VMware Site Recovery Manager 5.x guidelines for the IBM Storwize family

A step-by-step guide

IBM Systems and Technology Group ISV Enablement

February 2014

@IBMSystemsISVs

Table of contents

Abstract

This paper offers detailed configuration information regarding the IBM Storwize family in the VMware virtual infrastructure environment with VMware vCenter Site Recovery Manager. The purpose of this paper is to set appropriate expectations for customers with regard to high availability (HA) options and automated failover using VMware 5.x vCenter Site Recovery Manager.

Introduction

This paper describes how to implement VMware vCenter Site Recovery Manager when used with the IBM® Storwize® family. This paper helps readers to install and configure Site Recovery Manager with VMware HA and understand how to run a disaster recovery (DR) plan and a disaster recovery test. In addition, readers can have a conceptual understanding of what is required in a true disaster recovery scenario, which is much broader than just failover of the virtual infrastructure and storage environment.

Intended audience

This technical report is intended for:

- Customers and prospects looking to install and use VMware Site Recovery Manager with the IBM Storwize family storage system
- Users and management seeking information to install and use VMware Site Recovery Manager with the IBM Storwize family storage system

Scope

This technical report provides:

- Detailed deployment information on how to effectively install and use VMware Site Recovery Manager with the IBM Storwize V7000 storage system, which is applicable to the complete Storwize family
- Detailed design and implementation guide and configuration best practices

This technical report does not:

- Discuss any performance impact and analysis from a user perspective
- Replace any official manuals and documents from IBM and VMware on the products used in the solution

Prerequisites

This technical paper assumes familiarity with the following prerequisites:

- Basic knowledge of the following VMware virtualization technologies and products:
  - VMware vCenter Server 5.x
  - VMware vSphere ESXi 5.x
  - VMware vCenter Site Recovery Manager 5.x
- Basic knowledge of the IBM Storwize family storage system

Overview of a disaster recovery solution

As IT systems have become increasingly critical to the smooth operation of companies, the importance of ensuring their continued operation, including rapid recovery when subsystems fail, has increased.

Before selecting a disaster recovery strategy, a disaster recovery planner should refer to their organization's business continuity plan, which should indicate the key metrics of recovery point objective (RPO) and recovery time objective (RTO) for various business processes. The metrics specified for the business processes must then be mapped to the underlying IT systems and infrastructure that support those processes. After mapping the RTO and RPO metrics to IT infrastructure, the disaster recovery planner can determine the most suitable recovery strategy for each system.

RPO – This refers to the point in time to which data must be recovered as defined by the organization. This is generally a definition of what an organization determines as an acceptable loss in a disaster situation.

RTO – This refers to the duration of time and a service level within which a business process must be restored, following a disaster or disruption, to avoid unacceptable consequences associated with a break in business continuity.

The ideal solution will have both: a low RPO (in minutes) and RTO (ranges from minutes to hours). It is important to test a DR solution to find whether it is suitable and efficient for business continuity.

Benefits of implementing VMware vCenter Site Recovery Manager with the Storwize family

Implementing a virtualized environment using the VMware technology and VMware vCenter Site Recovery Manager on the Storwize family storage systems provides the infrastructure for unique opportunities to implement real working disaster recovery processes that are quick and easy to test, consume less additional storage, and significantly reduce RTO and RPO duration.

VMware vCenter Site Recovery Manager is a market leading disaster recovery management product. It ensures a simple, one of the most affordable and reliable disaster protections for all virtualized applications. Site Recovery Manager uses cost-efficient VMware vSphere Replication or storage-based replication to provide centralized management of recovery plans, enable non-disruptive testing, and automate site recovery and migration processes.

The disaster recovery solution combined with the IBM Storwize family and VMware vCenter Site Recovery Manager, helps to meet RTOs, reduces the costs traditionally associated with business continuance plans, and achieves low-risk and predictable results for recovery of a virtual environment. Site Recovery Manager integrates tightly with VMware vSphere, VMware vCenter Server, and storage replication software from IBM to process site failover to recover rapidly, reliably, and affordably. It eases disaster recovery risk and protects all business-critical systems and applications.

Key benefits

The key benefits of implementing VMware vCenter Site Recovery Manager are:

- Replace traditional and error-prone manual runbooks with simple and automated recovery plans.
- Enable frequent non-disruptive testing of recovery plans to ensure that they meet business requirements.
- Automate site recovery and migration processes to ensure fast and reliable recovery.
- Streamline planned migrations and preventive failovers.
- Choose among a broad set of replication options. Use VMware vSphere Replication for affordable replication, or storage-based replication for large, business-critical environments.

Overview of VMware vCenter Site Recovery Manager

VMware vCenter Site Recovery Manager is a business continuity and disaster recovery solution that helps you plan, test, and run a scheduled migration or emergency failover of data center services from one site to another.

Site Recovery Manager is an extension to VMware vCenter that enables integration with array-based replication, discovery, and management of replicated data stores, and automated migration of inventory from one vCenter to another. Site Recovery Manager servers coordinate the operations of the replicated storage arrays and vCenter servers at two sites so that when virtual machines (VMs) at the protected site are shut down, virtual machines at the recovery site are started up and use the data replicated from the protected site to assume responsibility for providing the same services. Transfer of 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 and the allocation of computer resources and the networks that might be accessed. Site Recovery Manager allows you to test a recovery plan, using a temporary copy of the replicated data, in a manner that does not disrupt the ongoing operations at site. Figure 1 shows the VMware vCenter Site Recovery Manager architecture.

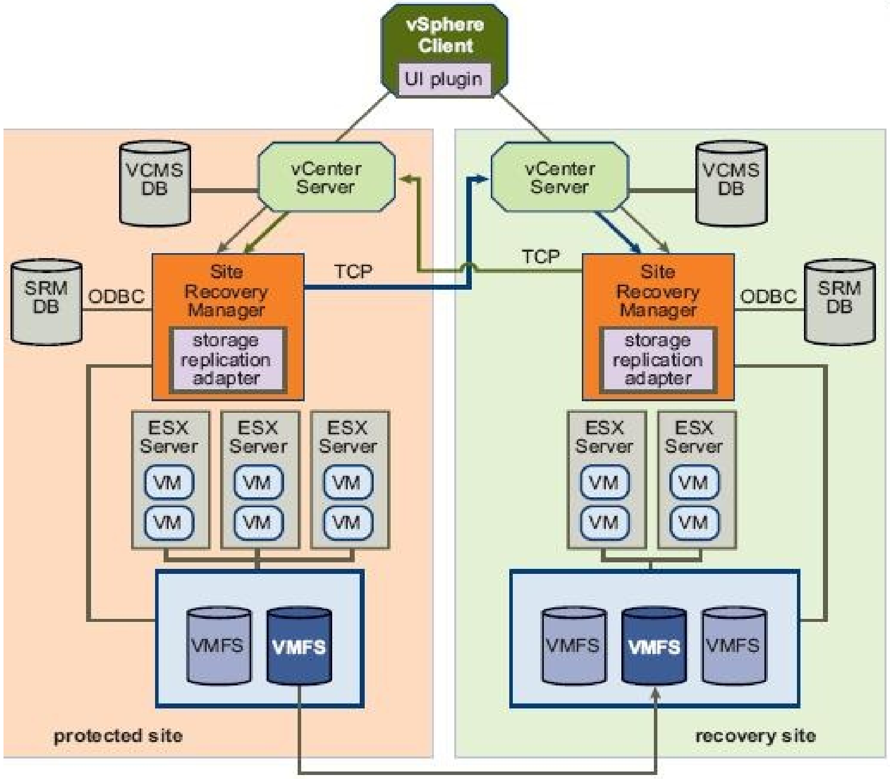

Figure 1: VMware vCenter Site Recovery Manager architecture

Protected and recovery sites

In a typical Site Recovery Manager installation, the protected site provides business-critical data center services, and the recovery site provides an alternative facility to which these services can be migrated.

The protected site can be any site where virtual infrastructure supports a critical business need. The recovery site can be located thousands of miles away, or in the same site. In the typical case, the recovery site is in a facility that is unlikely to be affected by any environmental, infrastructure, or other disturbances that affect the protected site. Site Recovery Manager has the following requirements for the vSphere configuration at each site:

- Each site must include at least one vSphere data center.
- The recovery site must support array−based replication with the protected site, and must have hardware and network resources that can support the same virtual machines and workloads as the protected site.
- One virtual machine must be located on a replicated data store at the protected site. This data store must be supported by a storage array that is compatible with Site Recovery Manager.
- The protected and recovery sites should be connected by a reliable IP network. Storage arrays might have additional network requirements.
- The recovery site should have access to the same public and private networks as the protected site, though not necessarily the same range of network addresses.

Storage Replication Adapter

Storage Replication Adapter (SRA) is an interface between the storage subsystem and Site Recovery Manager. IBM Storwize family SRA provides a means for VMware Site Recovery Manager to control the IBM Storwize family features, such as replication and IBM FlashCopy® without VMware Site Recovery Manager requiring any awareness of the storage system. The Storwize family SRA can be downloaded from the VMware website and installed on the server running Site Recovery Manager.

Site Recovery Manager database

The Site Recovery Manager requires its own database for storing data such as recovery plans and inventory information. Depending upon the version of Site Recovery Manager, this can be either Microsoft® SQL Server database or Oracle Server database or IBM DB2® Server database. You can find specific interoperability information at:
http://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php?

The Site Recovery Manager server requires its own database to store recovery plans, inventory information, and other associated data. Before installing the Site Recovery Manager server, you must configure and initialize the databases for the Site Recovery Manager server.

The Site Recovery Manager database at each site holds information about virtual machine configurations, protection groups, and recovery plans. Site Recovery Manager cannot use the vCenter Server database because it has different database schema requirements. The database must exist before Site Recovery Manager can be installed. If the Site Recovery Manager database at either site becomes corrupted, the Site Recovery Manager servers at both the sites will shut down. Each Site Recovery Manager site requires its own instance of the Site Recovery Manager database.

For the lab solution testing purpose, Microsoft SQL Server has been configured for the Site Recovery Manager database.

IBM storage systems

IBM provides a range of storage systems designed to meet challenges. IBM storage systems support many advanced storage functionalities such as IBM Real-time Compression™, automated tiering, storage virtualization, and thin provisioning. These advanced functions, combined with the performance and reliability expected of an IBM solution, result in better system performance and lower IT costs.

IBM Storwize family

The IBM Storwize family consists of a set of virtualized storage systems designed to reduce cost, simplify management, provide cost saving advanced functionality, and provide high scalability. To meet this challenge, IBM offers a range of storage systems running on the same core Storwize code base.

- **IBM System Storage SAN Volume Controller**
  IBM System Storage® SAN Volume Controller (SVC) provides a modular and highly scalable storage virtualization solution. SVC allows customers a single point of management for all the storage in their IT infrastructure. SVC also augments the virtualized storage systems by providing advanced storage features such as Real-time Compression and automated tiering, to name a few. For additional information about IBM SVC, refer to the following URL:
  ibm.com/systems/storage/software/virtualization/svc/

- **IBM Storwize V7000**
  The IBM Storwize V7000 storage system provides block storage enhanced with enterprise-class features to midrange customer environments. The Storwize V7000 system can scale up to 240 drivers per control enclosure. Additionally, up to four control enclosures can be clustered, allowing the Storwize V7000 system to scale up to 906 drives. For additional information about the IBM Storwize V7000 system, refer to the following URL:
  ibm.com/systems/storage/disk/storwize_v7000/index.html

- **IBM Storwize V7000 Unified**
  The IBM Storwize V7000 storage system combines the block storage capabilities of the Storwize V7000 system with advanced file storage capabilities to form a single system for greater ease of management and efficiency. The IBM Storwize V7000 Unified system contains the same 2U drive enclosures of the Storwize V7000 system and two 2U file modules. For additional information about the IBM Storwize V7000 Unified, refer to the following URL:
  ibm.com/systems/storage/disk/storwize_v7000/index.html

- **IBM Flex System V7000 Storage Node**
  IBM Flex System™ V7000 Storage Node is a high-performance block-storage solution that has been designed to integrate directly with IBM Flex System. IBM Flex System V7000 Storage Node provides advanced storage capabilities, such as IBM System Storage Easy Tier®, IBM Real-Time Compression, thin provisioning, and more. For more information the IBM Flex System V7000 Storage Node, refer to the following URL:
  ibm.com/systems/flex/storage/v7000/index.html

- **IBM Storwize V5000**
  The IBM Storwize V5000 system provides cost-efficient midrange storage. Built from the same

technology as the IBM SAN Volume Controller and IBM Storwize V7000 systems, the IBM Storwize V5000 system offers advanced storage features in a cost-conscious package. For additional information about the IBM Storwize V5000 system, refer to the following URL: **ibm.com**/systems/hk/storage/disk/storwize_v5000/index.html

- **IBM Storwize V3700**
  The IBM Storwize V3700 system is an entry-level storage system designed for ease of use and affordability. Built from the same technology used in all of the Storwize family, the Storwize V3700 system offers some of the advanced features that can be found in other Storwize models. For more information about the IBM Storwize V3700 system, refer to the following URL: **ibm.com**/systems/storage/disk/storwize_v3700/index.html

Using IBM Storwize family and VMware Site Recovery Manager for disaster recovery

VMware Site Recovery Manager accelerates and ensures reliable recovery and simplifies disaster recovery through the following functions:

- Automation
- Non-disruptive testing
- Eliminating complex manual-recovery steps
- Centralizing management of recovery plans

With Site Recovery Manager, you can create, update, and document disaster recovery plans to meet RTOs, RPOs, and compliance requirements. Site Recovery Manager enables you to run automated tests of recovery plans without disrupting the IT environment.

IBM Storwize family Copy Services features

The IBM Storwize family includes the following Copy Services features used by Site Recovery Manager:

- **FlashCopy** - Creates an instant, point-in-time copy from a source to a target volume.
- **Metro Mirror** - Provides a consistent copy of a source volume on a target volume. Data is written to the target volume synchronously as it is written to the source volume, so that the copy is continuously updated.

Note: You need to set up FlashCopy on the recovery site and Metro Mirror on the IBM Storwize family at both sites for use with VMware Site Recovery Manager.

FlashCopy

The FlashCopy function copies the contents of a source volume to a target volume. Any data that existed on the target volume is lost and is replaced by the copied data. After the copy operation has been completed, the target volumes contain the contents of the source volumes as they existed at a single point in time. The FlashCopy function is sometimes described as an instance of a **time-zero copy** (T 0) or **point-in-time copy technology**. Although a FlashCopy operation takes some time to complete, the resulting data on the target volume is presented so that the copy appears to have occurred immediately.

Metro Mirror

The Metro Mirror Copy Services features enable you to set up a relationship between two volumes, so that updates made to one volume are mirrored to the other volume. The volumes can be in the same cluster or on two different clusters.

An application needs to send only writes to a single volume and the system maintains two copies of the data. If the copies are separated by a significant distance, the Metro Mirror copies can be used as a backup for disaster recovery. A prerequisite for Metro Mirror operations between clusters is that the SAN fabric to which they are attached provides adequate bandwidth between the clusters.

For Metro Mirror, one volume is designated as the primary and the other volume is designated as the secondary. Host applications write data to the primary volume, and updates to the primary volume are copied to the secondary volume. Ordinarily, host applications do not perform I/O operations to the secondary volume.

The Metro Mirror feature provides a synchronous-copy process. When a host writes data to the primary volume, it does not receive confirmation of I/O completion until the write operation has completed for the write to both the primary volume and the secondary volume. This ensures that the secondary volume is always up-to-date with the primary volume in the event that a failover must be performed. However, the host I/O is subject to the latency and bandwidth limitations of the communication link to the secondary volume.

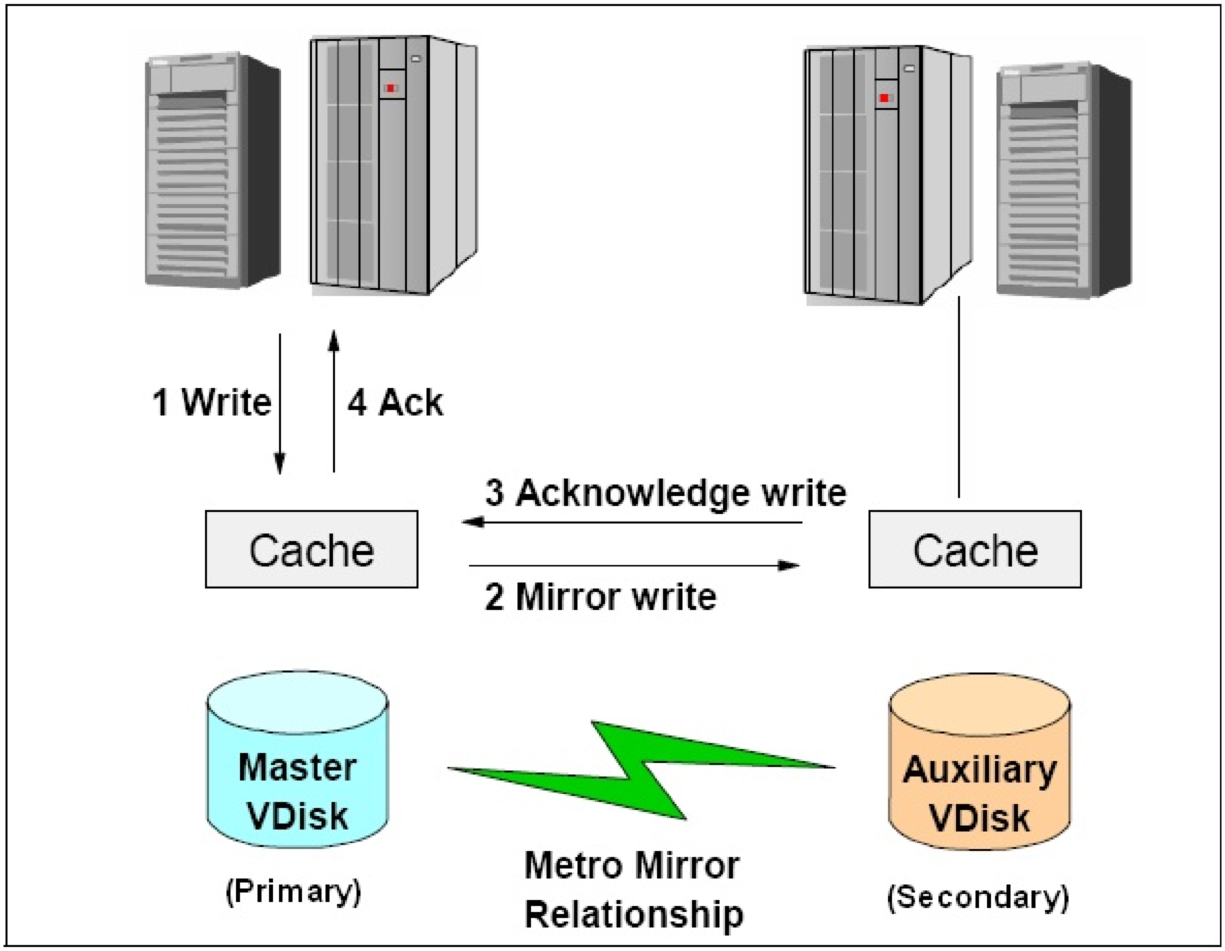

Figure 2: Write on volume in Metro Mirror relationship

Figure 2 depicts how a write operation to the master VDisk is mirrored to the cache of the auxiliary VDisk before an acknowledgement of the write operation is sent back to the host that issued the write. This process ensures that the auxiliary VDisk is synchronized in real time, in case it is needed in a failover situation.

The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km.

When the application performs a write update operation to a source volume at the production site, the following actions occur:

1. An application requests a write operation to the source volume. The write operation is written into cache at the production site.
2. Production site, send the write operation to the target cache and at the recovery site.
3. The Storwize family array at the recovery site signals that the write operation has been completed when the updated data is in its cache.
4. When Storwize family array at the production site receives notification from the target storage at the recovery site that the write operation has been completed, it returns the I/O complete status to your application.

IBM Storwize family partnership and replication

To use the Copy Services found in the IBM Storwize family, the Storwize storage systems need to be connected to each other in what is termed as a partnership. With the Storwize 7.2 code release, there are now two options for enabling this partnership.

Fibre Channel partnership for remote copy

The IBM Storwize family supports partnerships over Fibre Channel (FC). To facilitate long distance replication, Fibre Channel over IP (FCIP) bridging can be used. The lab demonstration uses this type of parternship.

IP Partnership for remote copy

Extending an FC network across remote distances can be very expensive. Storwize 7.2 and later offers native IP replication as an alternative. The Storwize storage system generates the IP frames that can be directly transmitted over Ethernet networks. This allows Storwize storage system to establish partnerships over native IP links without the use of an FCIP device. Native IP replication supports both Metro Mirror and Global Mirror.

The lab demonstration does not use IP replication. You can find information regarding the best practices in configuring a native IP partnership at:

ibm.com/partnerworld/wps/servlet/ContentHandler/stg_ast_sto_wp_configuring-system-storage-svc/lc=en_ALL_ZZ

Solution architecture

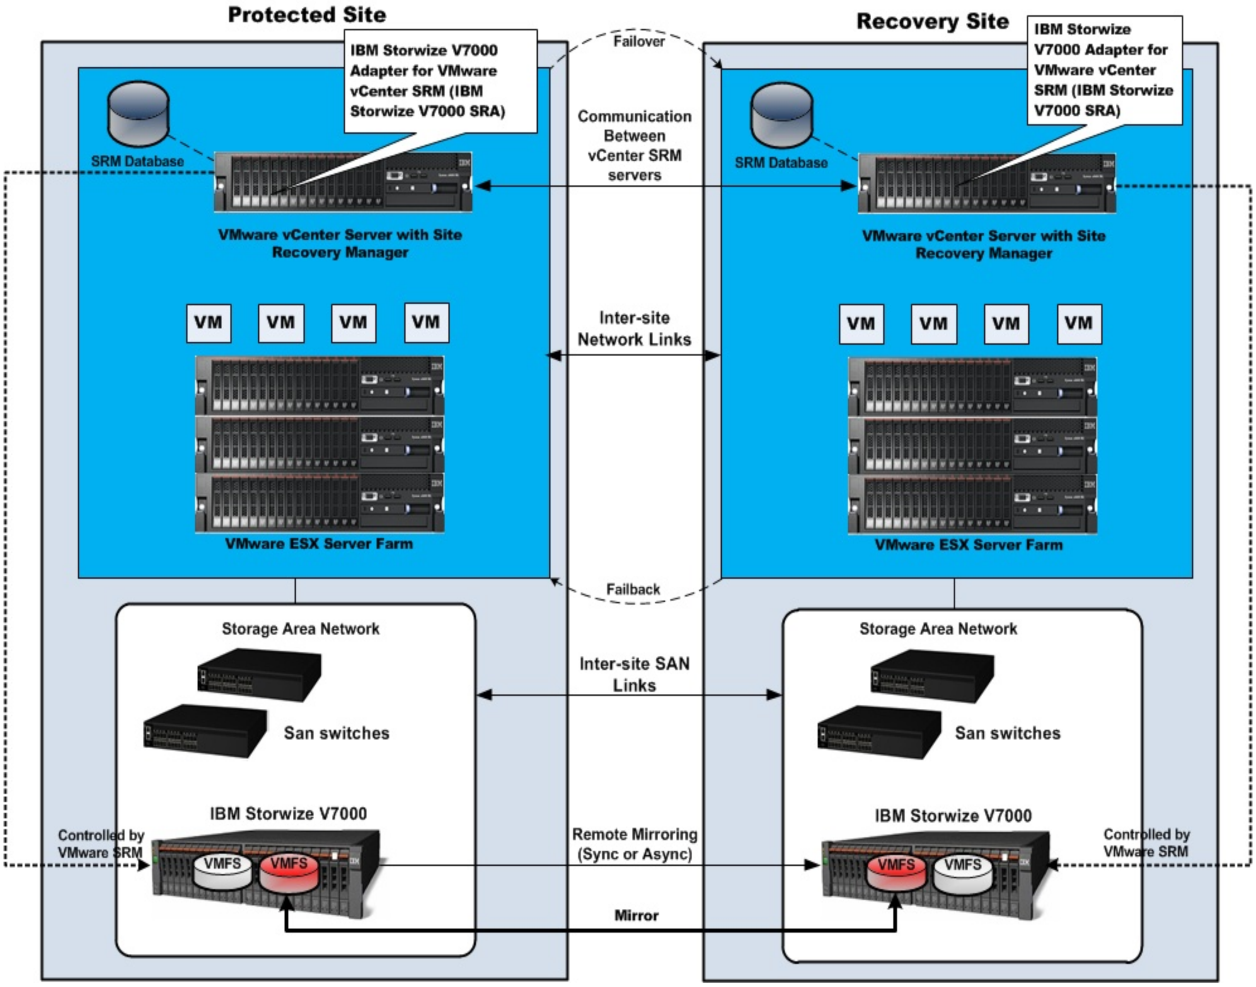

Figure 3: Solution architecture

Figure 3 depicts the basic architecture of the disaster recovery solution that is made up of IBM Storwize V7000 and VMware Site Recovery Manager for enterprise IT virtual infrastructure. This solution is build using VMware ESXi 5.x hosts, VMware vCenter servers, VMware Site Recovery Manager servers, and the IBM Storwize V7000 storage systems on both protected and recovery sites.

In the lab solution validation environment, for the solution validation purpose, the VMware ESXI 5.x hosts are configured with vSphere HA cluster. Protected and recovery sites are active with VMware ESXi hosts running Microsoft Windows® and Linux® VMs. The IBM Storwize V7000 Metro Mirror type Remote Copy replication has been configured and used for synchronous replication between protected site and recovery site.

Material list of solution setup in the lab

Table 1 lists the hardware and software used in the lab test solution setup.

Infrastructure component	**Vendor**	**Quantity**	**Details**
Server running VMware ESXi 5.x	IBM (IBM System x® 3650 M3)	4	
Server running VMware SRM server	IBM (IBM System x 3650 M3)	2	
Server running VMware vCenter server	IBM (IBM System x 3650 M3)	2	
Storage system	IBM	2	IBM Storwize V7000 Unified
			ibm.com/systems/storage/disk/storwize_v7000/index.html
SAN switch	IBM	4	IBM System Storage SAN24B-24
			ibm.com/systems/networking/switches/san/b-type/san24b-4/express/specifications.html
Network (Ethernet) switch	Cisco Catalyst 6509	1	
Software	IBM		IBM Storwize V7000 control enclosure version 7.1.0.1
	IBM		IBM System Storage SAN24B-24 software version 7.0.2b
	VMware		VMware vSphere ESXi 5.x
	VMware		VMware vSphere vCenter Server 5.x
	VMware		VMware vSphere Site Recovery Manager 5.x

Table 1: Hardware and software used in the lab test solution setup

Lab solution test configuration summary

This section provides a brief summary of the lab solution test setup:

1. Zone server to storage systems.

   Protected site:

   Servers: Isvsonas15 and isvsonas16 (IBM System x 3650 M3)

   Storage system: isv7k4 (IBM Storwize V7000)

   Recovery site:

   Servers: isvsonas19 and isvsonas34 (IBM System x 3650 M3)

   Storage system: isv7k5 (IBM Storwize V7000)

2. Ensure that a zone containing all the storage system host ports exists to support storage systems (isv7k4 and isv7k5) partnership.
3. Create server hosts on storage systems.
4. Create storage pools, managed disks (MDisks), and storage volumes on storage systems.

   Pools:

   Isv7k4: SRM_RAID5_Prot (protected site)

   Isv7k5: SRM_RAID5_Recov (recovery site)

   MDisks: RAID5, 2x7 HDDs 300 GB, 10 rpm 3.2TB (net)

Volumes (protected site):	isv7k4:	SRM_1_Prot 500 GB
		SRM_2_Prot 500 GB
		SRM_3_Prot 2TB
		SRM_4_Prot 100 GB
Volumes (Recovery site):	isv7k5:	SRM_1_Recov 500 GB
		SRM_2_ Recov 500 GB
		SRM_3_ Recov 2TB
		SRM_4_ Recov 100 GB

5. Map storage volumes to server hosts on storage systems.
6. Create partnership for storage systems.
7. Create Metro Mirror relationship for storage volumes between storage systems.
8. Place Metro Mirror volumes in a consistency group.

