Deployment Guide

StorMagic SvSAN on VMware vSphere 6.0 with Cisco UCS C240 M4

April 2016

Introduction

Virtualization in today's computer infrastructure requires storage that is accessible to multiple servers simultaneously. Concurrently shared storage supports capabilities such as VMware vMotion migration, high availability, and replication, which provide stability and application availability. Large data centers typically meet these requirements through the use of large-scale SAN devices. Frequently, these solutions do not address the needs of smaller deployment models, such as remote-office and branch-office (ROBO) deployments, because the SAN is costly to extend and is burdened by latency.

One of the greatest challenges for multisite organizations is ensuring the uptime and availability of business-critical applications in distributed IT environments. This goal becomes even more challenging as organizations strive to reduce capital expenditures and operating expenses.

This challenge now can be addressed by combining the leading, software-defined storage solution from StorMagic with the latest in server technology from Cisco.

This deployment guide provides an integrated infrastructure solution that supports shared storage at remote data centers, using capacity attached directly to the computing layer for a self-contained application-delivery platform. StorMagic SvSAN in combination with Cisco UCS® C-Series Rack Servers provides computing, networking, and storage resources in even the most remote locations. Using a proven and verified architecture, the solution lets you reliably deploy computing and storage support for hundreds of virtual desktops and the necessary infrastructure for virtual machines.

This joint solution allows edge enterprises to deploy data storage infrastructure that supports their multiple-site environments.

This document provides a reference architecture and deployment guide for StorMagic SvSAN on VMware vSphere 6.0 running in a Cisco UCS C240 M4 Rack Server environment.

Cisco UCS C-Series Rack Servers

Cisco UCS C-Series Rack Servers (Figure 1) are designed for both performance and expandability. They deliver unified computing in an industry-standard form factor to reduce total cost of ownership (TCO) and increase agility. Cisco offers multiple server models, each with different processing power, memory capacity, I/O, and internal storage capacities to address a wide variety of workload challenges.

Table of Contents

Cisco UCS C-Series Rack Servers provide the following benefits:

- Form-factor-independent entry point into Cisco Unified Computing System™ (Cisco UCS)
- Simplified and fast deployment of applications
- Extension of unified computing innovations and benefits to rack servers
- Increased customer choice, with unique benefits, in a familiar rack package
- Reduced TCO and increased business agility

Figure 1. Cisco UCS C-Series Rack Servers

Optimized for virtualized environments, Cisco C-Series servers are integral components of the next-generation data center. They unite computing, networking, storage access, and virtualization resources in a cohesive system designed to reduce TCO and increase business agility. With their integration and scalability features and unified management domain, they are particularly well suited to distributed infrastructure.

Cisco UCS C240 M4 Server

The Cisco UCS C240 M4 Rack Server (Figure 2) is an enterprise-class server designed to deliver exceptional performance, expandability, and efficiency for storage and I/O-intensive infrastructure workloads. Such workloads include big data analytics, virtualization, and graphics-intensive and bare-metal applications.

Figure 2. Cisco UCS C240 M4 Rack Server

The Cisco UCS C240 M4 delivers outstanding levels of expandability and performance in a 2-rack-unit (2RU) form factor for standalone systems as well as for environments managed by Cisco UCS (see http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c240-m4-rack-server/index.html).

It provides:

- Dual Intel® Xeon® processor E5-2600 v3 CPUs for improved performance that is suitable for nearly all 2-socket applications
- Next-generation double-data-rate 4 (DDR4) memory and 12-Gbps SAS throughput
- Innovative Cisco UCS virtual interface card (VIC) support in PCI Express (PCIe) and modular LAN on motherboard (mLOM) form factors
- Graphics-intensive experiences for more virtual users with support for the latest NVIDIA graphics processing units (GPUs)

Cisco UCS Virtual Interface Card 1227

The Cisco UCS VIC 1227 (Figure 3) offers a mLOM adapter. The mLOM slot, new to Cisco rack servers, can be used to install a Cisco VIC without consuming a PCIe slot, providing greater I/O expandability.

The Cisco UCS VIC 1227 is a dual-port, Enhanced Small Form-Factor Pluggable (SFP+), 10 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)–capable, PCIe mLOM adapter. It is designed exclusively for the M4 generation of Cisco UCS C-Series Rack Servers and the dense-storage Cisco UCS C3160 and C3260 Rack Servers (see http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1227/index.html).

Figure 3. Cisco UCS VIC 1227

VMware vSphere 6.0

VMware vSphere 6.0, the industry-leading virtualization platform, empowers users to virtualize any application with confidence. It redefines availability and simplifies the virtual data center. It offers a highly available, resilient, on-demand infrastructure that is an excellent foundation for any cloud environment. This release contains many new features and enhancements, many of which are industry-first features (see https://www.vmware.com/files/pdf/vsphere/VMW-WP-vSPHR-Whats-New-6-0.pdf).

StorMagic SvSAN

StorMagic SvSAN (Figure 4) is a software solution that enables enterprises to eliminate downtime for business-critical applications at the edge, where disruption directly leads to loss of service and revenue. SvSAN helps ensure high availability through a virtualized shared storage platform, so that business-critical edge applications remain operational.

StorMagic's typical customer has between 10 and 10,000 edge sites, at which local IT resources are not available, but application uptime is essential.

- SvSAN provides an intuitive, standardized management interface that allows multiple SvSANs, spread across remote sites, to be managed and provisioned quickly and simply, either locally or remotely, from a central location.
- SvSAN's efficient and flexible architecture and its modular approach enable it to meet the changing and increasingly demanding storage requirements of almost any organization.
- SvSAN enables organizations to create Small Computer System Interface over IP (iSCSI) SANs by deploying virtual storage appliances (VSA) using internal or direct-attached server storage. It can be used to provide highly available storage without the expense of an external SAN.

A VSA SAN behaves like a conventional iSCSI SAN, except that instead of running on an external physical storage system, the iSCSI storage appliance runs in a virtual machine on a hypervisor in a host server. Storage devices assigned to the VSA are treated as pool storage, from which iSCSI targets can be created. These iSCSI targets are usually assigned to hypervisor hosts, where they are formatted with a file system, but they can be mounted on any iSCSI initiator. Each hypervisor server can run a VSA, and iSCSI targets can be presented from the internal or direct-attached storage (DAS). These targets can also be mirrored between VSAs to create highly available storage. If one hypervisor host and its VSA become inactive or fail, the other hypervisor servers can continue to access the storage using the surviving VSA.

Figure 4. StorMagic SvSAN

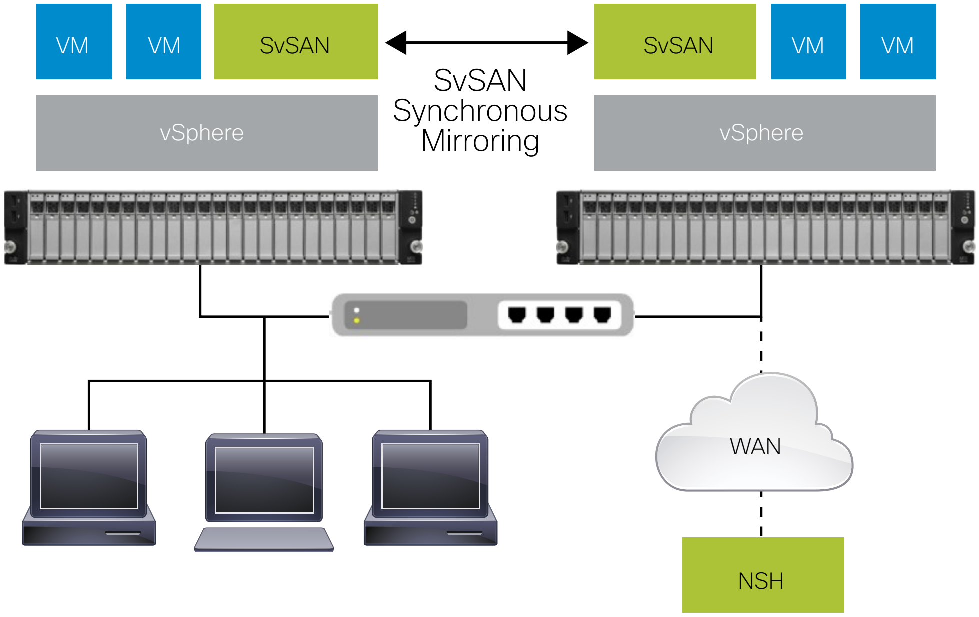

SvSAN benefits include:

- Abstraction of storage services away from traditional storage arrays, making it an important component of a software-defined storage strategy
- Elimination of the need for a physical SAN
- Virtualization of internal disk drives and external, DAS arrays to enable storage sharing among multiple servers
- High availability in a simple two-server solution
- At least 33 percent lower capital acquisition costs for distributed environments
- Over 40 percent lower TCO through reduced power and cooling costs, decreased sparing and maintenance costs, and reduced need for dedicated IT resources
- Greater application uptime
- Centralized storage management of entire multi-site infrastructure
- Rapid, scripted deployment and simultaneous updating of multiple sites through automated provisioning
- Optimal flexibility, because SvSAN is hardware and hypervisor independent and can scale as storage requirements grow
- Fast resynchronization through its restore capability, enabling users to replace a failed server with a new one and automatically rebuild the environment

StorMagic SvSAN with Cisco UCS C-Series Architecture

The architecture described in for this deployment guide is simple. It provides a highly available environment using a combination of features from StorMagic SvSAN, VMware vSphere, and Cisco UCS and using the local storage available on Cisco UCS C240 M4 Rack Servers.

Figure 5 shows the physical layout of the test environment. Each Cisco UCS C240 M4 server has an Intel Ethernet I350 1-Gbps network controller with two 1 Gigabit Ethernet LOM ports to connect to the network for VMware ESXi management, SvSAN management, and virtual machine management. A neutral storage host (NSH) running on a separate host is used to prevent a SvSAN cluster split-brain scenario. The NSH uses the management network for SvSAN heartbeat information only. The servers are connected back to back using a Cisco UCS VIC 1227 10-Gbps dual-port converged network adapter (CNA) SFP+ card for SvSAN iSCSI, mirroring, and vMotion traffic.

Figure 5. Cisco UCS C240 M4 Servers Physical Connectivity for SvSAN Deployment.

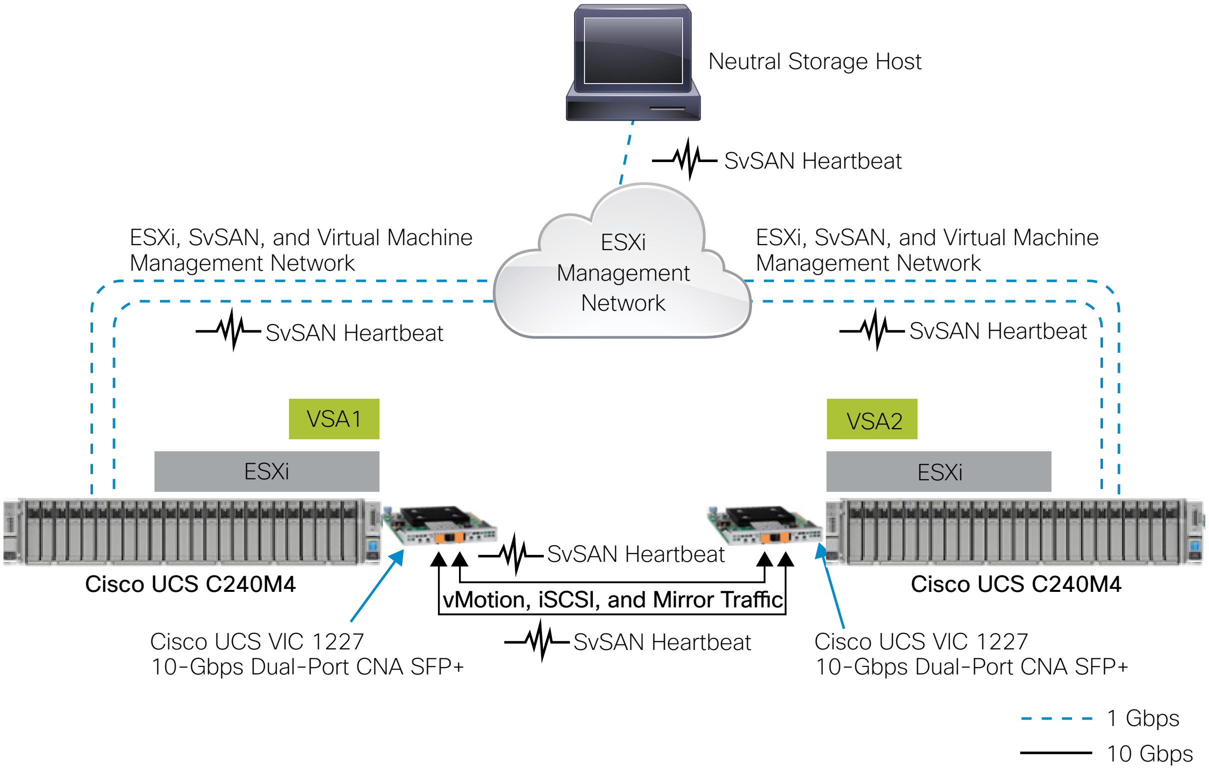

Figure 6 shows the Cisco UCS C240 M4 storage configuration for ESXi and SvSAN deployment in a VMware cluster environment. Each Cisco UCS C240 M4 server has two 32-GB Cisco Flexible Flash (FlexFlash) Secure Digital (SD) cards and eight 1.2-terabyte (TB) 6.0-Gbps 10,000-rpm SAS hard disk drives (HDDs). A hypervisor partition is created using two 32-GB FlexFlash SD cards in a mirrored (RAID 1) configuration to install ESXi. The eight 1.2-TB SAS drives are used to create a single RAID 5 drive group with two virtual drives. The first virtual drive, with a total capacity of 7.6-TB, is allocated explicitly to the VSA as a raw disk mapping (RDM). The VSA sees the virtual disk as a pool of storage resources from which iSCSI targets can be carved. These iSCSI targets can be mirrored with other VSAs to create highly available storage. The VSA deployment uses the second 30-GB virtual disk and must be configured as persistent storage on the ESXi host, and this data store requires 25-GB of free space.

