[**PACKT**] enterprise ✖
PUBLISHING
professional expertise distilled

Disaster Recovery Using VMware vSphere Replication and vCenter Site Recovery Manager

Abhilash GB

Chapter No. 1

"Installing and Configuring vCenter Site Recovery Manager (SRM) 5.5"

In this package, you will find:

A Biography of the author of the book

A preview chapter from the book, Chapter NO.1 "Installing and Configuring vCenter Site Recovery Manager (SRM) 5.5"

A synopsis of the book's content

Information on where to buy this book

About the Author

Abhilash GB is an author, a designer, and a VMware vExpert 2014, who specializes in the areas of Data Center Virtualization and cloud computing.

He is a VMware Certified Advanced Professional in Data Center Administration (VCAP4-DCA and VCAP5-DCA). He has been in the IT industry for nearly a decade, and he has been working on VMware products and technologies since the start of 2007. He is also the author of *VMware vSphere 5.1 Cookbook*, *Packt Publishing*.

He is a passionate author who's willing to contribute more titles to the VMware community, and also an aspiring engineer always ready to indulge in designing and creating great solutions for customers.

> I would like to dedicate this book to my wife and my parents. Without their patience and support, this book would not have been possible.
>
> Thanks to the technical reviewers Andy Grant, Kenneth van Ditmarsch, and Daniel Langenhan for their invaluable inputs.
>
> Special thanks to Nikhil Chinnari (Acquisition Editor), Harshal Ved (Project Coordinator), Govindan K (Content Development Editor), and Shashank Desai and Menza Mathew (Technical Editors) for their support during the course of writing this book.

Disaster Recovery Using VMware vSphere Replication and vCenter Site Recovery Manager

This book covers the use of vSphere Replication and VMware Site Recovery Manager for making your vSphere environment recoverable in the event of a disaster. All the concepts and tasks covered in this book are for vSphere Replication 5.5 and VMware vCenter Site Recovery Manager 5.5.

What This Book Covers

Chapter 1, *Installing and Configuring vCenter Site Recovery Manager (SRM) 5.5*, introduces you to the architecture of SRM and also guides you through the process of installing and configuring SRM to leverage array-based replication.

Chapter 2, *Creating Protection Groups and Recovery Plans*, teaches you how to configure protection for virtual machines by creating Protection Groups and creating an orchestrated runbook with the help of Recovery Plans.

Chapter 3, *Testing and Performing a Failover and Failback*, teaches you how to test the recovery plans that were created and also perform a Planned Migration, a Failover, and a Failback using them.

Chapter 4, *Deploying vSphere Replication 5.5*, guides you through the steps required in deploying vSphere Replication Appliances and vSphere Replication Servers.

Chapter 5, *Configuring and Using vSphere Replication 5.5*, teaches you how to add target sites and enable replication on virtual machines and recover them. It will also teach you to configure vCenter SRM to leverage vSphere Replication engine.

1
Installing and Configuring vCenter Site Recovery Manager (SRM) 5.5

In this chapter, we will cover the following topics:

- What is Site Recovery Manager (SRM)?
- Preparing storage for array-based replication
- Host presentation at protected and recovery sites
- Installing SRM on protected and recovery sites
- Installing SRM plugin for vSphere Client
- Pairing sites
- Installing Storage Replication Adapters (SRA)
- Adding an array manager
- Enabling an array pair
- Configuring placeholder datastores
- Creating resource, folder, and network mappings

Introduction

With today's IT infrastructures, be it virtual or physical, disaster recovery is of prime importance. Any business should be able to continue operating with reduced downtime for its sustainability amongst the competition. It also has a legal obligation towards customers to whom it sold its services. Two of the major factors used to market or sell a service are its High Availability and Recoverability.

Recoverability is the guarantee that the service offered and its data are protected against failures, and High Availability is the guarantee that the service offered would remain operational and the failures are handled in a way that the user of the service would not even know that there was a failure.

There are many ways in which businesses plan and implement disaster recovery. Although important, much of these decisions depend on the budgetary constraints. What turns out to be the most important is the existence of a disaster recovery plan. Gone are those days when you had to wait for a long period of time before all your critical applications were made available at a recovery site. With a lot of automation and scripting, businesses now expect better **Recovery Point Objective** (**RPO**) and **Recovery Time Objective** (**RTO**).

So what exactly are RPO and RTO?

RPO defines the amount of data an organization can afford to lose when measured against time.

RTO defines the amount of downtime the organization can afford for its services before it becomes operational again.

Both RPO and RTO are defined by time. For example, an organization can have an RPO set to 4 hours and RTO set to 1 hour. This means, it can afford to lose up to 4 hours of data, but it can only afford a service downtime up to 1 hour.

RTO only defines the amount of time a service can remain unavailable but doesn't account for the data loss. This is where RPO pitches in. It defines how much data loss can be afforded.

For example, if you were a company hosting an online document format conversion service, then setting a lower RTO value is very important because the customers will prefer access to the service, rather than to the historical data. The RPO value will determine how much historical data you will have to keep.

Both RPO and RTO help an organization to determine the type of backup and DR solution to meet the business requirements.

What is Site Recovery Manager (SRM)?

vCenter **Site Recovery Manager** (**SRM**) is an orchestration software that is used to automate disaster recovery testing and Failover. It can be configured to leverage either vSphere Replication or a supported array-based replication. With SRM, you can create Protection Groups and run Recovery Plans against them. The Recovery Plans can then be used to test the DR setup and perform a planned Failover, or it can be initiated during a disaster recovery. SRM is a not a product that performs an automatic Failover, which means that there is no intelligence built into SRM that would detect a disaster/outage and Failover the VMs. The disaster recovery process should be manually initiated. Hence, it is not a high availability solution; it is purely a tool that orchestrates a Recovery Plan.

Architecture

vCenter Site Recovery Manager is not a tool that works on its own. It needs to talk to other components in the vSphere environment. I will walk you through all the components involved in an environment protected using SRM.

The following are the components that will be involved in an SRM-protected environment:

Protected site	Recovery site
vCenter Server	vCenter Server
SRM instance	SRM instance
Array managers	Array managers
Storage Replication Adapter	Storage Replication Adapter

SRM requires both the protected and recovery sites to be managed by separate instances of the vCenter Server. It also requires an SRM instance at both the sites. SRM's functionalities are currently only available via the vSphere Client and not the vSphere Web Client. Hence, an SRM plugin needs to be installed on the same machine where the vSphere Client is installed. Refer to the following figure:
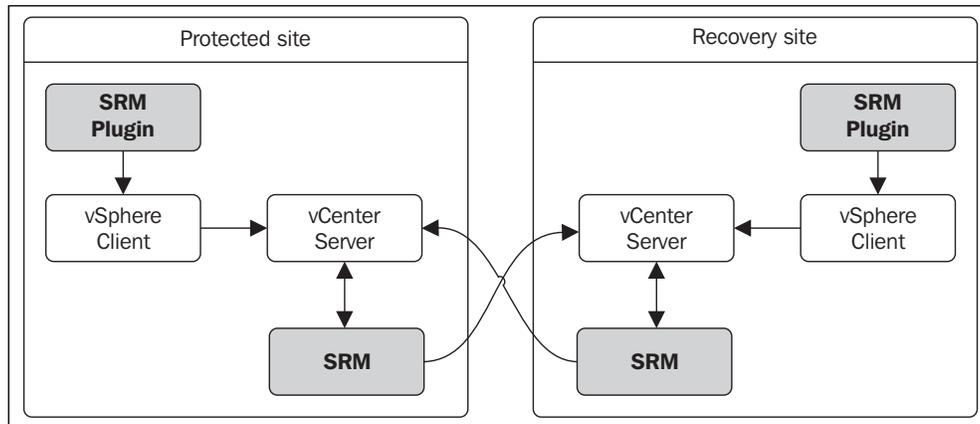
SRM as a solution cannot work on its own. This is because it is only an orchestration tool, and it does not include a replication engine. However, it can leverage either a supported array-based replication or VMware's proprietary replication engine, vSphere Replication. We have separate chapters covering the vSphere Replication.