Note: For lab solution validation purpose, the test team mounted iSCSI-based Virtual Machine File System (VMFS) volume. The raw device mapping (RDM) volumes are also supported.

Zone servers to storage systems

You need to perform the following steps in the lab to zone the servers to the IBM Storwize V7000 storage systems.

1. Assign appropriate alias to the host ports using the SAN management tool, as shown in Figure 4.

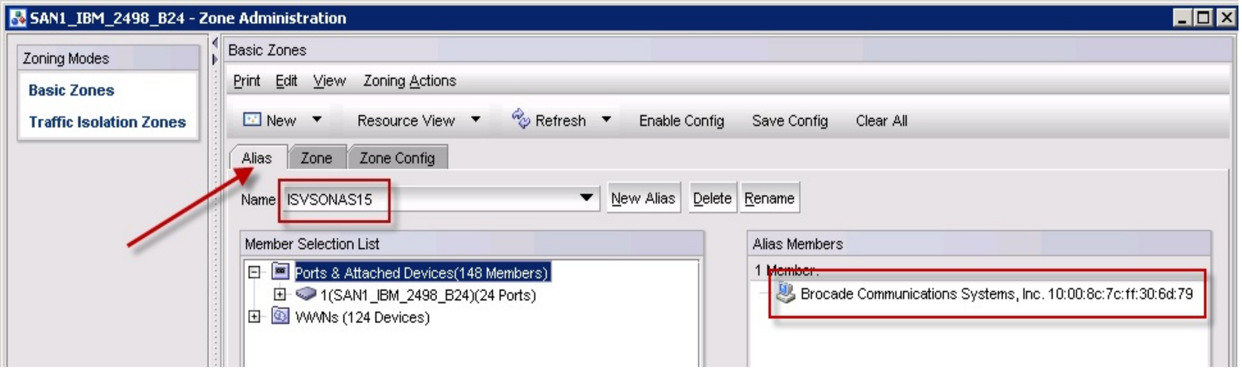

Name: ISVSONAS15

Member Selection List:
- Ports & Attached Devices(148 Members)
  - 1(SAN1_IBM_2498_B24)(24 Ports)
- WWNs (124 Devices)

Alias Members: 1 Member.
- Brocade Communications Systems, Inc. 10:00:8c:7c:ff:30:6d:79

Figure 4: Assign appropriate alias to the host ports

2. Assign appropriate alias to the IBM Storwize V7000 storage ports, as shown in Figure 5.

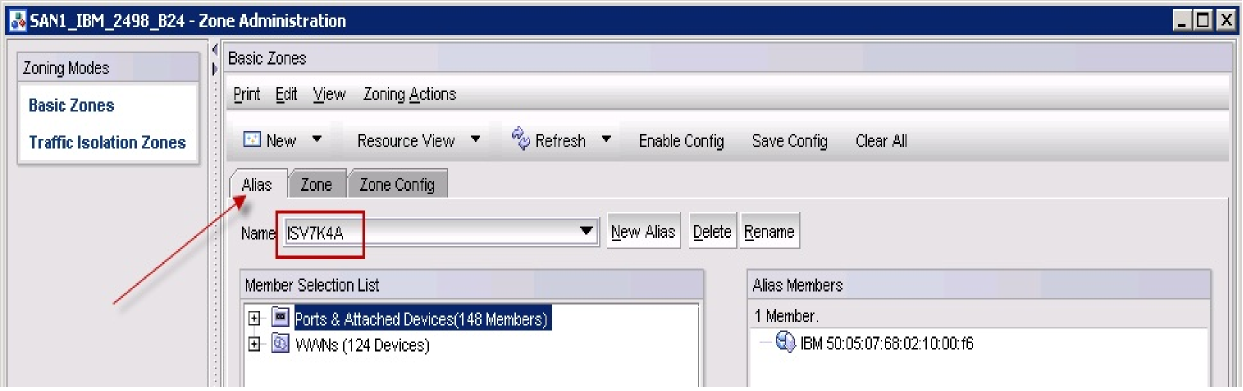

Name: ISV7K4A

Member Selection List:
- Ports & Attached Devices(148 Members)
- WWNs (124 Devices)

Alias Members: 1 Member.
- IBM 50:05:07:68:02:10:00:f6

Figure 5: Assign appropriate alias to the storage ports

3. Create a new zone using the SAN management tool, as shown in Figure 6.

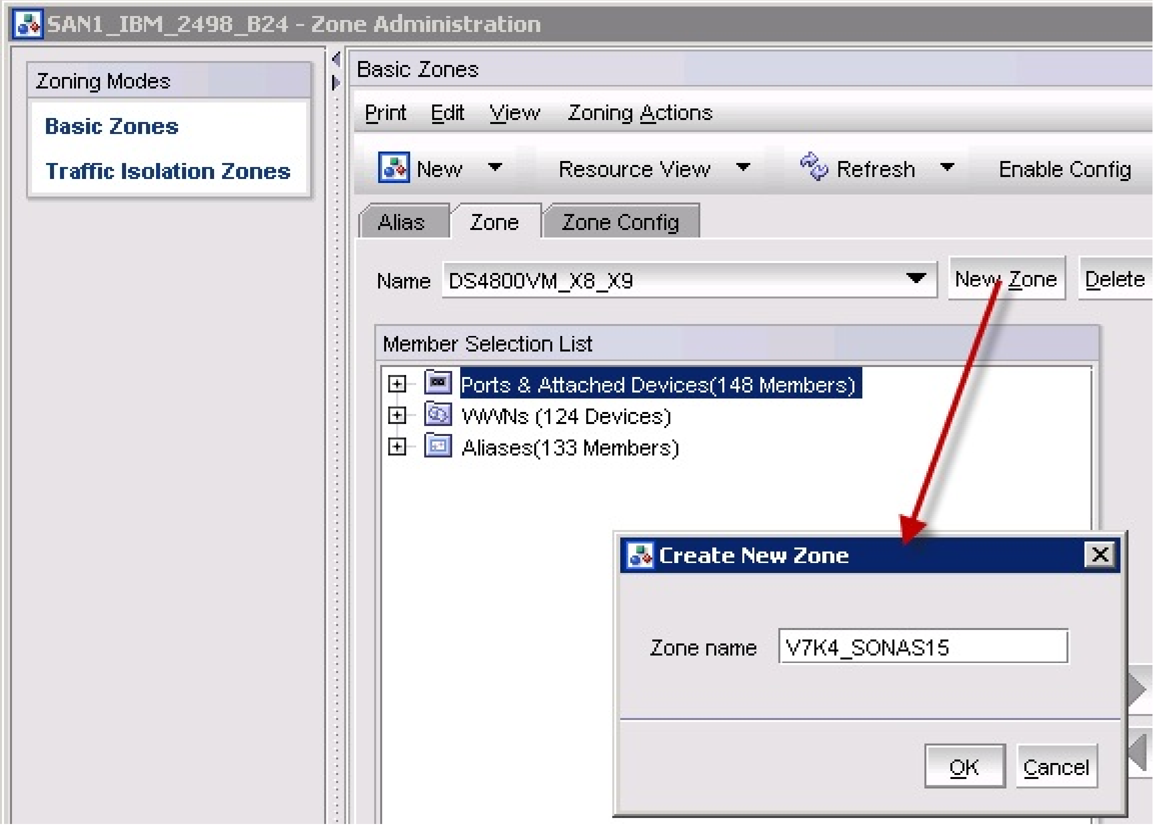

Name DS4800VM_X8_X9

Zone name V7K4_SONAS15

Figure 6: Create a new zone

4. Add the newly defined host port and IBM Storwize V7000 storage ports alias to the newly created zone as shown in Figure 7 and Figure 8.

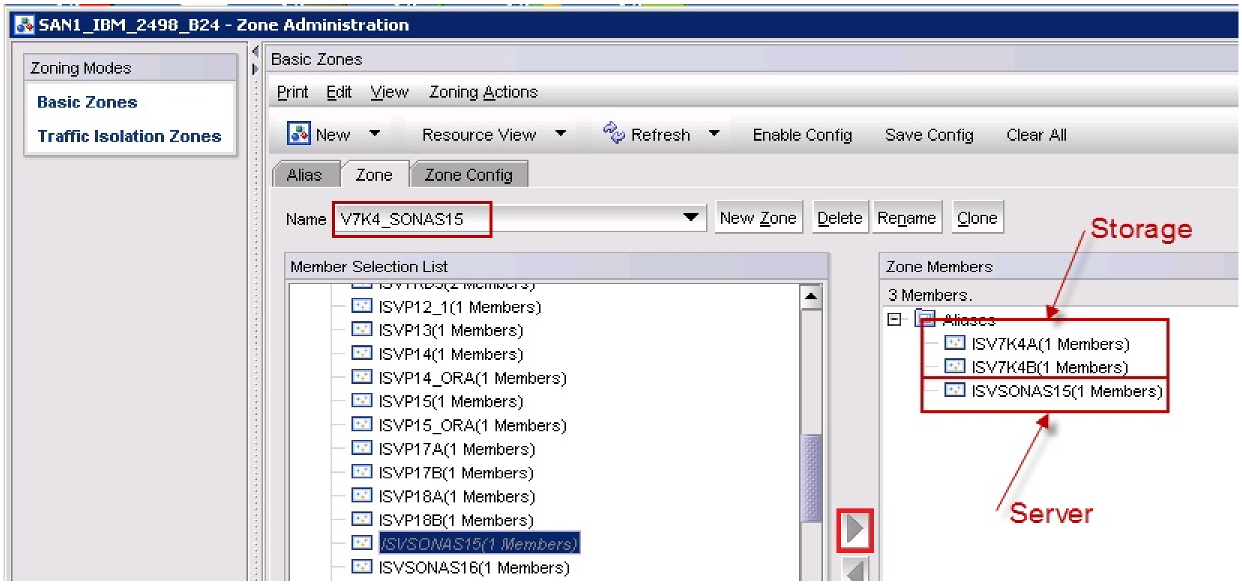

Figure 7: Add host ports and storage ports to the newly created zone

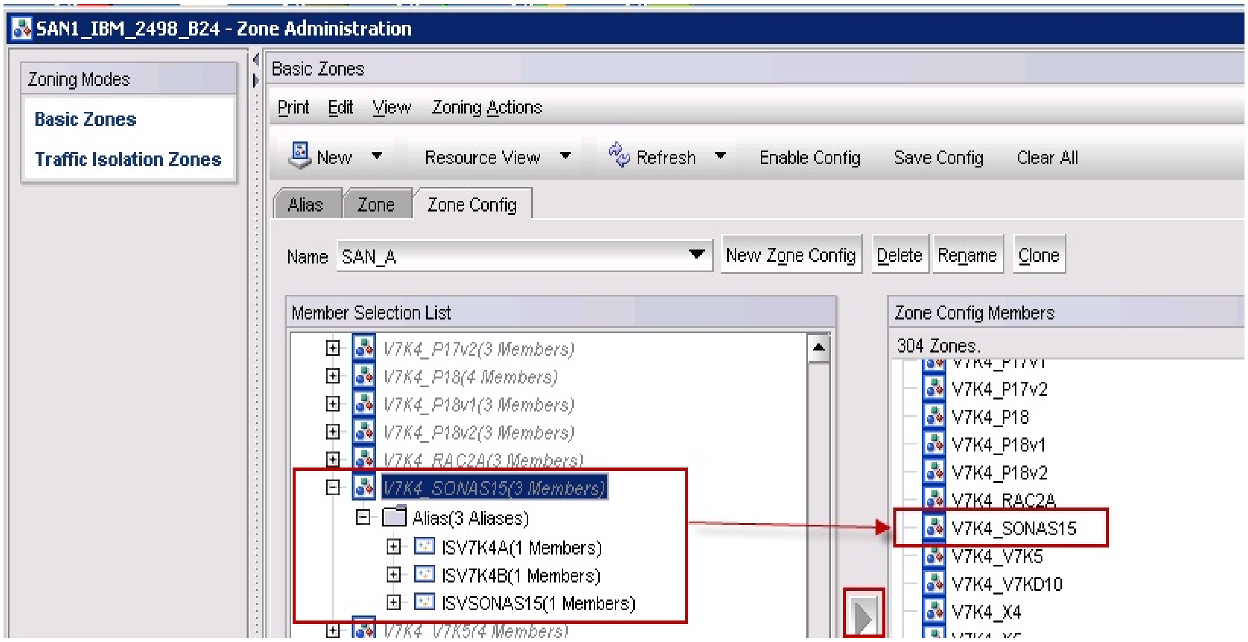

Figure 8: Newly configured host ports and storage ports in the new zone

5. Finally, save the configuration changes and then activate the zone set, as shown in Figure 9. Repeat all of the zoning steps for the redundant fabric.

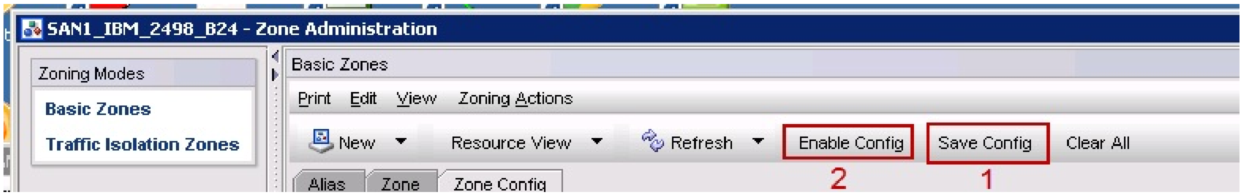

Figure 9: Save and enable the zoning configuration

Zoning IBM Storwize V7000 storage system ports for the partnership

Create a zone containing all the IBM Storwize V7000 storage system ports for the partnership, as shown in Figure 10. Repeat the steps for the redundant fabric.

Note: This step is required only when performing a Fibre Channel partnership. If you need to use an IP partnership, refer to the IP partnership configuration guide. The URL has been provided in appendix A.

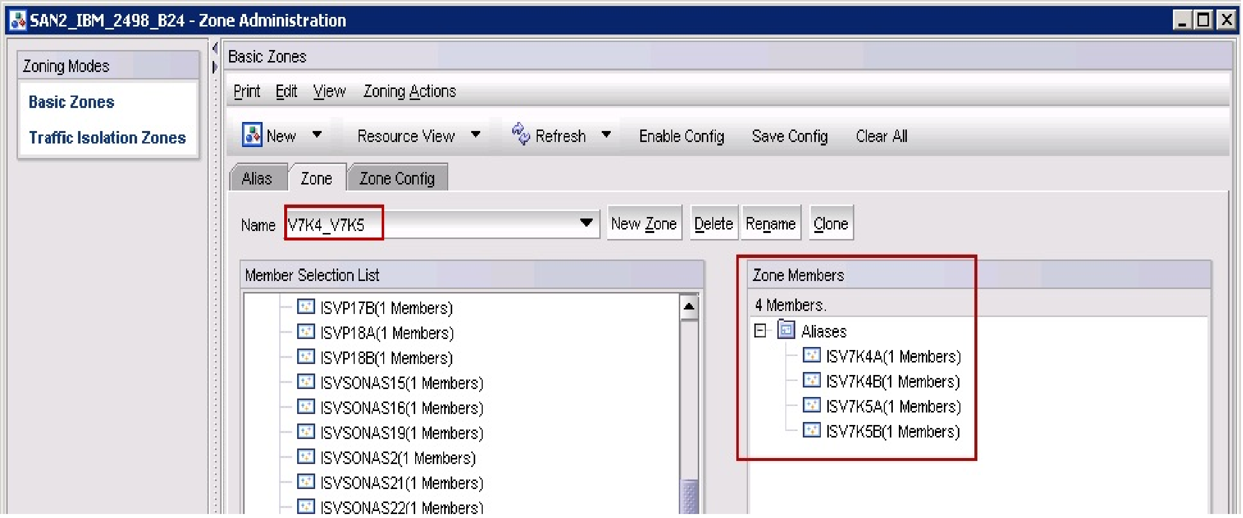

Name V7K4_V7K5

Member Selection List
- ISVP17B(1 Members)
- ISVP18A(1 Members)
- ISVP18B(1 Members)
- ISVSONAS15(1 Members)
- ISVSONAS16(1 Members)
- ISVSONAS19(1 Members)
- ISVSONAS2(1 Members)
- ISVSONAS21(1 Members)
- ISVSONAS22(1 Members)

Zone Members
4 Members.
- Aliases
  - ISV7K4A(1 Members)
  - ISV7K4B(1 Members)
  - ISV7K5A(1 Members)
  - ISV7K5B(1 Members)

Figure 10: Creating a zone containing all storage ports partnership

Create VMware ESXi hosts (Fibre Channel host) on the IBM Storwize V7000 storage system

In the solution test lab environment, the VMware ESXi servers are appropriately installed and configured with Fibre Channel host bus adapters (HBAs). This section describes how to create FC hosts using the IBM Storwize V7000 GUI.

1. Open the host configuration window by clicking **Hosts**, as shown in Figure 11.

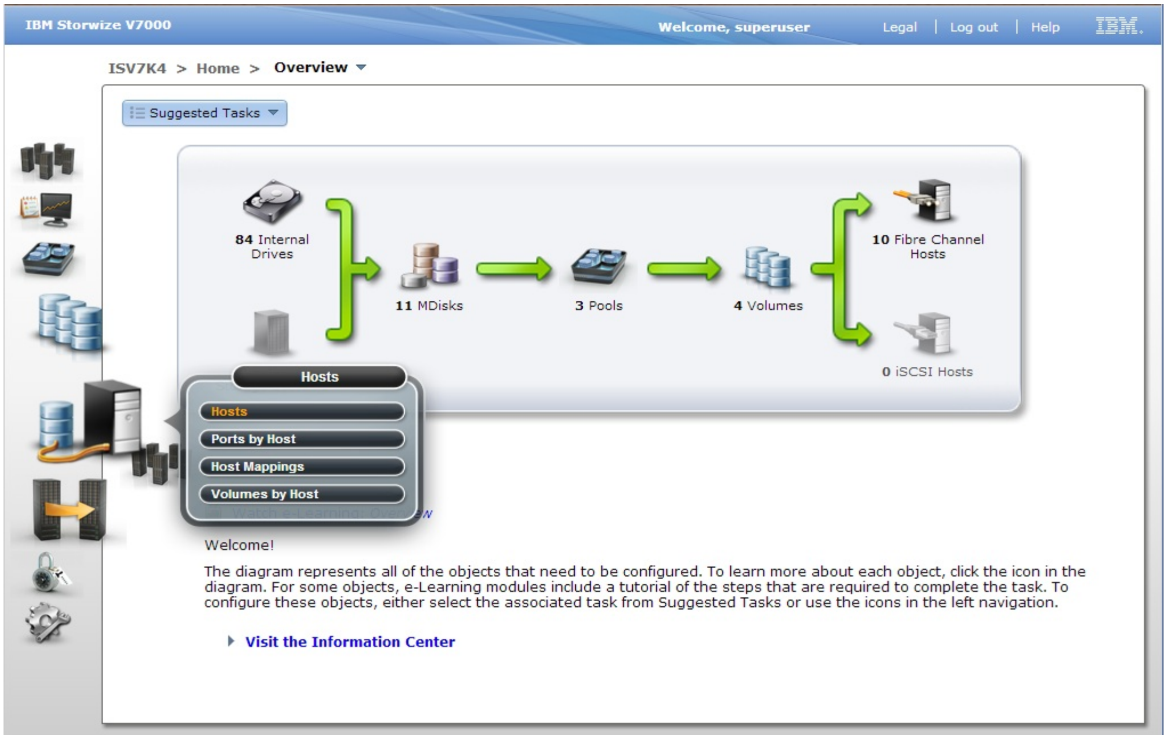

Welcome!

The diagram represents all of the objects that need to be configured. To learn more about each object, click the icon in the diagram. For some objects, e-Learning modules include a tutorial of the steps that are required to complete the task. To configure these objects, either select the associated task from Suggested Tasks or use the icons in the left navigation.

Figure 11: Open the host window

2. To create a new Fibre Channel host, click **New Host** and then click **Fibre Channel Host**, as shown in Figure 12.

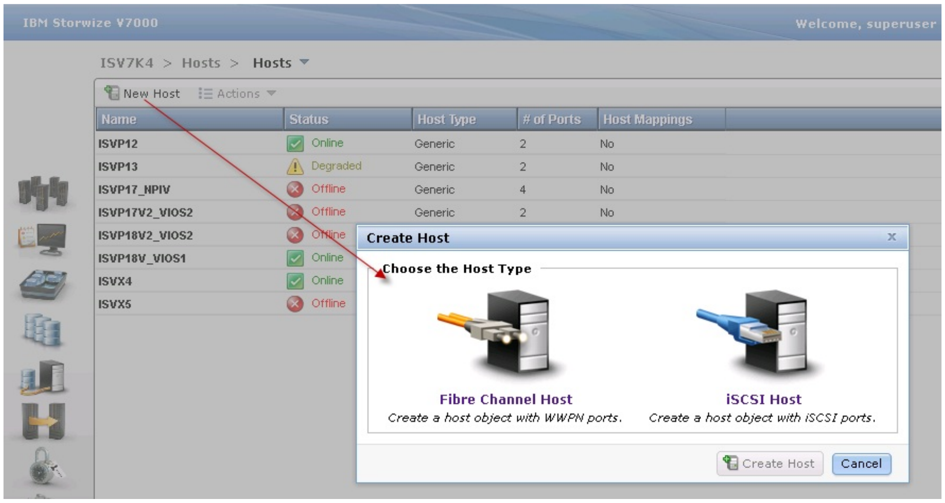

Name	Status	Host Type	# of Ports	Host Mappings
ISVP12	Online	Generic	2	No
ISVP13	Degraded	Generic	2	No
ISVP17_NPIV	Offline	Generic	4	No
ISVP17V2_VIOS2	Offline	Generic	2	No
ISVP18V2_VIOS2	Offline			
ISVP18V_VIOS1	Online			
ISVX4	Online			
ISVX5	Offline			

Create Host

Choose the Host Type

Fibre Channel Host
Create a host object with WWPN ports.

iSCSI Host
Create a host object with iSCSI ports.

Figure 12: Create a Fibre Channel host

3. Enter a host name and click the Fibre Channel Ports drop-down list to get a list of all known worldwide port names (WWPNs) (as shown in Figure 13). Validate your SAN zoning to select the appropriate WWPNs.

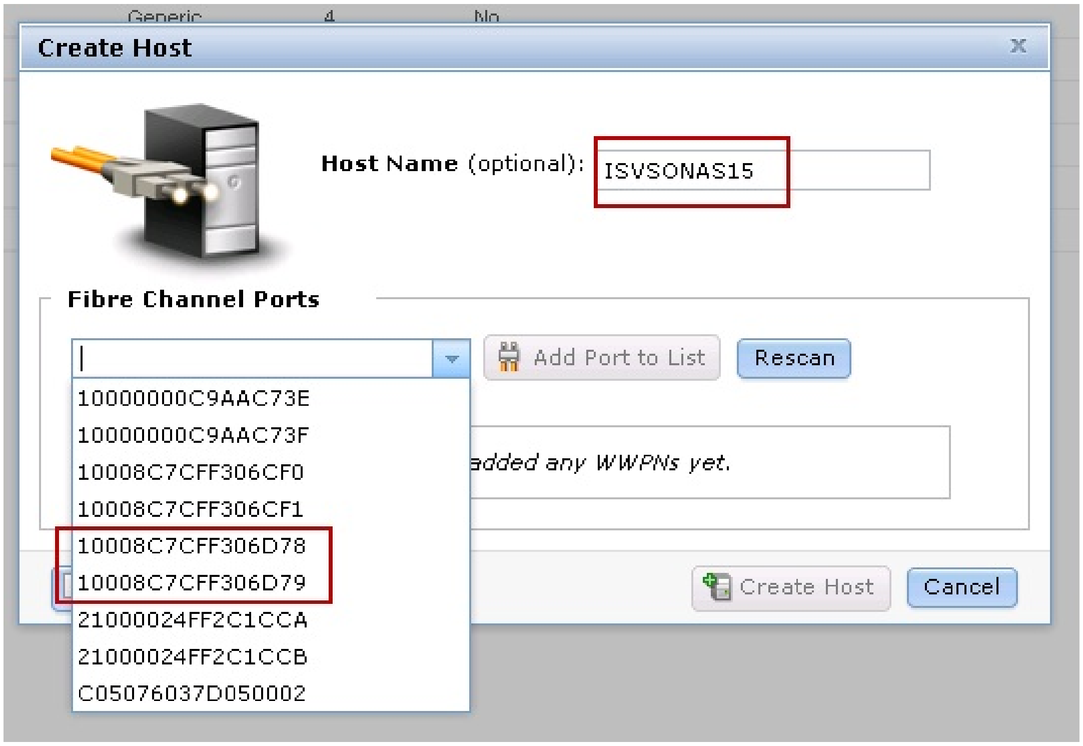

Create Host

Host Name (optional): ISVSONAS15

Fibre Channel Ports

10000000C9AAC73E
10000000C9AAC73F
10008C7CFF306CF0
10008C7CFF306CF1
10008C7CFF306D78
10008C7CFF306D79
21000024FF2C1CCA
21000024FF2C1CCB
C05076037D050002

added any WWPNs yet.