Figure 6. Cisco UCS C240 M4 Server Virtual Disk (RAID) Configuration for StorMagic SvSAN Deployment

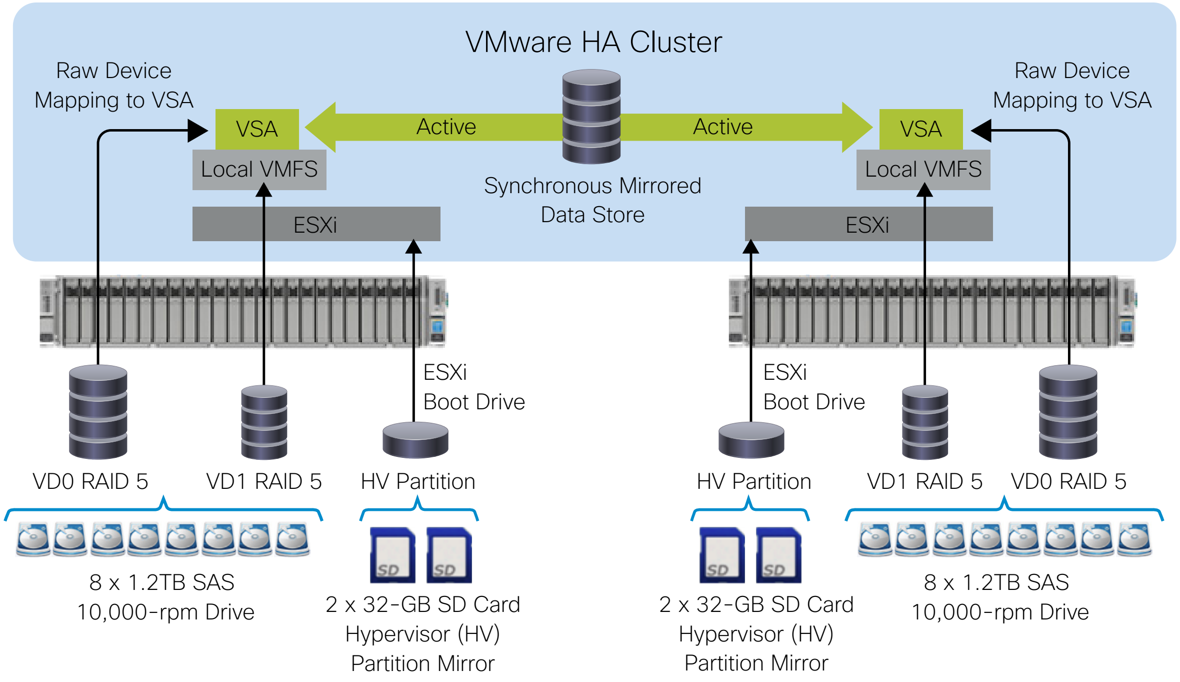
VMware HA Cluster
Raw Device Mapping to VSA
VSA
Active
Active
VSA
Raw Device Mapping to VSA
Local VMFS
Local VMFS
Synchronous Mirrored Data Store
ESXi
ESXi
ESXi Boot Drive
ESXi Boot Drive
VD0 RAID 5
VD1 RAID 5
HV Partition
HV Partition
VD1 RAID 5
VD0 RAID 5
SD
SD
SD
SD
8 x 1.2TB SAS 10,000-rpm Drive
2 x 32-GB SD Card Hypervisor (HV) Partition Mirror
2 x 32-GB SD Card Hypervisor (HV) Partition Mirror
8 x 1.2TB SAS 10,000-rpm Drive

Table 1 shows the virtual disk and RAID configuration for ESXi and SvSAN deployment on Cisco UCS C240 M4 servers.

Table 1. Virtual Disk Configuration

Virtual Disk Configuration	
FlexFlash hypervisor partition mirror	29.7-GB
	ESXi installation
Virtual Disk 0 RAID 5	7.6-TB
	Virtual drive allocated explicitly to the local VSA as an RDM
	VSA sees virtual drive as a single storage pool
	At least one iSCSI mirror should be carved from the storage pool
Virtual Disk 1 RAID 5	34.19-GB
	VSA deployment
	Space formatted with a VMware Virtual Machine File System (VMFS), creating a local data store for the host
	VSA deployed to this data store • Boot drive: 512MB • Journal drive: 20GB

This detailed architecture for the Cisco UCS with SvSAN deployment consists of the components listed in Table 2.

Table 2. Cisco UCS with StorMagic SvSAN Deployment Components

Component	Description
Cisco UCS	2 x Cisco UCS C240 M4 Rack Servers (x86 servers), each with: • 2 x Intel Xeon processor E5-2660 v3 CPUs • 16 x 16-GB DDR4 2133-MHz RDIMM (PC4-17000) dual rank x4 1.2V memory • 8 x Seagate 1.2-TB 6.0-Gbs 10,000-rpm SAS HDDs • Embedded (onboard) Intel I350 1-Gbps network controller, with 2 ports • 1 x Cisco 12G SAS Modular Raid Controller with 2-GB flash-based write cache (FBWC) (UCSC-MRAID12G-2GB) • 1 x Cisco UCS VIC 1227 10-Gbps dual-port CNA SFP+ (UCSC-MLOM-CSC-02) • 2 x 32-GB FlexFlash cards (UCS-SD-32G-S)
VMware software	• VMware vCenter 6.0.0b • VMware ESXi 6.0.0b
StorMagic SvSAN software	Release 5.3

Note: New software and firmware releases should be installed as listed in the Cisco UCS Hardware and Software Interoperability Matrix at http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html. For information about the StorMagic SvSAN certified and supported versions, see the VMware Hardware Compatibility List (HCL) at http://www.vmware.com/resources/compatibility/search.php.

Table 3 lists the virtual machine sizing assumptions used for the Cisco UCS with SvSAN deployment.

Table 3. Virtual Machine Configuration

Virtual machine instance size	• 4 virtual CPUs (vCPUs) with 6 GB of RAM • 50-GB OS logical unit number (LUN) and 100-GB raw LUN for I/O test

Deploying StorMagic SvSAN for VMware vSphere in a Cisco UCS Environment

Figure 7 shows the high-level steps for deploying SvSAN for vSphere in a Cisco UCS environment.

Figure 7. High-Level Steps for Deploying StorMagic SvSAN for VMware vSphere on Cisco UCS C240 M4 Servers

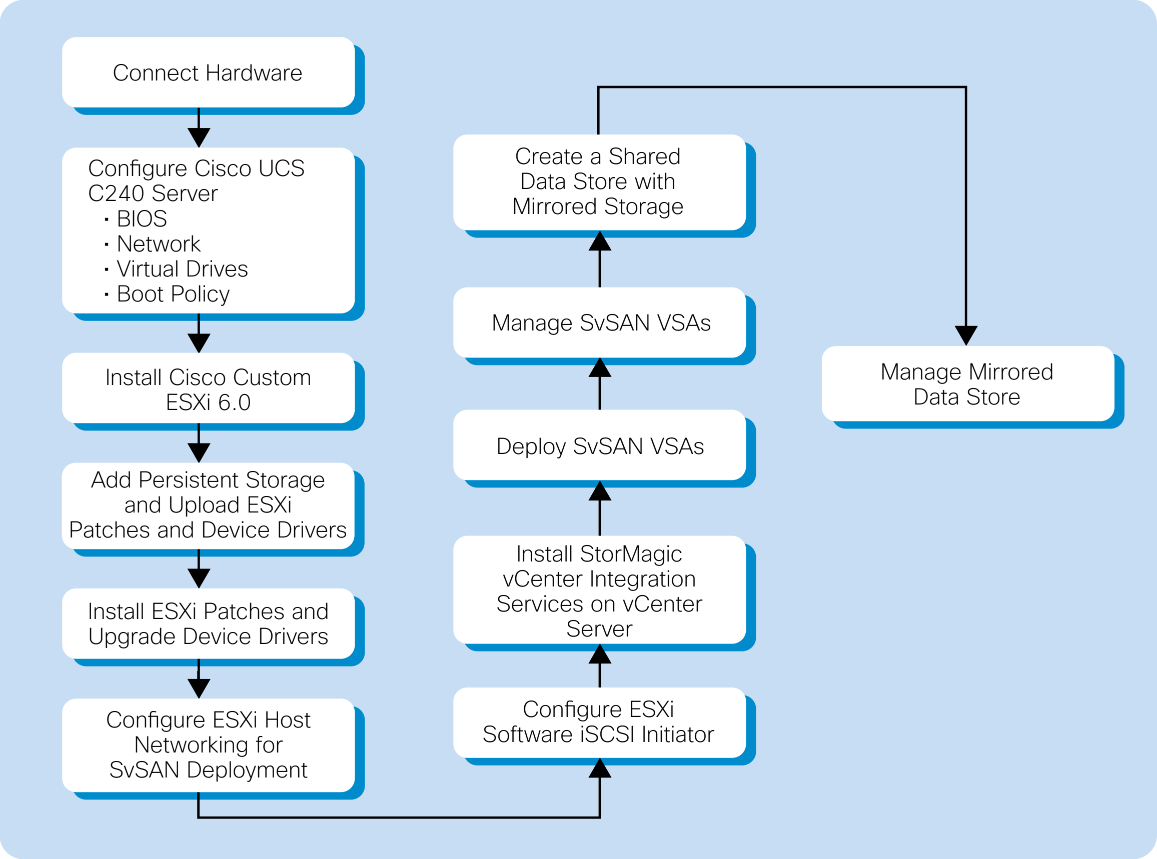
Connect Hardware
Configure Cisco UCS C240 Server
· BIOS
· Network
· Virtual Drives
· Boot Policy
Install Cisco Custom ESXi 6.0
Add Persistent Storage and Upload ESXi Patches and Device Drivers
Install ESXi Patches and Upgrade Device Drivers
Configure ESXi Host Networking for SvSAN Deployment
Configure ESXi Software iSCSI Initiator
Install StorMagic vCenter Integration Services on vCenter Server
Deploy SvSAN VSAs
Manage SvSAN VSAs
Create a Shared Data Store with Mirrored Storage
Manage Mirrored Data Store

Configuring Cisco UCS C240 M4 Rack Servers

This section provides detailed information about the Cisco UCS C240 M4 server configuration for ESXi and SvSAN deployment.

Configure BIOS Policy

The BIOS policy for the SvSAN environment is configured to achieve high performance, as shown Figure 8. For more information about the BIOS settings, refer to http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/whitepaper_C11-733130.html.

Figure 8. BIOS Policy

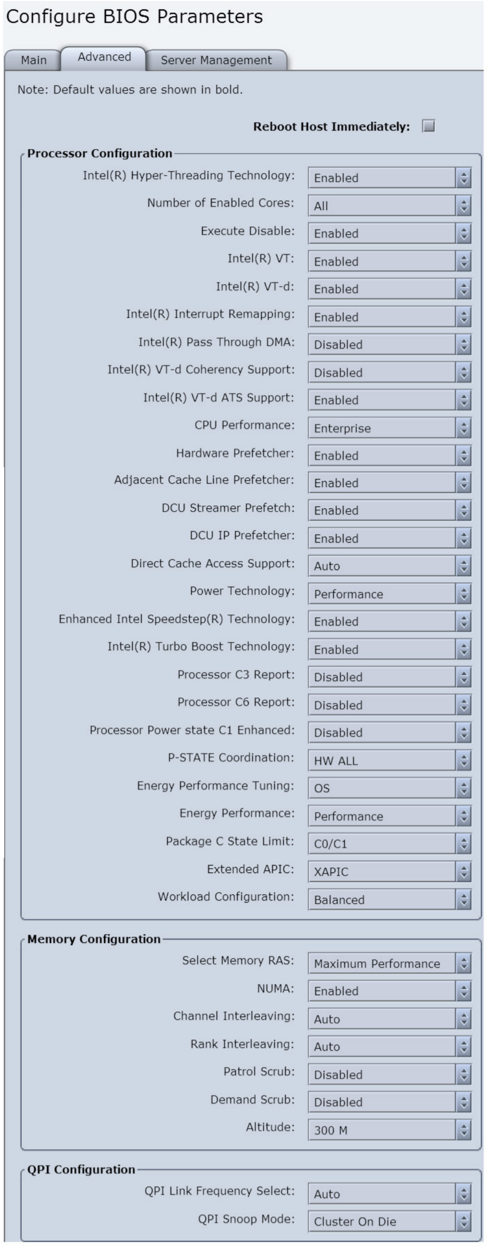

Configure BIOS Parameters

Note: Default values are shown in bold.

Reboot Host Immediately: □

Processor Configuration

Intel(R) Hyper-Threading Technology:	Enabled
Number of Enabled Cores:	All
Execute Disable:	Enabled
Intel(R) VT:	Enabled
Intel(R) VT-d:	Enabled
Intel(R) Interrupt Remapping:	Enabled
Intel(R) Pass Through DMA:	Disabled
Intel(R) VT-d Coherency Support:	Disabled
Intel(R) VT-d ATS Support:	Enabled
CPU Performance:	Enterprise
Hardware Prefetcher:	Enabled
Adjacent Cache Line Prefetcher:	Enabled
DCU Streamer Prefetch:	Enabled
DCU IP Prefetcher:	Enabled
Direct Cache Access Support:	Auto
Power Technology:	Performance
Enhanced Intel Speedstep(R) Technology:	Enabled
Intel(R) Turbo Boost Technology:	Enabled
Processor C3 Report:	Disabled
Processor C6 Report:	Disabled
Processor Power state C1 Enhanced:	Disabled
P-STATE Coordination:	HW ALL
Energy Performance Tuning:	OS
Energy Performance:	Performance
Package C State Limit:	C0/C1
Extended APIC:	XAPIC
Workload Configuration:	Balanced

Memory Configuration

Select Memory RAS:	Maximum Performance
NUMA:	Enabled
Channel Interleaving:	Auto
Rank Interleaving:	Auto
Patrol Scrub:	Disabled
Demand Scrub:	Disabled
Altitude:	300 M

QPI Configuration

QPI Link Frequency Select:	Auto
QPI Snoop Mode:	Cluster On Die

Configure Networking

Each Cisco UCS C240 M4 server has one onboard Intel I350 1-Gbps network controller, which has two 1-Gbps network ports and one Cisco UCS VIC 1227 10-Gbps dual-port CNA SFP+ adapter. Both the 1-Gbps network ports are configured for ESXi, SvSAN VSA, and virtual machine management. The servers are connected back to back using the Cisco UCS VIC 1227 for SvSAN iSCSI, mirroring, and vMotion traffic. Configure each 10-Gbps port with a maximum transmission unit (MTU) of 9000 through the Cisco Integrated Management Controller (IMC): choose **Cisco VIC Adapters > vNICs > Properties** (Figure 9).