Array manager

Each SRM instance needs to be configured with an array manager for it to communicate with the storage array. The array manager will detect the storage array using the information you supply to connect to the array. Before you add an array manager, you need to install an array-specific **Storage Replication Adapter** (**SRA**). This is because the array manager uses the SRA installed to collect replication information from the array. Refer to the following figure:
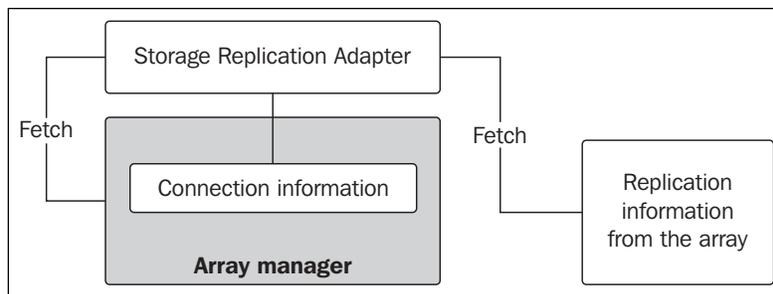
Storage Replication Adapter (SRA)

SRA is a storage vendor component that makes SRM aware of the replication configuration at the array. SRM leverages SRA's ability to gather information regarding the replicated volumes and direction of the replication from the array.

SRM also uses SRA for the following functions:

- Test Failover
- Recovery
- Reprotection

This is illustrated in the following figure:
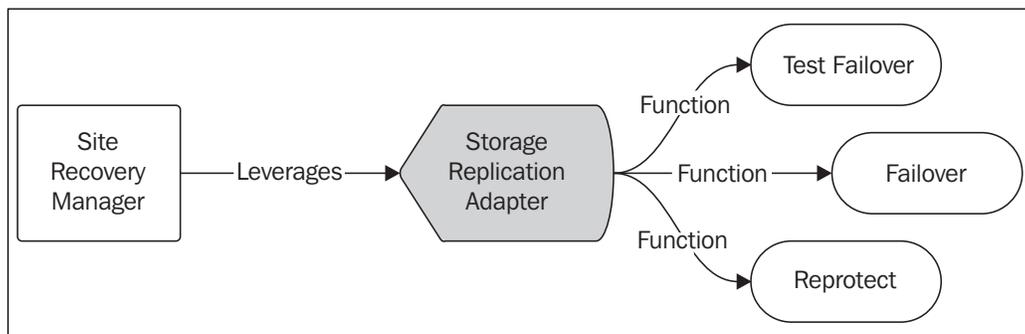
We will learn more about these functions in the next chapter. For now, it is important to understand that SRM requires SRA to be installed for all of its functions leveraging array-based replication.

When all these components are put together, a site protected by SRM will look as is shown in the following figure:
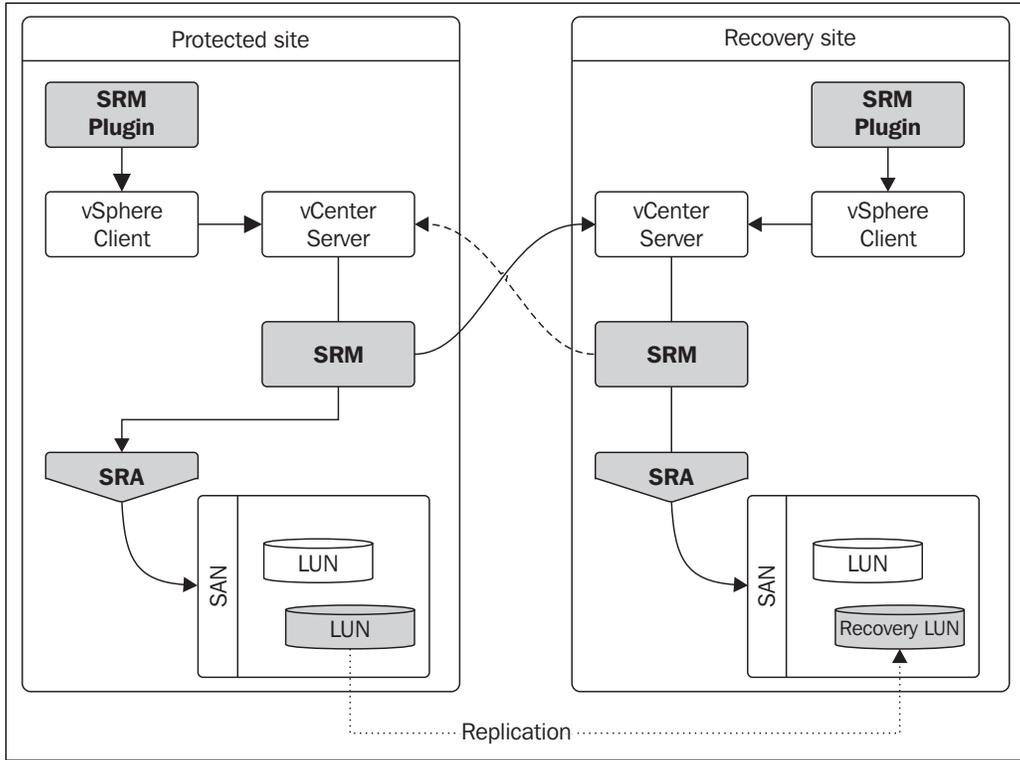
SRM conceptually assumes that both the protected and recovery sites are geographically separated. However, such a separation is not mandatory. You could use SRM to protect a chassis of servers and have another chassis in the same datacenter as the recovery site. Now that we have a brief understanding of the SRM architecture, it is time to learn how to set up these components.

Laying the groundwork for an SRM environment

You will need to perform a set of configuration activities to lay the groundwork for an SRM environment so that it can be used to test or execute the Recovery Plans.

Here is an outline of the tasks that need to be done to form an SRM environment:

- Preparing storage for an array-based replication
- Host presentation (zoning) at the protected and recovery sites
- Installing SRM on both the protected and recovery sites
- Installing the SRM plugin for vSphere Client
- Pairing the SRM instances
- Installing SRA
- Adding array managers
- Enabling array pairs
- Creating resource, folder, and network mappings
- Creating placeholder datastores

Preparing storage for an array-based replication

The first thing that you will need to do is make sure that your array is supported by VMware and licensed for an array-based replication by the array vendor. This is not a VMware license but a licensed feature from the storage vendor.

Now, to enable replication, you have a couple of approaches that you could employ, which are as follows:

Approach-1	Approach-2
• Identify the VMs that you want to protect	• Identify the VMs that you want to protect
• Identify the VMFS datastores the VMs have their files on	• Plan the sizing of a datastore large enough to hold all the identified VMs
• Identify the LUNs corresponding to the already identified datastores	• Create a LUN large enough to host the datastore
• Enable replication on the identified LUNs	• Present the new LUN to the hosts running the identified VM and create a new VMFS volume (datastore) on it
	• Migrate the VMs that you want to protect onto the new datastore
	• Enable replication on the new LUN that corresponds to the new datastore

Approach-1 is used in scenarios where the array does not have the spare capacity to provide a separate LUN to host-protected VMs. This approach adds an administrative overhead if the virtual machines are spread across multiple datastores. It also contributes to wastage of replication bandwidth and storage space since the LUNs that are replicated will contain unprotected virtual machine data.

Approach-2 is used in scenarios where you have ample spare capacity. This approach is the best as it reduces complexity, avoids replication bandwidth wastage, and reduces space wastage, as compared to Approach-1. However, this approach will have an impact on the size of the LUNs required at both the protected and replication sites.

Host presentation (zoning) at the protected and recovery sites

If you are involved in a new implementation, then you will have to plan how the ESXi hosts are zoned to the array at both the protected and recovery sites. This means that LUNs should be correctly zoned at the fabric. The details for protected site and recovery site arrays are as follows:

- At the protected site array, zone the ESXi hosts to communicate with the array and make sure that the LUNs housing the VMs to be protected are assigned to the ESXi hosts
- At the recovery site array, zone the ESXi hosts to the array, but do not map the replica LUNs to the hosts yet

Installing SRM on the protected and recovery sites

VCenter SRM has to be installed at both the protected and recovery sites for the disaster recovery setup to work. The installation process is identical regardless of the site it is being installed on; the only difference is that at each site, you will be registering the SRM installation to the vCenter Server managing that site.

SRM can either be installed on the same machine that has vCenter Server installed or on a different machine. The decision to choose either one of the installation models depends on how you want to size or separate the service-providing machines in your infrastructure. The most common deployment model is to have both vCenter and SRM on the same machine. The rationale behind this is that SRM will not work in a standalone mode; this means that if your vCenter Server goes down, there is no way you could access SRM. Like vCenter Server, SRM can be installed on a physical or virtual machine.

Another factor that you must take into account is the installation of SRA. SRAs have to be installed on the same machine where you already have SRM installed. Some SRAs need a reboot after installation. So, it is important to read through the storage vendor's documentation prior to proceeding to make a deployment choice for SRM. If the vCenter downtime is not feasible, then you will have to consider installing SRM on a separate machine.