Figure 13: Create a Fibre Channel host

4. Add all the ports that belong to the host, as shown in Figure 14.

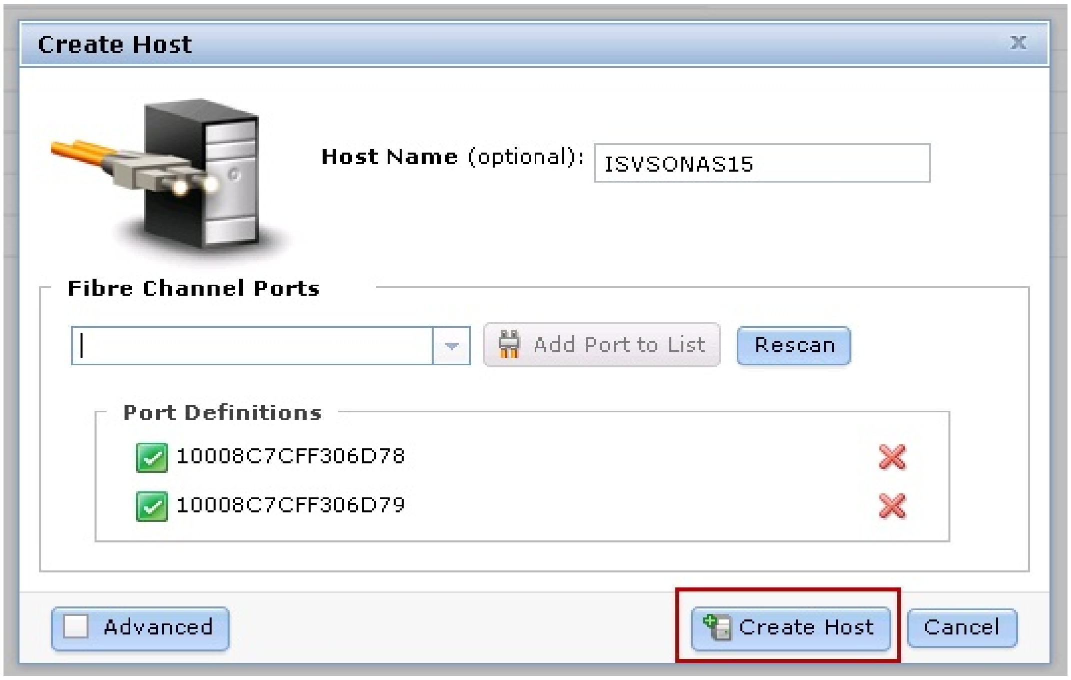

Create Host

Host Name (optional): ISVSONAS15

Fibre Channel Ports

Port Definitions

☑ 10008C7CFF306D78

☑ 10008C7CFF306D79

Figure 14: Add all WWPNs

5. Click **Create Host** and the wizard creates the host, as shown in Figure 15.

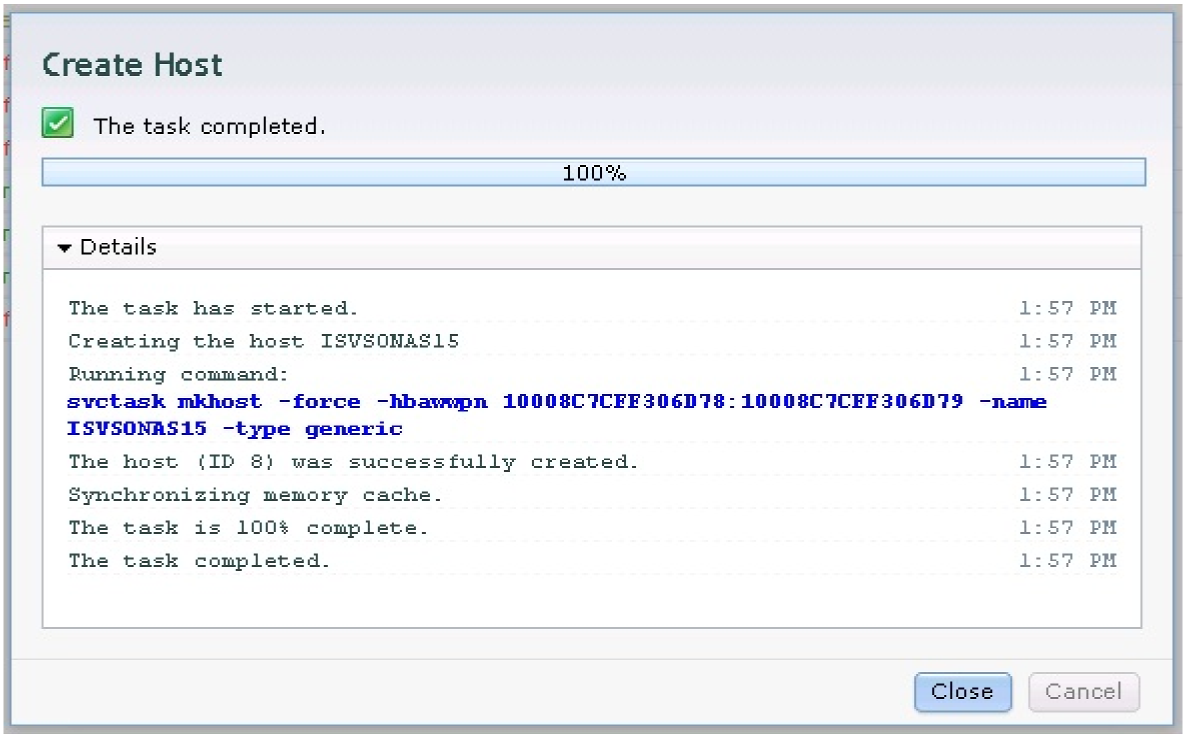

Create Host

☑ The task completed.

100%

Details

```
The task has started.                                        1:57 PM
Creating the host ISVSONAS15                                 1:57 PM
Running command:                                             1:57 PM
svctask mkhost -force -hbawwpn 10008C7CFF306D78:10008C7CFF306D79 -name
ISVSONAS15 -type generic
The host (ID 8) was successfully created.                    1:57 PM
Synchronizing memory cache.                                  1:57 PM
The task is 100% complete.                                   1:57 PM
The task completed.                                          1:57 PM
```

Figure 15: Host created successfully

6. Click **Close** to return to the host window, as shown in Figure 16.

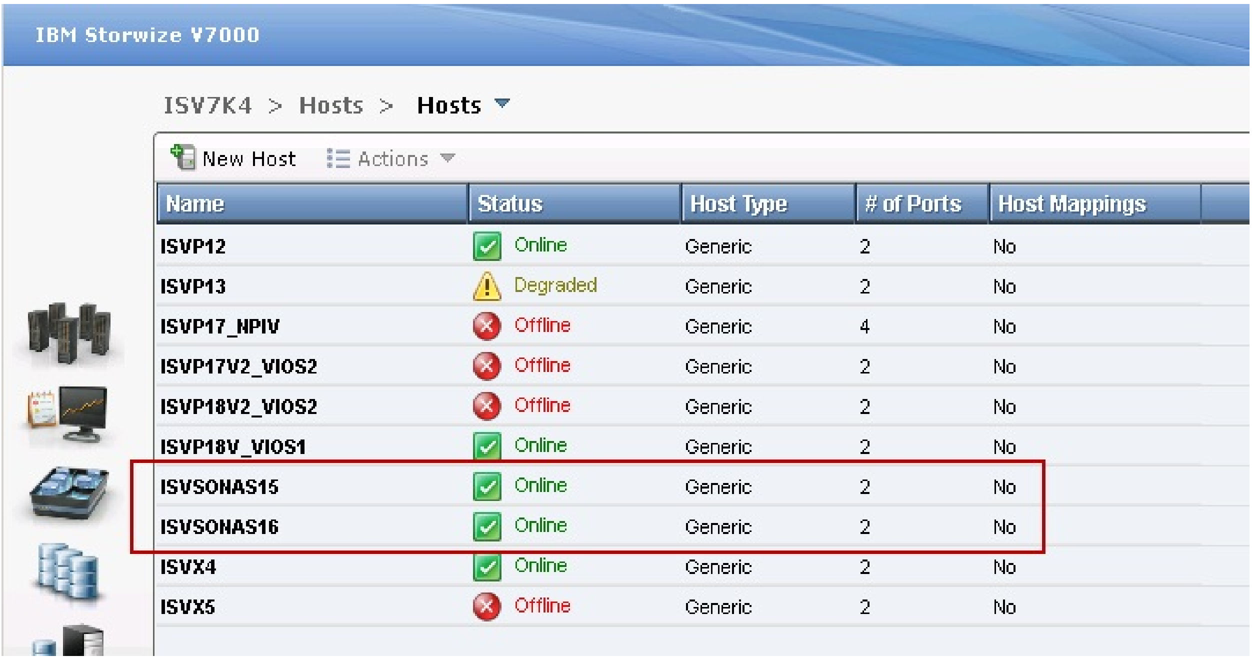

Name	Status	Host Type	# of Ports	Host Mappings
ISVP12	Online	Generic	2	No
ISVP13	Degraded	Generic	2	No
ISVP17_NPIV	Offline	Generic	4	No
ISVP17V2_VIOS2	Offline	Generic	2	No
ISVP18V2_VIOS2	Offline	Generic	2	No
ISVP18V_VIOS1	Online	Generic	2	No
ISVSONAS15	Online	Generic	2	No
ISVSONAS16	Online	Generic	2	No
ISVX4	Online	Generic	2	No
ISVX5	Offline	Generic	2	No

Figure 16: All hosts

Create storage pools, MDisks, and storage volumes on the storage systems

This section describes the fundamental steps involved in the internal storage configuration of the IBM Storwize V7000 system.

In the lab, for solution validation purpose, the test team created a new storage pools using the Basic RAID 5 configuration and validated the appropriate MDisks in the newly created storage pools. After that, new storage volumes are configured and appropriately mapped the newly created storage volumes to Fibre Channel hosts on the IBM Storwize V7000 storage systems. The steps involved creating storage pools, validation of MDisks, and configuring storage volumes (as described in this section).

1. Click the **Pools** icons at the left side of the window and then click **Internal Storage**, as shown in Figure 17.

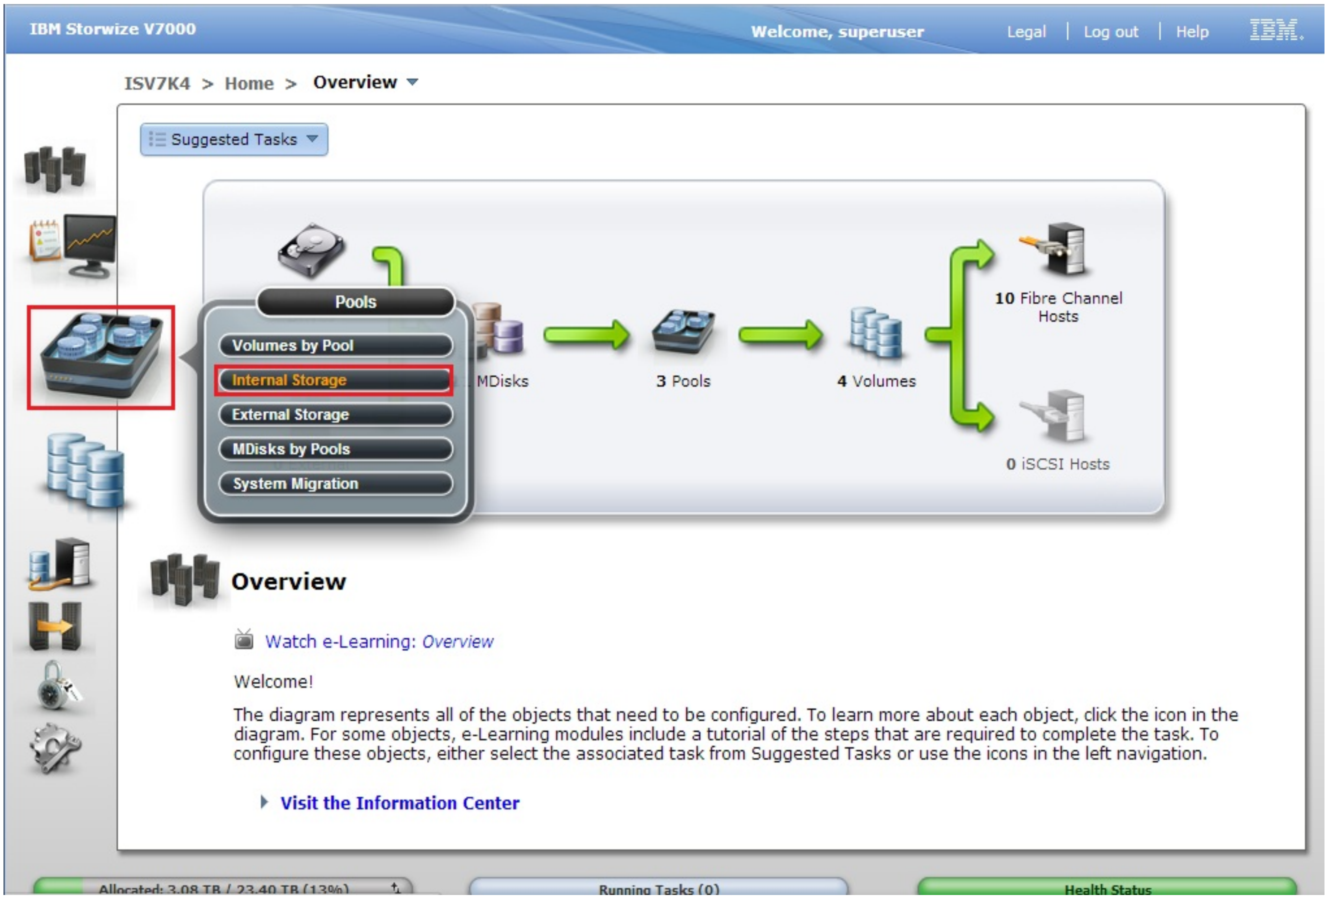

Welcome!

The diagram represents all of the objects that need to be configured. To learn more about each object, click the icon in the diagram. For some objects, e-Learning modules include a tutorial of the steps that are required to complete the task. To configure these objects, either select the associated task from Suggested Tasks or use the icons in the left navigation.

Figure 17: Accessing the Internal Storage window

2. Select an appropriate physical drive pool from the **Drive Class Filter** pane on the left side and click **Configure Storage** to launch the Configure Internal Storage wizard, as shown in Figure 18.

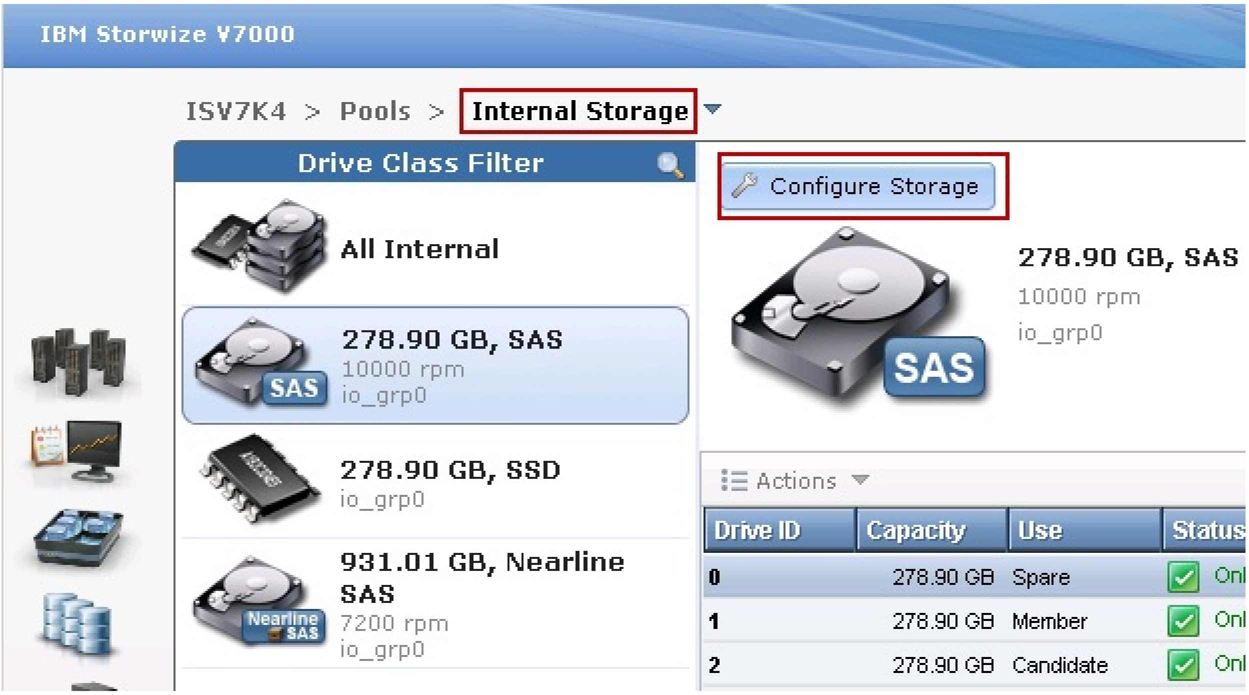

IBM Storwize V7000

ISV7K4 > Pools > Internal Storage

Drive Class Filter

All Internal

278.90 GB, SAS
10000 rpm
io_grp0

278.90 GB, SSD
io_grp0

931.01 GB, Nearline SAS
7200 rpm
io_grp0

Configure Storage

278.90 GB, SAS
10000 rpm
io_grp0

Actions

Drive ID	Capacity	Use	Status
0	278.90 GB	Spare	✓ Onl
1	278.90 GB	Member	✓ Onl
2	278.90 GB	Candidate	✓ Onl

Figure 18: Internal Storage window

3. In the solution lab test environment, the **Basic RAID-5** configuration is selected (as shown in Figure 19).

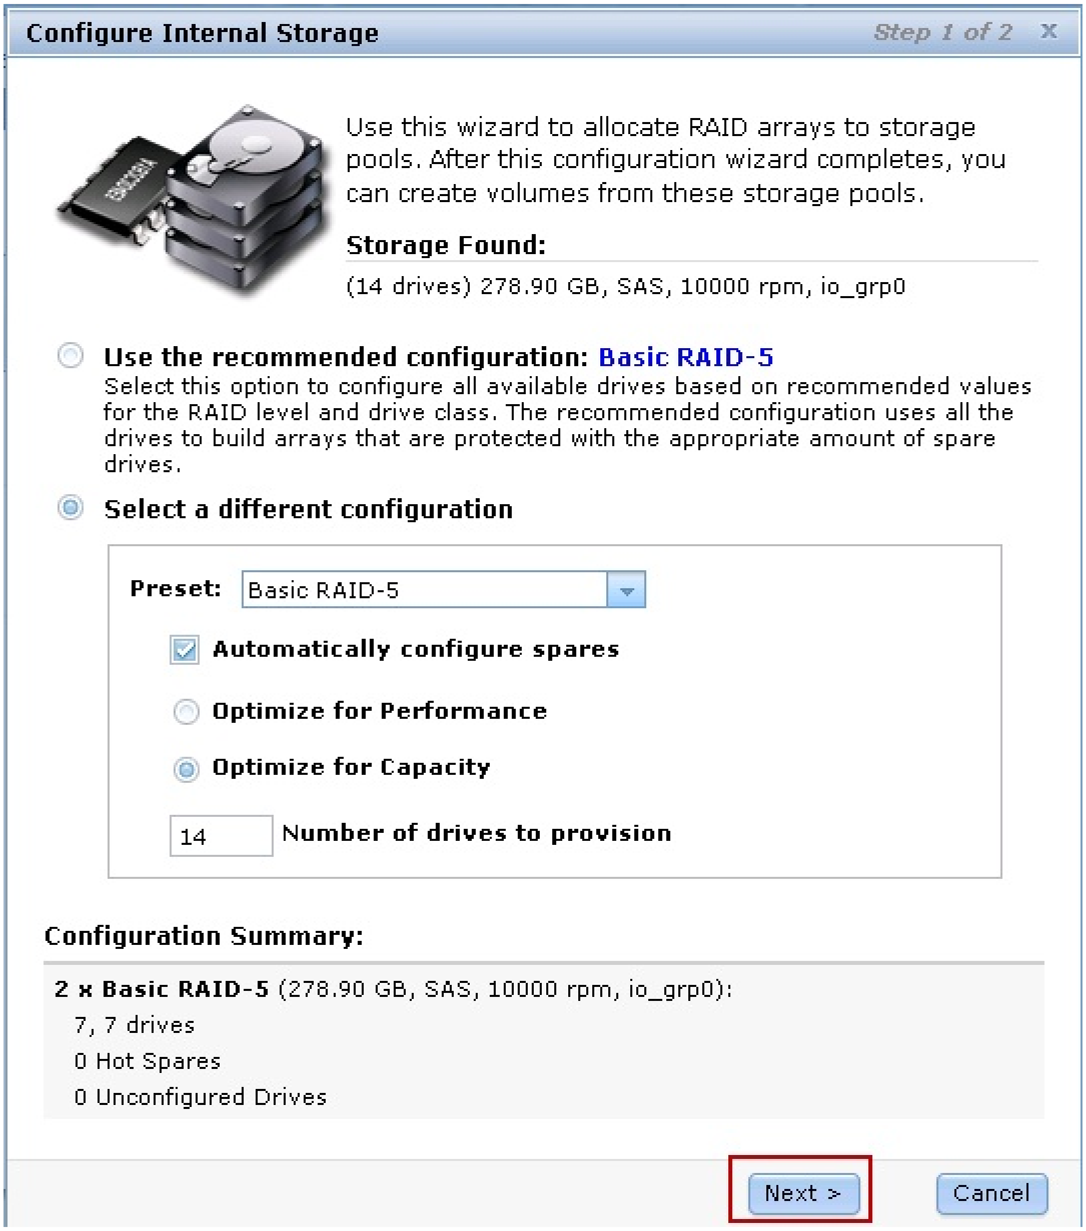

Configure Internal Storage

Use this wizard to allocate RAID arrays to storage pools. After this configuration wizard completes, you can create volumes from these storage pools.

Storage Found:

(14 drives) 278.90 GB, SAS, 10000 rpm, io_grp0

○ **Use the recommended configuration: Basic RAID-5**
Select this option to configure all available drives based on recommended values for the RAID level and drive class. The recommended configuration uses all the drives to build arrays that are protected with the appropriate amount of spare drives.

● **Select a different configuration**

Preset: Basic RAID-5

☑ **Automatically configure spares**

○ **Optimize for Performance**

● **Optimize for Capacity**

14 **Number of drives to provision**

Configuration Summary:

2 x Basic RAID-5 (278.90 GB, SAS, 10000 rpm, io_grp0):
7, 7 drives
0 Hot Spares
0 Unconfigured Drives

Figure 19: Select a RAID preset

4. Create a new storage pool (as shown in Figure 20) and click **Finish** to complete the wizard.

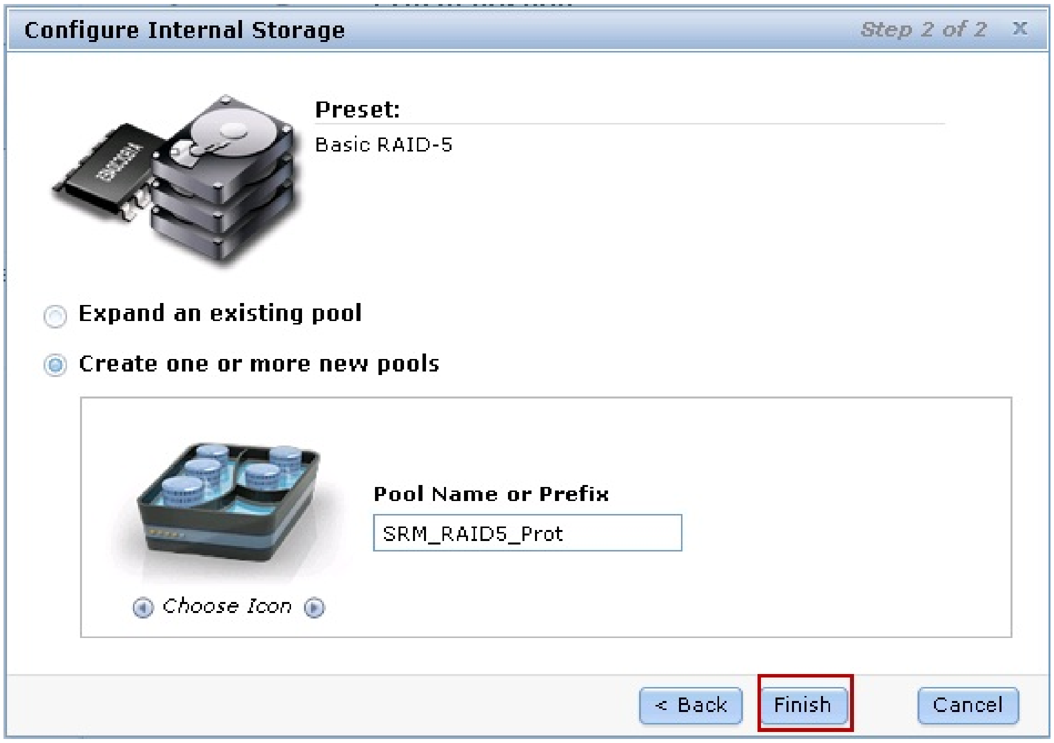

Configure Internal Storage

Step 2 of 2

Preset:
Basic RAID-5

○ Expand an existing pool

● Create one or more new pools

Pool Name or Prefix: SRM_RAID5_Prot

Figure 20: Create new storage pools

5. Click **Close** to return to the Internal Storage window, as show in Figure 21.

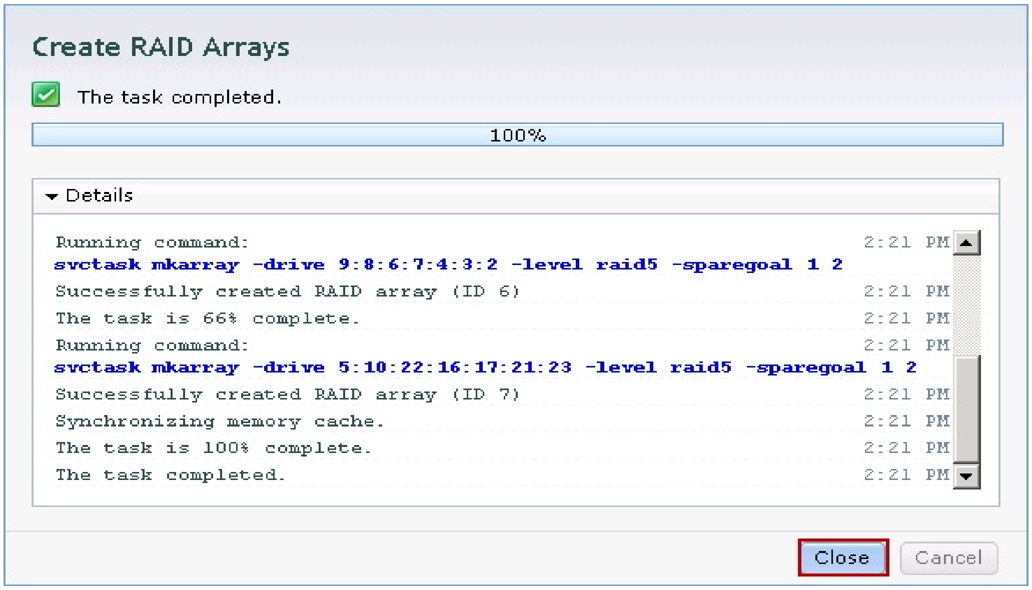

Create RAID Arrays

☑ The task completed.

100%

Details

```
Running command:                                                      2:21 PM
svctask mkarray -drive 9:8:6:7:4:3:2 -level raid5 -sparegoal 1 2
Successfully created RAID array (ID 6)                                2:21 PM
The task is 66% complete.                                             2:21 PM
Running command:                                                      2:21 PM
svctask mkarray -drive 5:10:22:16:17:21:23 -level raid5 -sparegoal 1 2
Successfully created RAID array (ID 7)                                2:21 PM
Synchronizing memory cache.                                           2:21 PM
The task is 100% complete.                                            2:21 PM
The task completed.                                                   2:21 PM
```

Figure 21: Complete the storage pool creation

6. Verify that the new storage pool and the MDisks were created properly by clicking the newly created storage pool, as shown inFigure 22

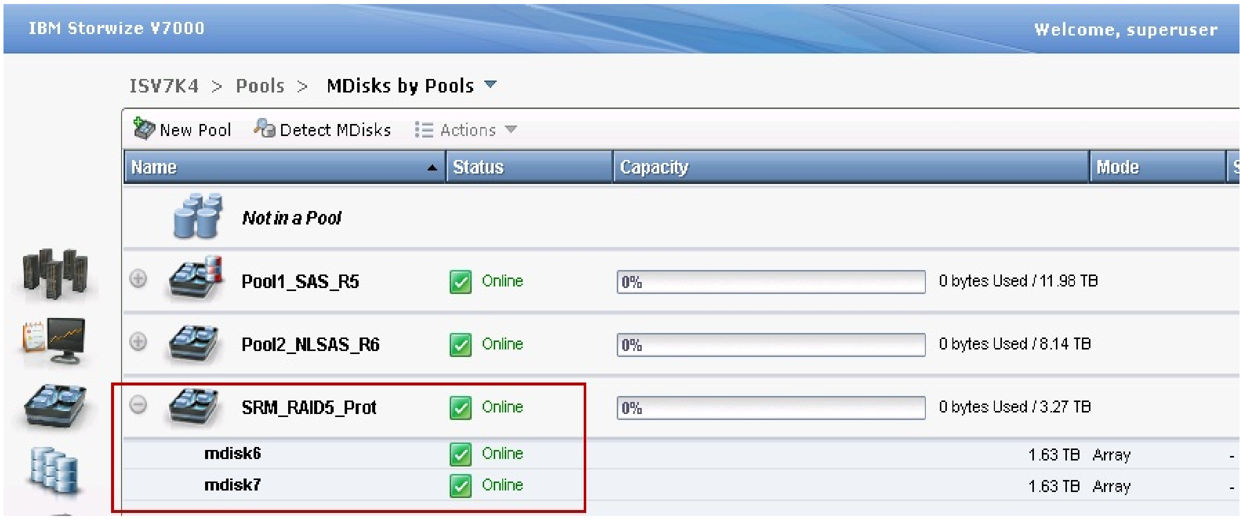

Name	Status	Capacity		Mode	S
Not in a Pool					
Pool1_SAS_R5	Online	0%	0 bytes Used / 11.98 TB		
Pool2_NLSAS_R6	Online	0%	0 bytes Used / 8.14 TB		
SRM_RAID5_Prot	Online	0%	0 bytes Used / 3.27 TB		
mdisk6	Online		1.63 TB	Array	-
mdisk7	Online		1.63 TB	Array	-

Figure 22 Completed storage pools

7. In the IBM Storwize V7000 GUI, click the **Volumes** icon and then click **Volumes** to open the Volumes window, as shown in Figure 23.