Figure 9. vNIC Details

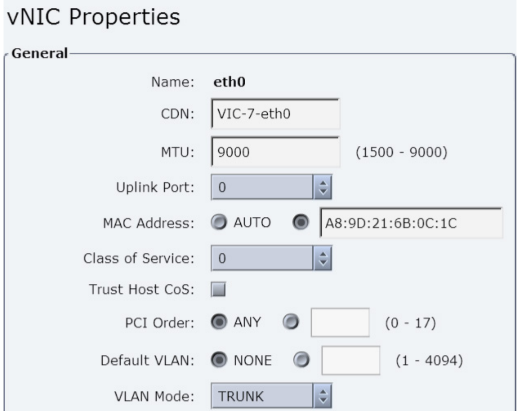

vNIC Properties

General

Name: **eth0**

CDN: VIC-7-eth0

MTU: 9000 (1500 - 9000)

Uplink Port: 0

MAC Address: ○ AUTO ● A8:9D:21:6B:0C:1C

Class of Service: 0

Trust Host CoS: □

PCI Order: ● ANY ○ (0 - 17)

Default VLAN: ● NONE ○ (1 - 4094)

VLAN Mode: TRUNK

Figure 10 shows the configuration.

Figure 10. Network Configuration

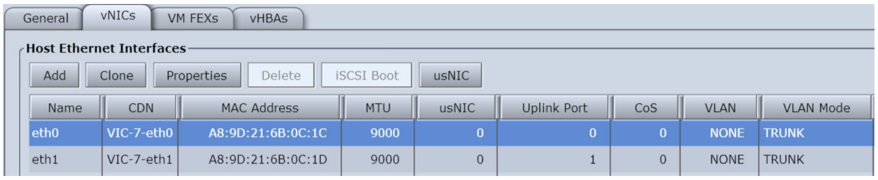

Host Ethernet Interfaces

Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode
eth0	VIC-7-eth0	A8:9D:21:6B:0C:1C	9000	0	0	0	NONE	TRUNK
eth1	VIC-7-eth1	A8:9D:21:6B:0C:1D	9000	0	1	0	NONE	TRUNK

Configure Storage

Each Cisco UCS C240 M4 server has two 32-GB FlexFlash SD cards and eight 1.2-TB SAS hard drives. In this configuration, ESXi is booted from the on-board mirrored (RAID 1) FlexFlash SD cards. For more information, see Cisco FlexFlash: Use and Manage Cisco Flexible Flash Internal SD Card for Cisco UCS C-Series Standalone Rack Servers. The FlexFlash SD card mirrored (RAID 1) configuration is performed through the Integrated Management Controller, as shown in the FlexFlash sample document at http://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/unified-computing/whitepaper_C11-718938.pdf.

Configuring RAID 5

Create a single RAID 5 drive group with two virtual drives using eight 1.2-TB SAS drives. In the Integrated Management Controller, create two virtual disks with sizes of 7.6-TB and 34.19-GB in a single underlying RAID 5 drive group as shown in the following steps and in Figures 11 through 14.

1. In the navigation pane, click the **Storage** tab.

2. On the Storage tab, select **Cisco 12G SAS Modular RAID Controller** and in the configuration pane, select the **Controller Info** tab.

3. Select **Create Virtual Drive from Unused Physical Drives** and create the virtual drive.

4. After virtual drive creation is complete, fully initialize the virtual drives.

Figure 11. Creating Virtual Drive from Unused Physical Drives

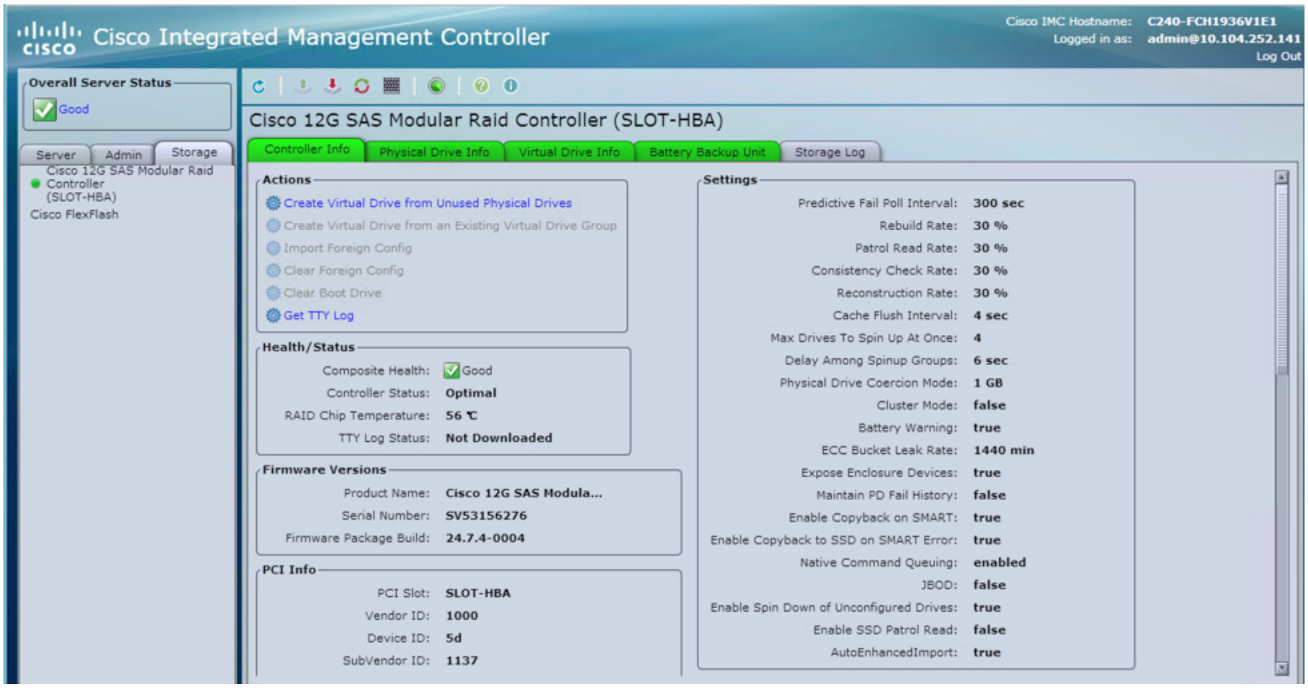

Cisco IMC Hostname: C240-FCH1936V1E1
Logged in as: admin@10.104.252.141

Overall Server Status: ✓ Good

Health/Status

Composite Health: ✓ Good
Controller Status: Optimal
RAID Chip Temperature: 56 ℃
TTY Log Status: Not Downloaded

Firmware Versions

Product Name: Cisco 12G SAS Modula...
Serial Number: SV53156276
Firmware Package Build: 24.7.4-0004

PCI Info

PCI Slot: SLOT-HBA
Vendor ID: 1000
Device ID: 5d
SubVendor ID: 1137

Settings

Predictive Fail Poll Interval: 300 sec
Rebuild Rate: 30 %
Patrol Read Rate: 30 %
Consistency Check Rate: 30 %
Reconstruction Rate: 30 %
Cache Flush Interval: 4 sec
Max Drives To Spin Up At Once: 4
Delay Among Spinup Groups: 6 sec
Physical Drive Coercion Mode: 1 GB
Cluster Mode: false
Battery Warning: true
ECC Bucket Leak Rate: 1440 min
Expose Enclosure Devices: true
Maintain PD Fail History: false
Enable Copyback on SMART: true
Enable Copyback to SSD on SMART Error: true
Native Command Queuing: enabled
JBOD: false
Enable Spin Down of Unconfigured Drives: true
Enable SSD Patrol Read: false
AutoEnhancedImport: true

Figure 12. Selecting Physical Drives to Create RAID and Set Virtual Drive Properties

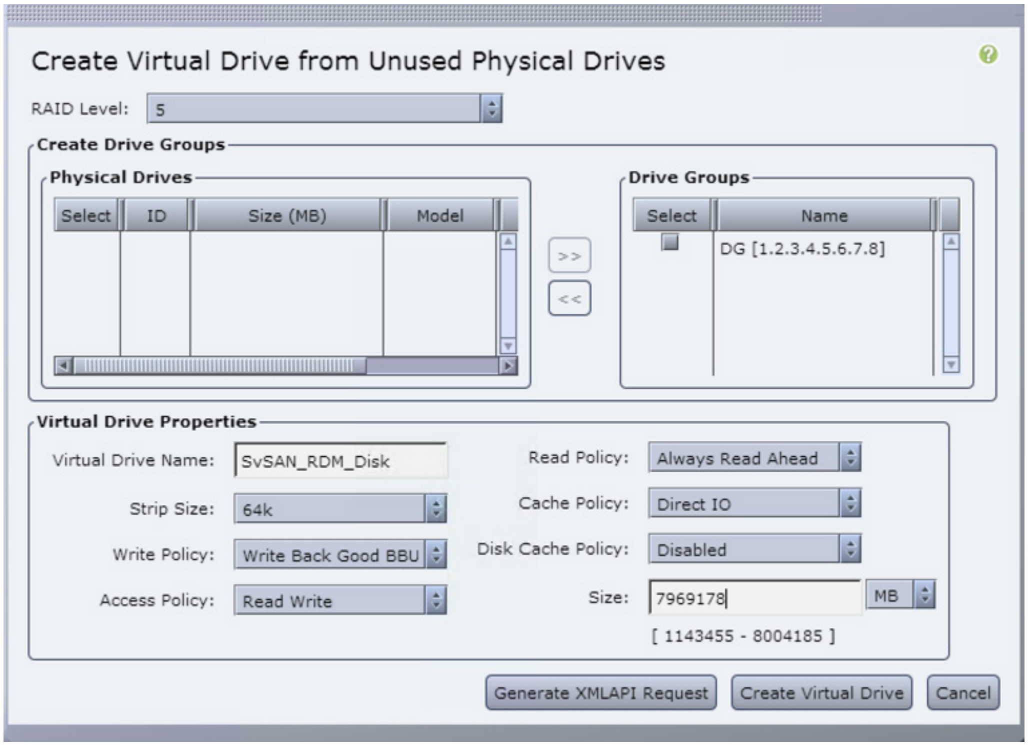

Create Virtual Drive from Unused Physical Drives

RAID Level: 5

Create Drive Groups

Physical Drives

Select	ID	Size (MB)	Model

Drive Groups

Select	Name
□	DG [1.2.3.4.5.6.7.8]

Virtual Drive Properties

Virtual Drive Name: SvSAN_RDM_Disk

Strip Size: 64k

Write Policy: Write Back Good BBU

Access Policy: Read Write

Read Policy: Always Read Ahead

Cache Policy: Direct IO

Disk Cache Policy: Disabled

Size: 7969178 MB

[1143455 - 8004185]

Figure 13. Creating a Virtual Drive from an Existing Virtual Drive Group (to Configure a StorMagic SvSAN Virtual Machine)

