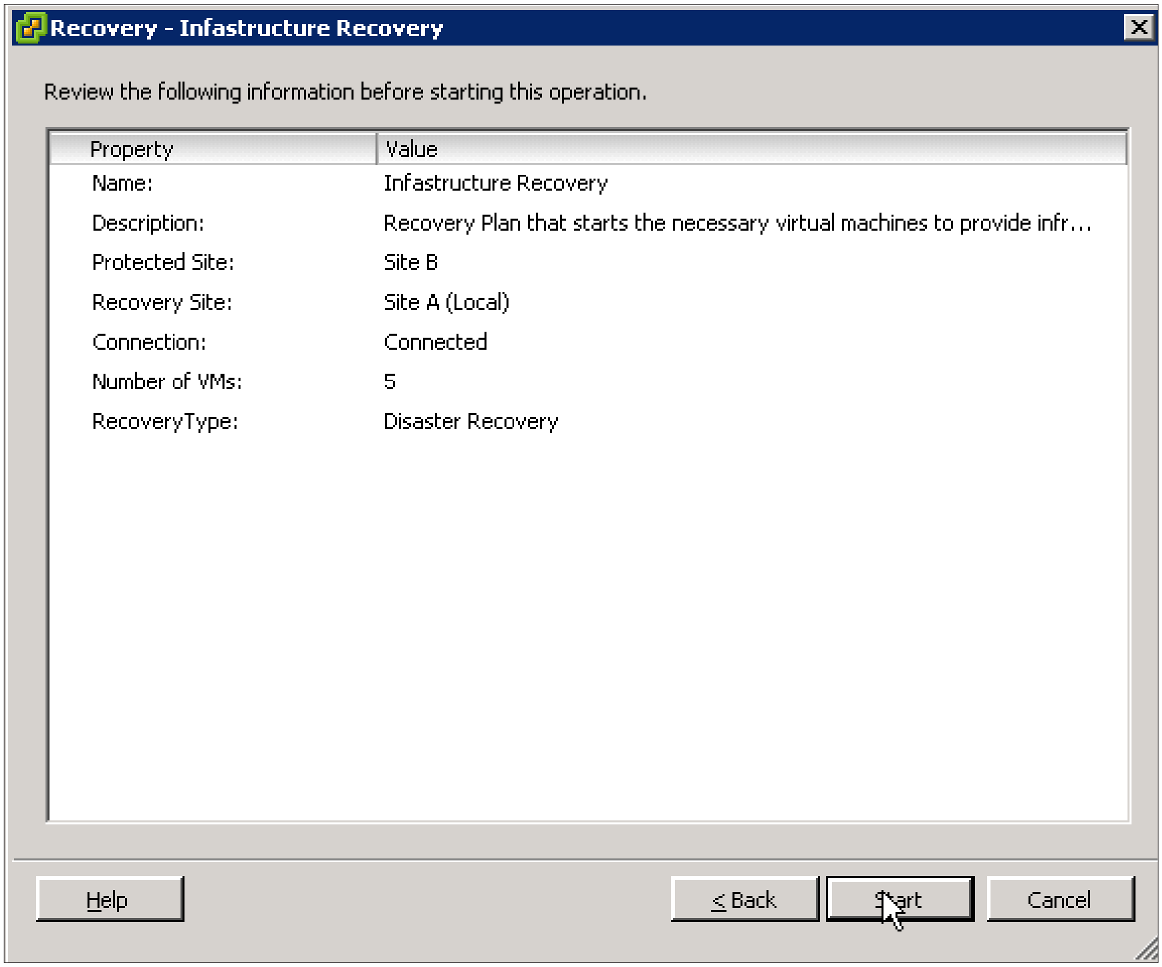

Recovery - Infastructure Recovery

Review the following information before starting this operation.

Property	Value
Name:	Infastructure Recovery
Description:	Recovery Plan that starts the necessary virtual machines to provide infr...
Protected Site:	Site B
Recovery Site:	Site A (Local)
Connection:	Connected
Number of VMs:	5
RecoveryType:	Disaster Recovery

Figure 122. Failback Is No More Than a Failover from Site B to Site A

When executing the failback of the recovery plan, it will only have the virtual machines and other information that were in the recovery plan for initial failover. If VMs have been added to the recovery site **after** failover, but **before** failback, they will not be automatically represented in the failback. If the environment has changed dramatically during the failed-over state, ensure that recovery plans are updated to reflect the new environment.