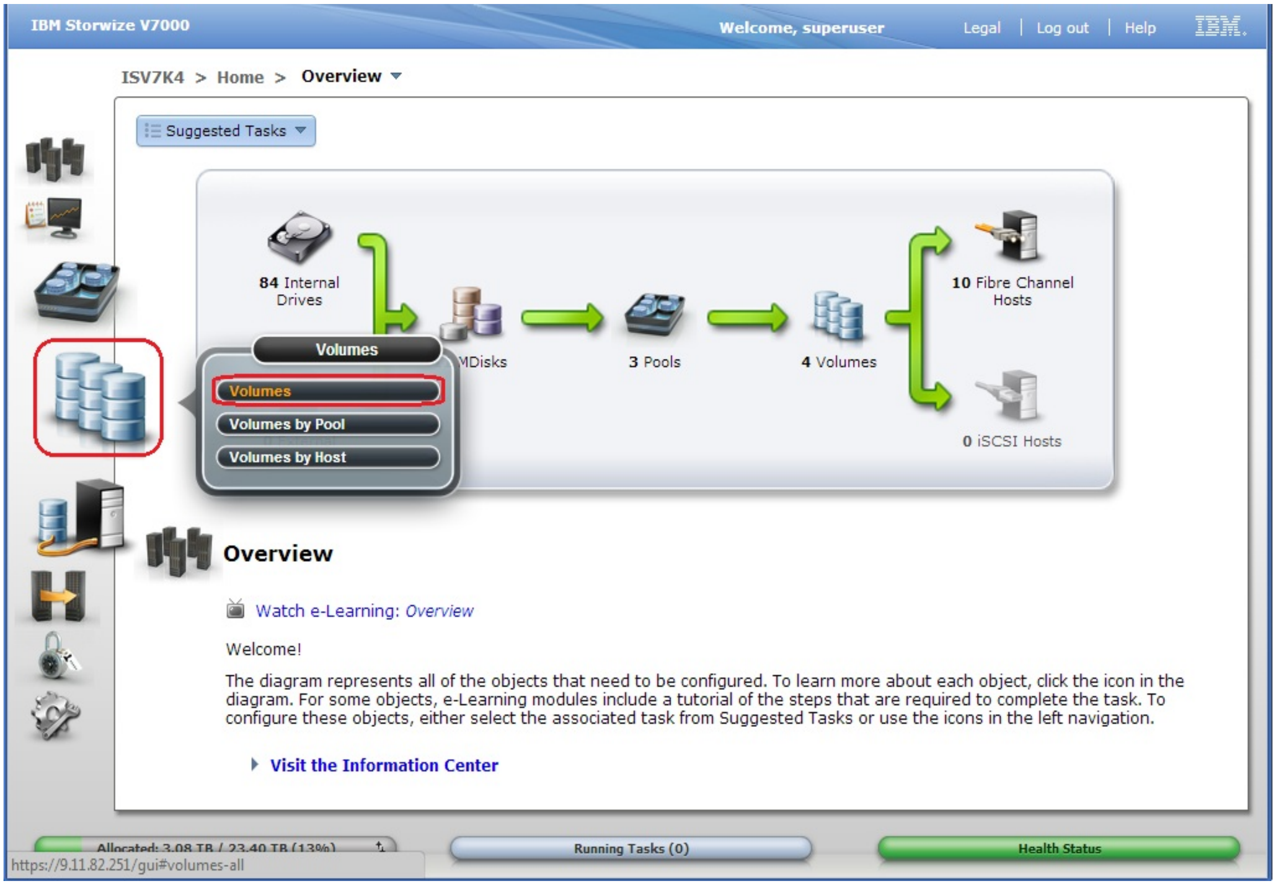

Welcome!

The diagram represents all of the objects that need to be configured. To learn more about each object, click the icon in the diagram. For some objects, e-Learning modules include a tutorial of the steps that are required to complete the task. To configure these objects, either select the associated task from Suggested Tasks or use the icons in the left navigation.

Figure 23: Volumes window

8. Click **New Volume** to start creating a new volume, as shown in Figure 24.

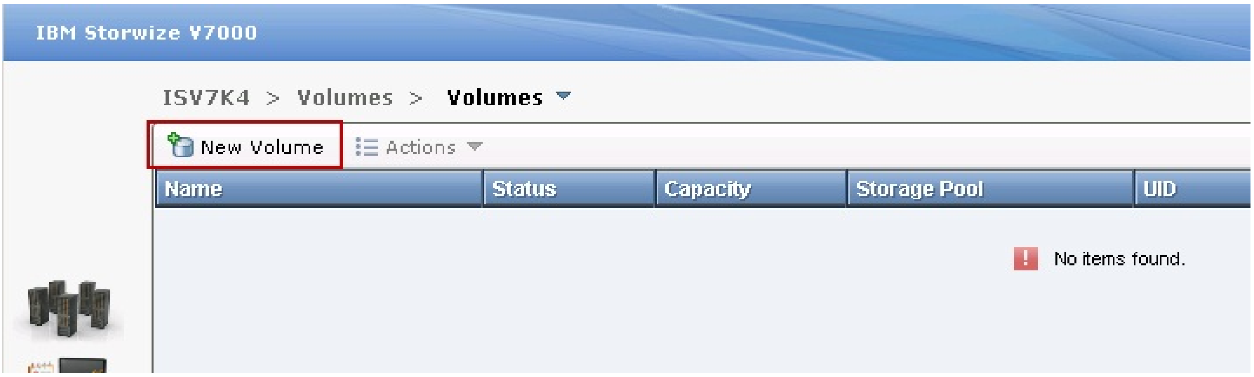

Name	Status	Capacity	Storage Pool	UID

No items found.

Figure 24: Create a new volume

9. For the solution validation purpose, click **Generic** and select the newly created storage pool, as shown in Figure 25.

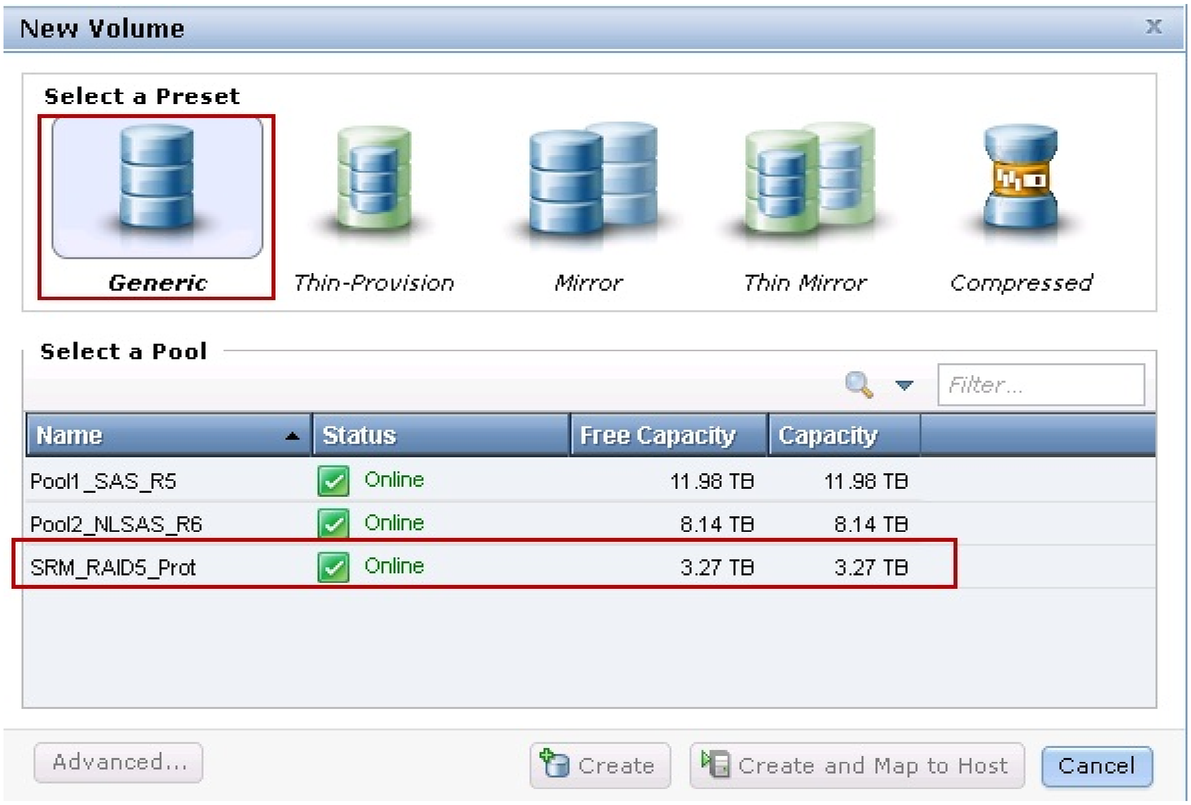

Name	Status	Free Capacity	Capacity
Pool1_SAS_R5	✓ Online	11.98 TB	11.98 TB
Pool2_NLSAS_R6	✓ Online	8.14 TB	8.14 TB
SRM_RAID5_Prot	✓ Online	3.27 TB	3.27 TB

Figure 25: Create a generic volume

10. In the solution test lab environment, four volumes have been configured by clicking **Create and Map to Host**, as shown in Figure 26.

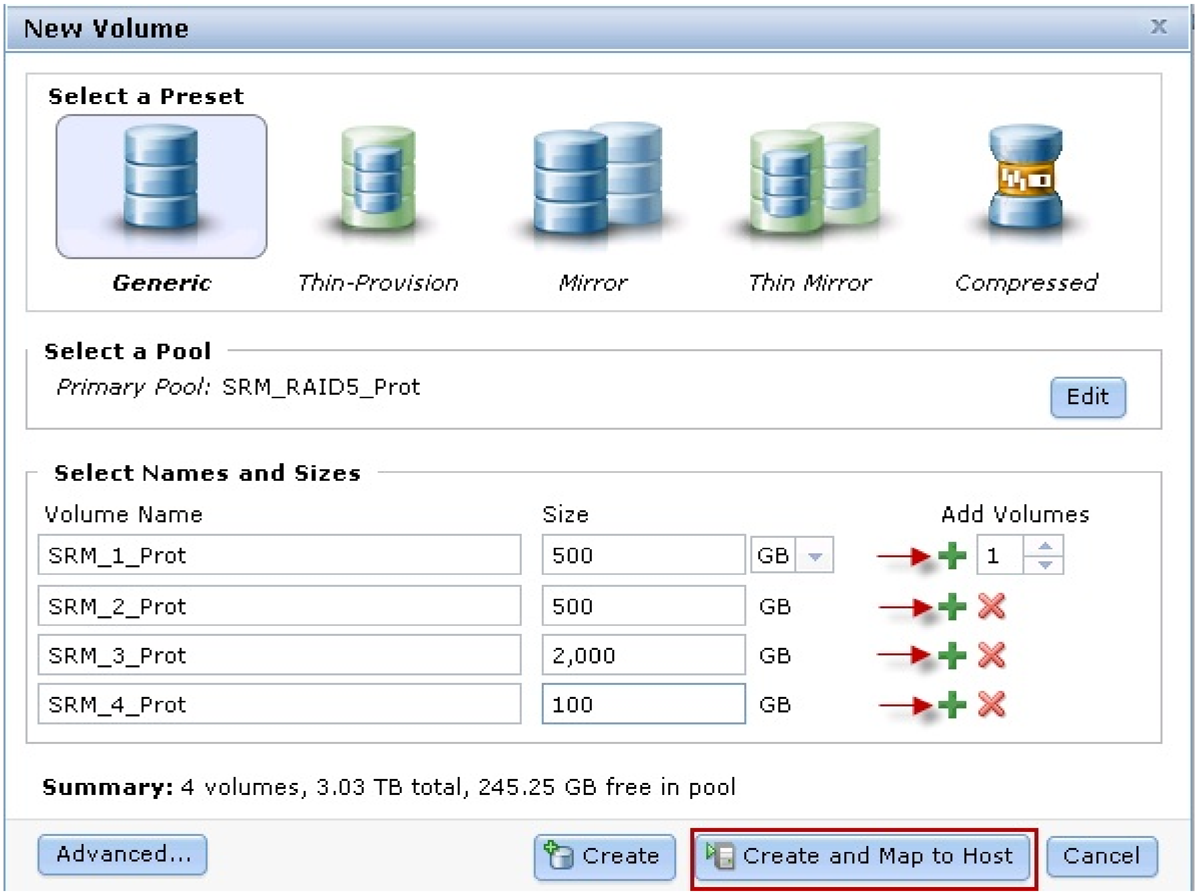

New Volume

Select a Preset

Generic, Thin-Provision, Mirror, Thin Mirror, Compressed

Select a Pool

Primary Pool: SRM_RAID5_Prot

Select Names and Sizes

Volume Name	Size		Add Volumes
SRM_1_Prot	500	GB	1
SRM_2_Prot	500	GB	
SRM_3_Prot	2,000	GB	
SRM_4_Prot	100	GB	

Summary: 4 volumes, 3.03 TB total, 245.25 GB free in pool

Figure 26: Create generic volumes

11. Click **Continue** to create volumes and map storage volumes to server hosts on the storage system, as shown in Figure 27.

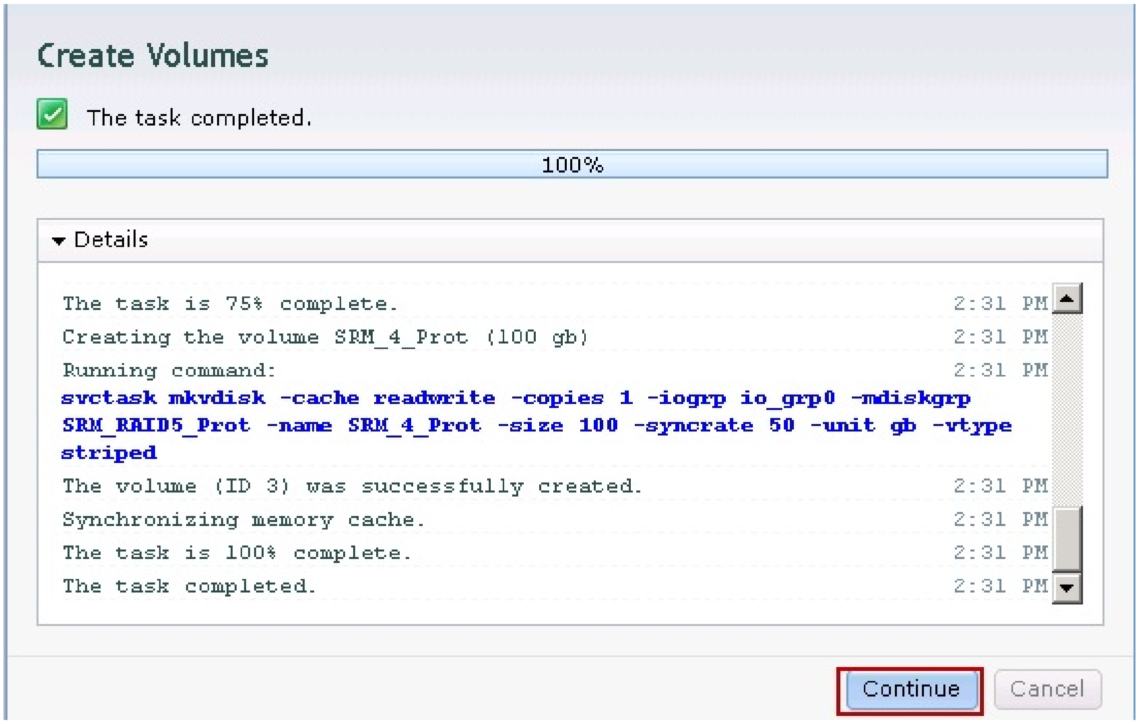

Create Volumes

☑ The task completed.

100%

▾ Details

```
The task is 75% complete.                                            2:31 PM
Creating the volume SRM_4_Prot (100 gb)                              2:31 PM
Running command:                                                     2:31 PM
svctask mkvdisk -cache readwrite -copies 1 -iogrp io_grp0 -mdiskgrp
SRM_RAID5_Prot -name SRM_4_Prot -size 100 -syncrate 50 -unit gb -vtype
striped
The volume (ID 3) was successfully created.                          2:31 PM
Synchronizing memory cache.                                          2:31 PM
The task is 100% complete.                                           2:31 PM
The task completed.                                                  2:31 PM
```

Figure 27: Create generic volumes

12. Select the appropriate server host to map to the newly created volumes, as shown in Figure 28.

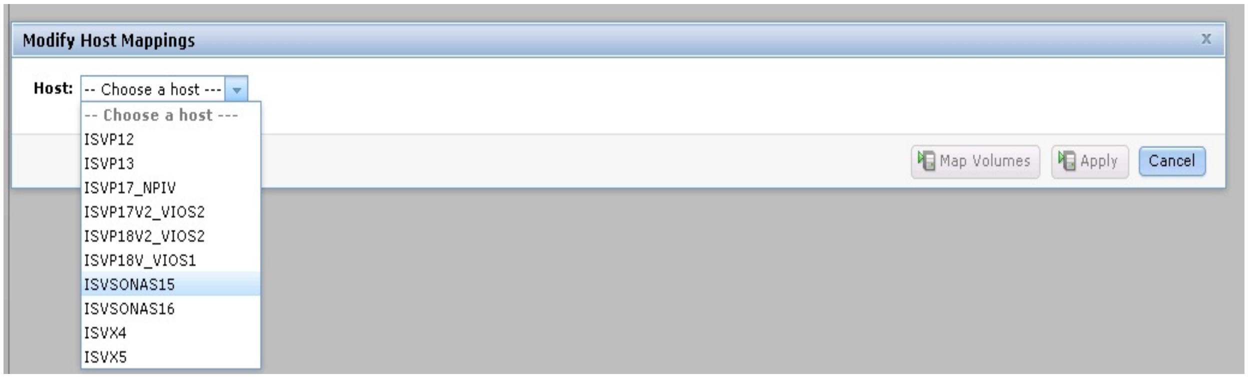

Modify Host Mappings

Host: -- Choose a host ---

Figure 28: Map storage volumes to the server host

13. Click **Apply** to finalize the storage volumes mapping to selected server host, as shown in Figure 29.

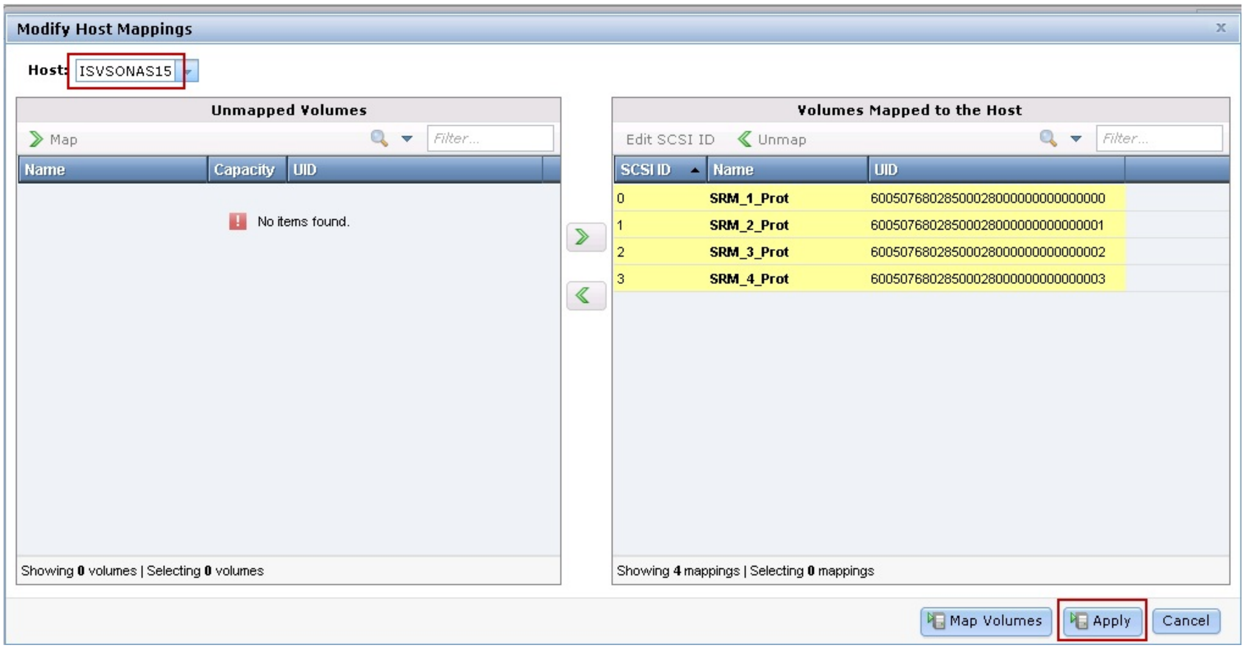

Modify Host Mappings

Host: ISVSONAS15

Unmapped Volumes

Name	Capacity	UID

No items found.

Showing **0** volumes | Selecting **0** volumes

Volumes Mapped to the Host

SCSI ID	Name	UID
0	**SRM_1_Prot**	60050768028500028000000000000000
1	**SRM_2_Prot**	60050768028500028000000000000001
2	**SRM_3_Prot**	60050768028500028000000000000002
3	**SRM_4_Prot**	60050768028500028000000000000003

Showing **4** mappings | Selecting **0** mappings

Figure 29: Map storage volumes to the selected server host

14. Click **Close** to complete the host mapping, as shown in Figure 30.

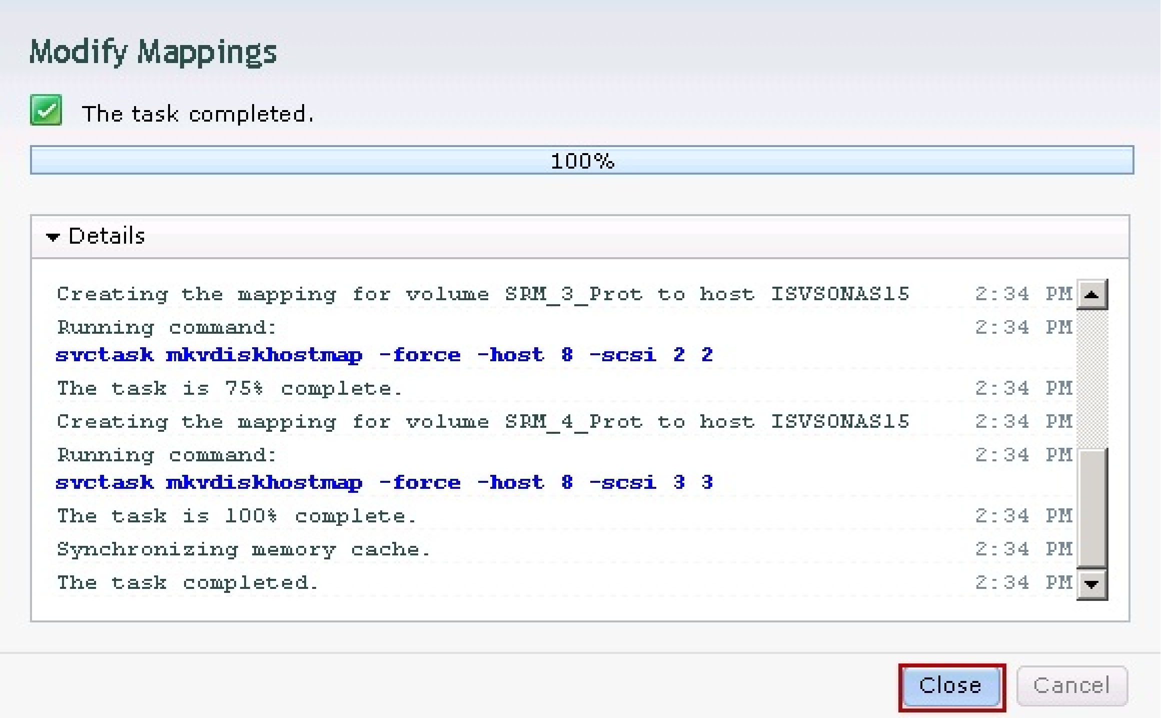

Modify Mappings

The task completed.

100%

Details

```
Creating the mapping for volume SRM_3_Prot to host ISVSONAS15     2:34 PM
Running command:                                                   2:34 PM
svctask mkvdiskhostmap -force -host 8 -scsi 2 2
The task is 75% complete.                                          2:34 PM
Creating the mapping for volume SRM_4_Prot to host ISVSONAS15     2:34 PM
Running command:                                                   2:34 PM
svctask mkvdiskhostmap -force -host 8 -scsi 3 3
The task is 100% complete.                                         2:34 PM
Synchronizing memory cache.                                        2:34 PM
The task completed.                                                2:34 PM
```

Figure 30: Modify mapping task completed

Managing remote copy using the GUI

This section describes the remote copy configuration using the GUI.

Remote copy consists of two methods for copying: Metro Mirror and Global Mirror. Metro Mirror is designed for metropolitan distances in conjunction with a synchronous copy requirement, while Global Mirror is designed for longer distances without requiring the hosts to wait for the full round-trip delay of the long distance link.

Metro Mirror and Global Mirror are IBM branded terms for the functions synchronous remote copy and asynchronous remote copy respectively.

Note: In the lab solution test setup, the team used Metro Mirror based remote copy method. Following steps are performed to configure the Metro Mirror based remote copy method.

Note: The partnership creating steps are specific to using a Fibre Channel based partnership. If you will be using a native IP partnership, refer to the IP partnership configuration guide in appendix A.

Create partnership for storage systems (protected site)

The first step is to configure the partnership for storage systems of the protected site and the recovery site. Following steps involves configuring the partnership for the storage systems.

Note: Start the partnership configuration on the IBM Storwize V7000 storage system at the protected site.

6. Click **Copy Services** and then click **Partnerships**, as shown in Figure 31.

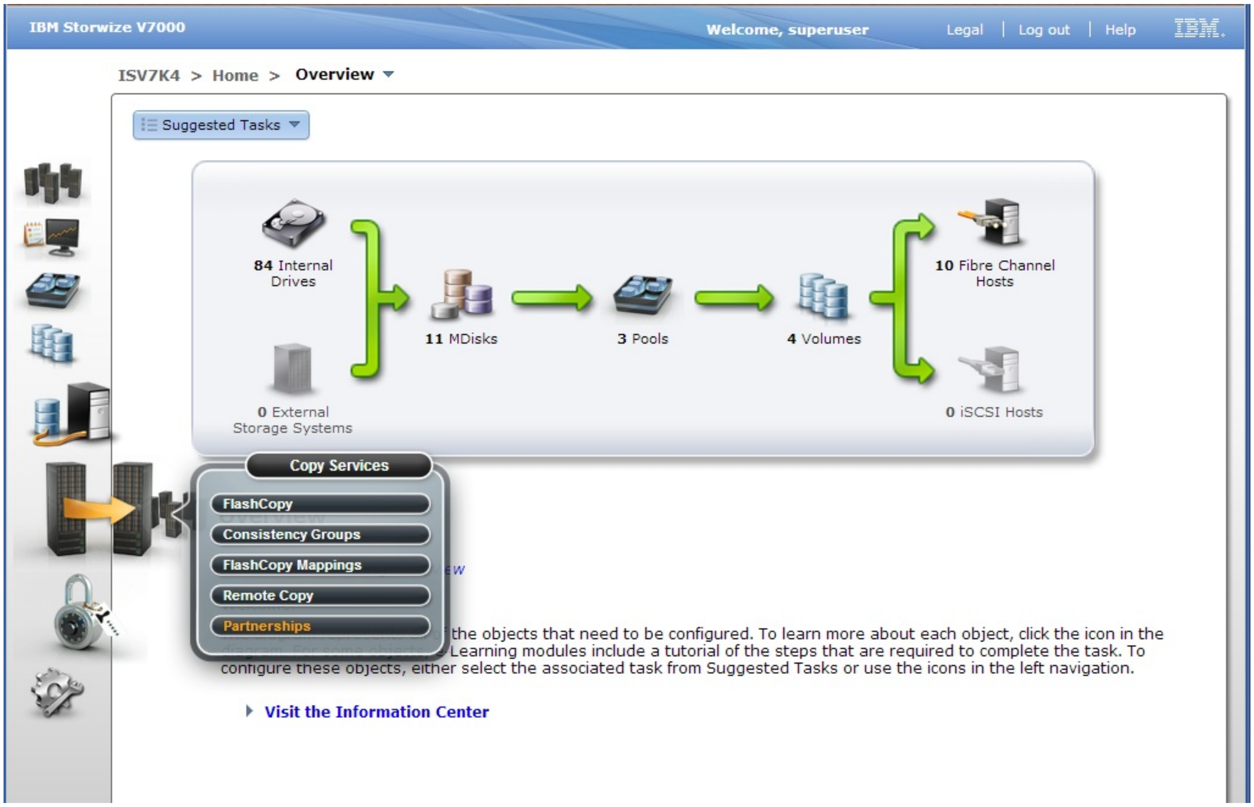

f the objects that need to be configured. To learn more about each object, click the icon in the diagram. For some objects, Learning modules include a tutorial of the steps that are required to complete the task. To configure these objects, either select the associated task from Suggested Tasks or use the icons in the left navigation.