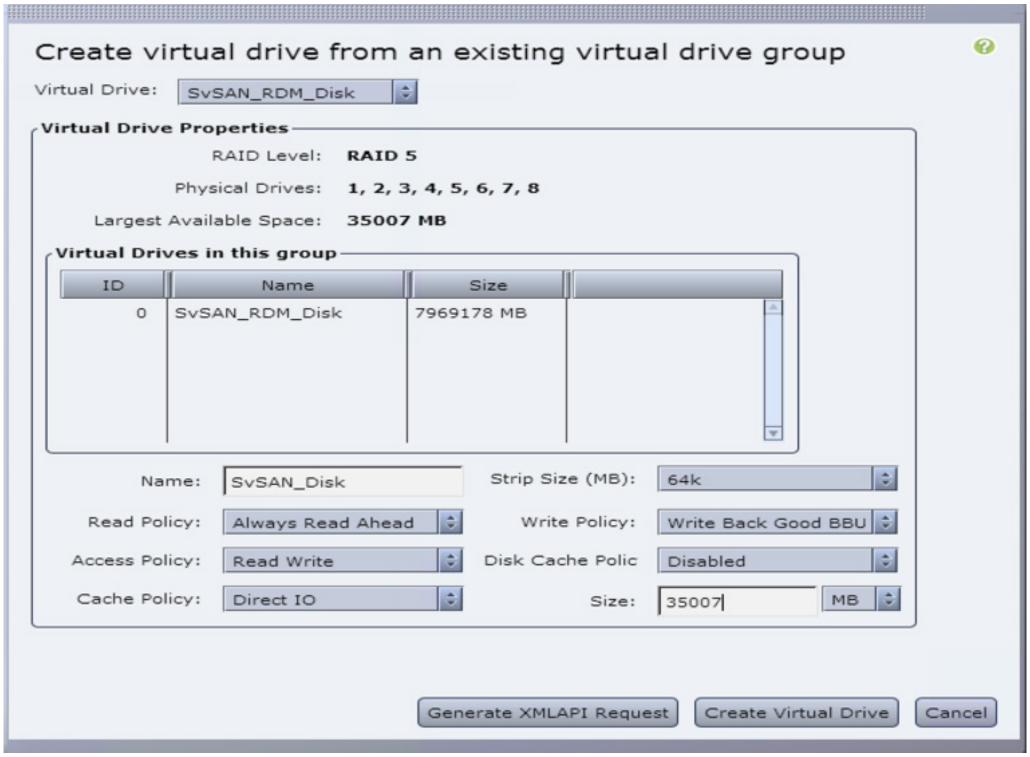

Create virtual drive from an existing virtual drive group

Virtual Drive: SvSAN_RDM_Disk

Virtual Drive Properties

RAID Level: **RAID 5**

Physical Drives: **1, 2, 3, 4, 5, 6, 7, 8**

Largest Available Space: **35007 MB**

Virtual Drives in this group

ID	Name	Size
0	SvSAN_RDM_Disk	7969178 MB

Name: SvSAN_Disk

Read Policy: Always Read Ahead

Access Policy: Read Write

Cache Policy: Direct IO

Strip Size (MB): 64k

Write Policy: Write Back Good BBU

Disk Cache Polic Disabled

Size: 35007 MB

Figure 14. Created Virtual Drives Are Listed on the Virtual Drive Info Tab

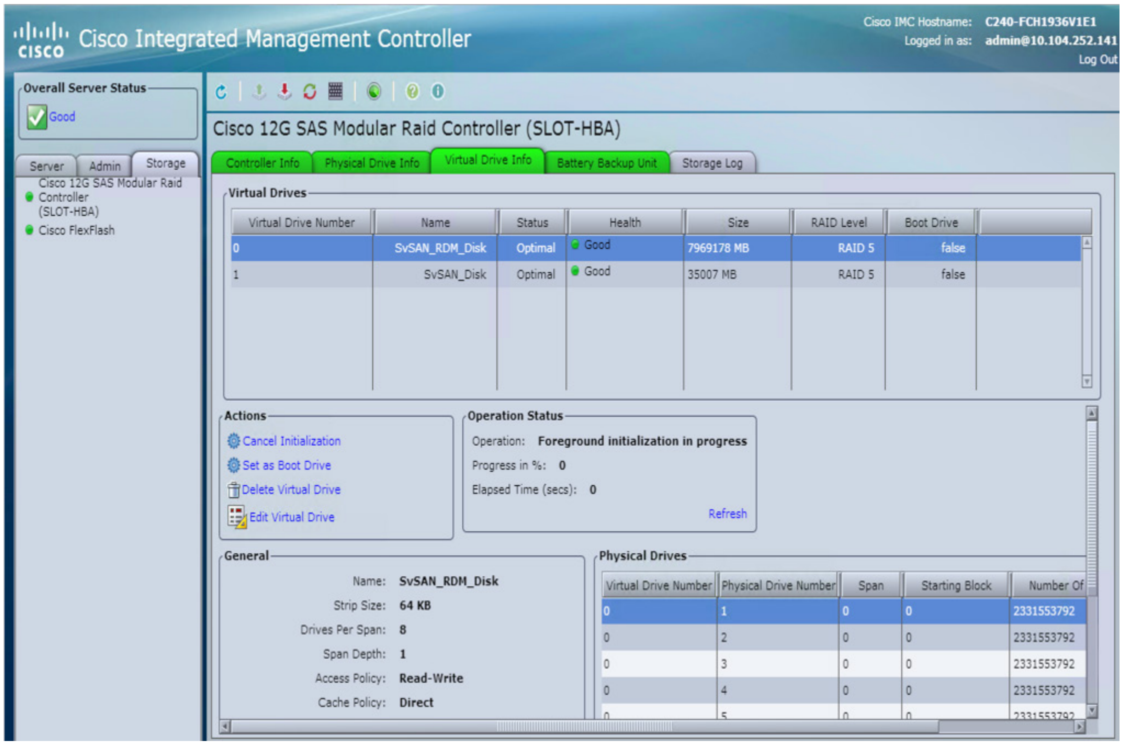

Cisco IMC Hostname: C240-FCH1936V1E1
Logged in as: admin@10.104.252.141

Cisco 12G SAS Modular Raid Controller (SLOT-HBA)

Virtual Drives

Virtual Drive Number	Name	Status	Health	Size	RAID Level	Boot Drive
0	SvSAN_RDM_Disk	Optimal	● Good	7969178 MB	RAID 5	false
1	SvSAN_Disk	Optimal	● Good	35007 MB	RAID 5	false

Operation Status

Operation: Foreground initialization in progress
Progress in %: 0
Elapsed Time (secs): 0

General

Name: SvSAN_RDM_Disk
Strip Size: 64 KB
Drives Per Span: 8
Span Depth: 1
Access Policy: Read-Write
Cache Policy: Direct

Physical Drives

Virtual Drive Number	Physical Drive Number	Span	Starting Block	Number Of
0	1	0	0	2331553792
0	2	0	0	2331553792
0	3	0	0	2331553792
0	4	0	0	2331553792
0	5	0	0	2331553792

Note: To create virtual drives for scenarios with additional expansion packs (either one or two expansion packs can be added, each with eight HDDs), follow the steps in the preceding section. For each expansion pack, one virtual drive is sufficient for storage allocation to the local VSA, because there is no need for additional SvSAN virtual machines.

Configure Boot Policy

Using the Integrated Management Controller, you can configure the order in which the server attempts to boot from available boot device types. In a traditional boot-order configuration, the Integrated Management Controller allows you to reorder the device types, but not the devices within the device types. The boot policy is created with an SD card and keyboard, video, and mouse (KVM) mapped DVD as the preferred boot device.

Configure the precision boot order as follows:

1. In the navigation pane, click the **Server** tab.
2. In the Server tab, click **BIOS**.
3. In the Actions area, click **Configure Boot Order**.
4. In the Configure Boot Order dialog box, select **Add SD Card** and **Add Virtual Media** for the preferred boot device. (Figure 15)
5. Click **Save**.

Figure 15. Configuring Boot Order

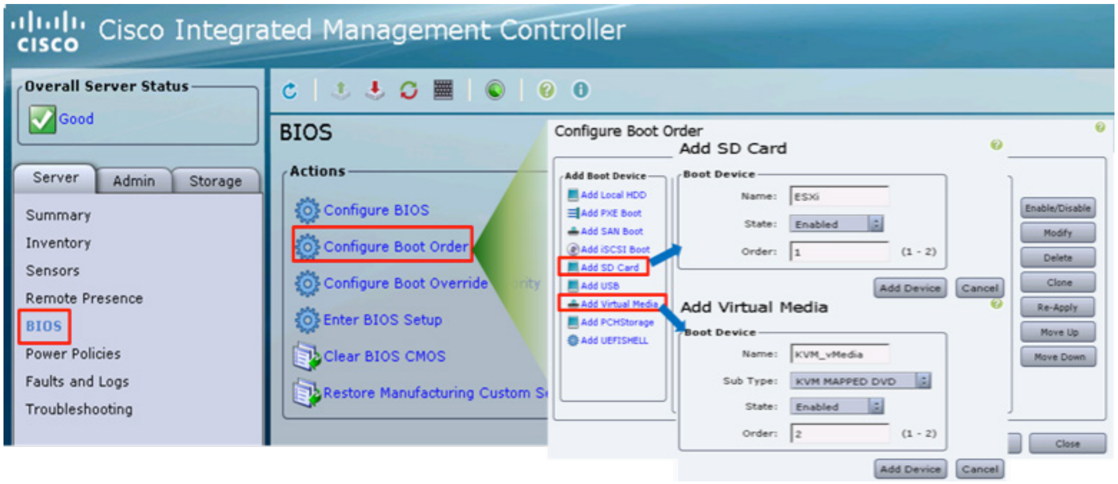

Add SD Card

Boot Device

Name: ESXi

State: Enabled

Order: 1 (1 - 2)

Add Virtual Media

Boot Device

Name: KVM_vMedia

Sub Type: KVM MAPPED DVD

State: Enabled

Order: 2 (1 - 2)

Installing VMware ESXi

Follow the steps here to install ESXi.

1. Map the Cisco custom ESXi 6.0 ISO image to the KVM virtual drive (Figure 16).

Figure 16. Mapping the Cisco Custom VMware ESXi Image to the KVM Virtual Drive

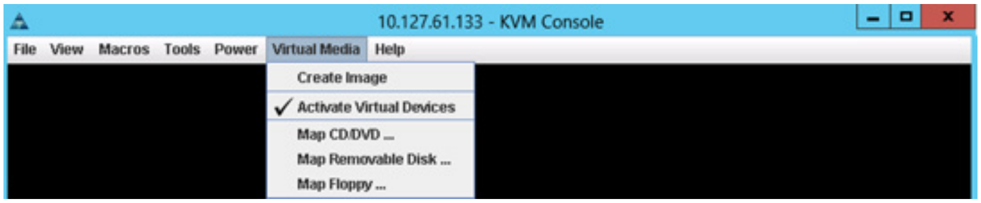

2. The system should detect a 32-GB SD card installed on the server. In the Integrated Management Controller WebGUI, verify that the SD card is detected and that virtual drive options are configured.
3. Boot to ESXi (Figure 17).

Figure 17. Booting to VMware ESXi

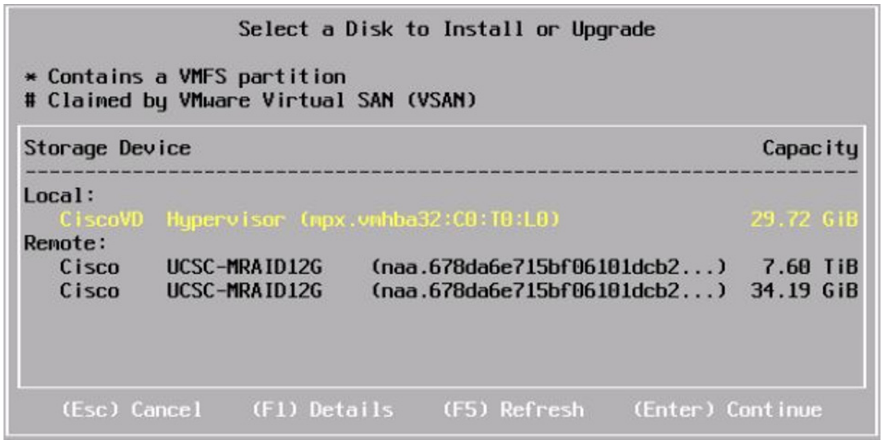

```
                 Select a Disk to Install or Upgrade

* Contains a VMFS partition
# Claimed by VMware Virtual SAN (VSAN)

Storage Device                                                  Capacity
------------------------------------------------------------------------
Local:
   CiscoVD  Hypervisor (mpx.vmhba32:C0:T0:L0)                  29.72 GiB
Remote:
   Cisco    UCSC-MRAID12G    (naa.678da6e715bf06101dcb2...)    7.60 TiB
   Cisco    UCSC-MRAID12G    (naa.678da6e715bf06101dcb2...)   34.19 GiB


   (Esc) Cancel     (F1) Details     (F5) Refresh     (Enter) Continue
```

4. Select the hypervisor partition for OS installation and install ESXi.
5. After ESXi is installed, reboot the system and configure the IP address.
6. Log in to the ESXi host using vSphere Client.
7. Add storage for VSA deployment and for uploading the ESXi patches and device drivers.

Note: See the section Create Persistent Storage for VMware ESXi Logs and VSA Deployment for information about adding the storage to the ESXi host.

8. Download ESXi patches and device drivers from the following link and upload using the data-store browser from vSphere Client: https://my.vmware.com/group/vmware/patch#search
9. Put the ESXi host into maintenance mode.
   `vim-cmd hostsvc/maintenance_mode_enter`
10. From the ESXi shell, install the latest ESXi version (Figure 18).

Figure 18. Installing VMware ESXi

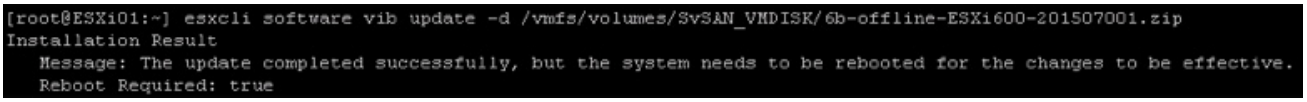

```
[root@ESXi01:~] esxcli software vib update -d /vmfs/volumes/SvSAN_VMDISK/6b-offline-ESXi600-201507001.zip
Installation Result
   Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
   Reboot Required: true
```

11. Reboot the ESXi host and install the latest drivers for the network interface card (NIC), storage, and other components if required from the Cisco driver ISO image (Figure 19)

    Get access to the latest drivers from this link:
    https://software.cisco.com/download/type.html?mdfid=286281356&flowid=71443.

Figure 19. Installing VMware ESXi Drivers

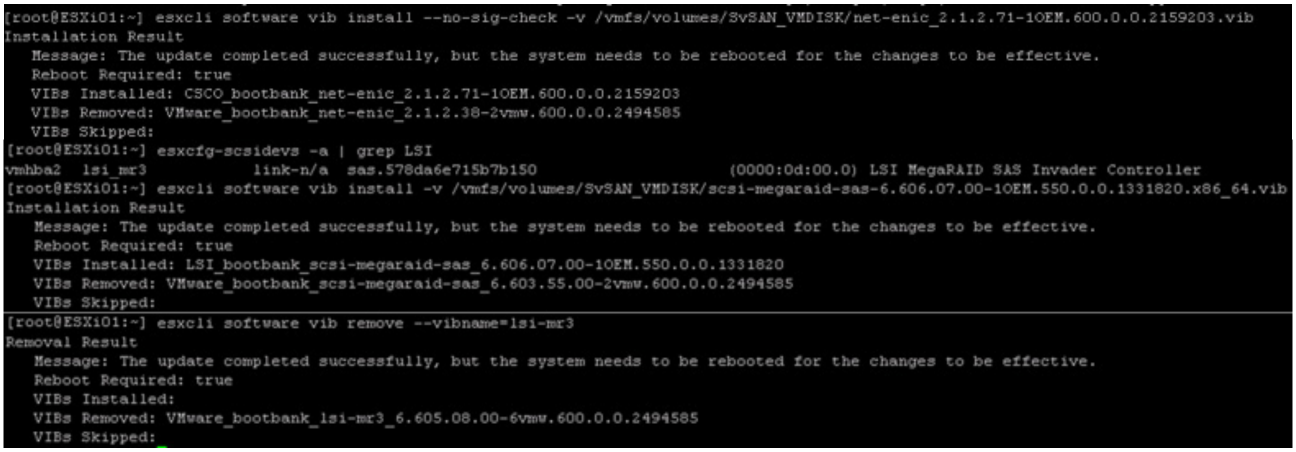

```
[root@ESXi01:~] esxcli software vib install --no-sig-check -v /vmfs/volumes/SvSAN_VMDISK/net-enic_2.1.2.71-1OEM.600.0.0.2159203.vib
Installation Result
   Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
   Reboot Required: true
   VIBs Installed: CSCO_bootbank_net-enic_2.1.2.71-1OEM.600.0.0.2159203
   VIBs Removed: VMware_bootbank_net-enic_2.1.2.38-2vmw.600.0.0.2494585
   VIBs Skipped:
[root@ESXi01:~] esxcfg-scsidevs -a | grep LSI
vmhba2  lsi_mr3           link-n/a  sas.578da6e715b7b150                  (0000:0d:00.0) LSI MegaRAID SAS Invader Controller
[root@ESXi01:~] esxcli software vib install -v /vmfs/volumes/SvSAN_VMDISK/scsi-megaraid-sas-6.606.07.00-1OEM.550.0.0.1331820.x86_64.vib
Installation Result
   Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
   Reboot Required: true
   VIBs Installed: LSI_bootbank_scsi-megaraid-sas_6.606.07.00-1OEM.550.0.0.1331820
   VIBs Removed: VMware_bootbank_scsi-megaraid-sas_6.603.55.00-2vmw.600.0.0.2494585
   VIBs Skipped:
[root@ESXi01:~] esxcli software vib remove --vibname=lsi-mr3
Removal Result
   Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
   Reboot Required: true
   VIBs Installed:
   VIBs Removed: VMware_bootbank_lsi-mr3_6.605.08.00-6vmw.600.0.0.2494585
   VIBs Skipped:
```

Note 1: Remove old drivers if they were not removed during the new driver installation and upgrade processes.