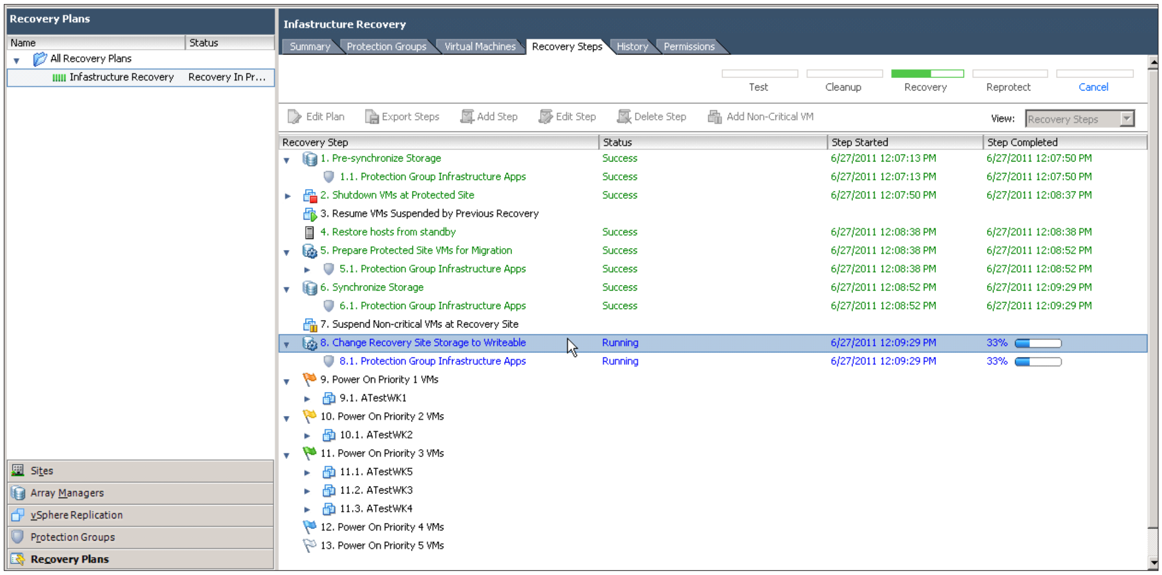

Recovery Step	Status	Step Started	Step Completed
1. Pre-synchronize Storage	Success	6/27/2011 12:07:13 PM	6/27/2011 12:07:50 PM
1.1. Protection Group Infrastructure Apps	Success	6/27/2011 12:07:13 PM	6/27/2011 12:07:50 PM
2. Shutdown VMs at Protected Site	Success	6/27/2011 12:07:50 PM	6/27/2011 12:08:37 PM
3. Resume VMs Suspended by Previous Recovery			
4. Restore hosts from standby	Success	6/27/2011 12:08:38 PM	6/27/2011 12:08:38 PM
5. Prepare Protected Site VMs for Migration	Success	6/27/2011 12:08:38 PM	6/27/2011 12:08:52 PM
5.1. Protection Group Infrastructure Apps	Success	6/27/2011 12:08:38 PM	6/27/2011 12:08:52 PM
6. Synchronize Storage	Success	6/27/2011 12:08:52 PM	6/27/2011 12:09:29 PM
6.1. Protection Group Infrastructure Apps	Success	6/27/2011 12:08:52 PM	6/27/2011 12:09:29 PM
7. Suspend Non-critical VMs at Recovery Site			
8. Change Recovery Site Storage to Writeable	Running	6/27/2011 12:09:29 PM	33%
8.1. Protection Group Infrastructure Apps	Running	6/27/2011 12:09:29 PM	33%
9. Power On Priority 1 VMs			
9.1. ATestWK1			
10. Power On Priority 2 VMs			
10.1. ATestWK2			
11. Power On Priority 3 VMs			
11.1. ATestWK5			
11.2. ATestWK3			
11.3. ATestWK4			
12. Power On Priority 4 VMs			
13. Power On Priority 5 VMs			

Figure 123. Failback Will Not Automatically Update Recovery Plans with VMs That Are Added to the Failover Site

At any time, you may wish to get a graphic reminder of the state of replication and learn what datastores are being protected to which site. This becomes especially important with regards to the process of reprotects and failbacks to ensure that data is being synchronized correctly. To ensure that the reprotect has successfully reversed the direction of replication and that a failback is going to be successful, you might choose to return to the **Array Managers** section of SRM and examine the direction of replication for each relevant device by clicking the appropriate **Array Manager** and selecting the **Devices** tab. Here, you can see in graphic detail the direction in which devices are replicating, and can ensure that the failover will be directed appropriately.