Nevertheless, it is important to be aware of the software and hardware requirements of a software installation before it is actually installed. This is to make sure that you don't run into compatibility or supportability issues during the course of using the product. To understand the requirement of SRM, refer to page number 23 in *Chapter 2*, *Site Recovery Manager System Requirements*, in the *Site Recovery Manager installation and configuration guide for SRM 5.5* document available at `http://pubs.vmware.com/ srm-55/topic/com.vmware.ICbase/PDF/srm-install-config-5-5.pdf`.

The following flowchart depicts the processes involved in installing vCenter SRM:
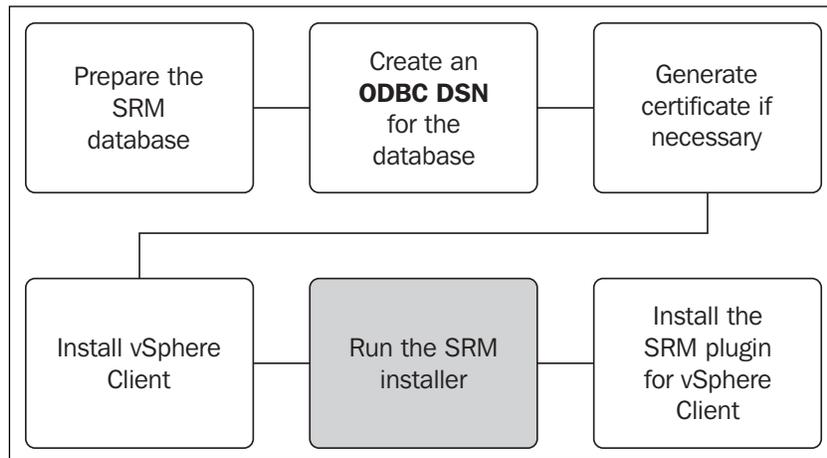
Performing the SRM installation

Let's assume that the SRM database and the 64-bit DSN have already been created. We will delve directly into the installation procedure using the SRM installer.

Before you begin, you will need to download the SRM installation bundle from the VMware website. It can be downloaded by navigating to `www.vmware.com` and then to the **vCenter Site Recovery Manager** option in the **Downloads** menu. You need to log in to your `my.vmware.com` account before you download the executable.

The following procedure will guide you through the SRM installation wizard:

1. Double-click on the downloaded executable to load the installer.
2. On the welcome screen of the installation wizard, click on **Next** to continue.

3. Choose a destination folder for the installer to place the files. The default location is `C:\Program Files\VMware\VMware vCenter Site Recovery Manager\`. You can change this by clicking on the **Change** button. For now, I have chosen to leave the default in place. Click on **Next** to continue.

4. On the next screen, you will be prompted to install the vSphere Replication UI plugin for SRM. You can choose to install or not install the plugin at this stage. Since we are not discussing vSphere Replication in this chapter, I have selected **Do not install vSphere Replication** as the option. Click on **Next** to continue.

5. On the next screen, provide the FQDN/IP and the credentials of the vCenter Server; the SRM instance should also be registered. Use of a separate service account, instead of the built-in administrator account, is recommended. In most cases where SRM is currently being installed, it is the local vCenter Server managing the site. Click on **Next** to continue.

6. You will then be prompted to a certificate source. Here, you can either let the installer generate a certificate, or you can supply a certificate file generated by a certificate authority.

7. The options available are as follows:
    ° **Automatically generate a certificate**
    ° **Use a PKCS#12 certificate file**

8. Make a selection of your choice and click on **Next** to continue.

    Here, we have chosen to let the installer generate a new certificate. Use the second option if you already have a certificate file from your certificate authority. VMware recommends using CA-signed certificates for all its products.
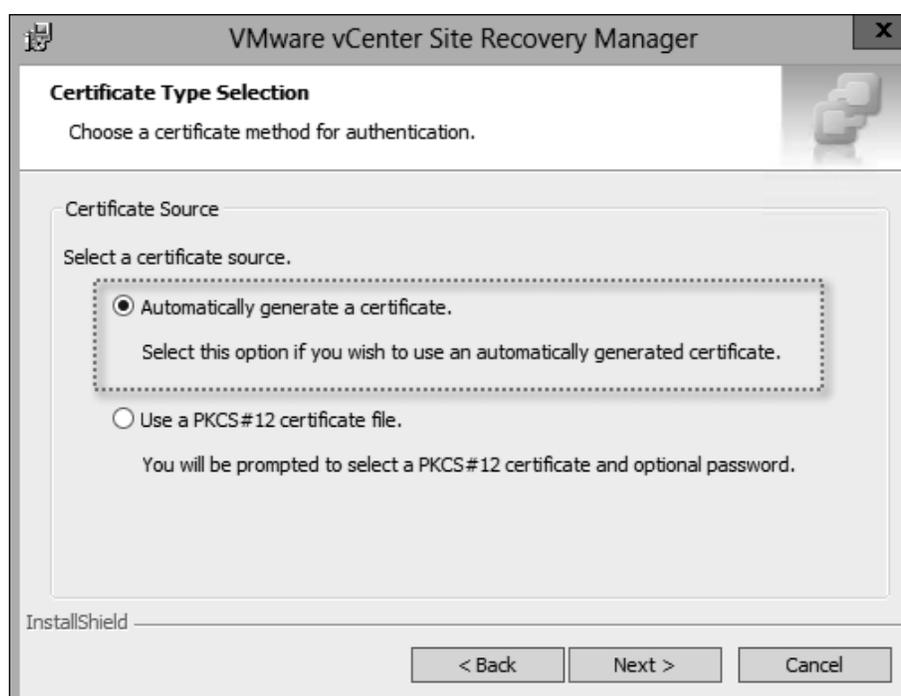
9. On the next screen, supply the details (organization and organization unit) for the certificate generation, and then click on **Next** to continue. You will be prompted for this information only if you have chosen to automatically generate a certificate.

10. Supply a local site name, two e-mail addresses of the administrators who need to be notified of any event, and the IP address of the machine where we will install SRM. The local site name can be any name that you supply. In this case, I have used the name of the vCenter Server that is managing the site. The site name can be changed by going to the **Advanced Settings** tab of the site post installation. Click on **Next** to continue.

11. Now, you will be prompted to supply the details of the database that you previously created for SRM.

12. The following details are needed to proceed further:

    ° The name of the 64-bit DSN that is configured to connect to the SRM database.

    ° The database user credentials, which could be a user that you created manually at the database server for the SRM database. Although, in this example, I have used the `sa` credentials, it is not a recommended practice to expose the `sa` credentials. In most environments, the `sa` account is used by the database administrator. Consider using a separate service account.

13. Supply the details and click on **Next** to continue.

14. On the **Ready to install the Program** screen, click on **Install** to begin the installation.

15. Once the installation is complete, click on **Finish** to exit the installer.

Installing the SRM plugin for vSphere Client

The SRM functions are exposed in the vSphere Client UI with the help of the SRM plugin. The SRM installer does not install this plugin. This is because it is a plugin for the vSphere Client and not the vCenter Server. The plugin needs to be installed separately on the machine where you have vSphere Client installed.

This is how you do it:

1. Connect to vCenter Server using vSphere Client.

2. Navigate to the **Manage Plug-ins** option in the **Plug-ins** tab, as shown in the following screenshot:
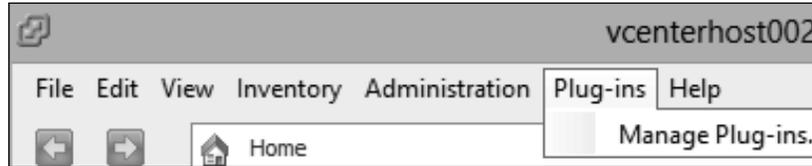
3. In the **Plug-in Manager** window, the SRM plugin should be listed under the **Available Plug-ins** category.
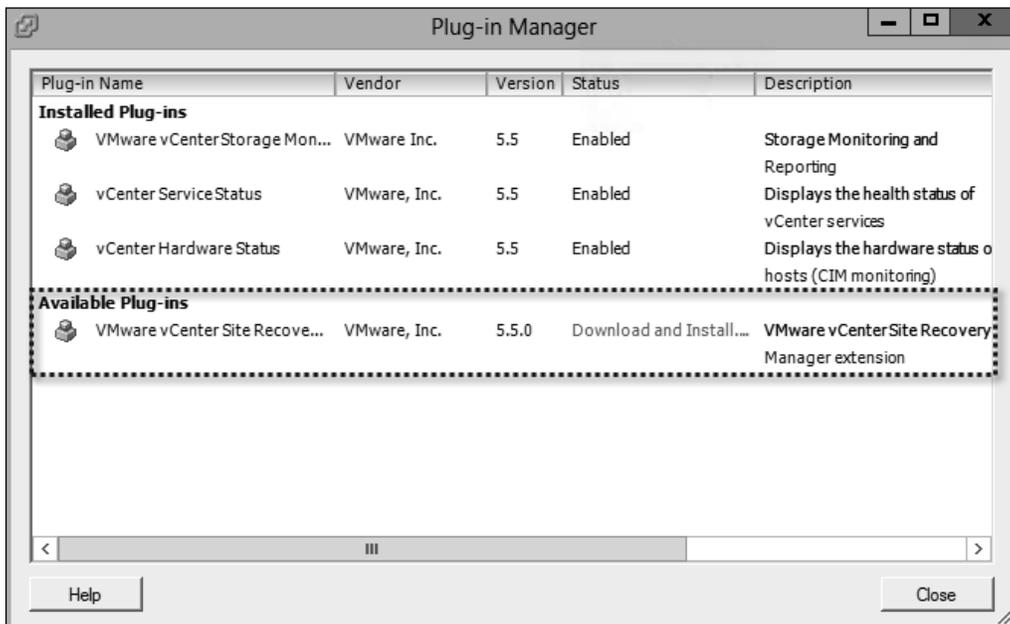
4. Click on the **Download and Install** link to fetch and install the SRM Plugin. The plugin installation is pretty straightforward.