Note 2: You can update ESX drivers using a ZIP file or a VIB file. For a ZIP file, use the -d flag, and for a VIB file, use the -v flag.

12. Reboot the ESXi host and exit from maintenance mode:
    `vim-cmd hostsvc/maintenance_mode_exit`

Create Persistent Storage for VMware ESXi Logs and VSA Deployment

Repeat the following steps on each ESXi host to add the storage for the SvSAN VSA deployment.

1. Log in to vSphere Client and select the host from the **Inventory** pane.
2. Click the **Configuration** tab and click **Storage** in the **Hardware** pane.
3. Click **Datastores** and click **Add Storage**.
4. Select the **Disk/LUN** storage type (Figure 20). Then click **Next**.

Figure 20. Selecting the Storage Type

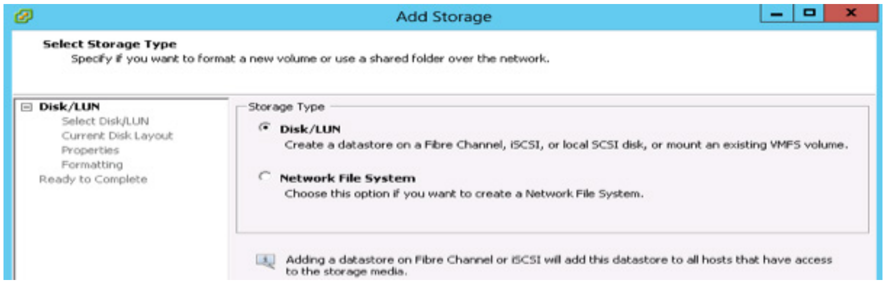

Select Storage Type
Specify if you want to format a new volume or use a shared folder over the network.

Storage Type

● **Disk/LUN**
Create a datastore on a Fibre Channel, iSCSI, or local SCSI disk, or mount an existing VMFS volume.

○ **Network File System**
Choose this option if you want to create a Network File System.

Adding a datastore on Fibre Channel or iSCSI will add this datastore to all hosts that have access to the storage media.

5. On the **Select Disk/LUN** page, select the 34.19-GB storage device (Figure 21). Then click **Next**.

Figure 21. Selecting the Storage Device

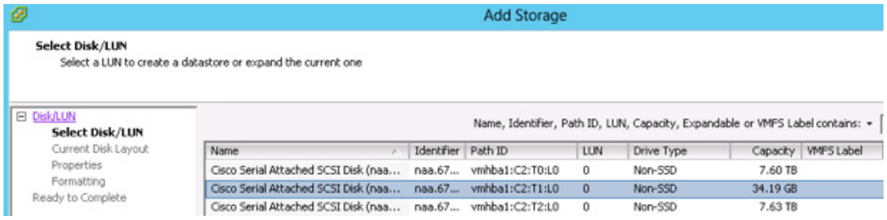

Select Disk/LUN
Select a LUN to create a datastore or expand the current one

Name, Identifier, Path ID, LUN, Capacity, Expandable or VMFS Label contains:

Name	Identifier	Path ID	LUN	Drive Type	Capacity	VMFS Label
Cisco Serial Attached SCSI Disk (naa...	naa.67...	vmhba1:C2:T0:L0	0	Non-SSD	7.60 TB	
Cisco Serial Attached SCSI Disk (naa...	naa.67...	vmhba1:C2:T1:L0	0	Non-SSD	34.19 GB	
Cisco Serial Attached SCSI Disk (naa...	naa.67...	vmhba1:C2:T2:L0	0	Non-SSD	7.63 TB	

6. On the **Current Disk Layout** page, click **Next**.
7. On the **Properties** page, type the data-store name and click **Next**.
8. Specify the storage capacity. Then click **Next**.
9. Click **Finish** to complete the process (Figure 22).

Figure 22. Ready to Complete the Add Storage Process

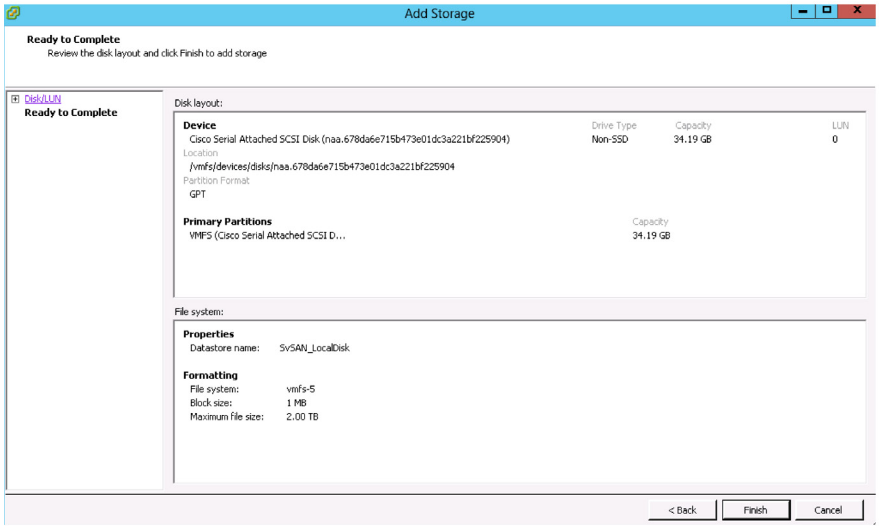

Ready to Complete
Review the disk layout and click Finish to add storage

Disk layout:

Device
Cisco Serial Attached SCSI Disk (naa.678da6e715b473e01dc3a221bf225904)
Drive Type: Non-SSD
Capacity: 34.19 GB
LUN: 0
Location
/vmfs/devices/disks/naa.678da6e715b473e01dc3a221bf225904
Partition Format
GPT

Primary Partitions
VMFS (Cisco Serial Attached SCSI D...
Capacity: 34.19 GB

File system:

Properties
Datastore name: SvSAN_LocalDisk

Formatting
File system: vmfs-5
Block size: 1 MB
Maximum file size: 2.00 TB

10. Click the **Configuration** tab. Under **Software**, click **Advanced Settings**.

11. Verify that Syslog.global.logDir points to a persistent location. The directory should be specified as [datastorename] path_to_file, where the path is relative to the data store: for example, [datastore1] /systemlogs.

12. If the Syslog.global.logDir field is empty or explicitly points to a scratch partition, verify that the field ScratchConfig.CurrentScratchLocation shows a location on persistent storage (Figure 23).

Figure 23. Verifying That Storage Uses a Persistent Location

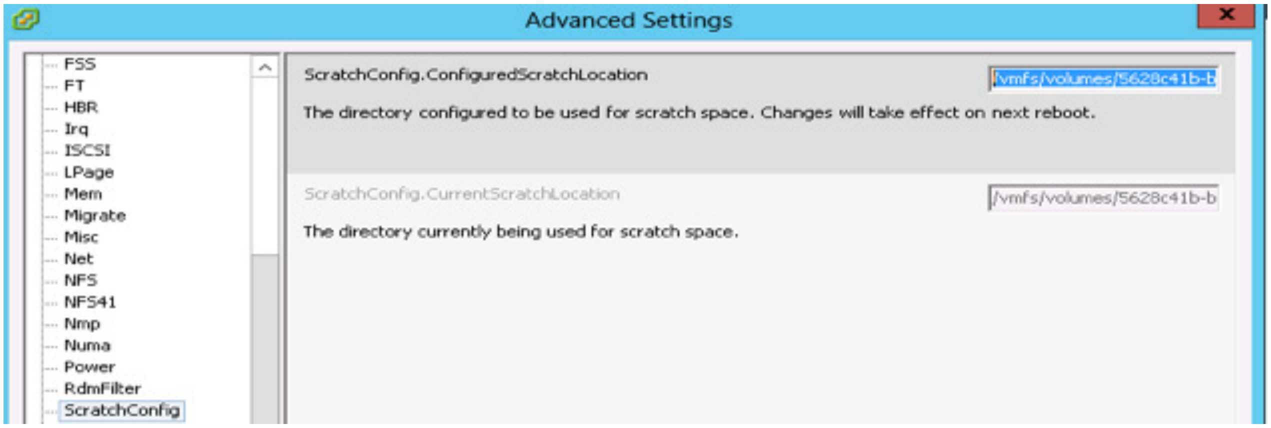

ScratchConfig.ConfiguredScratchLocation: /vmfs/volumes/5628c41b-b
The directory configured to be used for scratch space. Changes will take effect on next reboot.

ScratchConfig.CurrentScratchLocation: /vmfs/volumes/5628c41b-b
The directory currently being used for scratch space.

Note: For more information about configuring the scratch location for system logs, see the VMware knowledgebase article at http://kb.vmware.com/kb/2032823.

13. Reboot the ESXi host to make the changes take effect.

Configuring the VMware ESXi Network for StorMagic SvSAN VSA Deployment

SvSAN uses three types of logical interfaces for traffic:

- **Management:** This interface is used for accessing the WebGUI, plug-in, or CLI. At least one management interface is required.
- **iSCSI:** This interface listens for incoming connection requests from iSCSI initiators. At least one iSCSI is required.
- **Mirror:** This interface is used by VSA nodes to communicate data and metadata associated with mirrored volumes.

StorMagic recommends the use of four or more physical NICs between hosts and VSAs to provide the recommended level of redundancy and load balancing.

Each Cisco UCS C240 M4 server has two 1-Gbps network ports and two 10-Gbps network ports. The test environment reported in this document used four network ports along with a neutral storage host (NSH) running outside the VMware High Availability (HA) cluster to prevent a split-brain scenario.

The aim of this configuration is to eliminate the need for a physical switch on the iSCSI and mirror networks, allowing the use of crossover cables by deploying three virtual (vSwitches). This configuration also splits the iSCSI and mirroring traffic onto separate links. If the iSCSI network fails because of a cable pull or a faulty NIC, the mirror network is automatically used for iSCSI traffic by the ESXi hosts. If the mirror network fails, mirroring traffic automatically goes through the iSCSI network to help ensure that the mirrors stay synchronized. After the fault has been resolved, traffic automatically fails back. To achieve the highest possible networking throughput, the VSAs can be configured to use jumbo frames.

Configure the vSwitches

Each host is configured with three vSwitches and configured as shown here.

vSwitch0 is composed of two 1-Gbps network adapters (teamed in ESXi) and used for ESXi, SvSAN, and virtual machine management. The VSA uses vSwitch0 for management, administration, and monitoring purposes only. The NSH uses the management network to monitor the SvSAN heartbeat. vSwitch1 uses one 10-Gbps network port and is used for vMotion, SvSAN iSCSI, and mirror traffic. vSwitch2 also uses one 10-Gbps network port and is used for vMotion, SvSAN iSCSI, and mirror traffic.

vSwitch0

- VMkernel: Management network
- Virtual machine port group: Production virtual machines and SvSAN management

vSwitch1

- VMkernel: Port network: vMotion, iSCSI traffic
- Virtual machine port group: SvSAN iSCSI and mirror traffic

vSwitch2

- VMkernel: Port network; vMotion and iSCSI traffic
- Virtual machine port group: SvSAN iSCSI and mirror traffic

Edit all the vSwitches and VMkernel ports used for iSCSI, mirror, and vMotion traffic and set the MTU size to 9000, as shown in the Figures 24 and 25.

Figure 24. Configuring vSwitch Ports

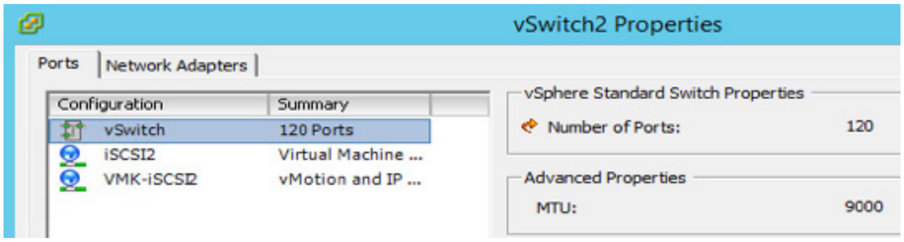

Configuration	Summary
vSwitch	120 Ports
iSCSI2	Virtual Machine ...
VMK-iSCSI2	vMotion and IP ...

vSphere Standard Switch Properties

Number of Ports: 120

Advanced Properties

MTU: 9000

Figure 25. Configuring vSwitches

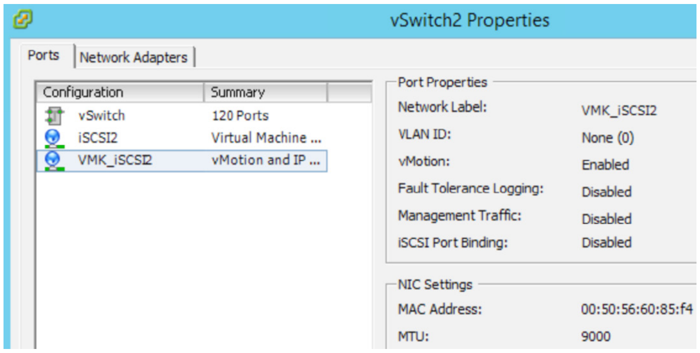

Configuration	Summary
vSwitch	120 Ports
iSCSI2	Virtual Machine ...
VMK_iSCSI2	vMotion and IP ...