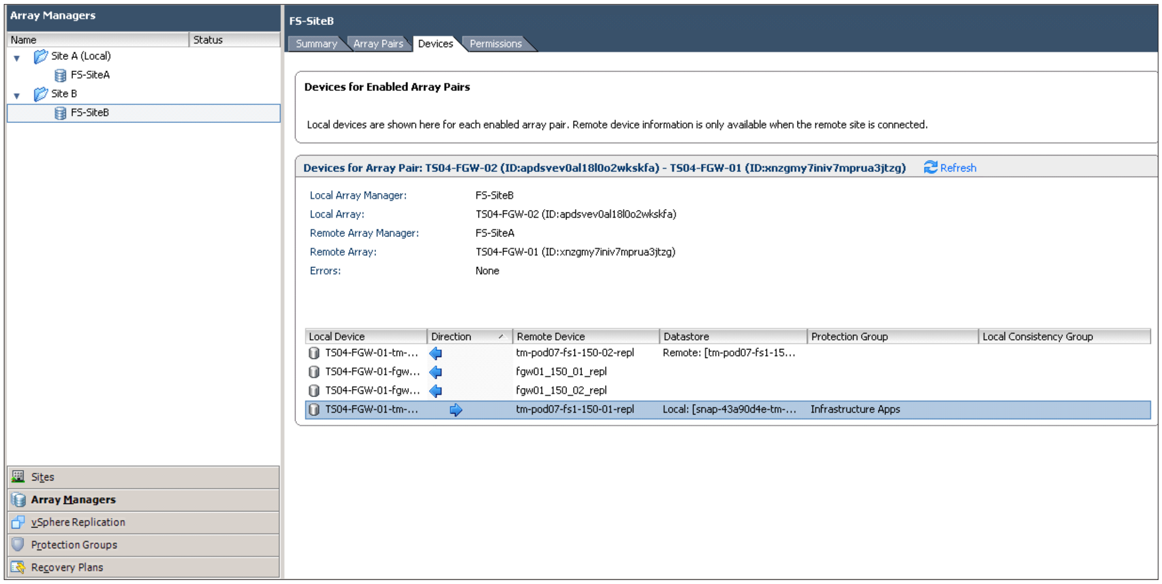

Devices for Enabled Array Pairs

Local devices are shown here for each enabled array pair. Remote device information is only available when the remote site is connected.

Devices for Array Pair: TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa) - TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg)

Local Array Manager:	FS-SiteB
Local Array:	TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa)
Remote Array Manager:	FS-SiteA
Remote Array:	TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg)
Errors:	None

Local Device	Direction	Remote Device	Datastore	Protection Group	Local Consistency Group
TS04-FGW-01-tm-…	⬅	tm-pod07-fs1-150-02-repl	Remote: [tm-pod07-fs1-15…		
TS04-FGW-01-fgw…	⬅	fgw01_150_01_repl			
TS04-FGW-01-fgw…	⬅	fgw01_150_02_repl			
TS04-FGW-01-tm-…	➡	tm-pod07-fs1-150-01-repl	Local: [snap-43a90d4e-tm-…	Infrastructure Apps	

Figure 124. Device-Specific Replication Information

*NOTE: Following a successful failback, you should **remember** to run **reprotect** once again to reverse the replication of the now failed-back environment. This will ensure that the environment is once more protected and ready for a failover. Consider a failover and failback a four-step process – failover, reprotect, failback, reprotect.*