5. Once the installation is complete, the vCenter inventory home page should list **Site Recovery** under **Solutions and Applications**.
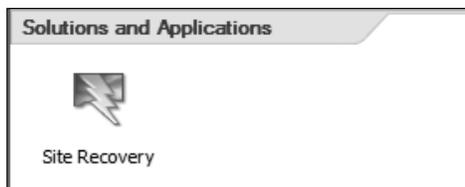
Pairing sites

Once SRM is installed on both the sites, the next step is to pair the sites together. The pairing process establishes a connection between the vCenter Servers at the protected and recovery sites, which in turn makes the SRM instances at both the sites aware of its counterpart at the other site (protected/recovery). Without the sites being paired, we can't proceed further with the configuration of the DR setup.

This is how the sites are paired:

1. Connect to the protected/recovery site vCenter Server using vSphere Client.
2. Navigate to the inventory home page and click on **Site Recovery**.
3. Click on **Sites** in the left pane.
4. Right-click on the local site listed, and click on **Configure Connection** to bring up the configure connection wizard. Refer to the following screenshot:
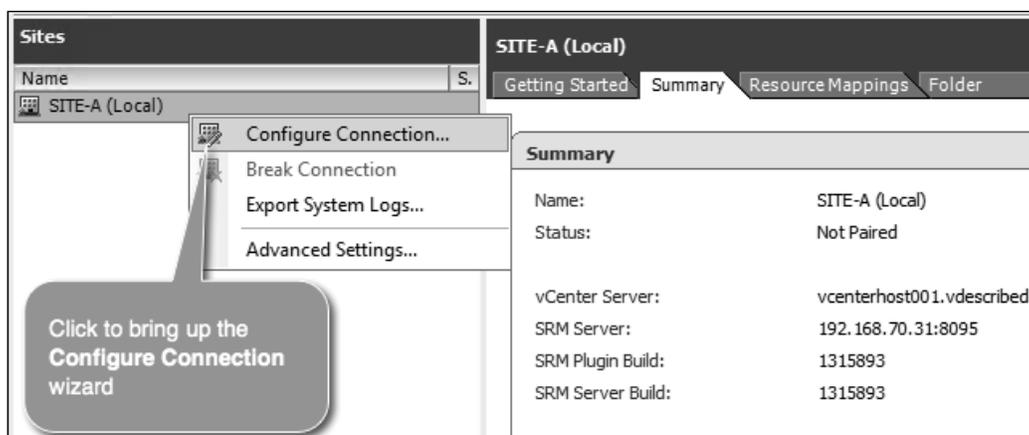
5. In the **Configure Connection** wizard, supply the FQDN of the remote vCenter Server and click on **Next** to continue. Accept any subsequent certificate warnings.

6. Furnish the administrator credentials, and click on **Next**. Accept any subsequent certificate warnings.

7. This will begin the pairing process, which will establish a connection with the remote vCenter Server and SRM, and also establish reciprocity. Click on **Finish** to exit the wizard.
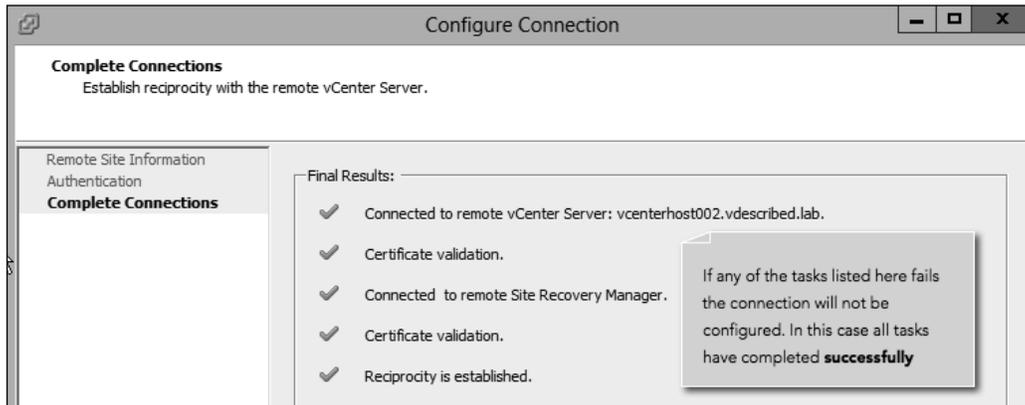
8. Once you exit the wizard, you will be prompted again for the administrator credentials of the remote vCenter Server. Enter the credentials and click on **OK**. Ignore any subsequent certificate warnings. Refer to the following screenshot:
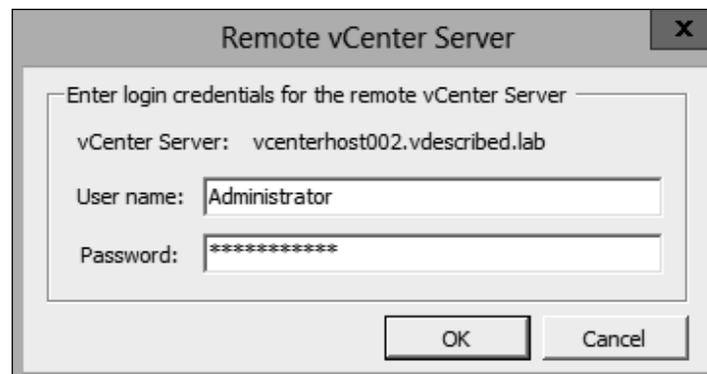
9. You should see both the sites listed under the **Sites** pane now:
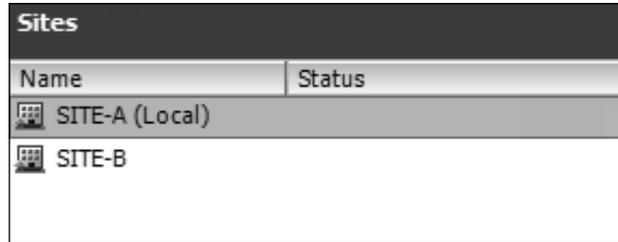
Keep in mind that the pairing is done only from one of the sites. This is because the pairing process establishes reciprocity by configuring the connection in the reverse direction as well. However, when you open the site recovery solution at the remote vCenter Server, you will be prompted to enter the administrator credentials of the other site.

Installing Storage Replication Adapters

Once you have the SRM instances installed and paired, the next step is to install the Storage Replication Adapters. SRAs are coded and provided by the storage vendors. VMware certifies the SRAs and posts them as compatible with the SRM.

Downloading SRA

The certified versions of SRA can be downloaded directly from VMware's website. Keep in mind that most vendors publish the updated versions of SRA at their website before it is certified by VMware. Since SRA is a vendor-supported component, you can choose to install the latest version available from the vendor, if that is known to fix a problem you are dealing with.

This is how you can download an SRA:

1. Go to VMware's website at `www.vmware.com`.

2. Navigate to the **vCenter Site Recovery Manager** option in **Downloads** under the **Product Downloads** category.

3. Once you are at the download page for vCenter SRM, click on the **Go to Downloads** hyperlink listed against SRA.

4. At the **Download Storage Replication Adapters for VMware vCenter Site Recovery Manager** page, you will see a list of all the certified SRAs. Click on the **Download Now** button corresponding to the needed SRA.