Port Properties

Network Label:	VMK_iSCSI2
VLAN ID:	None (0)
vMotion:	Enabled
Fault Tolerance Logging:	Disabled
Management Traffic:	Disabled
iSCSI Port Binding:	Disabled

NIC Settings

MAC Address:	00:50:56:60:85:f4
MTU:	9000

Figure 26 shows all the vSwitches and the VMkernel and virtual machine port groups created on all the ESXi hosts.

Figure 26. VMware ESXi Network Configuration: Three vSwitches

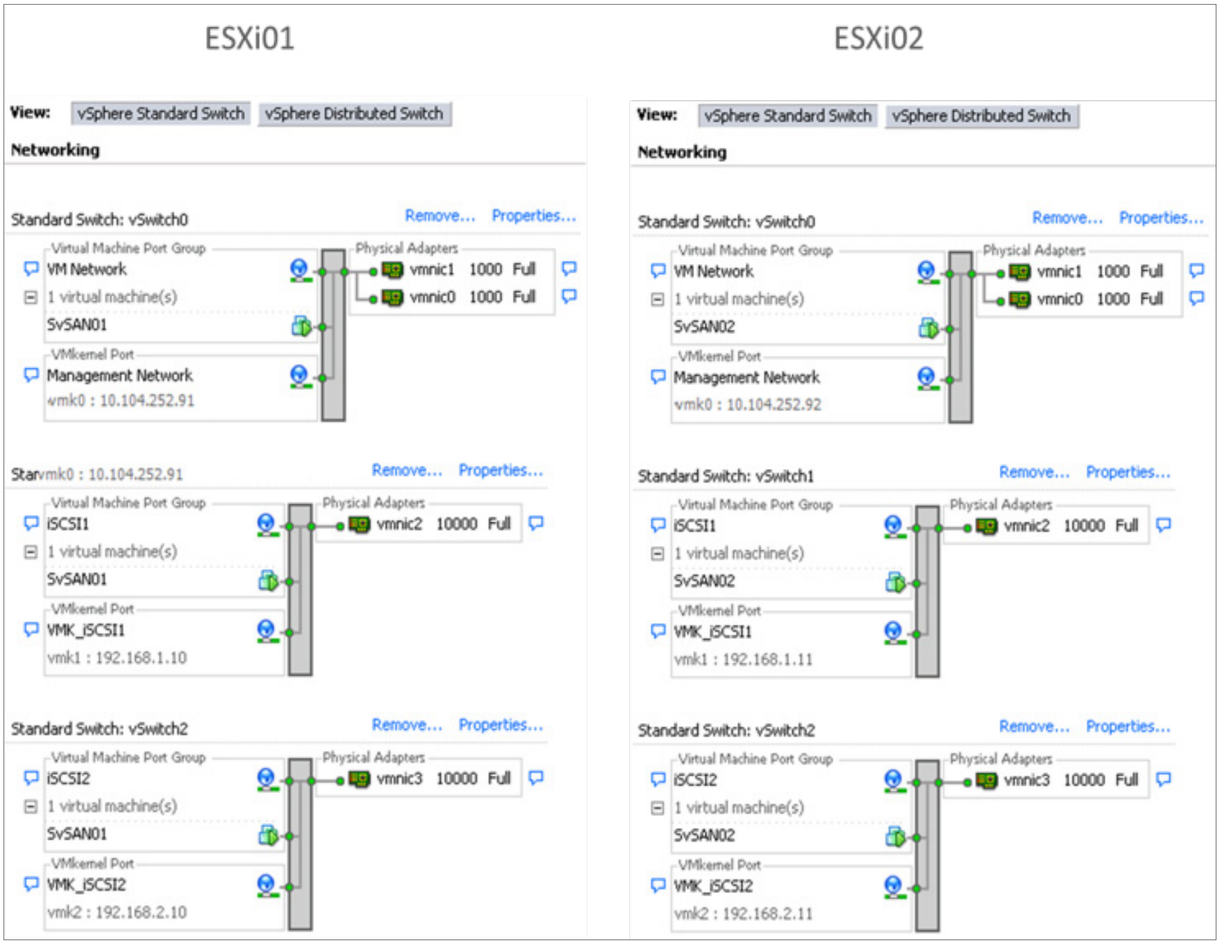

ESXi01

View: vSphere Standard Switch | vSphere Distributed Switch

Networking

Standard Switch: vSwitch0 — Remove... Properties...

- Virtual Machine Port Group: VM Network — 1 virtual machine(s): SvSAN01
- VMkernel Port: Management Network — vmk0 : 10.104.252.91
- Physical Adapters: vmnic1 1000 Full; vmnic0 1000 Full

Starvmk0 : 10.104.252.91 — Remove... Properties...

- Virtual Machine Port Group: iSCSI1 — 1 virtual machine(s): SvSAN01
- VMkernel Port: VMK_iSCSI1 — vmk1 : 192.168.1.10
- Physical Adapters: vmnic2 10000 Full

Standard Switch: vSwitch2 — Remove... Properties...

- Virtual Machine Port Group: iSCSI2 — 1 virtual machine(s): SvSAN01
- VMkernel Port: VMK_iSCSI2 — vmk2 : 192.168.2.10
- Physical Adapters: vmnic3 10000 Full

ESXi02

View: vSphere Standard Switch | vSphere Distributed Switch

Networking

Standard Switch: vSwitch0 — Remove... Properties...

- Virtual Machine Port Group: VM Network — 1 virtual machine(s): SvSAN02
- VMkernel Port: Management Network — vmk0 : 10.104.252.92
- Physical Adapters: vmnic1 1000 Full; vmnic0 1000 Full

Standard Switch: vSwitch1 — Remove... Properties...

- Virtual Machine Port Group: iSCSI1 — 1 virtual machine(s): SvSAN02
- VMkernel Port: VMK_iSCSI1 — vmk1 : 192.168.1.11
- Physical Adapters: vmnic2 10000 Full

Standard Switch: vSwitch2 — Remove... Properties...

- Virtual Machine Port Group: iSCSI2 — 1 virtual machine(s): SvSAN02
- VMkernel Port: VMK_iSCSI2 — vmk2 : 192.168.2.11
- Physical Adapters: vmnic3 10000 Full

Figure 27 shows the SvSAN architecture with three vSwitches and an NSH.

Figure 27. StorMagic SvSAN Architecture with Neutral Storage Host

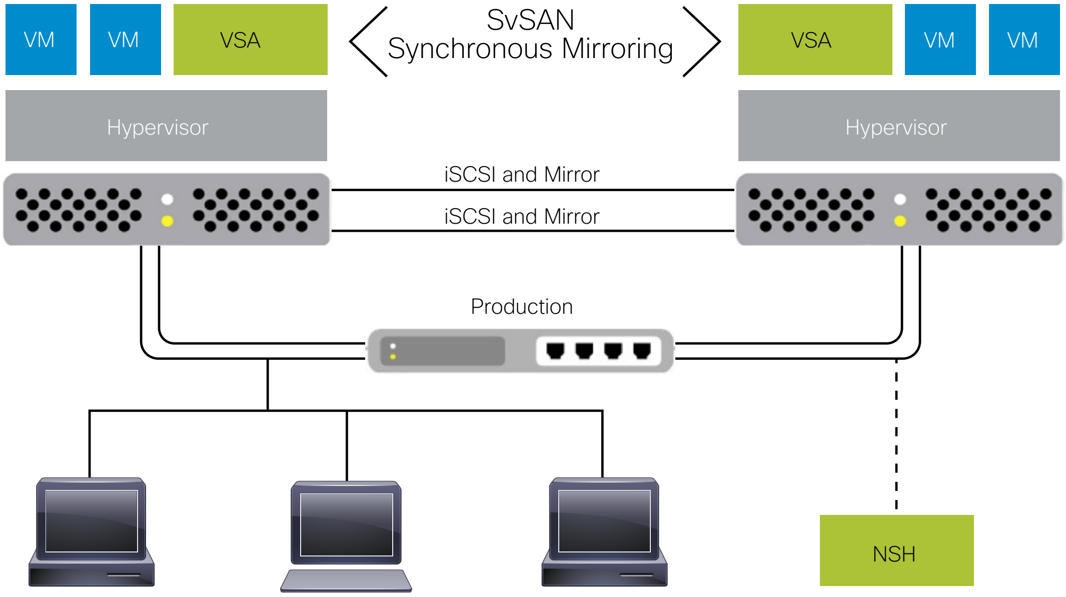

Configure the VMware ESXi Software iSCSI Initiator

Repeat the following steps on each ESXi host to configure the ESXi software iSCSI initiator.

1. In vSphere Client or vSphere Web Client, select the ESXi host.
2. On the **Configuration** tab, click **Storage Adapters**.
3. Click **Add** and add the software iSCSI adapter.

Note: SAN connectivity is configured for the ESXi host's iSCSI software adapter. This connectivity requires at least one vSwitch configured with VMkernel ports.

4. Enable the iSCSI port (TCP port 3260) on the VMware firewall.
5. Enable the iSCSI software initiator on ESXi hosts.
6. Enable Secure Shell (SSH) on the ESXi hosts to allow VSAs to be deployed.

Table 4 lists the network ports that SvSAN uses.

Table 4. Network Ports

Component or Service	Protocol or Port
Discovery	UDP 4174
XMLRPC server	TCP 8000
Inter-VSA communications	TCP 4174
SMDXMLRPC	TCP 43127
Web Access	TCP 80 and 443
Management services	TCP 8990
vSphere Client plug-in	TCP 80 and 443
Miscellaneous VMware services	TCP 16961, 16962, and 16963

Additional port numbers used by SvSAN are listed at http://www.stormagic.com/manual/SvSAN_5-3/en/Content/port-numbers-used-by-SvSAN.htm.

Installing StorMagic vCenter Integration Services Package and Neutral Storage Service on VMware vCenter Server

SvSAN management components are installed directly on vCenter Server. This package includes a number of services that are required to orchestrate VSA deployment and storage provisioning.

Before you begin the installation, verify that your system meets the requirements listed in Table 5. You can use Microsoft Windows or VMware vCenter Server Virtual Appliance (vCSA).

Table 5. VMware vCenter and StorMagic SvSAN Version Compatibility

VMware vCenter	StorMagic SvSAN
VMware vCenter version	Release 5.3
VMware vCenter Server 5.1 VMware vCenter Server 5.1a VMware vCenter Server 5.1b VMware vCenter Server 5.1 Update 1 VMware vCenter Server 5.1 Update 1a VMware vCenter Server 5.1 Update 1b VMware vCenter Server 5.1 Update 1c VMware vCenter Server 5.1 Update 2 VMware vCenter Server 5.1 Update 2a VMware vCenter Server 5.1 Update 3 VMware vCenter Server 5.1 Update 3a	Compatible
VMware vCenter 5.5 VMware vCenter 5.5.0a VMware vCenter 5.5.0b VMware vCenter 5.5.0c VMware vCenter 5.5.0 Update 1 VMware vCenter 5.5.0 Update 1a VMware vCenter 5.5.0 Update 1b VMware vCenter 5.5.0 Update 1c VMware vCenter 5.5.0 Update 2 VMware vCenter 5.5.0 Update 2d VMware vCenter 5.5.0 Update 2e	Compatible
VMware vCenter Server 6.0.0 VMware vCenter Server 6.0.0 Update 1	Compatible

Visual C++ runtime libraries and Microsoft .NET FX 3.5 SP1 are required. These are automatically installed (if they are not already present) when the SvSAN setup executable file (setup.exe) is run.

Note: If you are using Microsoft Windows Server 2012, you must perform this operation manually by choosing **Server Manager** > **Add Roles and Features**.

This section describes how to install the StorMagic VCenter Integration Services package and Neutral Storage Service in vCenter. The NSH is recommended to prevent a split-brain scenario in the event of a network fault.

Before you install SvSAN vCenter integration Service on vCenter Server, create a data center and a cluster and add ESXi hosts to the cluster.

Note: Disable the HA and DRS features while configuring the VMware cluster. After the SvSAN mirrored storage and workload is configured, you can enable the HA and DRS features and configure the VMware cluster according to your requirements.

In the Cisco UCS C-Series configuration, NSH is installed on a vCenter Server that is running on a separate ESXi host. This setup helps prevent a split-brain scenario in the event of a network fault.

1. Log in to the vCenter Server.
2. Select **Run as Administrator** and run the setup.exe file provided in the downloaded zip file.
3. Click **Next**. The end-user license agreement (EULA) opens.

Note: There are different versions of the EULA for the USA and for the rest of the world. Scroll down to find the appropriate agreement.

4. To accept the agreement, select the box. Then click **Next**.
5. Click **Custom**.
6. Verify that the **StorMagic Neutral Storage Service** and **StorMagic vCenter Integration Services** options are selected (Figure 28). The plug-in components will be installed on vCenter Server with the Neutral Storage Service.