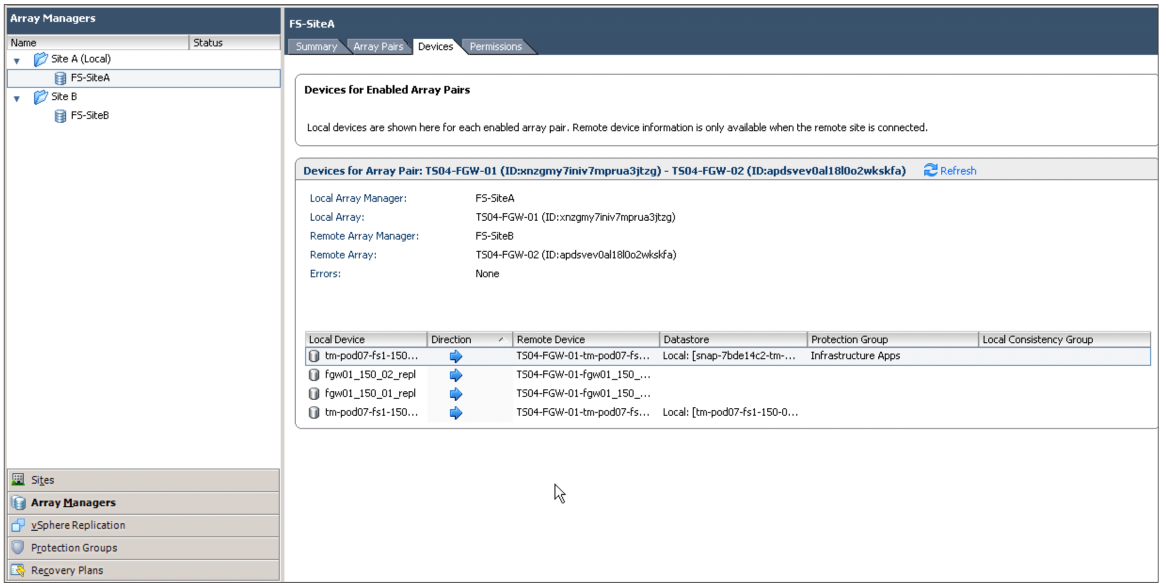

Devices for Enabled Array Pairs

Local devices are shown here for each enabled array pair. Remote device information is only available when the remote site is connected.

Devices for Array Pair: TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg) - TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa)

Local Array Manager:	FS-SiteA
Local Array:	TS04-FGW-01 (ID:xnzgmy7iniv7mprua3jtzg)
Remote Array Manager:	FS-SiteB
Remote Array:	TS04-FGW-02 (ID:apdsvev0al18l0o2wkskfa)
Errors:	None

Local Device	Direction	Remote Device	Datastore	Protection Group	Local Consistency Group
tm-pod07-fs1-150…	➡	TS04-FGW-01-tm-pod07-fs…	Local: [snap-7bde14c2-tm-…	Infrastructure Apps	
fgw01_150_02_repl	➡	TS04-FGW-01-fgw01_150_…			
fgw01_150_01_repl	➡	TS04-FGW-01-fgw01_150_…			
tm-pod07-fs1-150…	➡	TS04-FGW-01-tm-pod07-fs…	Local: [tm-pod07-fs1-150-0…		

Figure 125. Ensure That You Reprotect After Failback and That Devices Are Protected Correctly Once More

Summary

VMware vCenter Site Recovery Manager (SRM) leverages your vCenter and vSphere platform to improve disaster recovery in the following ways:

- **Rapid** —automating the disaster recovery process for your virtual machines by eliminating the complexities of traditional physical disaster recovery.
- **Reliable**—ensuring the proper execution of the recovery plan, enabling easier, more frequent tests in an isolated environment without impacting services in the protected site.
- **Manageable**—centrally managing recovery plans and making plans dynamic to match a dynamic virtualized environment.
- **Affordable**—utilizing appropriate replication technology for your needs, while safely increasing utilization of recovery site infrastructure and reducing management costs associated with DR practices.

Site Recovery Manager enables you to do the following:

- **Expand disaster recovery protection**—now any workload in a virtual machine can be protected with minimal incremental effort and cost.
- **Reduce time to recovery**—as soon as a disaster is declared, Site Recovery Manager allows for the recovery of protected virtual machines with a few mouse clicks to the designated recovery site.
- **Increase reliability of recovery**—replication of the system state ensures that your protected virtual machines have all they need to start up in the protected site. Hardware independence that is realized through your VMware Infrastructure eliminates failures due to different hardware.
- **Enable easier and more frequent testing**—Site Recovery Manager enables you to test your recovery plan in an isolated environment without impacting services in the protected site while using the actual failover sequence that will be executed during a real disaster.

Site Recovery Manager 5.0 provides additional features—vSphere Replication, automatic reprotection and failback, new means of handling dependencies and priorities, and a simpler user interface that you can leverage to extend your disaster recovery plan to cover even more of your business-continuity needs.

This guide provides you with step-by-step instructions on how to set up automated disaster recovery workflows using Site Recovery Manager, as well as information on other cutting-edge DR features in Site Recovery Manager. With Site Recovery Manager, you can design and implement a comprehensive disaster recovery plan for your virtual environment. After going through the evaluation exercises in this guide, you should be able to make the right choice to implement your disaster-recovery solutions in your virtual datacenter.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com