Installing SRA

Once downloaded, the SRA component has to be installed on both the sites. In most cases, the SRA installation is pretty simple and straightforward, but this can be different from vendor to vendor. You need to refer to the vendor documentation for the installation procedure.

Once the installation is complete, follow this procedure to discover the installed SRA component:

1. Connect to the protected/recovery site of vCenter Server using vSphere Client.
2. Navigate to the inventory home page and click on **Site Recovery**.
3. Click on **Array Managers** in the left pane and navigate to the **SRAs** tab.
4. In the **SRAs** tab, click on **Rescan SRAs** to discover the installed SRA.
5. Repeat the procedure at the recovery site as well.

Adding an array manager

Once you have the SRA installed and discovered at both the sites, you will need to add an array manager at both the sites. An array manager is required to discover the replicated LUNs and perform other storage operations initiated by SRM.

This is how you add an array manager:

1. Connect to the protected/recovery site of vCenter Server using vSphere Client.
2. Navigate to the vCenter Server's inventory home page and click on **Site Recovery**.

3. Click on **Array Managers** in the left pane, select the site, and click on
    Add Array Manager to bring up the **Add Array Manager** wizard,
    as shown in the following screenshot:
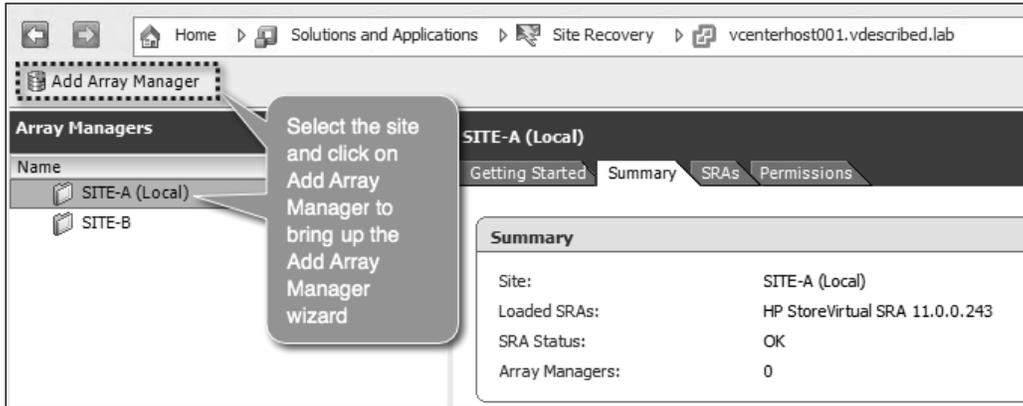
4. In the wizard, supply a **Display Name** for the array manager and an **SRA**
    Type. The SRA type value field will be prepopulated with the SRA that is
    already installed. Click on **Next** to continue.
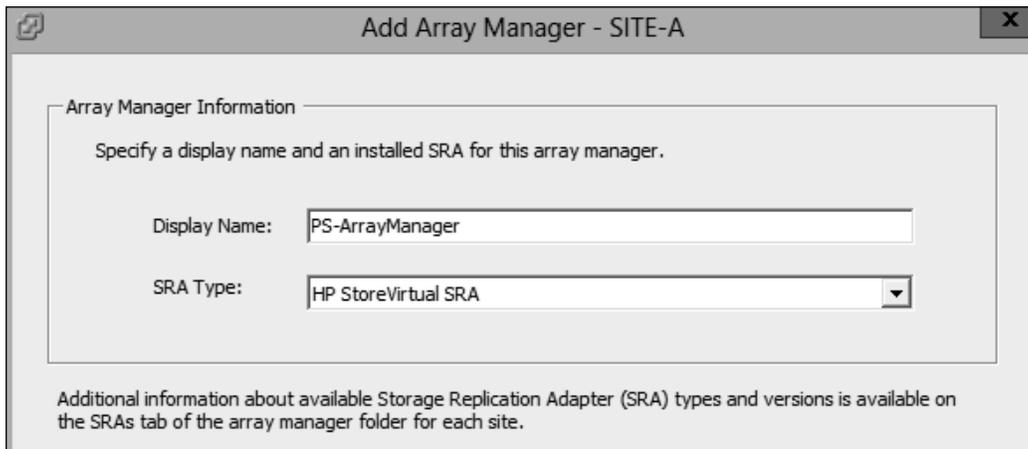
5. On the next screen, you will be prompted to enter the IP address of the storage node or a management server, which does the array management. Keep in mind that the information this screen prompts for differs from array to array and from vendor to vendor. It is purely dependent on what SRA is being used. Here in this example, I use the HP StoreVirtual (LeftHand) SRA and what I have entered as the Virtual IP (VIP) of the cluster the **Node Storage Module** (**NSM**) is part of. If I don't use this, all the NSMs in the cluster are involved in the replication of SRM, and then I can supply the IP addresses of the involved NSMs separated by a comma.

6. Supply the details and click on **Next** to continue. Refer to the following screenshot:

7. The next screen will read **Success** if the array managers are added successfully. Click on **Finish** to exit the wizard.

8. Repeat the same procedure on the recovery site as well.

9. Once done, the **Array Managers** for both the sites should be listed, as shown in the following screenshot:

Enabling an array pair

An array pair shows the replication relationship between two arrays. Before you enable an array pair, you need SRA installed and the array manager added at both the sites. For the array manager to detect an array pair, there should be a replication schedule already created between the arrays. Refer to the vendor documentation to understand what a replication schedule means for the vendor's array and the procedure to create it.

This is how you enable an array pair:

1. Make sure that there is a replication schedule enabled between the two arrays.

2. Navigate to the vCenter Server's inventory home page and click on **Site Recovery**.

3. Click on **Array Managers** in the left pane.

4. Select an added **Array Manager** (local or remote), and click on **Refresh** to discover an array pair, as shown in the following screenshot:
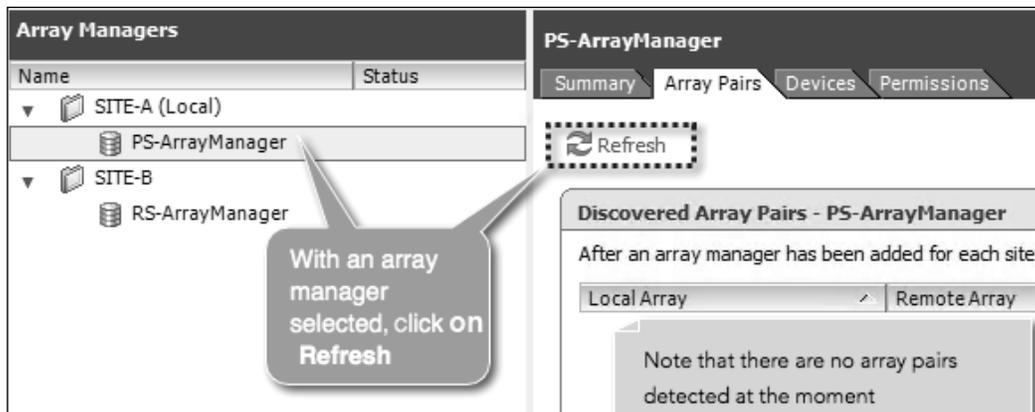
5. If the **Refresh** action discovers an array pair, the array pair will be listed. The **Refresh** action has to be done at both of the sites.