Figure 28. Verifying That StorMagic Neutral Storage Service and StorMagic vCenter Integration Services Are Selected

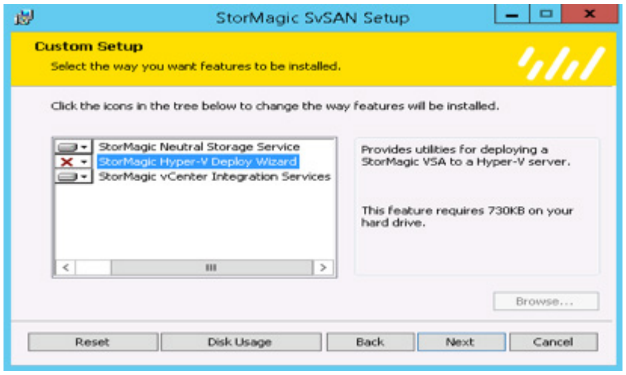

StorMagic SvSAN Setup

Custom Setup

Select the way you want features to be installed.

Click the icons in the tree below to change the way features will be installed.

- StorMagic Neutral Storage Service
- ☒ StorMagic Hyper-V Deploy Wizard
- StorMagic vCenter Integration Services

Provides utilities for deploying a StorMagic VSA to a Hyper-V server.

This feature requires 730KB on your hard drive.

7. Click **Next** and complete the final steps of the wizard. The StorMagic tab will appear in the data center on the Manage tab (Figure 29).

Figure 29. Completed Installation

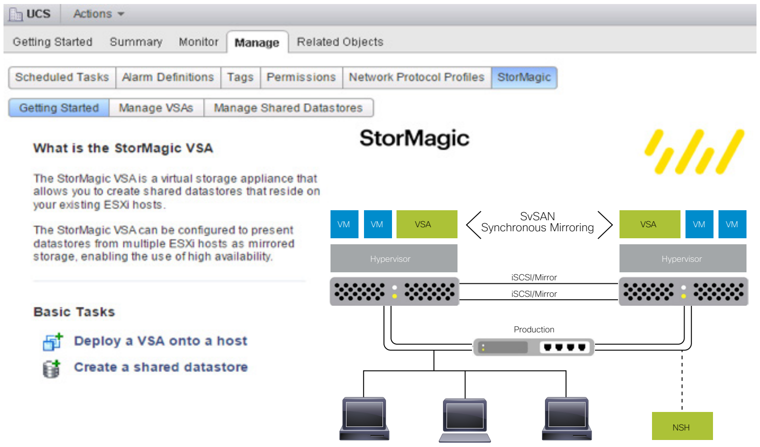

What is the StorMagic VSA

The StorMagic VSA is a virtual storage appliance that allows you to create shared datastores that reside on your existing ESXi hosts.

The StorMagic VSA can be configured to present datastores from multiple ESXi hosts as mirrored storage, enabling the use of high availability.

Basic Tasks

- Deploy a VSA onto a host
- Create a shared datastore

Deploying and Managing Virtual Storage Appliances

Before you deploy VSAs, verify that your system meets the following requirements:

- Configure vSwitch1 and vSwitch2, with one VMkernel port for SAN connectivity configured for the vCenter iSCSI software adapter.
- If you are using RDMs to assign storage, enable SSH on the host. However, SSH needs to be enabled only for the duration of the deployment of the VSA. You can disable it immediately after deployment is complete.
- Add the virtual disk that was created for the SvSAN VSA virtual machines in ESXi and create a VMFS data store.
- Enable the iSCSI software initiator on ESXi hosts.
- Enable the iSCSI port (TCP port 3260) on the VMware firewall.

A typical SvSAN environment consists of two hypervisor hosts. You can deploy a VSA by using the plug-in. You can install a VSA on each host in a data store residing on server internal or direct-attached storage. A 512-MB boot disk is required along with a 20-GB disk for storing metadata. Any additional internal or direct-attached storage assigned to the SvSAN VSA is used to create synchronously mirrored volumes. The VSA requires a CPU reservation of 2 GHz and a memory reservation of 1 GB.

Deploy a VSA to a Host

Repeat the following steps for each VSA to be deployed.

1. Open the plug-in: in vSphere Web Client, select your data center and then navigate to **Manage > StorMagic** (Figure 30).

Figure 30. Opening the Plug-in

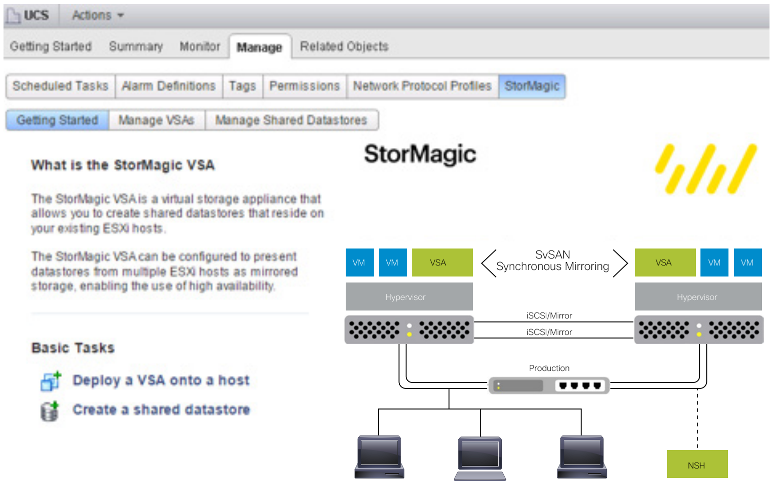

What is the StorMagic VSA

The StorMagic VSA is a virtual storage appliance that allows you to create shared datastores that reside on your existing ESXi hosts.

The StorMagic VSA can be configured to present datastores from multiple ESXi hosts as mirrored storage, enabling the use of high availability.

Basic Tasks

2. Click **Deploy a VSA onto a host**. The deployment wizard opens. Click **Next**.
3. Select a host to which to deploy a VSA. Click **Next**.
4. Read and accept the terms and conditions. Click **Next**.

5. Enter a hostname for the VSA (unique on your network) and domain name and choose the data store on which the VSA virtual machine is to reside (Figure 31). The data store should be on internal or direct-attached server storage. Two virtual machine disks (VMDKs) are created in this data store: a 512-MB boot disk and a 20-GB journal disk.

Figure 31. Deploying a VSA

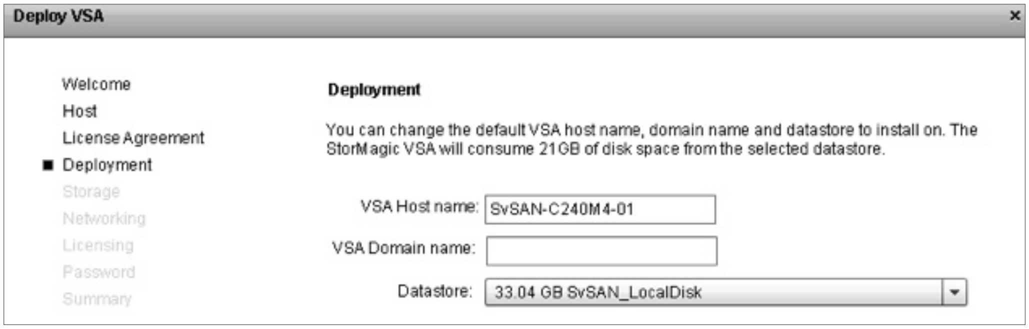

Deployment

You can change the default VSA host name, domain name and datastore to install on. The StorMagic VSA will consume 21GB of disk space from the selected datastore.

VSA Host name: SvSAN-C240M4-01

VSA Domain name:

Datastore: 33.04 GB SvSAN_LocalDisk

6. Choose the desired storage allocation technique (Figure 32). Raw device mappings offer the best performance; however, SSH must be enabled on the host, but only for the duration of the deployment (SSH can be disabled immediately after deployment is complete).

Warning: If you select an RDM device that has data stored on it, that data will be permanently deleted during deployment.

Figure 32. Choosing the Storage Allocation Technique

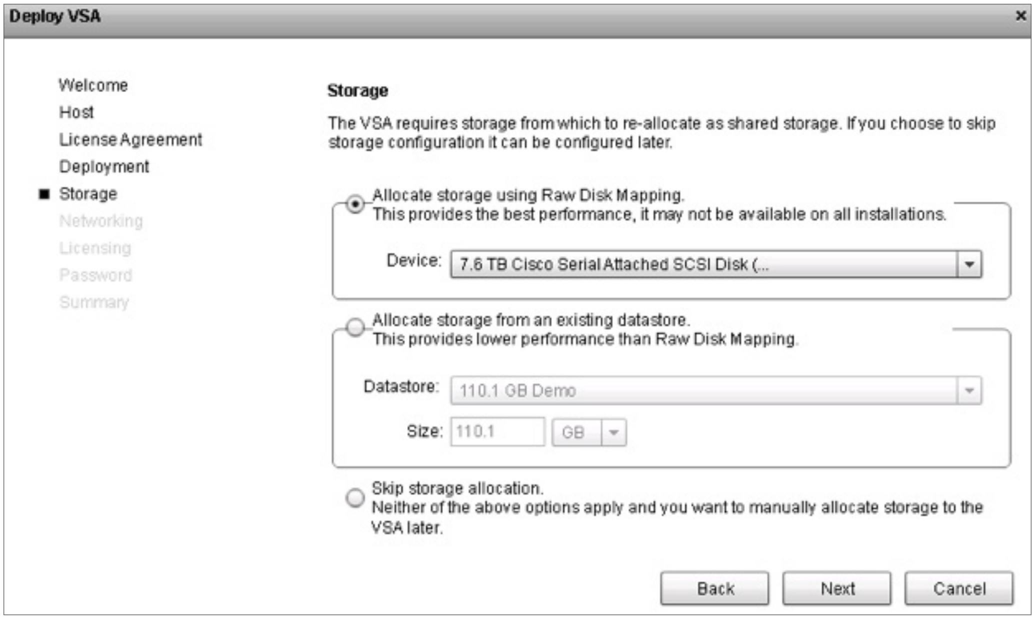

Storage

The VSA requires storage from which to re-allocate as shared storage. If you choose to skip storage configuration it can be configured later.

● Allocate storage using Raw Disk Mapping.
This provides the best performance, it may not be available on all installations.

Device: 7.6 TB Cisco Serial Attached SCSI Disk (...

○ Allocate storage from an existing datastore.
This provides lower performance than Raw Disk Mapping.

Datastore: 110.1 GB Demo

Size: 110.1 GB

○ Skip storage allocation.
Neither of the above options apply and you want to manually allocate storage to the VSA later.

7. If you want to allocate multiple RDMs as a pool, select the **Advanced options** checkbox. If the pool you want is listed, choose it. Otherwise, click **Add**. The Create Storage Pool window opens.

Note: Complete the fields to create the pool; then click OK. For more information, see the section "Creating a Pool in the WebGUI" in the StorMagic documentation at : http://www.stormagic.com/manual/SvSAN_5-3/en/.

8. Click **Next**. In the Caching configuration window, select the **Skip cache allocation** radio button.

Note: Optionally, you can enable an SSD cache device on the VSA. RAID 1 with two SSD devices is recommended (check the box). Select the SSD volumes to be used as the cache storage. When the VSA is deployed, cache storage with the default chunk size is created automatically and is ready for use by simple and mirrored volumes. For more information about this licensed feature, see http://www.stormagic.com/manual/SvSAN_5-3/en/#caching.htm%3FTocPath%3DWeb%20GUI%7CTargets%7C_____6.

9. You can let the Networking page acquire IP addresses through Domain Host Configuration Protocol (DHCP), as shown in Figure 33, or you can set the IP addresses statically. Select the network interface and click **Configure** to select the interface traffic types. In the setup discussed here, the virtual machine network interface is configured for management and iSCSI1 traffic, and the iSCSI2 interface is configured for iSCSI and mirroring traffic, as shown in Figure 34.

Figure 33. Configuring Network Interfaces Automatically with DHCP

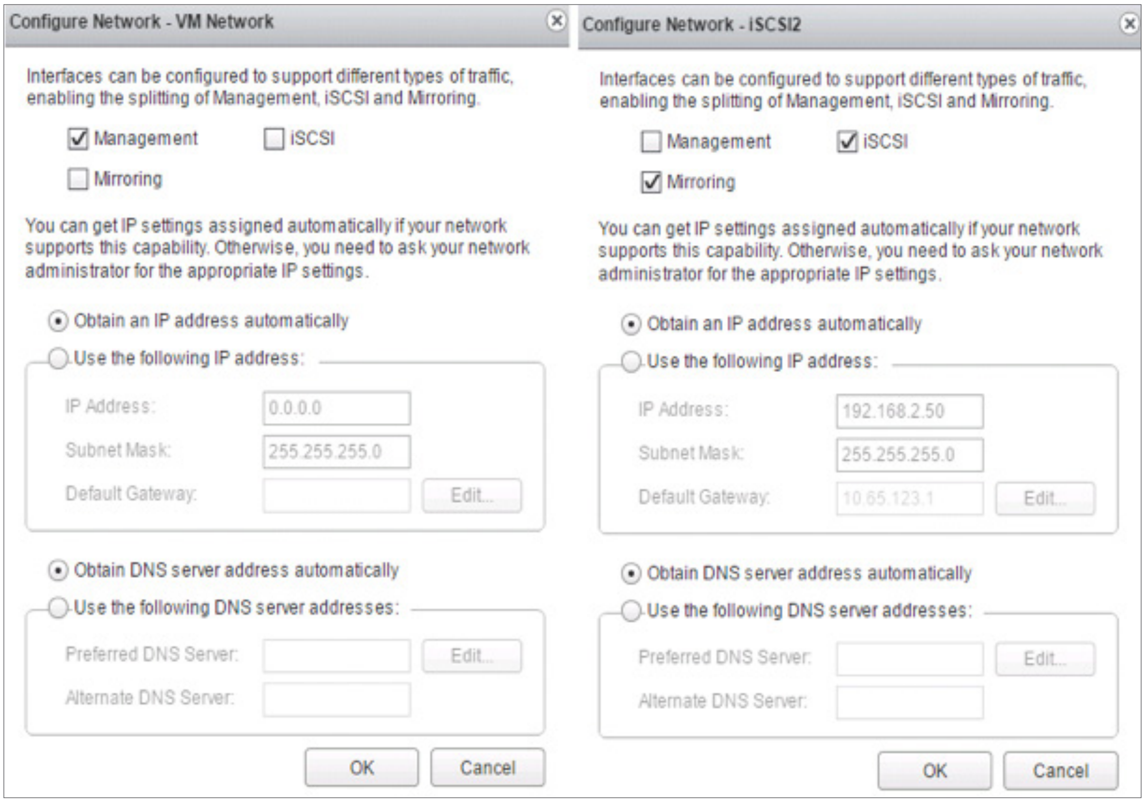

Configure Network - VM Network

Interfaces can be configured to support different types of traffic, enabling the splitting of Management, iSCSI and Mirroring.