Figure 31: Partnership window

7. Click **New Partnership** to create a new partnership, as shown in Figure 32.

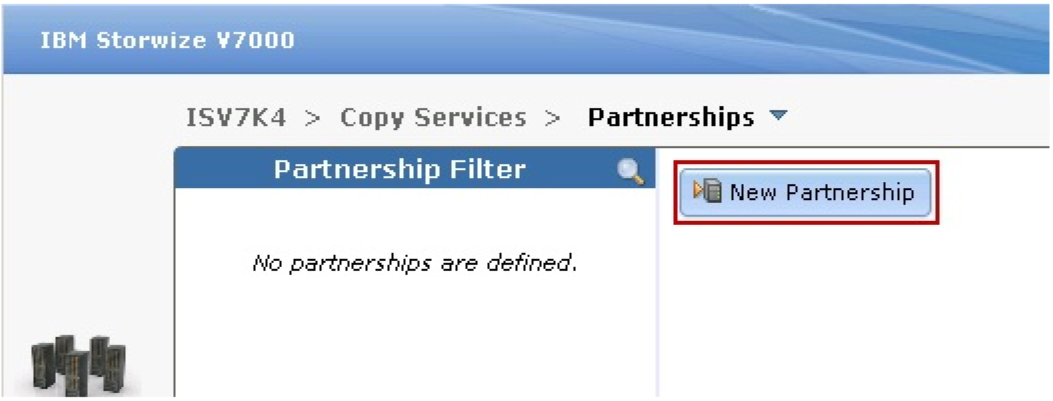

No partnerships are defined.

Figure 32: Create a new partnership

8. Provide the recovery site IBM Storwize V7000 storage system information (as shown in Figure 33) and click **Create.**

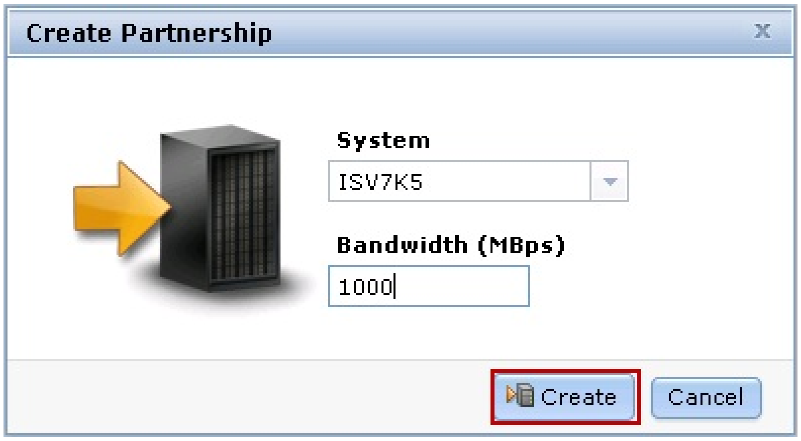

System

ISV7K5

Bandwidth (MBps)

1000

Figure 33: Create a new partnership

9. Click **Close** after the new partnership is created, as shown in Figure 34.

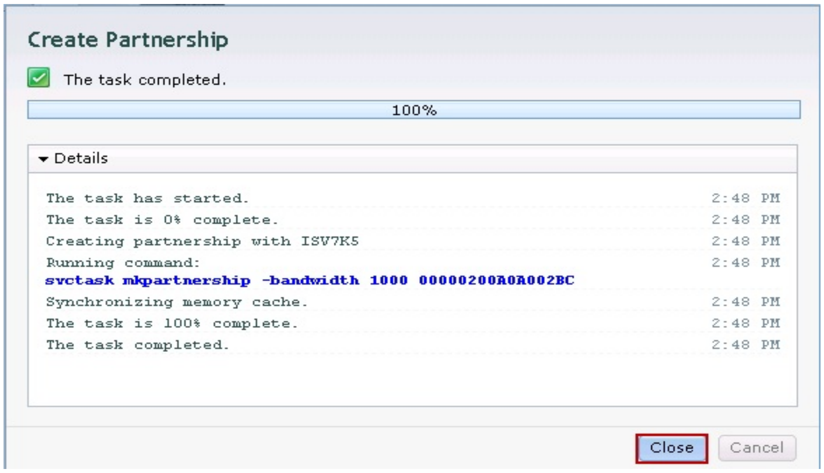

Create Partnership

☑ The task completed.

100%

Details

```
The task has started.                                      2:48 PM
The task is 0% complete.                                   2:48 PM
Creating partnership with ISV7K5                           2:48 PM
Running command:                                           2:48 PM
svctask mkpartnership -bandwidth 1000 00000200A0A002BC
Synchronizing memory cache.                                2:48 PM
The task is 100% complete.                                 2:48 PM
The task completed.                                        2:48 PM
```

Figure 34: Partnership task complete

Note: The recovery site's IBM Storwize V7000 storage system health status shows degraded until partnership is created for recovery site, as shown in Figure 35.

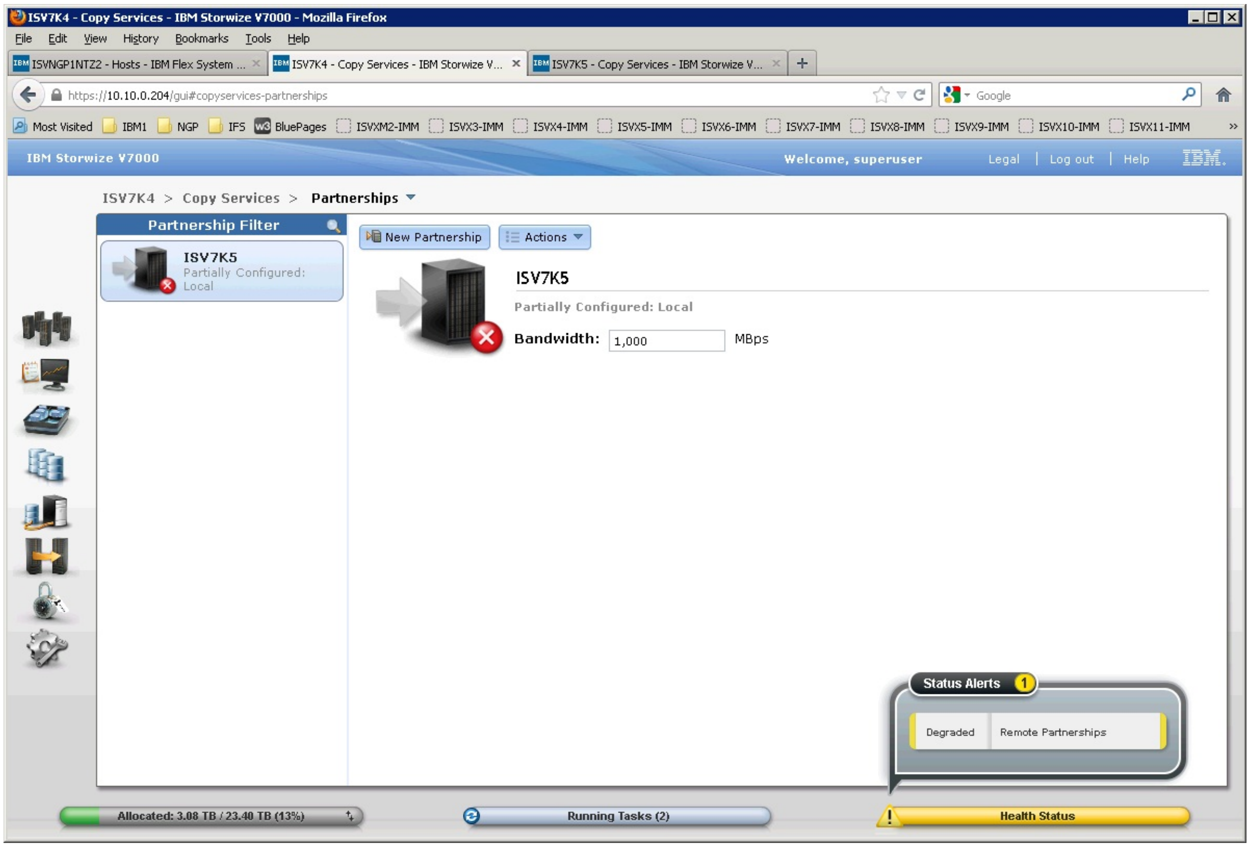

ISV7K4 - Copy Services - IBM Storwize V7000 - Mozilla Firefox
File Edit View History Bookmarks Tools Help
ISVNGP1NTZ2 - Hosts - IBM Flex System ...
ISV7K4 - Copy Services - IBM Storwize V...
ISV7K5 - Copy Services - IBM Storwize V...
https://10.10.0.204/gui#copyservices-partnerships
Google
Most Visited IBM1 NGP IFS BluePages ISVXM2-IMM ISVX3-IMM ISVX4-IMM ISVX5-IMM ISVX6-IMM ISVX7-IMM ISVX8-IMM ISVX9-IMM ISVX10-IMM ISVX11-IMM
IBM Storwize V7000
Welcome, superuser
Legal | Log out | Help
IBM.
ISV7K4 > Copy Services > Partnerships
Partnership Filter
ISV7K5 Partially Configured: Local
New Partnership
Actions

ISV7K5

Partially Configured: Local

Bandwidth: 1,000 MBps

Status Alerts 1

Degraded | Remote Partnerships

Allocated: 3.08 TB / 23.40 TB (13%)

Running Tasks (2)

Health Status

Figure 35: Partnership health status

Create partnership for storage systems (recovery site)

This section describes the partnership configuration on the IBM Storwize V7000 system at the recovery site.

Click **New Partnership** to create a new partnership on the IBM Storwize V7000 system at the recovery site as shown in Figure 36.

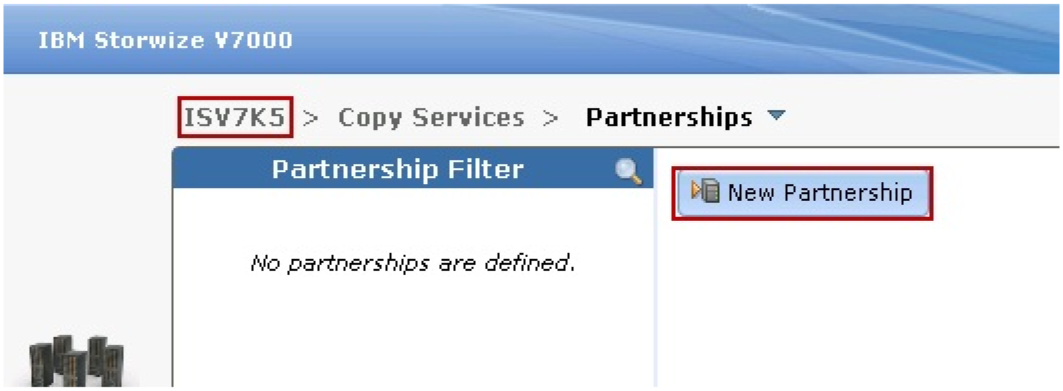

No partnerships are defined.

Figure 36: Create a new partnership

Figure 37 shows the new partnership on protected site storage system has been established with protected site storage system

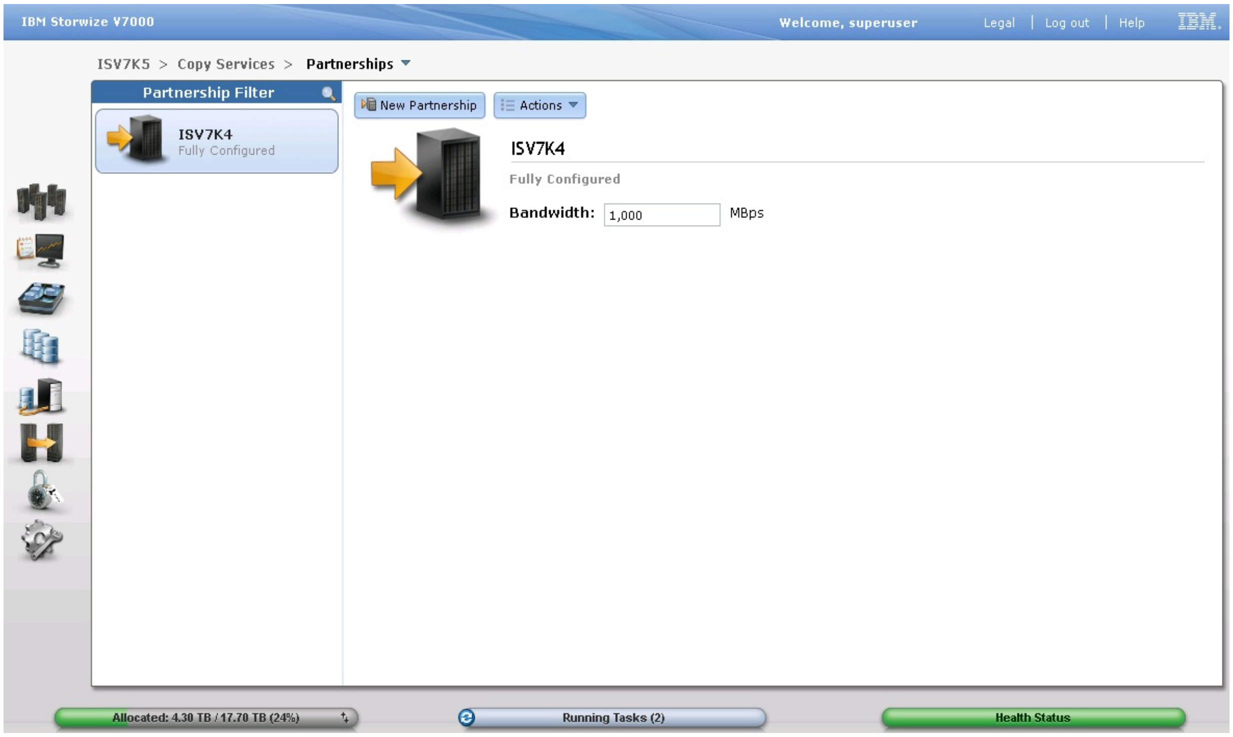

ISV7K4
Fully Configured

ISV7K4
Fully Configured
Bandwidth: 1,000 MBps

Allocated: 4.30 TB / 17.70 TB (24%)
Running Tasks (2)
Health Status

Figure 37: New partnership is established with the protected site storage system

Configure remote copy (Metro Mirror) relationship for storage volumes between protected site and recovery site IBM Storwize V7000 storage systems

To create a new remote copy relationship, click **Copy Services** and then click **Remote Copy** to open the Remote Copy window (as shown in Figure 38).

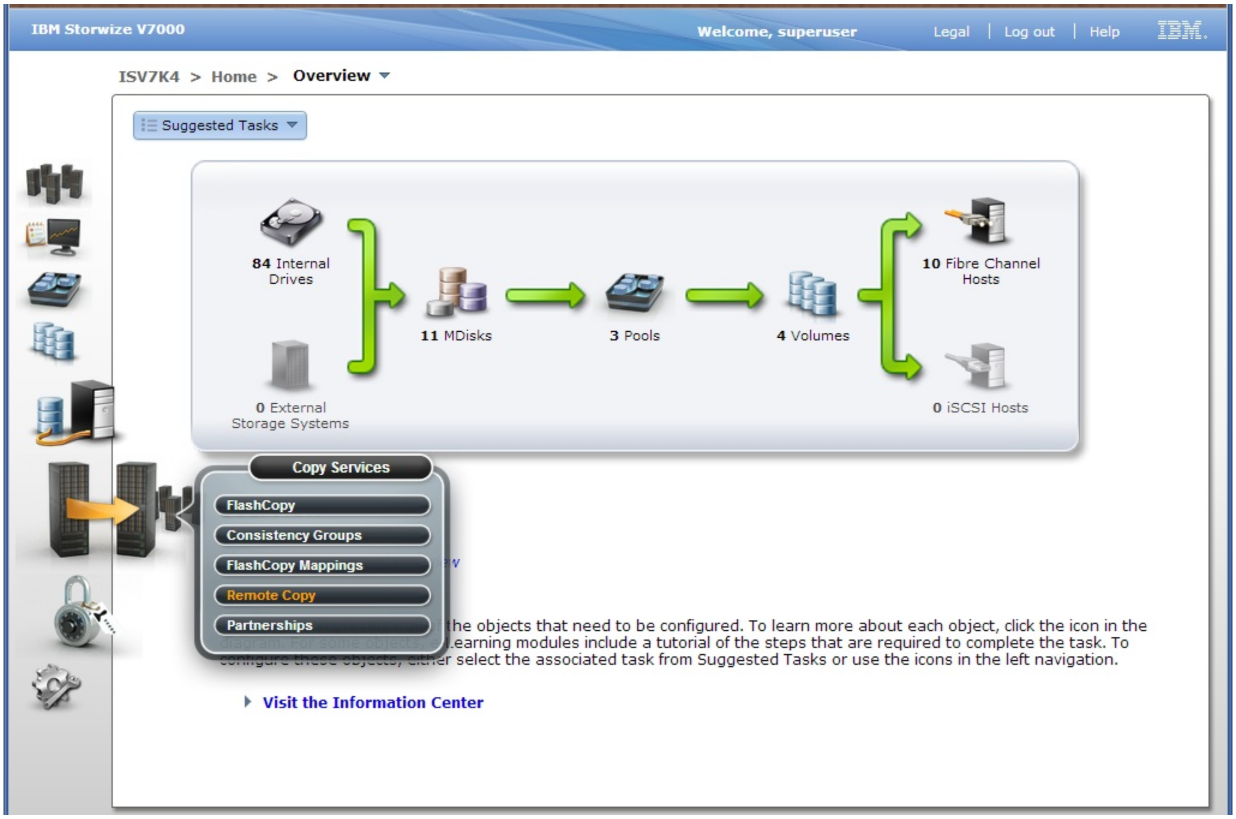

…he objects that need to be configured. To learn more about each object, click the icon in the diagram. For some objects, …earning modules include a tutorial of the steps that are required to complete the task. To configure these objects, either select the associated task from Suggested Tasks or use the icons in the left navigation.

Figure 38: Open the Remote Copy window

1. Click **New Relationship** to create a Metro Mirror relationship for the storage volumes between the IBM Storwize V7000 storage systems (as shown in Figure 39).

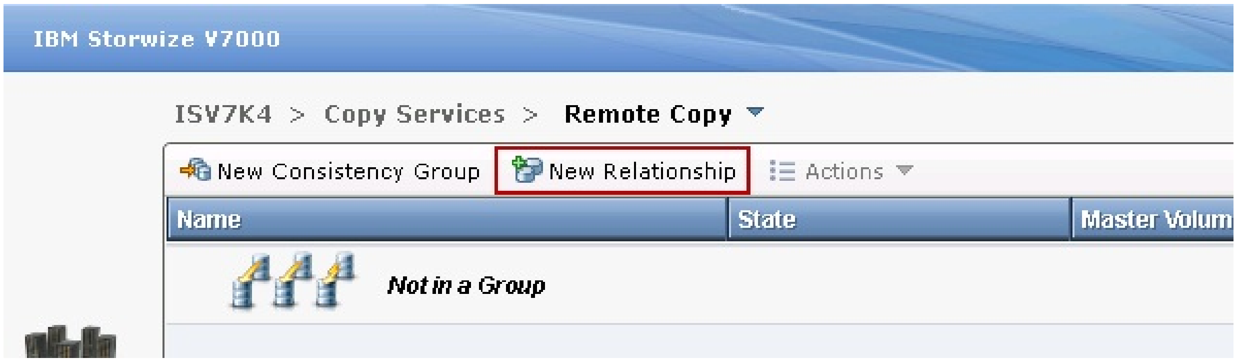

Figure 39: Remote Copy window

2. Select **Metro Mirror** as the relationship type (as shown in Figure 40) and click **Next**.

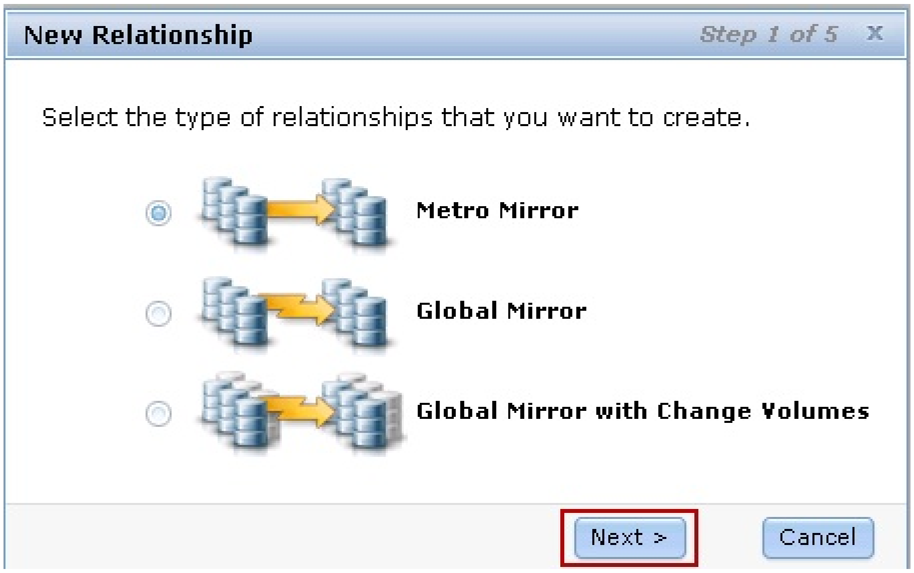

New Relationship — Step 1 of 5

Select the type of relationships that you want to create.

● Metro Mirror

○ Global Mirror

○ Global Mirror with Change Volumes

Figure 40: Metro Mirror relationship

3. Specify the location of the auxiliary volumes and click **Next**, as shown in Figure 41.

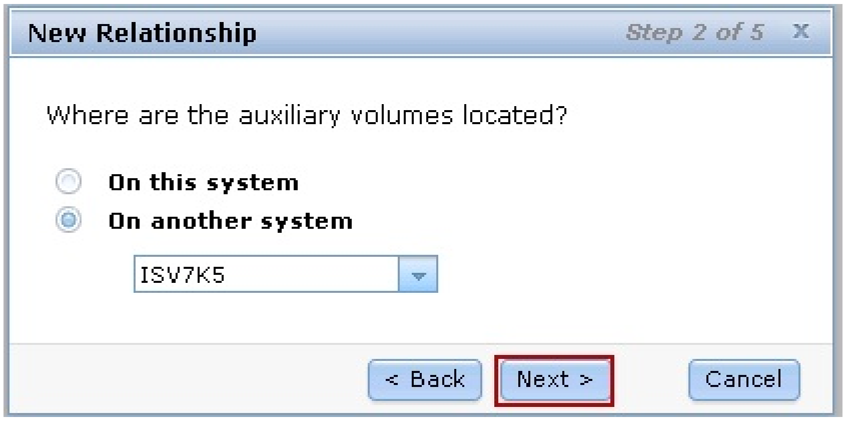

New Relationship — Step 2 of 5

Where are the auxiliary volumes located?

○ On this system

● On another system

ISV7K5

Figure 41: Setting the auxiliary volume location

4. Configure the master and auxiliary volumes relationships by clicking **Add** (as shown in Figure 42, Figure 43, Figure 44, and Figure 45, then click **Next**.

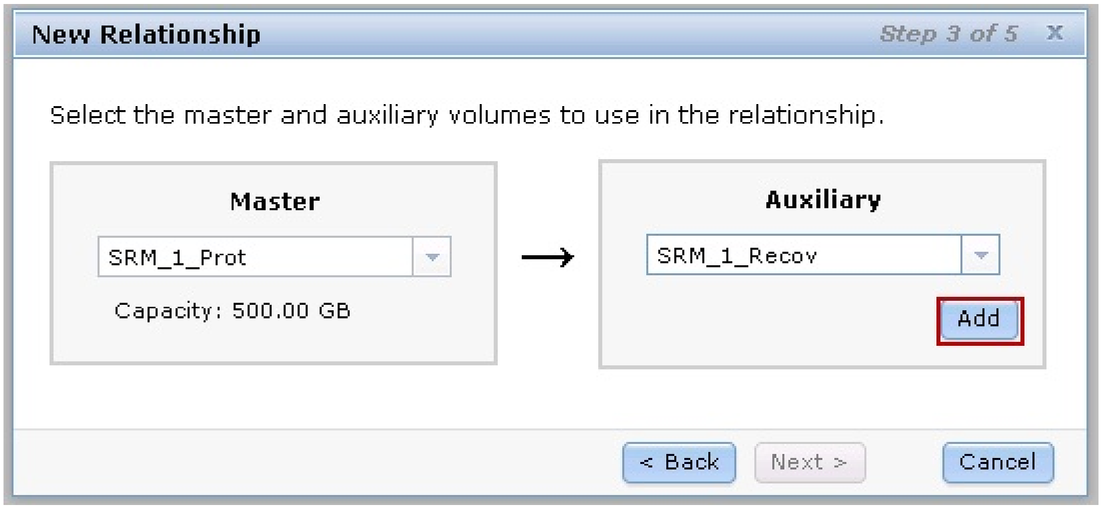

New Relationship — Step 3 of 5

Select the master and auxiliary volumes to use in the relationship.

Master: SRM_1_Prot

Capacity: 500.00 GB

Auxiliary: SRM_1_Recov

Figure 42: Setting remote copy relationship

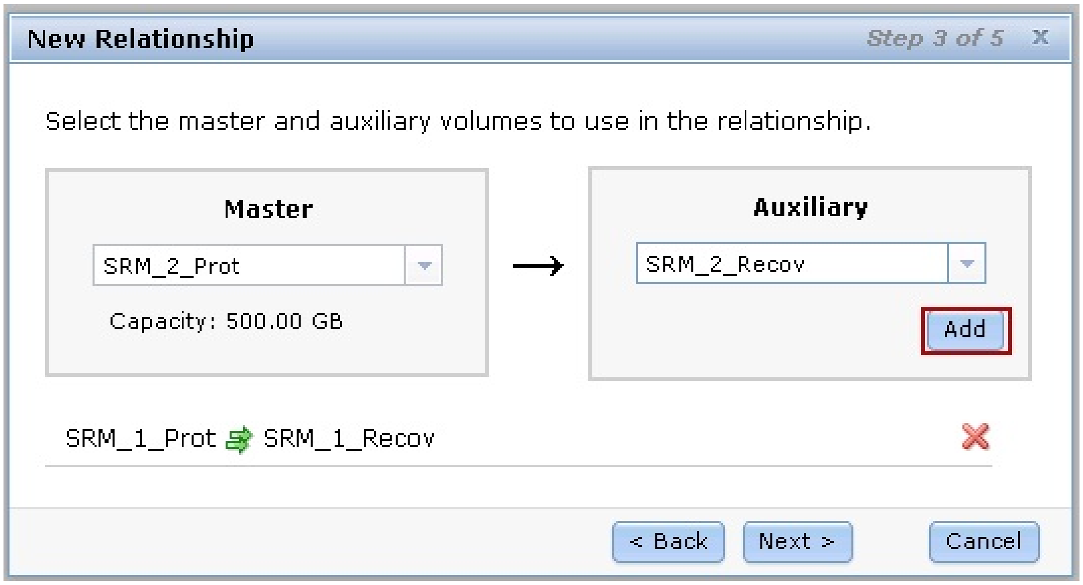

New Relationship — Step 3 of 5

Select the master and auxiliary volumes to use in the relationship.