6. Array pairs discovered are not enabled by default. To enable an array pair, select an array pair and click on **Enable**. This operation has to done at only one of the sites.
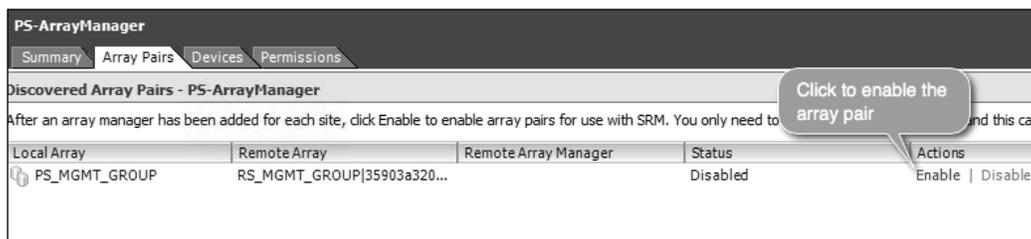
7. When an array pair is enabled, it tries to discover the devices (LUNs) for which a replication schedule is enabled at the array. Keep in mind that not all devices with a replication schedule are displayed as a device for the array pair; only the ones that are presented to a host at the protected site are displayed. To view the detected and filtered replication-enabled devices with the array manager selected, navigate to the **Devices** tab and click on **Refresh**, as shown in the following screenshot:
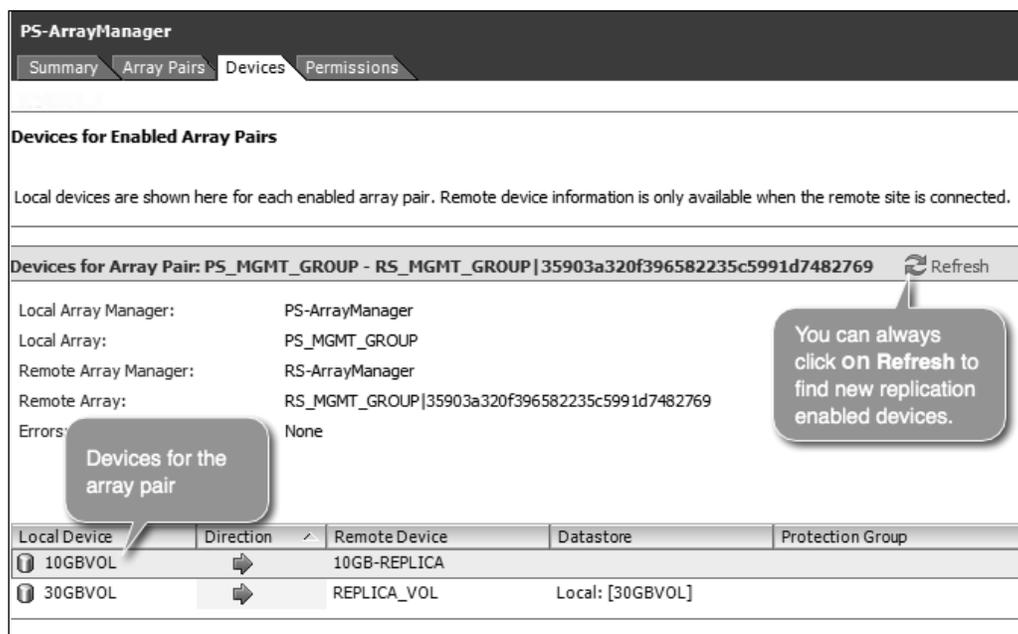
Configuring placeholder datastores

In the next chapter, you will learn how to create Protection Groups. For every virtual machine that becomes part of a Protection Group, SRM creates a shadow virtual machine. A placeholder datastore is used to store the files for the shadow virtual machines. The datastore used for this purpose should be accessible to all the hosts in the datacenter/cluster serving the role of a recovery-host. We will learn more about Protection Groups and shadow virtual machines in the next chapter. For now, understand that configuring placeholder datastores is an essential step in forming an SRM environment.

Assuming that each of these paired sites is geographically separated, each site will have its own placeholder datastore. The following figure shows the site and placeholder datastore relationship:
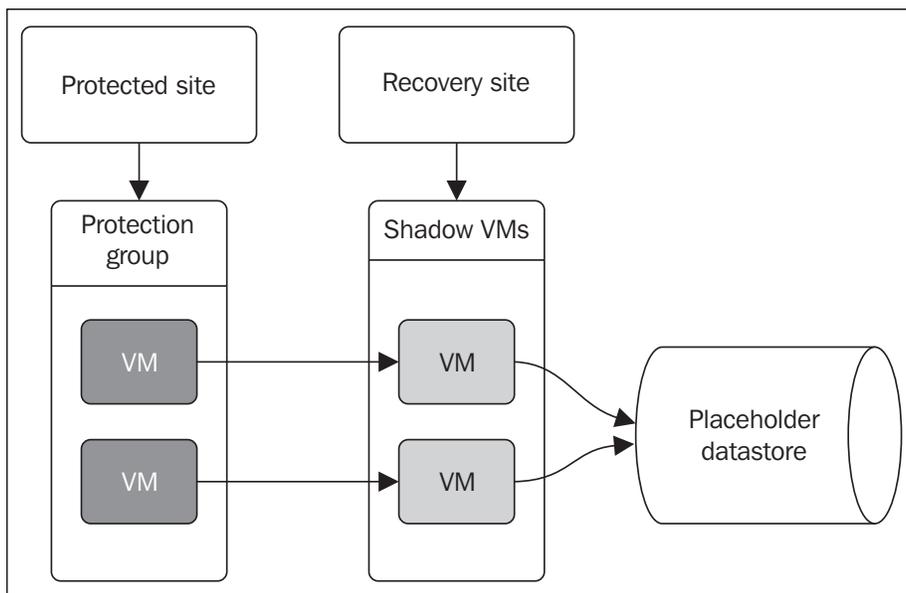
This is how you configure placeholder datastores:

1. Navigate to vCenter Server's inventory home page and click on **Site Recovery**.

2. Click on **Sites** in the left pane and select a site. Navigate to the **Placeholder Datastores** tab and click on **Configure Placeholder Datastore**, as shown in the following screenshot:
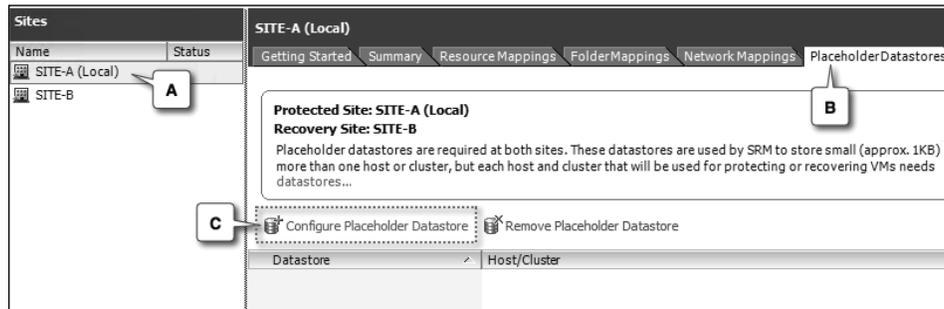
3. In the **Configure Placeholder Datastore** window, select an appropriate datastore and click on **OK**. To confirm the selection, exit the window.

4. Now, the **Placeholder Datastores** tab should show the configured placeholder. Refer to the following screenshot:
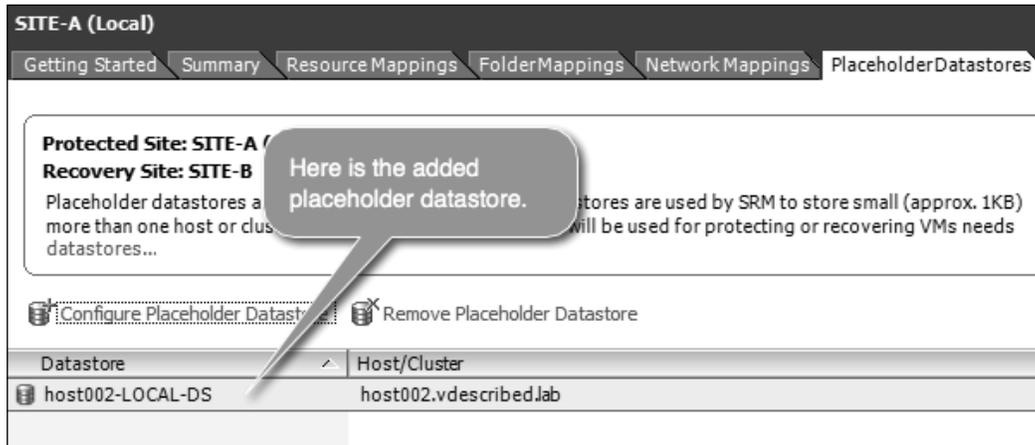
5. If you plan to configure a Failback, repeat the procedure in the recovery site.

Creating resource, folder, and network mappings

Creating resource, folder, and network mappings facilitates further orchestration of the Recovery Plan that will be executed for either a Planned Migration or a Failover. Without these mappings, you won't be able to configure protection on the virtual machines, and the protection status will indicate that these mapping are missing. We will learn more about Protection Groups in the next chapter.

Apart from being a requirement for creating Protection Groups, there are other use cases as well. The following table shows a few common ones:

Use cases	Mapping to use
If the designated recovery site runs other workload, then you may want to create a separate folder for the VMs from the protected site.	Folder mappings
If there is a separate cluster/resource pool at the recovery site to host, the VMs are recovered from the protected site.	Resource mappings
If there are vSwitch/DSwitch port groups at the recovery site for the recovered VMs, we use the network mapping.	Network mappings

Resource mappings

We need to provide a correlation between the compute resource containers on both the sites. The compute resource containers are cluster, resource pool, and ESXi host. This is achieved with the help of resource mappings.

Resource mappings respect the presence of these containers, which means that if there is a cluster or resource pool at the site, the ESXi hosts are not made available as a selectable compute container.

This is how you configure resource mappings:

1. Navigate to vCenter Server's inventory home page and click on **Site Recovery**.

2. Click on **Sites** in the left pane, select a site, and navigate to the **Resource Mappings** tab. Select the resource container (a cluster, resource pool, or host) you want to map, and click on **Configure Mapping** to bring up the **Mapping** window. Refer to the following screenshot:
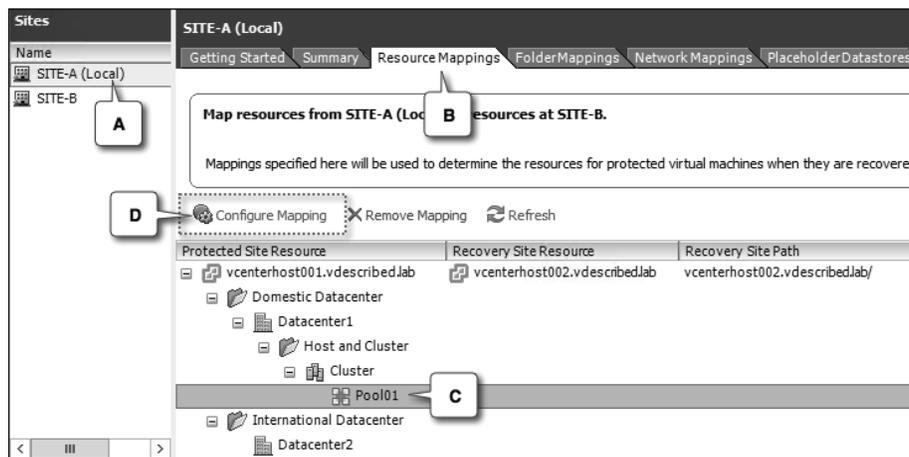
3. In the **Mapping** window, browse the resource inventory of the recovery site, select the destination resource container (a cluster, resource pool, or host), and click on **OK** to confirm.

---
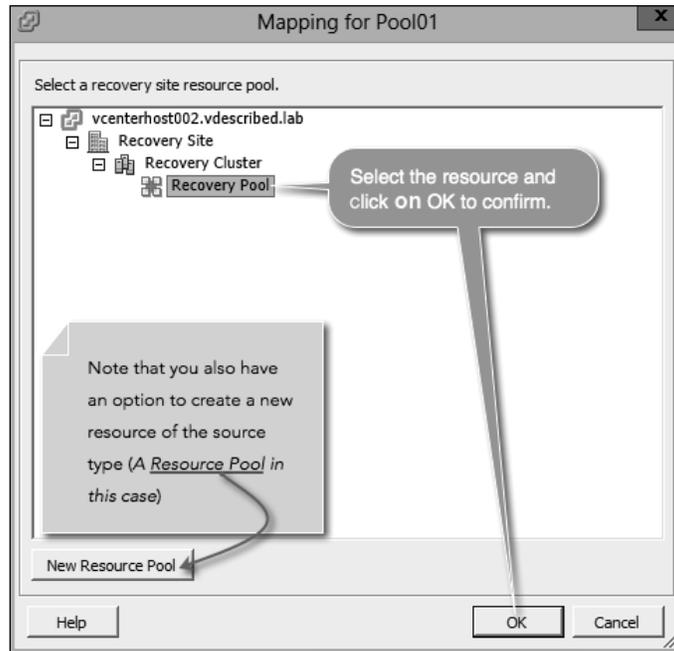
4. The **Resource Mapping** tab should now show the mapped **Recovery Site Resource**.
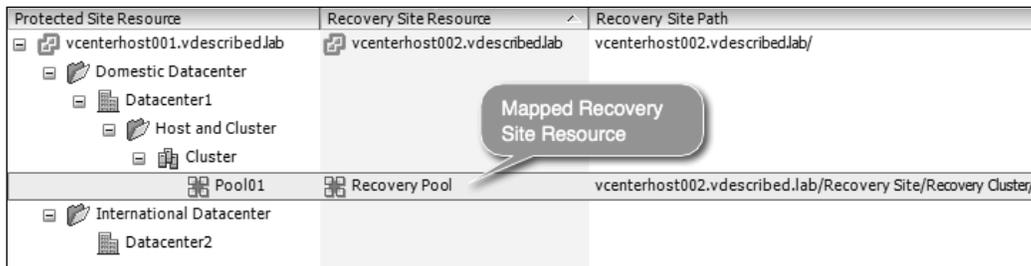
Folder mappings

Folders are inventory containers that can only be created using vCenter Server. They are used to group inventory objects of the same type for easier management.

There are different types of folders. The folder type is determined by the inventory-hierarchy level they are created at. The folder names are as follows:

- Datacenter folder
- Hosts and clusters folder
- Virtual machine and template folder
- Network folder
- Storage folder

The vSphere Web Client provides UI menu options to create a folder of the following types, without needing to navigate to an appropriate inventory-hierarchy level to create them:

- Hosts and clusters folder
- Network folders
- Storage folder
- Virtual machine and template folder

In the case of SRM folder mappings, we will deal with only virtual machine folders and its parent datacenter. You will not be able to configure mapping for any of the other folder types.

This is how you configure folder mapping:

1. Navigate to vCenter Server's inventory home page and click on **Site Recovery**.

2. Click on **Sites** in the left pane, select a site, and navigate to the **Folder Mappings** tab. Select the virtual machine folder that you want to map, and click on **Configure Mapping** to bring up the **Mapping** window, as shown in the following screenshot:
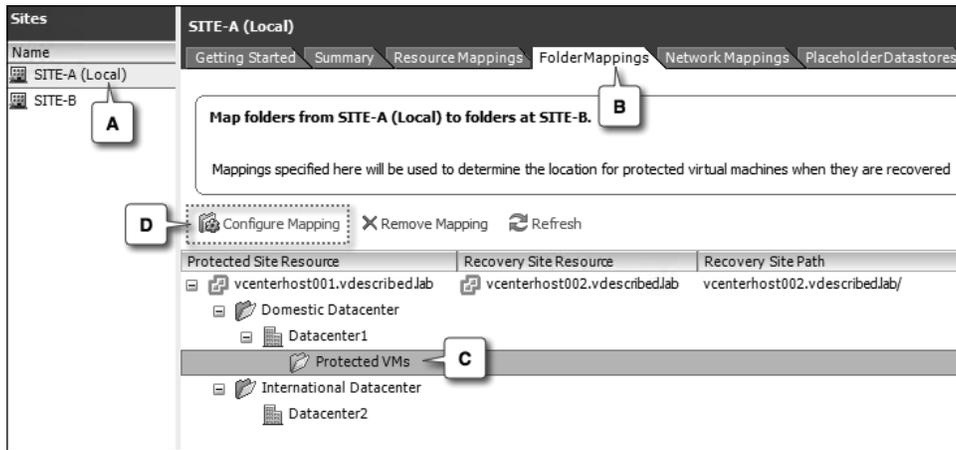
3. In the **Mapping for Protected VMs** window, browse the virtual machine folder inventory of the recovery site, select the destination folder, and click on **OK** to confirm.
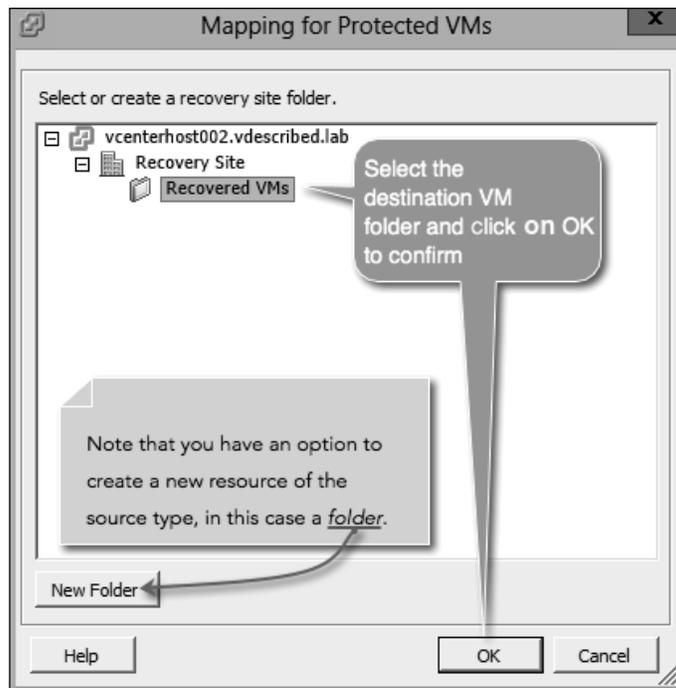
> It is important to make sure that the parent datacenter of which the virtual machine folder is part of is mapped as well. The procedure to configure mapping is the same.