Master: SRM_2_Prot

Capacity: 500.00 GB

Auxiliary: SRM_2_Recov

SRM_1_Prot → SRM_1_Recov

Figure 43: Setting remote copy relationship (continued)

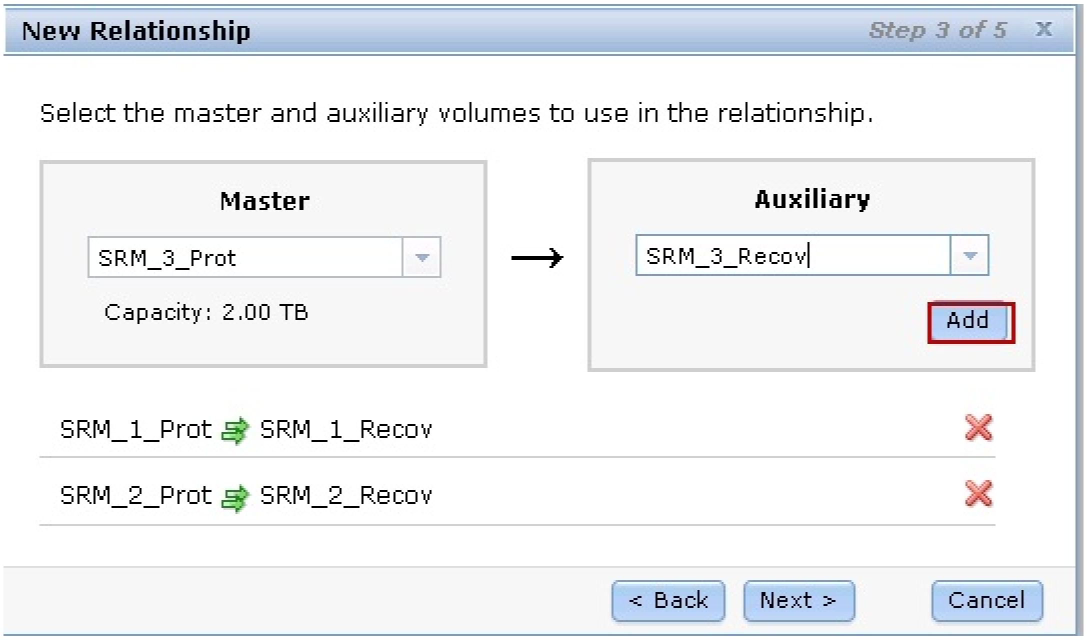
New Relationship

Step 3 of 5

Select the master and auxiliary volumes to use in the relationship.

Master: SRM_3_Prot

Capacity: 2.00 TB

Auxiliary: SRM_3_Recov

SRM_1_Prot → SRM_1_Recov

SRM_2_Prot → SRM_2_Recov

Figure 44: Setting remote copy relationship (continued)

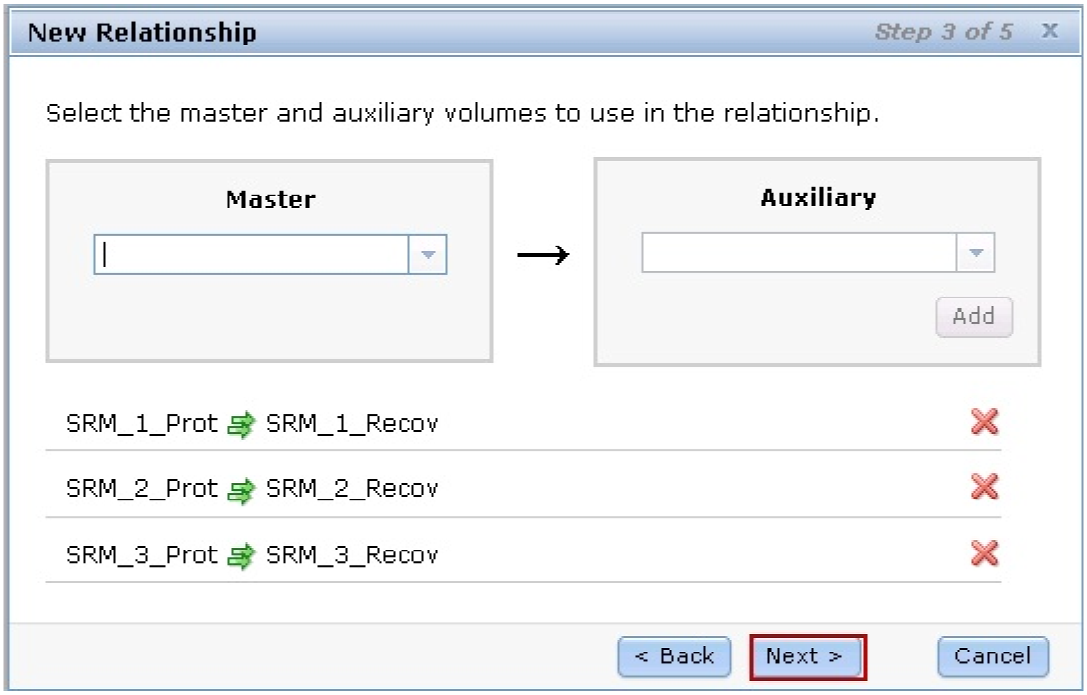
New Relationship

Step 3 of 5

Select the master and auxiliary volumes to use in the relationship.

Master

Auxiliary

SRM_1_Prot → SRM_1_Recov

SRM_2_Prot → SRM_2_Recov

SRM_3_Prot → SRM_3_Recov

Figure 45: Setting remote copy relationship (continued)

5. A window opens and asks if the volumes in the relationship are already synchronized. In most situations, the data on the master volume and on the auxiliary volume are not identical. So, select **No, the volumes are not synchronized** and click **Next** to enable an initial copy, as shown in Figure 46.

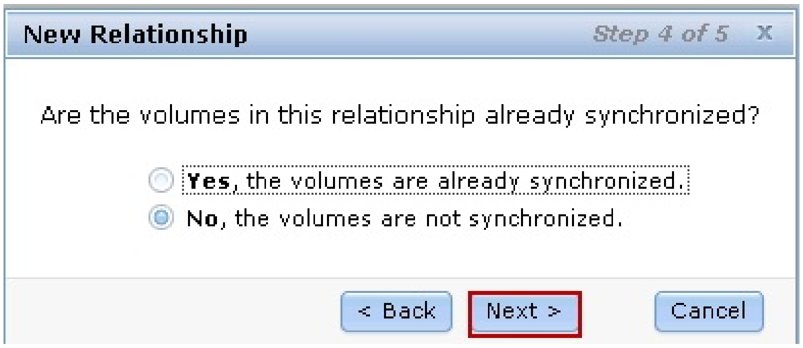

New Relationship — Step 4 of 5

Are the volumes in this relationship already synchronized?

○ **Yes**, the volumes are already synchronized.

● **No**, the volumes are not synchronized.

Figure 46: Activate an initial data copy

6. Select **Yes, start copying now** and click **Finish**, as shown in Figure 47.

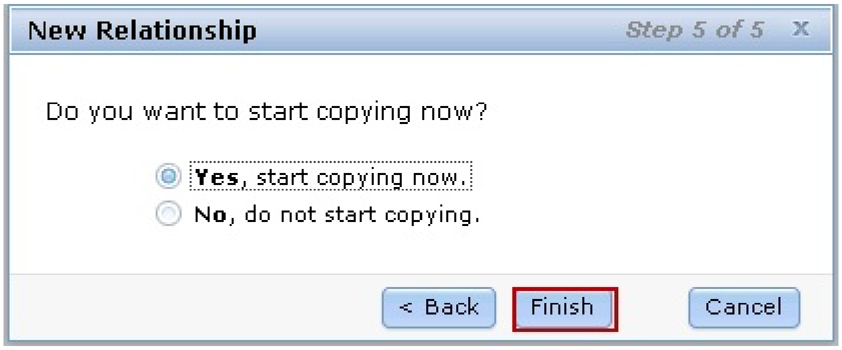

New Relationship — Step 5 of 5

Do you want to start copying now?

● **Yes**, start copying now.

○ **No**, do not start copying.

Figure 47: Specifying if you want to start copying now or later

Create a consistency group and place the Metro Mirror volumes in a consistency group

A consistency group is a logical entity that groups copy relationships. By grouping the relationships, you can ensure that these relationships are managed in unison and the data within the group is in a consistent state. The following steps show how the test team configured the consistency group for the Global Mirror volumes in the lab solution test setup.

1. In the IBM Storwize V7000 window, click **Copy Services** and then click **Remote Copy**.
2. Click **New Consistency Group**, as shown in Figure 48.

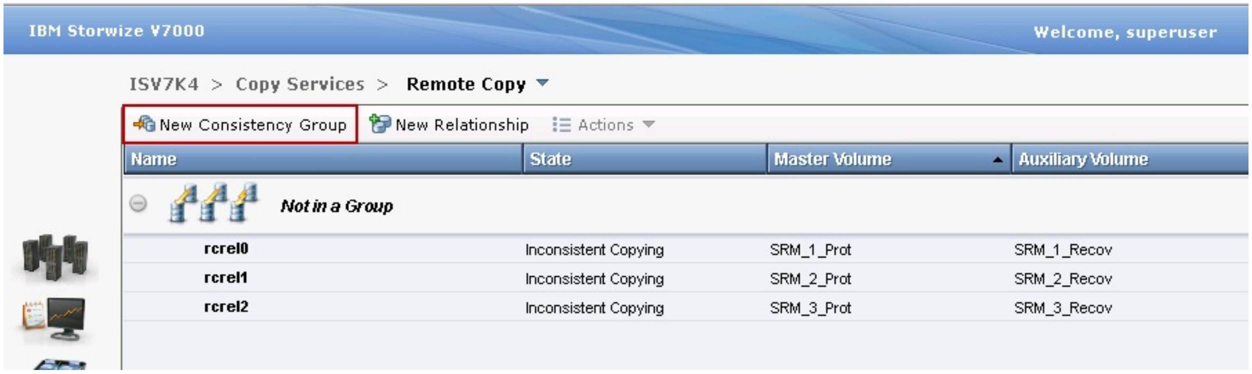

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
rcrel0	Inconsistent Copying	SRM_1_Prot	SRM_1_Recov
rcrel1	Inconsistent Copying	SRM_2_Prot	SRM_2_Recov
rcrel2	Inconsistent Copying	SRM_3_Prot	SRM_3_Recov

Figure 48: New consistency group

3. Enter a name for the consistency group, and then click **Next (**as shown in Figure 49).

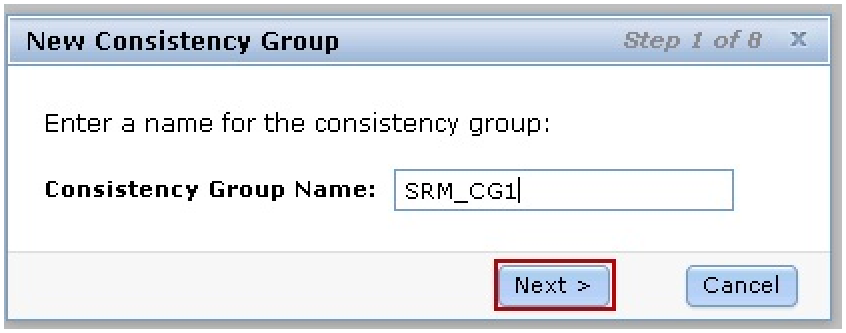

Enter a name for the consistency group:

Consistency Group Name: SRM_CG1

Figure 49: Enter a name for the new consistency group

4. Select **On another system**, to specify the location of the auxiliary volumes, as shown in Figure 50.

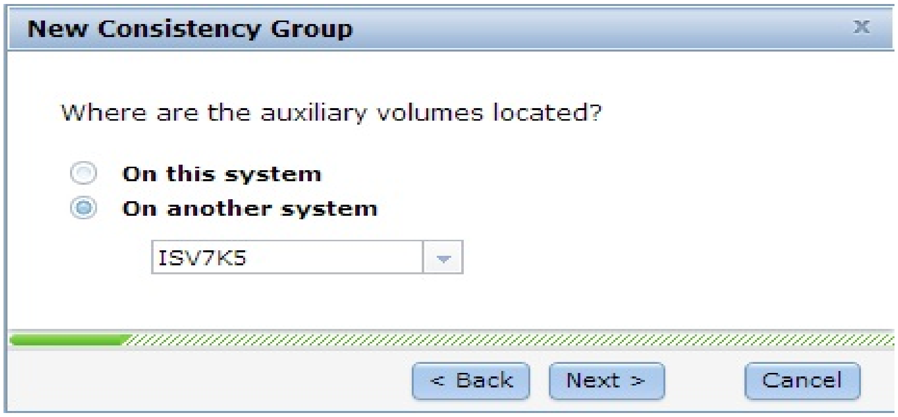

New Consistency Group

Where are the auxiliary volumes located?

○ On this system

● On another system

ISV7K5

Figure 50: Auxiliary volumes location

5. Select **Metro Mirror** as the relationship type and click **Next**, as shown in Figure 51.

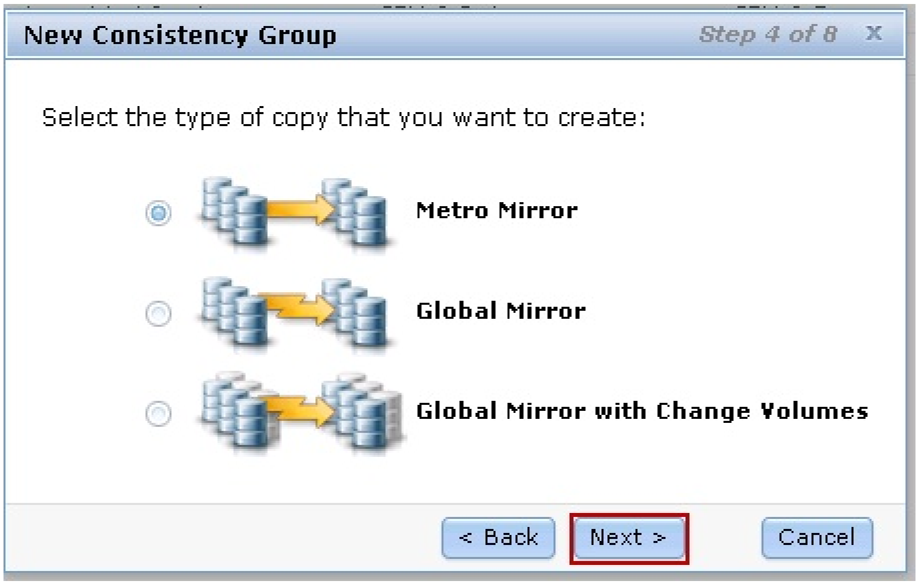

New Consistency Group — *Step 4 of 8*

Select the type of copy that you want to create:

● Metro Mirror

○ Global Mirror

○ Global Mirror with Change Volumes

Figure 51: Select Metro Mirror relationship

6. As shown in Figure 52, select the existing relationships to add to the group and then click **Next**.

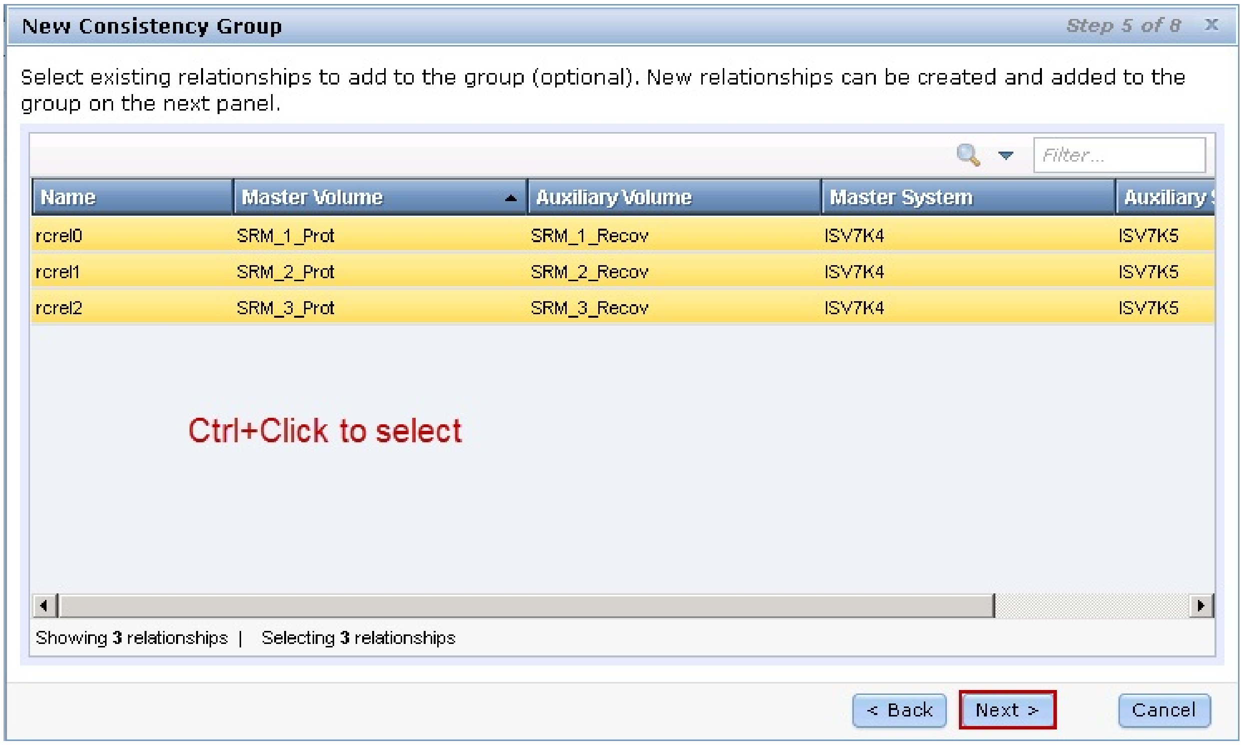

Select existing relationships to add to the group (optional). New relationships can be created and added to the group on the next panel.

Name	Master Volume	Auxiliary Volume	Master System	Auxiliary S
rcrel0	SRM_1_Prot	SRM_1_Recov	ISV7K4	ISV7K5
rcrel1	SRM_2_Prot	SRM_2_Recov	ISV7K4	ISV7K5
rcrel2	SRM_3_Prot	SRM_3_Recov	ISV7K4	ISV7K5

Showing **3** relationships | Selecting **3** relationships

Figure 52: Select the existing relationships to add to the group

7. Click **Next** to continue or add additional master and auxiliary volume to create new relationships, as shown in Figure 53.

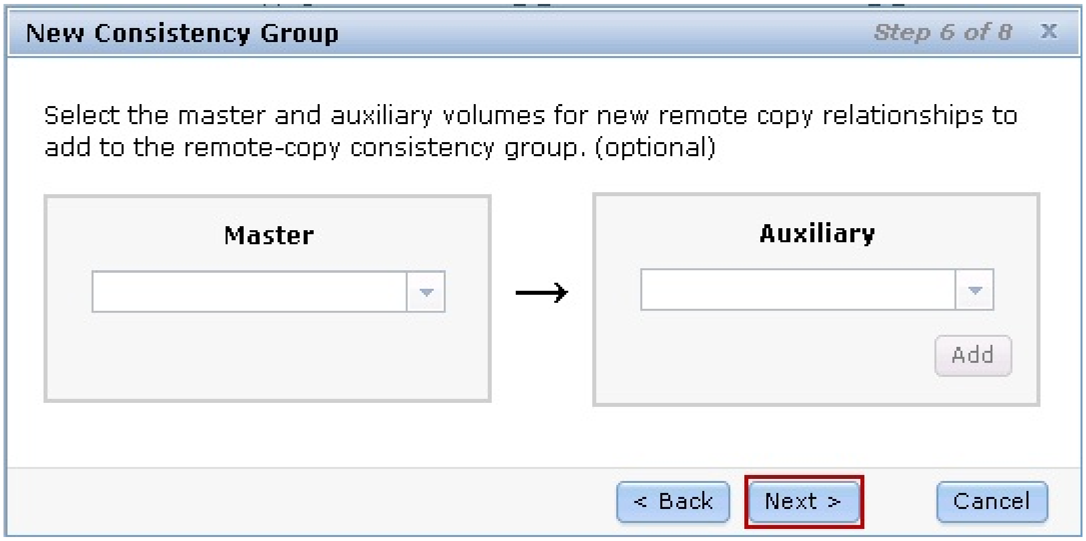

Select the master and auxiliary volumes for new remote copy relationships to add to the remote-copy consistency group. (optional)

Figure 53: Create relationships between the master and auxiliary volumes

8. Select **No, the volumes are not synchronized** and click **Next**, as shown in Figure 54.

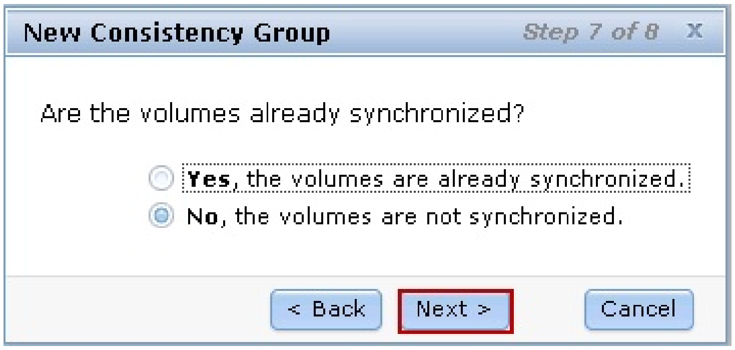

New Consistency Group

Step 7 of 8

Are the volumes already synchronized?

○ **Yes**, the volumes are already synchronized.

● **No**, the volumes are not synchronized.

Figure 54: Volume synchronization option

9. Select **Yes, start copying now** and click **Finish** to complete the consistency group configuration, as shown in Figure 55.

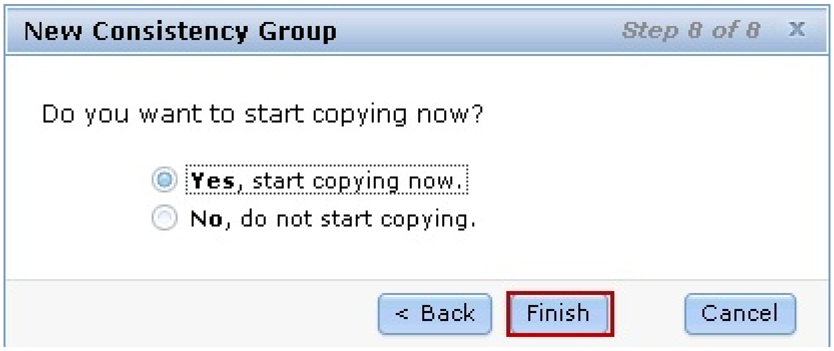

New Consistency Group

Step 8 of 8

Do you want to start copying now?

● **Yes**, start copying now.

○ **No**, do not start copying.

Figure 55: Start copying now after configuring new consistency group

10. Validate that the consistency group is successfully created and then click **Close** to return to the remote copy window, as shown in Figure 56.

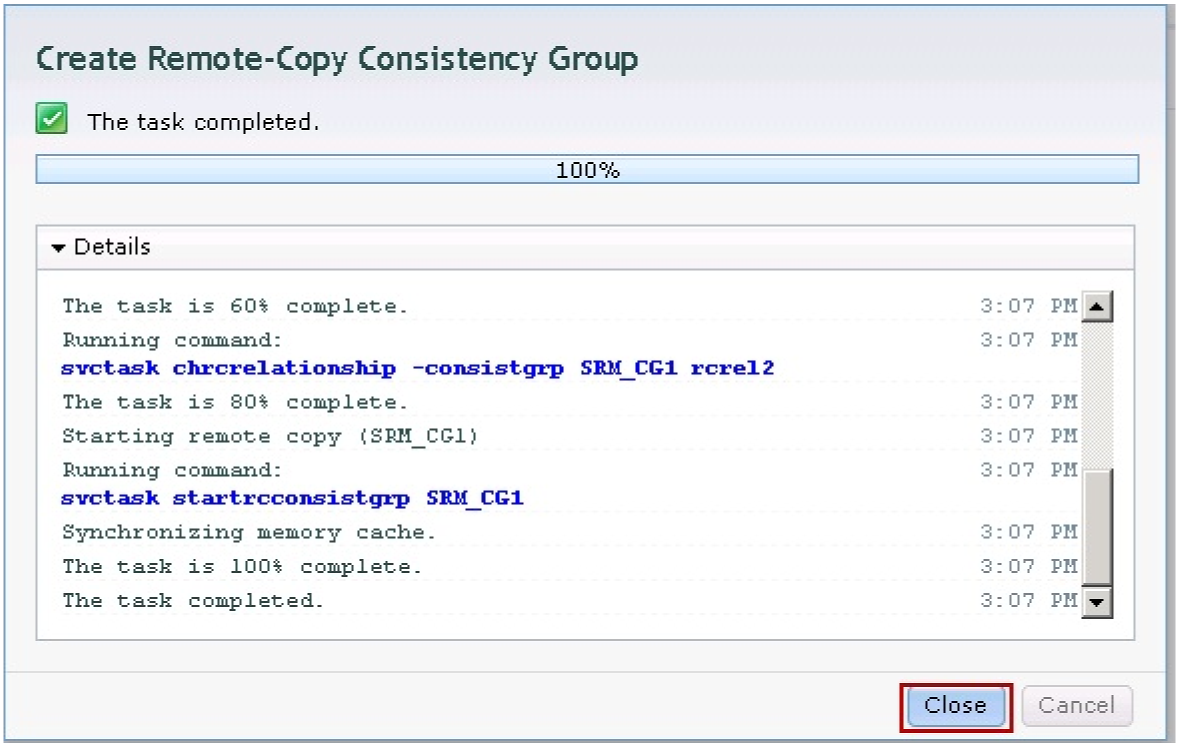

```
The task is 60% complete.                        3:07 PM
Running command:                                 3:07 PM
svctask chrcrelationship -consistgrp SRM_CG1 rcrel2
The task is 80% complete.                        3:07 PM
Starting remote copy (SRM_CG1)                   3:07 PM
Running command:                                 3:07 PM
svctask startrcconsistgrp SRM_CG1
Synchronizing memory cache.                      3:07 PM
The task is 100% complete.                       3:07 PM
The task completed.                              3:07 PM
```

Figure 56: Consistency group created successfully

Note: Metro Mirror logical unit numbers (LUNs) will not be available until copying is complete.

Installing and configuring disaster recovery solution using VMware Site Recovery Manager

This section illustrates important configuration and installation tasks to be performed for the enterprise virtual infrastructure disaster recovery solution using VMware Site Recovery Manager.

- Configure vSphere HA.
- Install the Site Recovery Manager plug-in.
- Install SRA.
- Configure the Site Recovery Manager plug-in.

Requirements for vSphere HA cluster

This section explains the prerequisites and the procedure for configuring vSphere HA.

Prerequisites

The following prerequisites must be met for configuring vSphere HA.

- All hosts must be licensed for vSphere HA.
- You need at least two hosts in the cluster.
- All hosts need to be configured with static IP addresses. If you are using Dynamic Host Configuration Protocol (DHCP), you must ensure that the address for each host persists across restarts.
- There should be at least one management network in common among all hosts and the best practice is to have at least two. Management networks differ depending on the version of the host that you are using.

For more detailed information about configuring the vSphere HA, refer:

http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-51-availability-guide.pdf

Network path redundancy

Network path redundancy between cluster nodes is important for vSphere HA reliability. A single service console network ends up being a single point of failure, and can result in failover, although only the network might have failed.