4. The **Folder Mapping** tab should now show the mapped recovery site folder. Refer to the following screenshot:
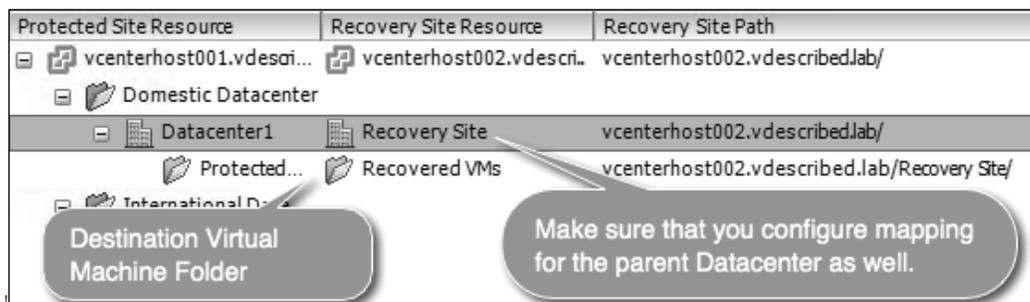
Network mappings

Network configuration at the protected and recovery sites need not be identical. Network mappings provide a method to form a correlation between the port groups (standard or distributed) of the protected and recovery steps.

Let's say we have a port group with the name VM Network at the protected site, and it is mapped to a port group with the name Recovery Network at the recovery site. In this case, a virtual machine that is connected to VM Network will be reconfigured to use the Recovery Network when failed over.

This is how you configure network mappings:

1. Navigate to vCenter Server's inventory home page and click on **Site Recovery**.

2. Click on **Sites** in the left pane, select a site, and navigate to the **Network Mappings** tab. Select the port group (standard/distributed) that you want to map, and click on **Configure Mapping** to bring up the **Mapping** window, as shown in the following screenshot:
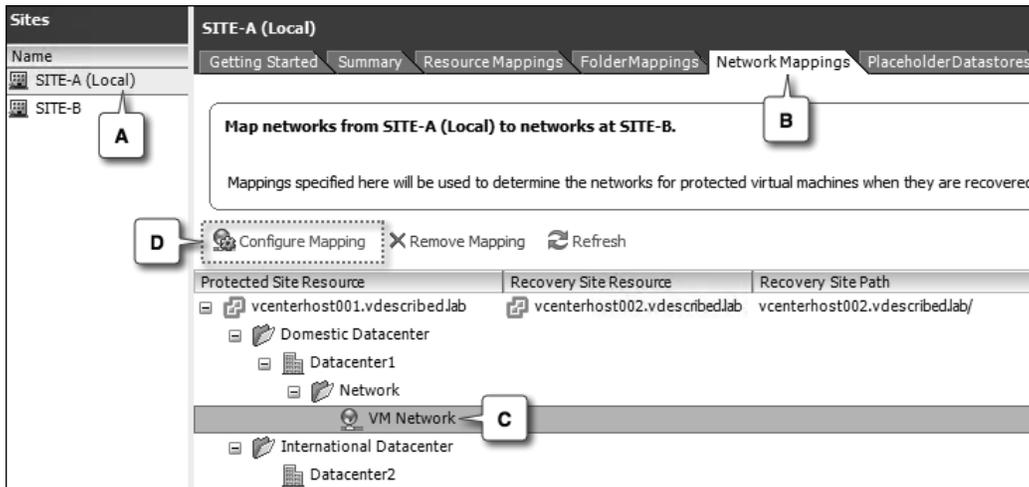
3. In the **Mapping for VM Network** window, browse the network inventory of
    the recovery site, select the destination port group, and click on **OK** to confirm.
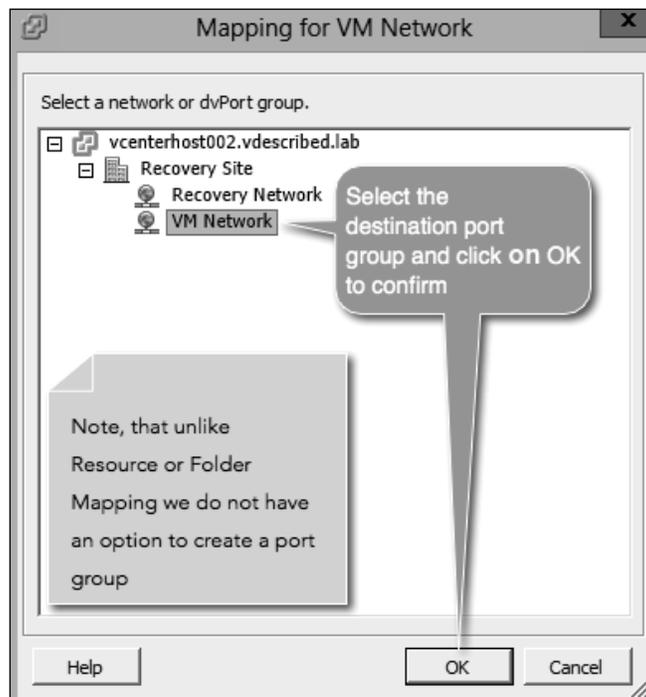
4. The **Network Mappings** tab should now show the mapped recovery site port group, as shown in the following screenshot:

Virtual machine swap file location

With SRM implementations, there is a common argument about the placement of the virtual machine swap files. Some would suggest maintaining a separate datastore for the virtual machine swap files, while some are against it. Before we try to understand the rationale behind either of the design choices, it is important to know what a virtual machine swap file is.

Every virtual machine will have a swap file (`.vswp`). This swap file is created every time a virtual machine is powered on. The size of the swap file is equal to the size of the memory assigned to the virtual machine, unless there is a reservation. If there is a memory reservation, then the size of the swap file will be equal to the size of the unreserved memory. Although rare, some environments use limits on memory as well.

So, the ideal formula to calculate the size of the swap file is as follows:

Swap file size = memory limit – memory reservation

The default memory reservation is 0 MB, and the default limit is equal to the configured size of the memory. By default, the swap file is stored along with the virtual machine in its working directory.

Design choice 1 – separate datastore for the swap files

The rationale is that the swap file is created every time a virtual machine is powered on. Since the VM will be powered on at the recovery site, the swap file will be created at that time. Hence, there is no need to replicate the swap files. The following table illustrates the pros and cons of this:

Pros	Cons
Swap file replication, if avoided, can reduce the bandwidth utilization for storage replication.	Single point of failure.
Reduces the need for the storage space at the recovery site, which otherwise would be needed for the swap files.	The swap location should be chosen at a per-host level; this would mean a lot of manual work in a large environment.
	Need to accommodate a separate large LUN; this could affect the available spare capacity of the array.

Design choice 2 – store the swap files in the virtual machines' working directory

The rationale is that apart from reduced replication bandwidth usage, there is no real advantage of maintaining a separate datastore for the swap files. Most SRM implementations will already have made sure that there will be more than enough bandwidth to make storage replication feasible. Also, not all virtual machines frequently use the swap files, unless the vSphere environment is oversubscribed and the virtual machines are frequently contending for memory resources. In most cases, the swap files will be replicated during the initial sync. Subsequent synchronizations will include swap files created consequent to power off and power on operations. Keep in mind that a guest OS reboot will not trigger the recreation of the swap files. The following table illustrates the pros and cons of this:

Pros	Cons
No administrative overhead, which would otherwise be needed to configuring a swap datastore per host.	Bandwidth wastage, due to the replication of the swap files.
No single point of failure.	Space wastage at the recovery site, which can otherwise be avoided if the swap files are not duplicated onto the replica LUNs.

> The design choices and the rationale behind them can vary depending on the environment you are dealing with. The rationales are only guidelines.

Summary

In this chapter, we learned what VMware vCenter SRM is and how it can be installed and configured to lay the groundwork for any SRM environment. In the next chapter, we will learn how to enable protection of the virtual machine workload by creating Protection Groups and Recovery Plans.

Where to buy this book

You can buy Disaster Recovery Using VMware vSphere Replication and vCenter Site Recovery Manager from the Packt Publishing website:
`http://www.packtpub.com/vsphere-disaster-recovery-using-srm-and-vsphere-replication/book`.

Free shipping to the US, UK, Europe and selected Asian countries. For more information, please read our shipping policy.

Alternatively, you can buy the book from Amazon, BN.com, Computer Manuals and most internet book retailers.

www.PacktPub.com

For More Information:
www.packtpub.com/vsphere-disaster-recovery-using-srm-and-vsphere-replication/book