If you have only one management network, any failure between the host and the cluster can cause an unnecessary (or false) failover activity if heartbeat data store connectivity is not retained during the network failure. Possible failures include network interface card (NIC) failures, network cable failures, network cable removal, and switch resets. Consider these possible sources of failure between hosts and try to minimize them, typically by providing network redundancy.

You can implement network redundancy at the NIC level with NIC teaming, or at the management network level. In most implementations, NIC teaming provides sufficient redundancy, but you can use or add management network redundancy, if required. Redundant management networking allows reliable detection of failures and prevents isolation or partition conditions from occurring, because heartbeats can be sent over multiple networks.

Configure the fewest possible number of hardware segments between the servers in a cluster. The goal is to limit single points of failure. Additionally, routes with too many hops can cause networking packet delays for heartbeats and increase the possible points of failure. Figure 57 shows the network path redundancy configuration.

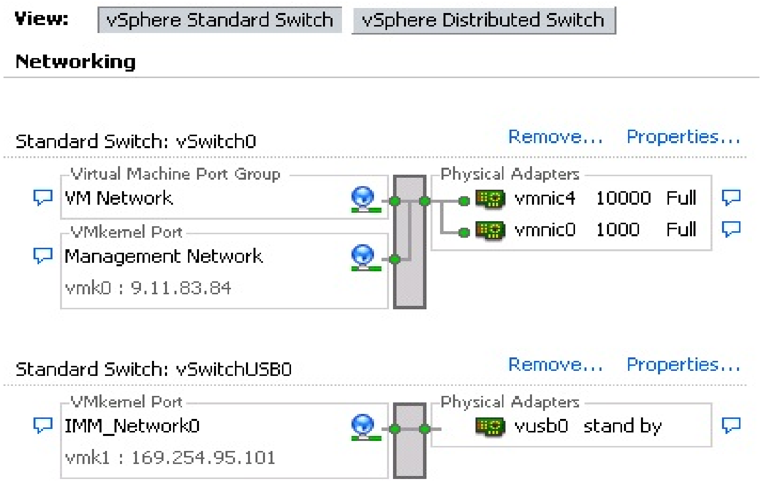

Figure 57: Network path redundancy configuration

Steps to configure vSphere HA

Perform the following steps to configure vSphere HA.

1. Connect to vCenter Server 5.x using vSphere client 5.x with appropriate user credentials.
2. Select the **Hosts and Clusters** view.
3. Right-click the data center in the Inventory tree and click **New Cluster**.
4. Complete the **New Cluster** wizard. Do not enable vSphere HA (or DRS) at this time.
5. Click **Finish** to close the wizard and create the cluster.
6. Based on your plan for the resources and networking architecture of the cluster, use the vSphere Client to add hosts to the cluster.
7. Right-click the cluster and click **Edit Settings**.
   The cluster's **Settings** dialog box is where you can modify the vSphere HA (and other) settings for the cluster.
8. On the **Cluster Features** page, select **Turn On vSphere HA**.
9. Configure the following vSphere HA settings as appropriate for your cluster.
   - Host Monitoring Status
   - Admission Control
   - Virtual Machine Options
   - VM Monitoring
10. Click **OK** to close the cluster's **Settings** dialog box.

For more information about configuring vSphere HA, refer to the *vSphere Availability Guide* at: http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-51-availability-guide.pdf

Figure 58 shows the solution test lab protection site vSphere HA configuration.

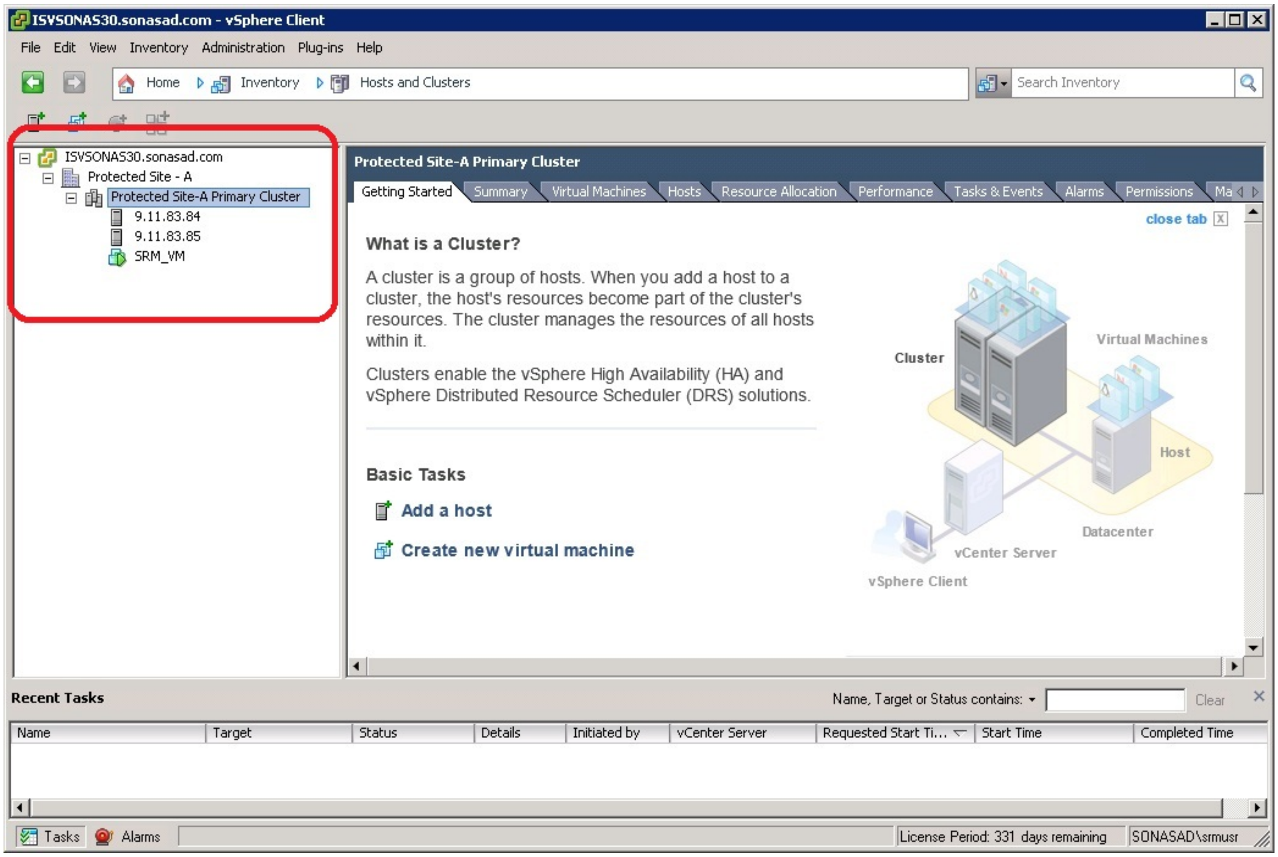

What is a Cluster?

A cluster is a group of hosts. When you add a host to a cluster, the host's resources become part of the cluster's resources. The cluster manages the resources of all hosts within it.

Clusters enable the vSphere High Availability (HA) and vSphere Distributed Resource Scheduler (DRS) solutions.

Basic Tasks

Add a host

Create new virtual machine

Figure 58: Protection site vSphere HA configuration

Install the Site Recovery Manager server and the Site Recovery Manager plug-in

This section provides the procedures used to configure and install the VMware Site Recovery Manager in the lab solution validation setup.

1. Create a Site Recovery Manager database. In the solution validation setup, Microsoft SQL Server 2008 R2 has been configured to create the Site Recovery Manager database.
2. Download the Site Recovery Manager server from http://downloads.vmware.com/ and install the Site Recovery Manager Server at the protected site and at the recovery site.
3. Start the vSphere Client and connect to the vCenter Server at either the protected or the recovery site.
4. On the vSphere Client menu bar, click **Plug-ins** and then click **Manage Plug-ins**.
5. In the Plug-In Manager window, in the **Available Plug-ins** section, locate **VMware vCenter Site Recover Manager Extension** and click **Download and Install**.
6. Review and accept the certificate.
   Note: This step only occurs if you use certificate-based authentication.
7. After the download is complete, click **Run** to start the installation wizard, select the installation language, and click **OK**.
8. Click **Next** to start the installation and then click **Next** again on the VMware Patents page.
9. Select **I accept the terms in the license agreement** and then click **Next**.
10. Click **Install**.
11. After the installation is complete, click **Finish**.

A new Site Recovery icon appears on the Home page in the vSphere Client, as shown in Figure 59.

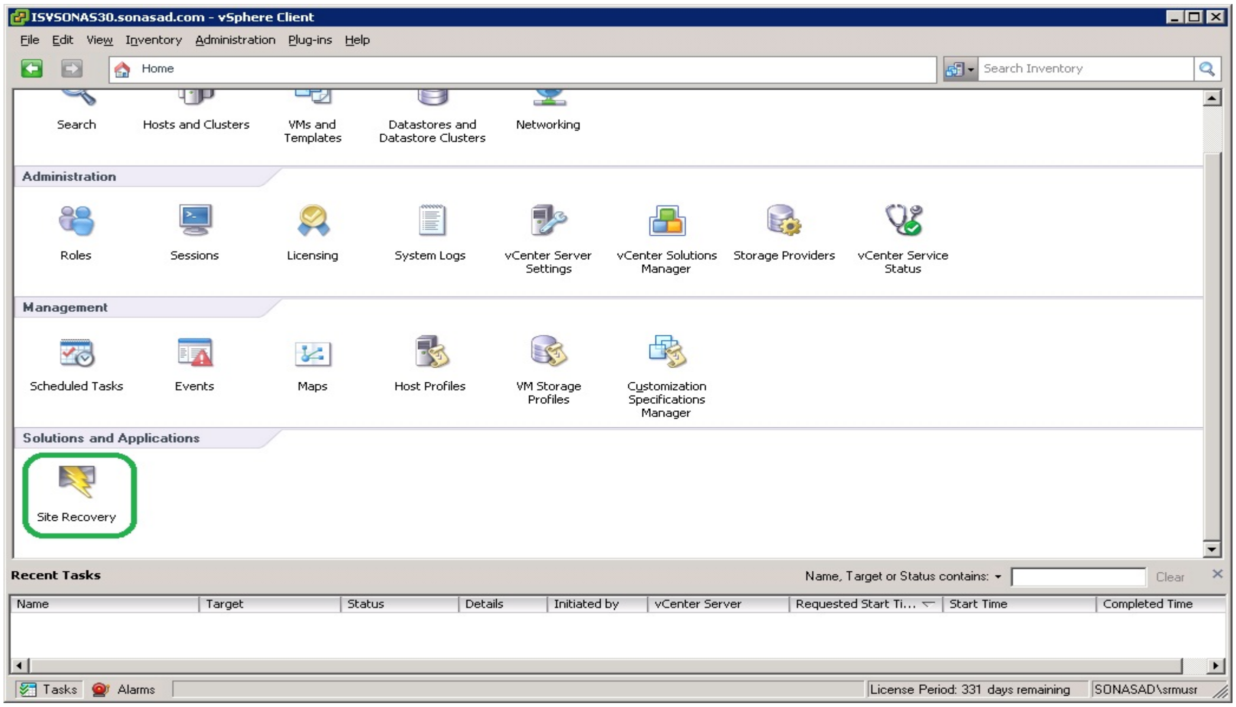

Figure 59: Site Recovery Manager plug-in

Connect the protected and recovery sites

It is important to connect the protected and the recovery sites before using VMware Site Recovery Manager. The sites must authenticate with each other. This is known as site pairing.

Note: Before pairing the protected and recovery sites, make sure that you have correctly installed the appropriate VMware vCenter Site Recovery Manager license keys.

For more information on pairing protected and recovery sites, refer to the VMware *Site Recovery Manager Installation and Configuration* guide at:

http://pubs.vmware.com/srm-51/topic/com.vmware.ICbase/PDF/srm-install-config-5-1.pdf

Figure 60 shows the site connection (pairing) configuration of the solution lab validation setup.

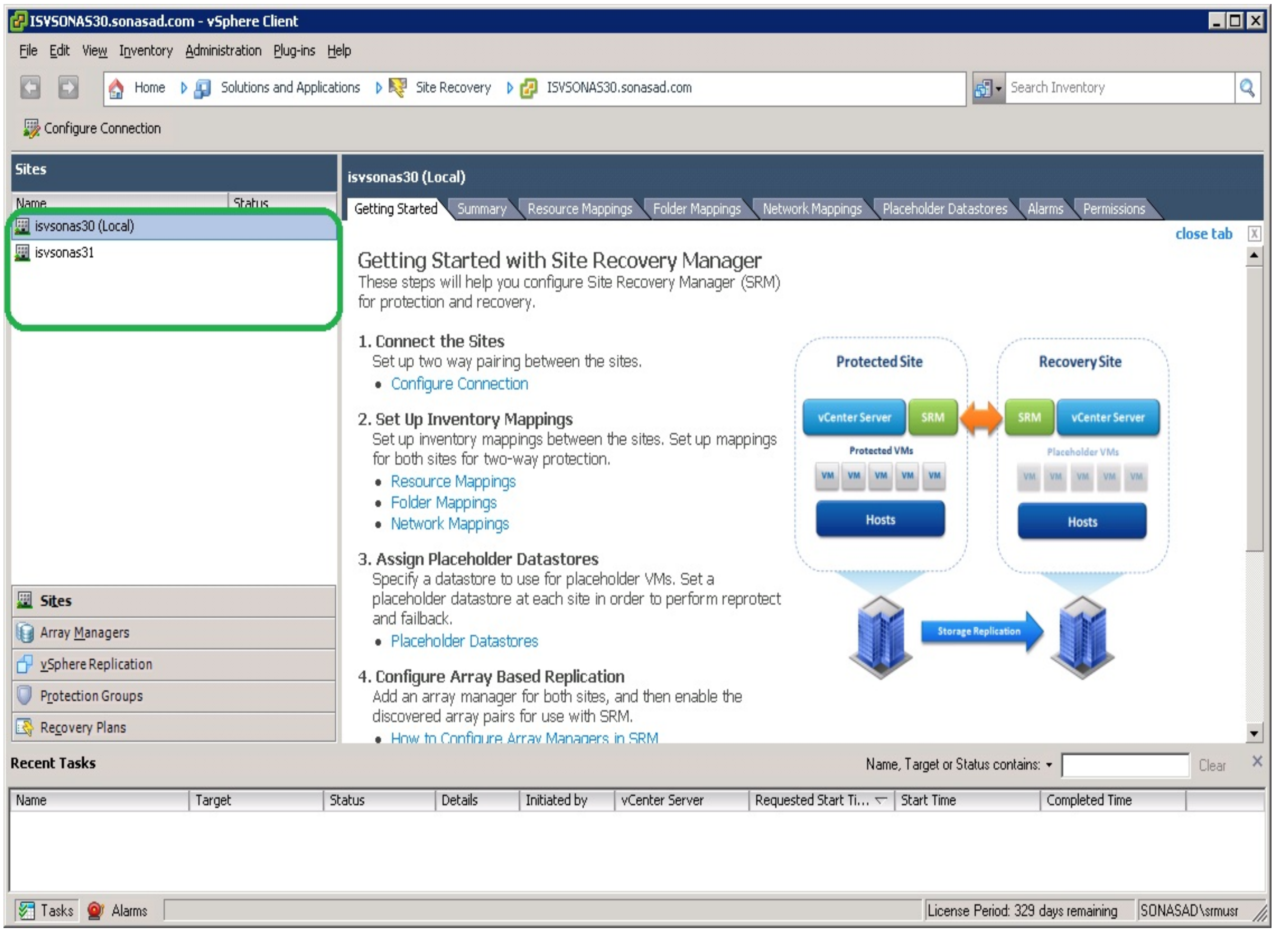

Getting Started with Site Recovery Manager

These steps will help you configure Site Recovery Manager (SRM) for protection and recovery.

1. Connect the Sites
   Set up two way pairing between the sites.
   - Configure Connection

2. Set Up Inventory Mappings
   Set up inventory mappings between the sites. Set up mappings for both sites for two-way protection.
   - Resource Mappings
   - Folder Mappings
   - Network Mappings

3. Assign Placeholder Datastores
   Specify a datastore to use for placeholder VMs. Set a placeholder datastore at each site in order to perform reprotect and failback.
   - Placeholder Datastores

4. Configure Array Based Replication
   Add an array manager for both sites, and then enable the discovered array pairs for use with SRM.
   - How to Configure Array Managers in SRM

Figure 60: Site connection configuration

Installing and configuring SRA for IBM Storwize V7000

This section illustrates the procedure for installing and configuring SRA for IBM Storwize V7000.

1. Download the SRA from http://www.vmware.com/download/srm/.
2. Install SRA on both protection and recovery sites. **Note:** It is important to install and configure Site Recovery Manager before installing SRA for IBM Storwize V7000. For further information, refer to the user guide that comes with SRA for IBM Storwize V7000.
3. After successful installation of SRA for IBM Storwize V7000, run the configuration utility on both protection and recovery sites of the Site Recovery Manager server. Figure 61 shows the SRA configuration utility. Select **Standard** as the volume type and enter the appropriate value in the **Test MDisk Group ID** field.
   Note: The Test MDisk Group ID option instructs SRA to which MDisk group on the recovery site IBM System Storage SAN Volume Controller must be used to create FlashCopy (target) volumes during test recovery operation. For more information, refer: http://pic.dhe.ibm.com/infocenter/strhosts/ic/topic/com.ibm.help.strghosts.doc/SVC%20SRA/2.1.0/PDFs/SVC_Adapter_for_VMware_VC_SRM_2.1.0_UG.pdf

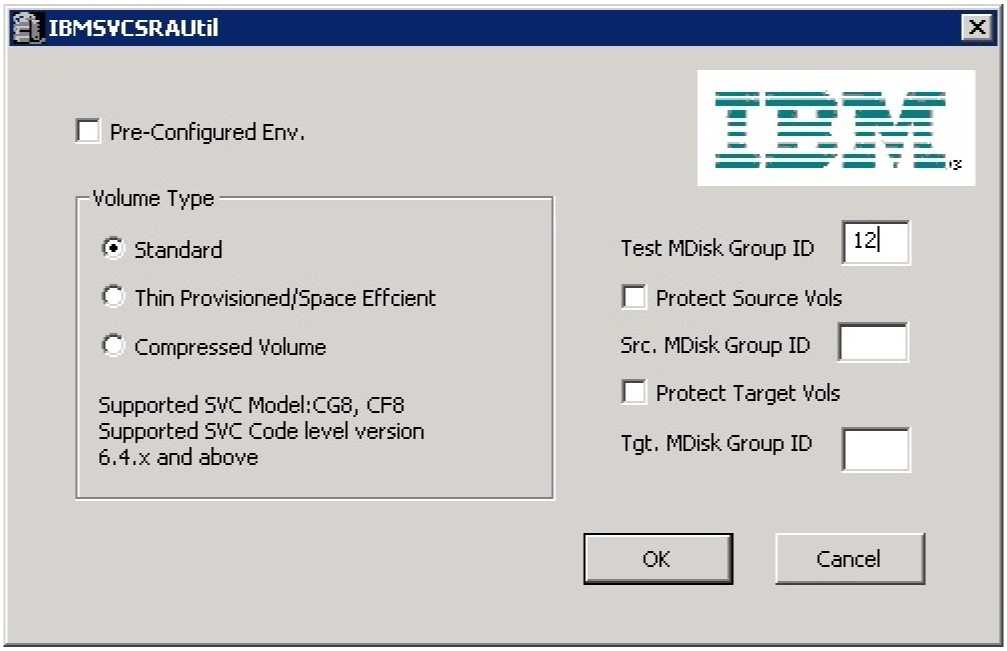

IBMSVCSRAUtil

□ Pre-Configured Env.

Volume Type

● Standard

○ Thin Provisioned/Space Effcient

○ Compressed Volume

Supported SVC Model:CG8, CF8
Supported SVC Code level version
6.4.x and above

Test MDisk Group ID: 12

□ Protect Source Vols

Src. MDisk Group ID:

□ Protect Target Vols

Tgt. MDisk Group ID:

Figure 61: SRA configuration utility

4. Using vSphere client, connect to Site Recovery Manager and select the array managers in the left pane. Click the **SRAs** tab and click **Rescan SRAs**. This refreshes the SRA information, allowing Site Recovery Manager to discover the SRA (as shown in Figure 62).

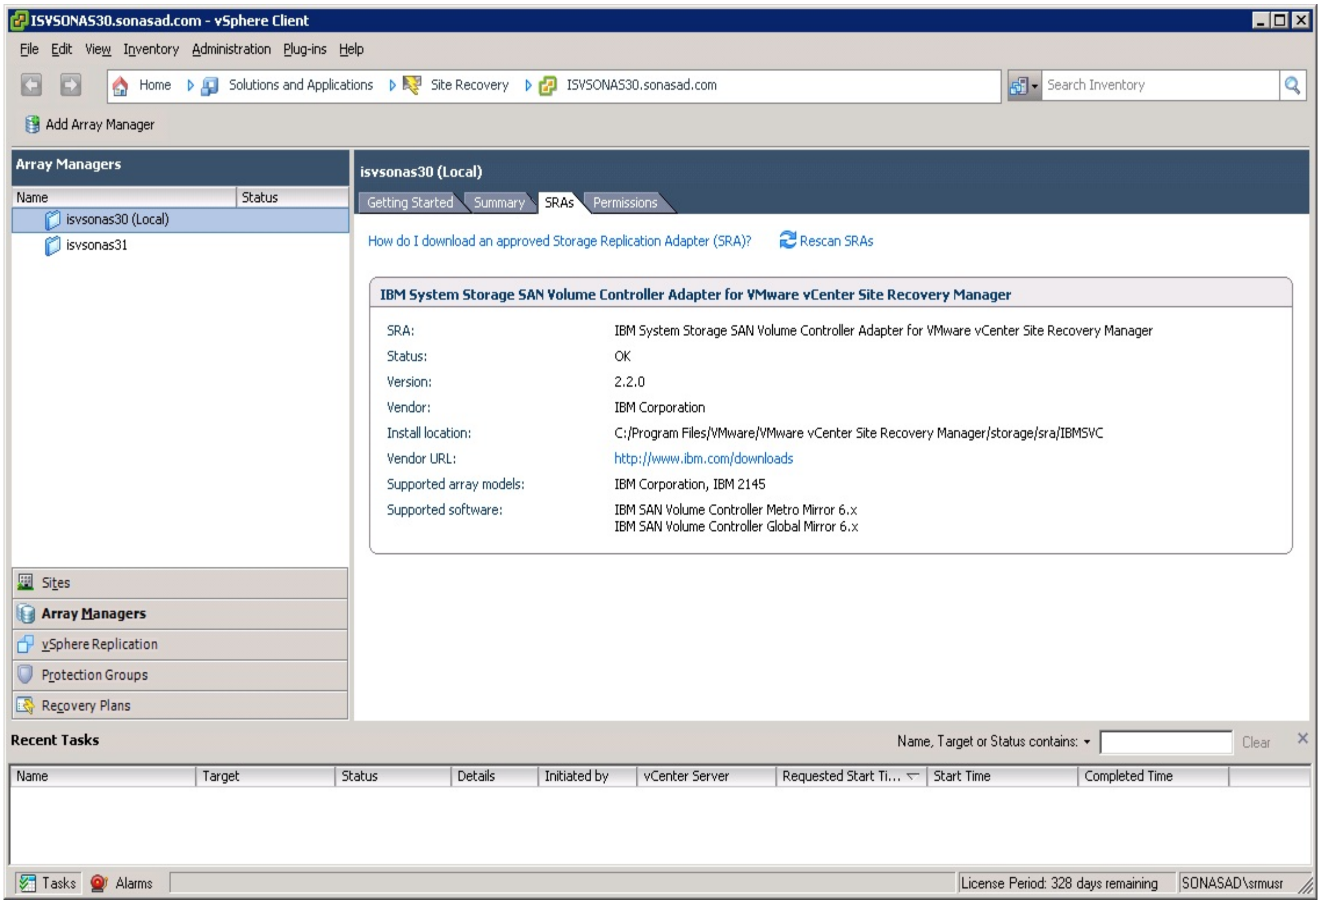

IBM System Storage SAN Volume Controller Adapter for VMware vCenter Site Recovery Manager

SRA:	IBM System Storage SAN Volume Controller Adapter for VMware vCenter Site Recovery Manager
Status:	OK
Version:	2.2.0
Vendor:	IBM Corporation
Install location:	C:/Program Files/VMware/VMware vCenter Site Recovery Manager/storage/sra/IBMSVC
Vendor URL:	http://www.ibm.com/downloads
Supported array models:	IBM Corporation, IBM 2145
Supported software:	IBM SAN Volume Controller Metro Mirror 6.x IBM SAN Volume Controller Global Mirror 6.x

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Completed Time

Figure 62: SRA for IBM Storwize V7000 installation status

Configuring array managers

After you pair the protected site and the recovery site as shown earlier, configure their respective array managers so that Site Recovery Manager can discover replicated devices, compute data store groups, and initiate storage operations.

Note: Before you configure array managers, it is mandatory to:

- Connect both protection and recovery sites.
- Install and configure SRA at both protection and recovery sites.

For more information about how to configure array managers, refer to the VMware *Site Recovery Manager Installation and Configuration* guide at: http://pubs.vmware.com/srm-51/topic/com.vmware.ICbase/PDF/srm-install-config-5-1.pdf

Note: It is important to provide the primary and remote IBM Storwize V7000 storage systems while configuring array managers (as shown in Figure 63).

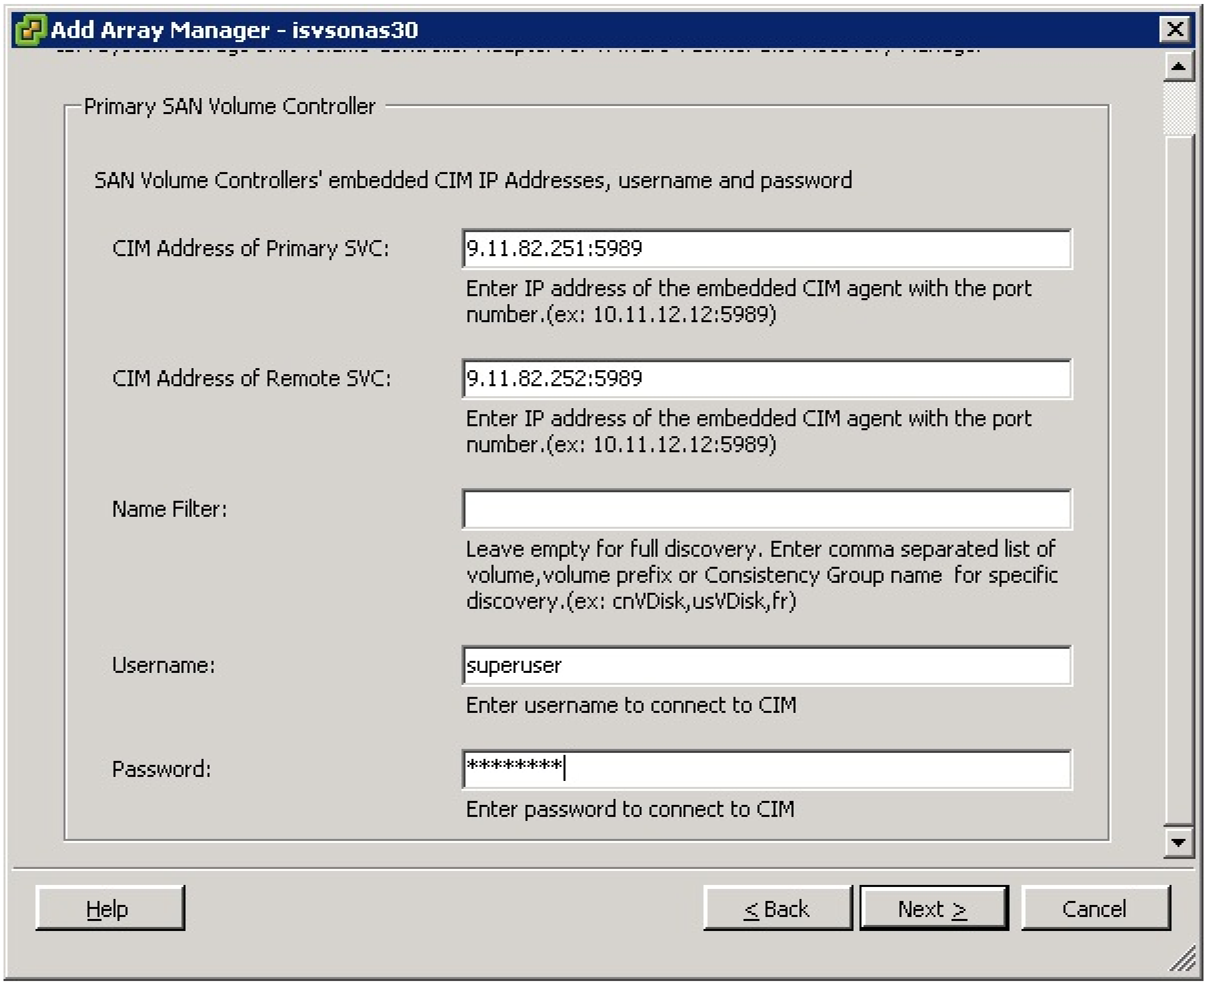

Primary SAN Volume Controller

SAN Volume Controllers' embedded CIM IP Addresses, username and password

CIM Address of Primary SVC: 9.11.82.251:5989

Enter IP address of the embedded CIM agent with the port number.(ex: 10.11.12.12:5989)

CIM Address of Remote SVC: 9.11.82.252:5989

Enter IP address of the embedded CIM agent with the port number.(ex: 10.11.12.12:5989)

Name Filter:

Leave empty for full discovery. Enter comma separated list of volume,volume prefix or Consistency Group name for specific discovery.(ex: cnVDisk,usVDisk,fr)

Username: superuser

Enter username to connect to CIM

Password: ********

Enter password to connect to CIM

Figure 63: Configuring array managers

Select an array pair in the **Discovered Array Pairs** section and click **Enable**. Figure 64 shows the solution lab setup with an array pair enabled.

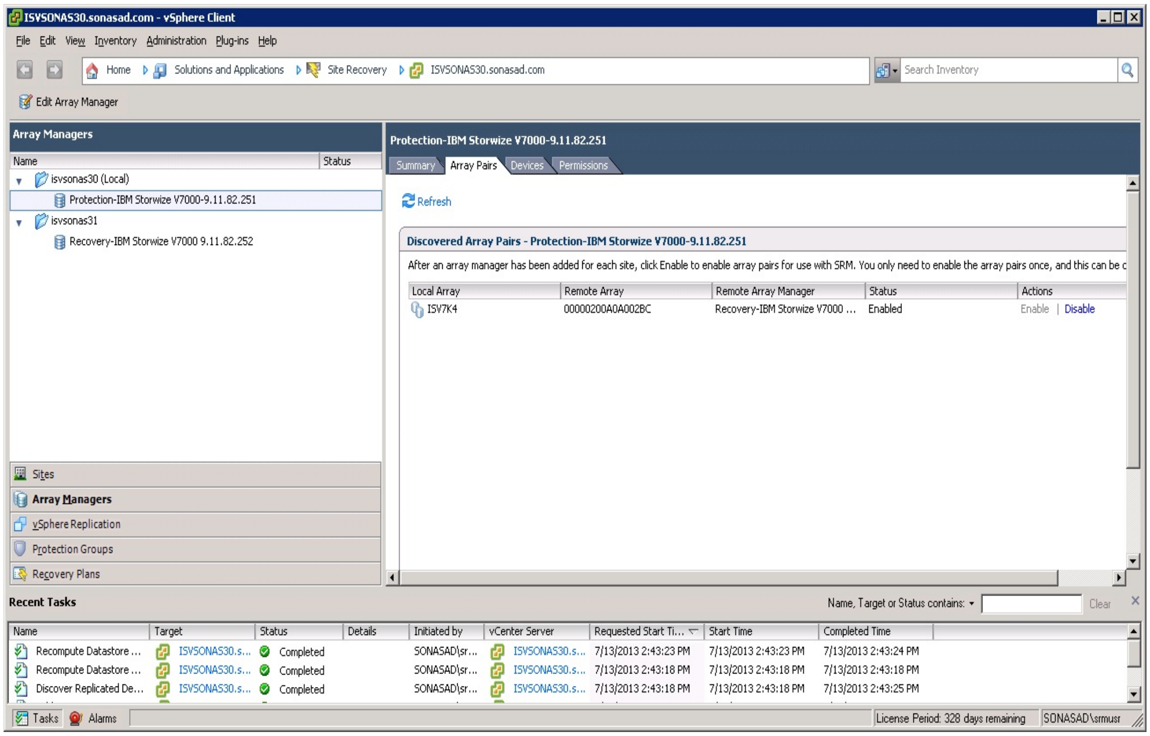

Name	Status
isvsonas30 (Local)	
Protection-IBM Storwize V7000-9.11.82.251	
isvsonas31	
Recovery-IBM Storwize V7000 9.11.82.252	

Protection-IBM Storwize V7000-9.11.82.251

Discovered Array Pairs - Protection-IBM Storwize V7000-9.11.82.251

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 c

Local Array	Remote Array	Remote Array Manager	Status	Actions
ISV7K4	00000200A0A002BC	Recovery-IBM Storwize V7000 ...	Enabled	Enable | Disable

Recent Tasks

Name, Target or Status contains:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Recompute Datastore ...	ISVSONAS30.s...	Completed		SONASAD\sr...	ISVSONAS30.s...	7/13/2013 2:43:23 PM	7/13/2013 2:43:23 PM	7/13/2013 2:43:24 PM
Recompute Datastore ...	ISVSONAS30.s...	Completed		SONASAD\sr...	ISVSONAS30.s...	7/13/2013 2:43:18 PM	7/13/2013 2:43:18 PM	7/13/2013 2:43:18 PM
Discover Replicated De...	ISVSONAS30.s...	Completed		SONASAD\sr...	ISVSONAS30.s...	7/13/2013 2:43:18 PM	7/13/2013 2:43:18 PM	7/13/2013 2:43:25 PM

License Period: 328 days remaining | SONASAD\srmusr

Figure 64: Array pair configuration

Finally, click the **Devices** tab to validate the array manager configuration (as show in Figure 65).

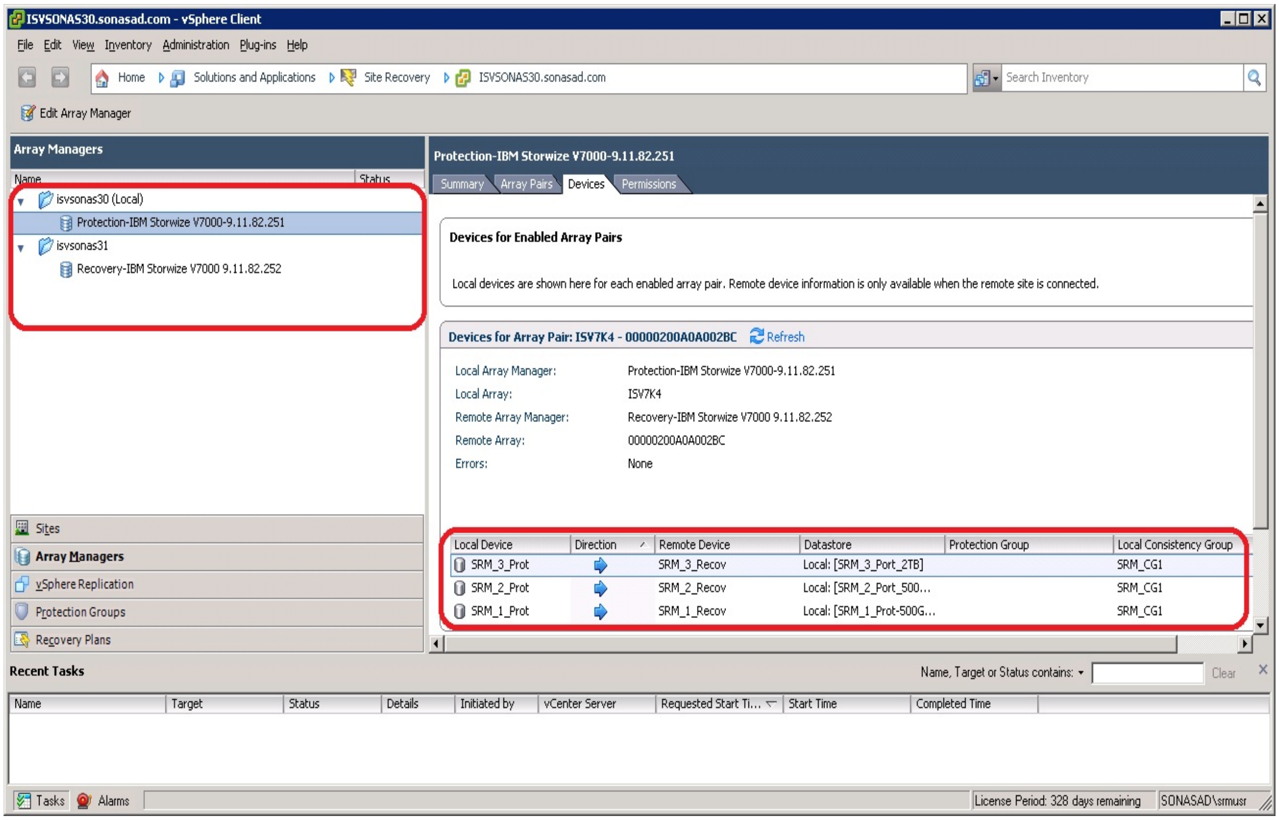
ISVSONAS30.sonasad.com - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home Solutions and Applications Site Recovery ISVSONAS30.sonasad.com
Search Inventory
Edit Array Manager
Array Managers
Name Status
isvsonas30 (Local)
Protection-IBM Storwize V7000-9.11.82.251
isvsonas31
Recovery-IBM Storwize V7000 9.11.82.252
Sites
Array Managers
vSphere Replication
Protection Groups
Recovery Plans
Protection-IBM Storwize V7000-9.11.82.251
Summary Array Pairs Devices Permissions
Refresh
Recent Tasks
Name, Target or Status contains:
Clear
Name Target Status Details Initiated by vCenter Server Requested Start Ti... Start Time Completed Time
Tasks Alarms
License Period: 328 days remaining
SONASAD\srmusr

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: ISV7K4 - 00000200A0A002BC

Local Array Manager:	Protection-IBM Storwize V7000-9.11.82.251
Local Array:	ISV7K4
Remote Array Manager:	Recovery-IBM Storwize V7000 9.11.82.252
Remote Array:	00000200A0A002BC
Errors:	None

Local Device	Direction	Remote Device	Datastore	Protection Group	Local Consistency Group
SRM_3_Prot	➡	SRM_3_Recov	Local: [SRM_3_Port_2TB]		SRM_CG1
SRM_2_Prot	➡	SRM_2_Recov	Local: [SRM_2_Port_500...		SRM_CG1
SRM_1_Prot	➡	SRM_1_Recov	Local: [SRM_1_Prot-500G...		SRM_CG1

Figure 65: Array manager device configuration on the protected site

Note: Repeat the same steps on the recovery site also (as shown in Figure 66).

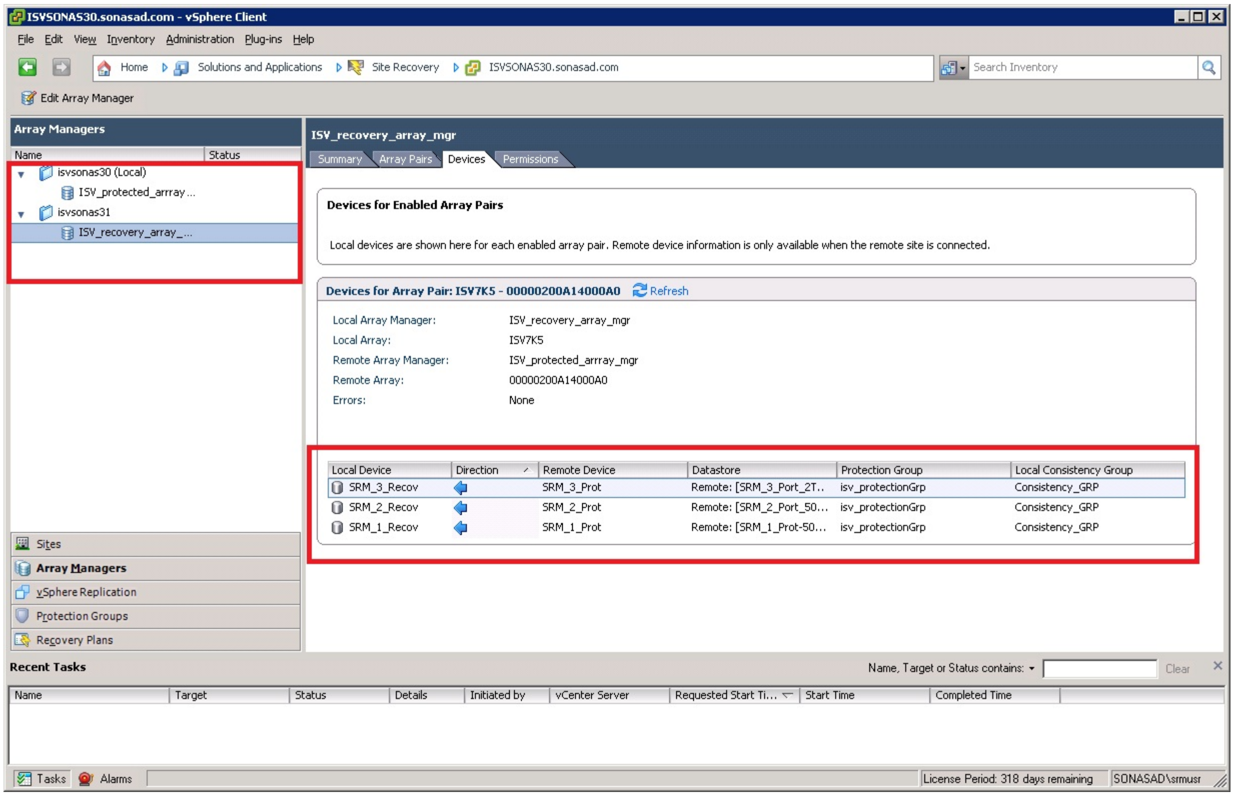

Devices for Enabled Array Pairs

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: ISV7K5 - 00000200A14000A0

Local Array Manager:	ISV_recovery_array_mgr
Local Array:	ISV7K5
Remote Array Manager:	ISV_protected_arrray_mgr
Remote Array:	00000200A14000A0
Errors:	None

Local Device	Direction	Remote Device	Datastore	Protection Group	Local Consistency Group
SRM_3_Recov	⬅	SRM_3_Prot	Remote: [SRM_3_Port_2T..	isv_protectionGrp	Consistency_GRP
SRM_2_Recov	⬅	SRM_2_Prot	Remote: [SRM_2_Port_50...	isv_protectionGrp	Consistency_GRP
SRM_1_Recov	⬅	SRM_1_Prot	Remote: [SRM_1_Prot-50...	isv_protectionGrp	Consistency_GRP

Figure 66: Array manager device configuration on the recovery site

Configuring inventory mappings

Inventory mappings establish recovery site defaults for the folders, networks, and resource pools to which the recovered virtual machines are assigned. These mappings are created at the protection site, and they should apply to all virtual machines in all protection groups at that site.

Ensure that resources at the protection site are mapped to the 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 mappings are not created, the VDisks must be specified individually for each virtual machine that is added to a protection group. A virtual machine cannot be protected unless it has valid inventory mappings for networks, folders, and computer resources. Inventory mappings are not required for resources that will not be used by protected virtual machines.

You need to perform the following steps to configure inventory mappings in the solution test lab.

1. Click **Sites** in the left pane and select the site for which you need to configure inventory mappings.
2. Click a tab for a type of inventory object to configure.
   You can choose the **Resource Mappings** tab, **Folder Mappings** tab, or the **Network Mappings** tab.
3. Select an inventory object and click **Configure Mapping**.
4. Expand the inventory items and navigate to the network, folder, or resource pool on the recovery site to which you need to map the protected site resource.
5. Optionally, you can specify how to create the mapping:
   - Select an existing resource from the list
   - Click **New Folder** or **New Resource Pool** to create a folder or a resource pool on the recovery site to which you need to map.

The selected resource appears in the Recovery Site Resources column, and its path relative to the root of the vCenter Server on the recovery site appears in the Recovery Site Path column.

Creating protection groups

A protection group is a group of VMs that are intended to fail over together. Grouping such VMs together on a single VMFS LUN is convenient. This procedure is also preferred because a one-to-one relationship between a protection group and a data store exists.

Before creating a protection group, the protected site must be connected to the recovery site, and the array managers must be configured. To be protected, a virtual machine must have the folder, network connection, and resource pool assignments that are also valid at the recovery site. Configure inventory mappings before creating protection groups, unless intending to configure these mappings individually for each member of the group.

You need to perform the following steps to create a protection group in the solution lab environment.

1. Click **Protection Groups** and click **Create Protection Group**.
2. On the Select Site and Protection Group Type page, select which site to protect and select **Array based replication (SAN)** as the protection group type.
3. Select an array pair and click **Next**. Figure 67 shows the selected configuration for the protection group.

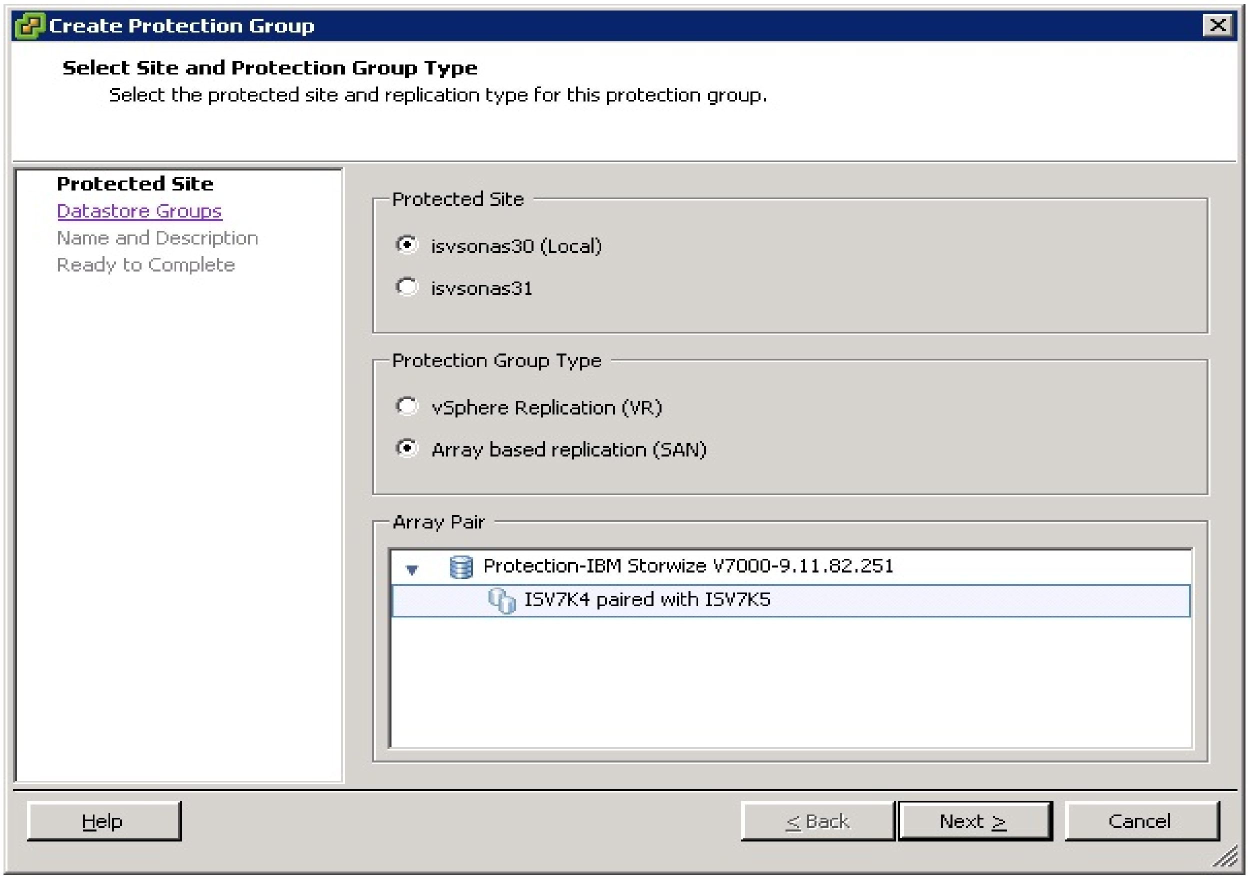

Create Protection Group

Select Site and Protection Group Type

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

Protected Site
Datastore Groups
Name and Description
Ready to Complete

Protected Site

● isvsonas30 (Local)
○ isvsonas31

Protection Group Type

○ vSphere Replication (VR)
● Array based replication (SAN)

Array Pair

Protection-IBM Storwize V7000-9.11.82.251
ISV7K4 paired with ISV7K5

Figure 67: Select the site and protection group type

4. Select the data store group from the list and click **Next** (as shown in Figure 68).

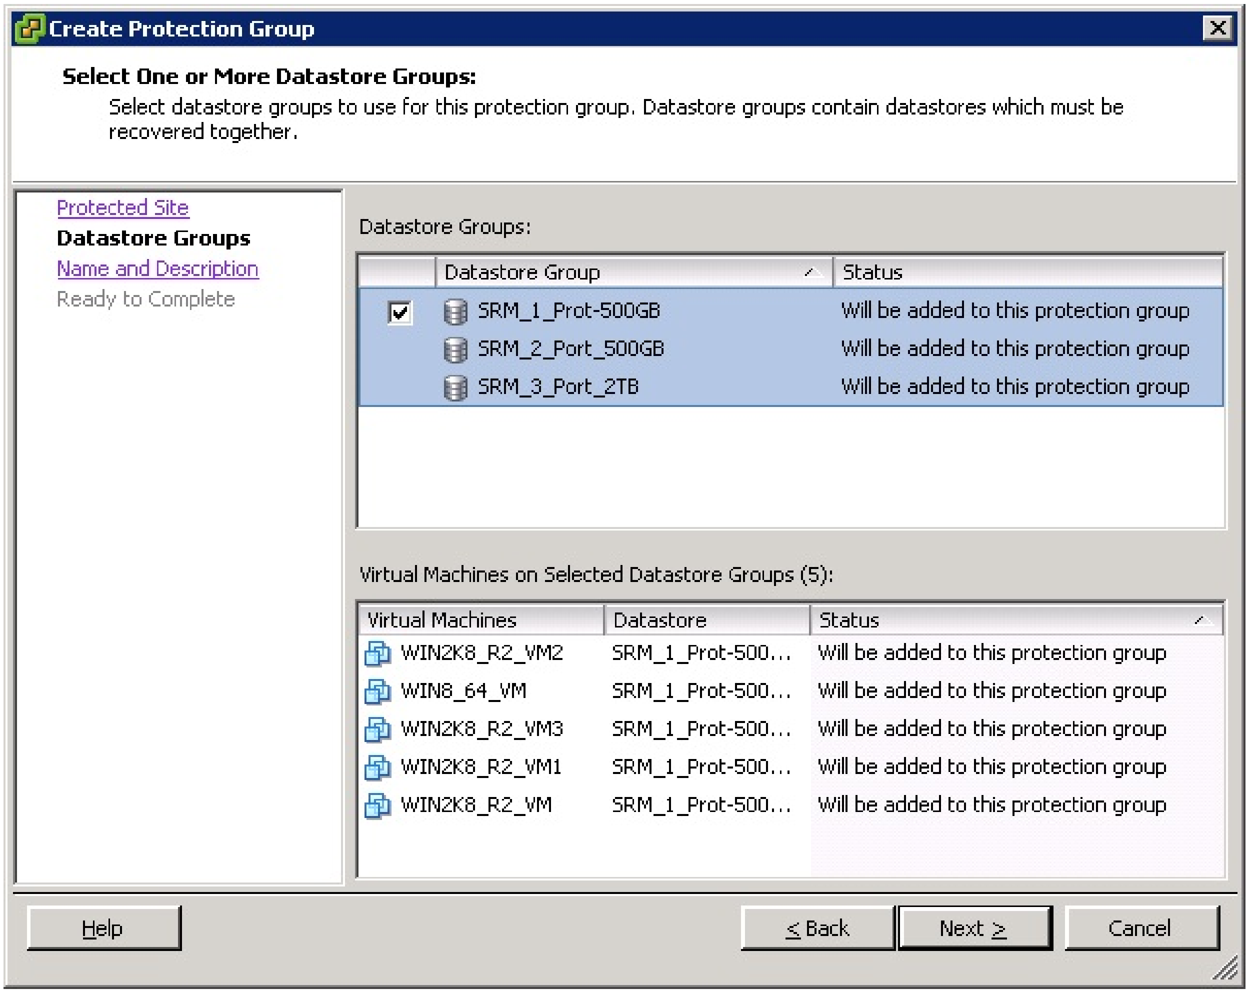

Select One or More Datastore Groups:

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

Datastore Groups:

	Datastore Group	Status
☑	SRM_1_Prot-500GB	Will be added to this protection group
	SRM_2_Port_500GB	Will be added to this protection group
	SRM_3_Port_2TB	Will be added to this protection group

Virtual Machines on Selected Datastore Groups (5):

Virtual Machines	Datastore	Status
WIN2K8_R2_VM2	SRM_1_Prot-500...	Will be added to this protection group
WIN8_64_VM	SRM_1_Prot-500...	Will be added to this protection group
WIN2K8_R2_VM3	SRM_1_Prot-500...	Will be added to this protection group
WIN2K8_R2_VM1	SRM_1_Prot-500...	Will be added to this protection group
WIN2K8_R2_VM	SRM_1_Prot-500...	Will be added to this protection group

Figure 68: Select One or More Datastore Groups page

5. Enter a name and optionally, a description for the protection group and click **Next**.
6. Finally, click **Finish** to create the protection group and begin the protection of the specified virtual machines.

For more information on configuring, refer to *IBM System Storage SAN Volume Controller Adapter for VMware vCenter Site Recovery Manager* at:

http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strghosts.doc%2FSVC_SRA-homepage.html

Creating, testing, and running recovery plans

Finally, after you configure Site Recovery Manager at the protected and recovery sites, you can create, test, and run a recovery plan. For more information about how to create, test, and run a recovery plan, refer to the *Site Recovery Manager Administration* guide at:

http://pubs.vmware.com/srm-51/topic/com.vmware.ICbase/PDF/srm-admin-5-1.pdf

Summary

This paper describes all the steps required to install and configure disaster recovery solution with IBM Storwize V7000 and VMware Site Recovery Manager, and demonstrates that Site Recovery Manager helps to easily manage complex VMware environments. It also describes how to configure Storwize family with Copy Services features, and how the SRA works together with Site Recovery Manager to implement the disaster recovery solution.

VMware vCenter Site Recovery Manager, combined with IBM Storwize family Copy Services features, FlashCopy, and Metro Mirror, provides a simplified disaster recovery solution. An easy-to-use GUI is all that is required to manage the task of defining, testing, and running a disaster recovery plan. This contrasts well with the manual scripting and complexity of more traditional disaster recovery solutions.

Appendix A: Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld®

  ibm.com/partnerworld/systems

- IBM Redbooks®

  ibm.com/redbooks

- IBM Publications Center

  www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US

- IBM Storwize family

  ibm.com/systems/storage/storwize/

- IBM Storwize family IP Partnership Configuration Guide
  ibm.com/partnerworld/wps/servlet/ContentHandler/stg_ast_sto_wp_configuring-system-storage-svc/lc=en_ALL_ZZ

- VMware vCenter Site Recovery Manager Documentation
  https://www.vmware.com/support/pubs/srm_pubs.html

@IBMSystemsISVs

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.