☑ Management ☐ iSCSI
☐ Mirroring

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

● Obtain an IP address automatically
○ Use the following IP address:

IP Address: 0.0.0.0
Subnet Mask: 255.255.255.0
Default Gateway:

● Obtain DNS server address automatically
○ Use the following DNS server addresses:

Preferred DNS Server:
Alternate DNS Server:

Configure Network - iSCSI2

Interfaces can be configured to support different types of traffic, enabling the splitting of Management, iSCSI and Mirroring.

☐ Management ☑ iSCSI
☑ Mirroring

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

● Obtain an IP address automatically
○ Use the following IP address:

IP Address: 192.168.2.50
Subnet Mask: 255.255.255.0
Default Gateway: 10.65.123.1

● Obtain DNS server address automatically
○ Use the following DNS server addresses:

Preferred DNS Server:
Alternate DNS Server:

Figure 34. Configuring Network Interfaces Statically

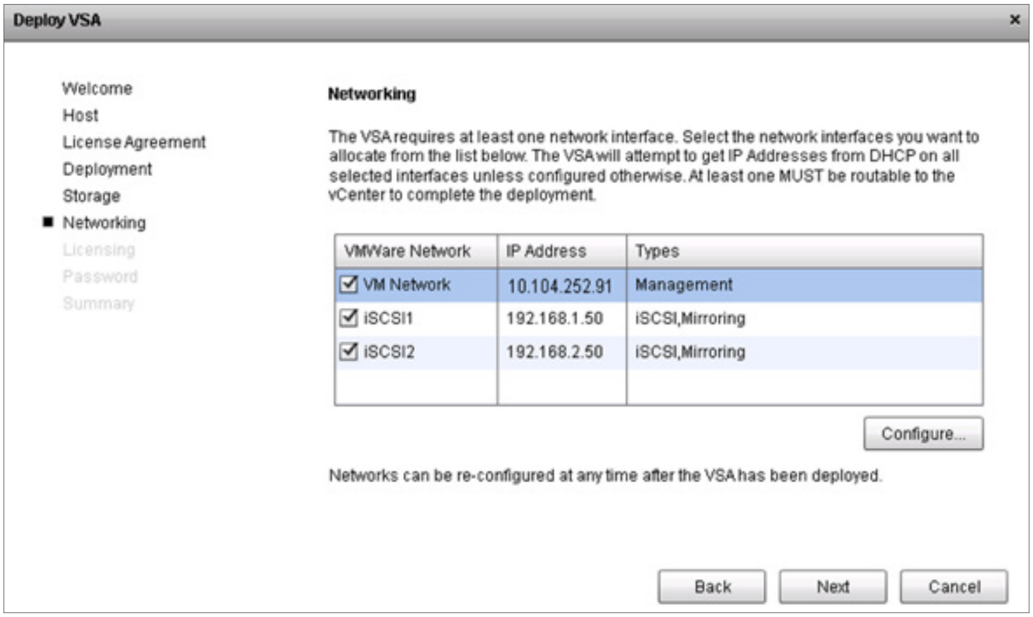

Networking

The VSA requires at least one network interface. Select the network interfaces you want to allocate from the list below. The VSA will attempt to get IP Addresses from DHCP on all selected interfaces unless configured otherwise. At least one MUST be routable to the vCenter to complete the deployment.

VMWare Network	IP Address	Types
☑ VM Network	10.104.252.91	Management
☑ iSCSI1	192.168.1.50	iSCSI,Mirroring
☑ iSCSI2	192.168.2.50	iSCSI,Mirroring

Networks can be re-configured at any time after the VSA has been deployed.

Note: Multiple networks are listed if multiple vSwitches are configured on the ESXi host. The VSA creates an interface on all vSwitches by default. If you do not want the VSA to create an interface on specific vSwitches, clear the box associated with the virtual machine port group. For each interface, you can choose its type.

10. Enter the VSA license information and click **Next**.

Note: During deployment, the VSA attempts to connect to StorMagic's license server to validate the license. If it needs to use a proxy server to do this, supply the proxy server details. If the VSA does not have Internet connectivity, license it later using an offline activation mechanism.

11. Enter the VSA management password, then click **Next**. Summary information about the VSA is displayed.
12. Click **Finish** to deploy the VSA.
13. You can monitor the deployment process in the VMware Recent Tasks window. Use the **StorMagic deploy OVF** task to view the overall percentage of the deployment completed. When this task finishes, the VSA will be booted and become operational.

Note: A VSA is deployed from an Open Virtualized Format (OVF) file. An OVF file is a platform-independent, efficient, extensible, and open packaging and distribution format for virtual machines, making it an excellent format for packaging StorMagic VSAs for distribution and deployment.

14. After the VSA has been deployed, you can use it to create data stores

Note: The VSA requires at least one network interface to be routable to vCenter; otherwise, the VSA deployment will fail.

Configure Jumbo Frames on a StorMagic SvSAN VSA Logical NIC

Repeat the following steps on each SvSAN VSA to configure jumbo frames on VSA logical NICs.

1. Log in to the vSphere Web Client and Click **Hosts and Clusters**.
2. Select the data center at which the VSAs are deployed. Navigate to **Manage > StorMagic > Manage VSAs**.
3. Select the VSA. A management URL will appear. Click the URL.
4. A new tab will open connecting directly to the VSA. The default username is admin. Supply the password entered at VSA deployment.
5. Click **Network** and then select the network device you want to edit.
6. From the **Actions** menu, choose **Edit**.
7. Edit the MTU value to the desired size to match the rest of the environment network configuration (Figure 35). Then click **Apply** (Figure 36).

Figure 35. Editing the Network Device

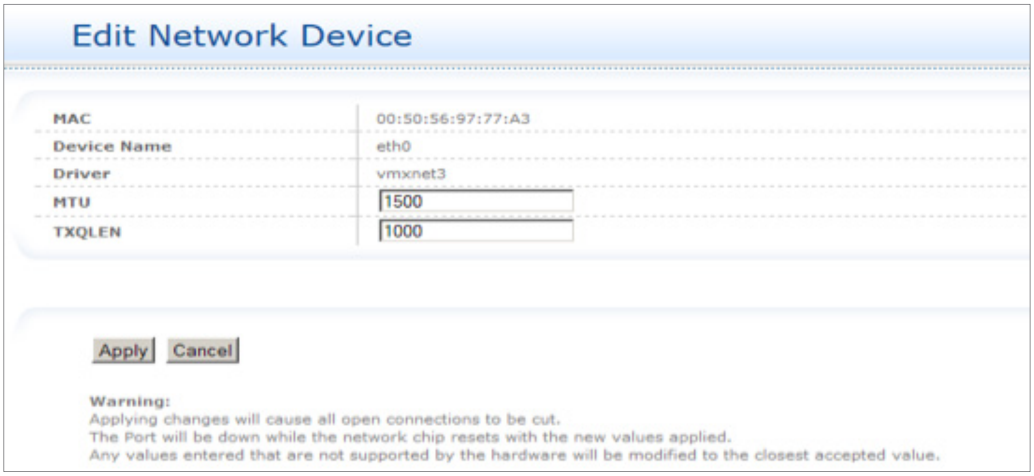

Edit Network Device

MAC	00:50:56:97:77:A3
Device Name	eth0
Driver	vmxnet3
MTU	1500
TXQLEN	1000

Warning:
Applying changes will cause all open connections to be cut.
The Port will be down while the network chip resets with the new values applied.
Any values entered that are not supported by the hardware will be modified to the closest accepted value.

Figure 36. Edited Devices

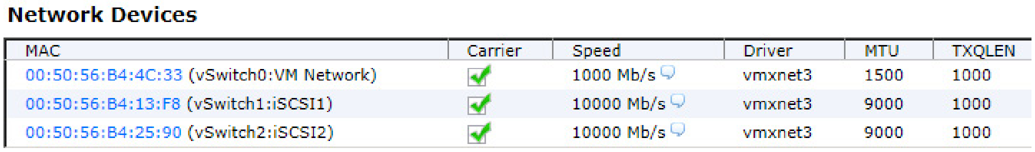

Network Devices

MAC	Carrier	Speed	Driver	MTU	TXQLEN
00:50:56:B4:4C:33 (vSwitch0:VM Network)	✓	1000 Mb/s	vmxnet3	1500	1000
00:50:56:B4:13:F8 (vSwitch1:iSCSI1)	✓	10000 Mb/s	vmxnet3	9000	1000
00:50:56:B4:25:90 (vSwitch2:iSCSI2)	✓	10000 Mb/s	vmxnet3	9000	1000

Manage VSAs

You can manage VSAs in the plug-in.

1. To check the status of your VSAs, in the plug-in, click **Manage VSAs** (Figure 37).

Figure 37. Checking the Status of VSAs

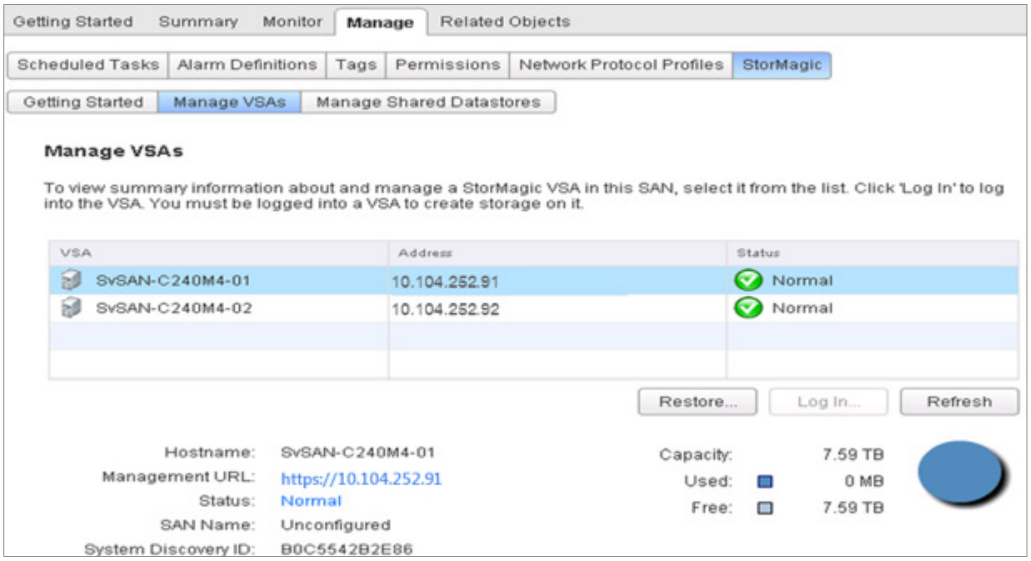

Manage VSAs

To view summary information about and manage a StorMagic VSA in this SAN, select it from the list. Click 'Log In' to log into the VSA. You must be logged into a VSA to create storage on it.

VSA	Address	Status
SvSAN-C240M4-01	10.104.252.91	✓ Normal
SvSAN-C240M4-02	10.104.252.92	✓ Normal

Hostname: SvSAN-C240M4-01
Management URL: https://10.104.252.91
Status: Normal
SAN Name: Unconfigured
System Discovery ID: B0C5542B2E86

Capacity:		7.59 TB
Used:	■	0 MB
Free:	□	7.59 TB

2. The screen lists the VSAs with their IP addresses and system status. Select a VSA to see more information about it, including the amount of pool storage used and the amount available, the system serial number, and the firmware version. If you want to change the VSA virtual machine name, right-click the VSA's virtual machine and then choose **Rename**.

Creating a Shared Data Store with Mirrored Storage

Use the steps here to create a shared data store with mirrored storage.

1. In vSphere Web Client, select your data center and then choose **Manage > StorMagic**. The plug-in opens (Figure 38).

Figure 38. Opening the Plug-in

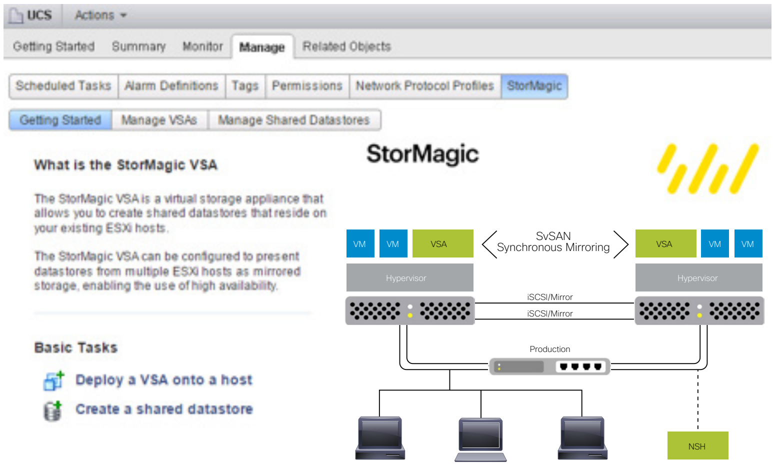

What is the StorMagic VSA

The StorMagic VSA is a virtual storage appliance that allows you to create shared datastores that reside on your existing ESXi hosts.

The StorMagic VSA can be configured to present datastores from multiple ESXi hosts as mirrored storage, enabling the use of high availability.

Basic Tasks

Deploy a VSA onto a host

Create a shared datastore

2. Click **Create a shared datastore**. The Create Datastore wizard opens. Click **Next**.
3. Enter the data store name.
4. Set the provisioned size. To set the size so that all the available space is used, select **Use all**.
5. To create mirrored storage, select two VSAs.
6. By default, the pool to be used on each VSA is selected automatically (Figure 39). To specify the pool, select the **Advanced options** checkbox.

Figure 39. Creating a Shared Data Store

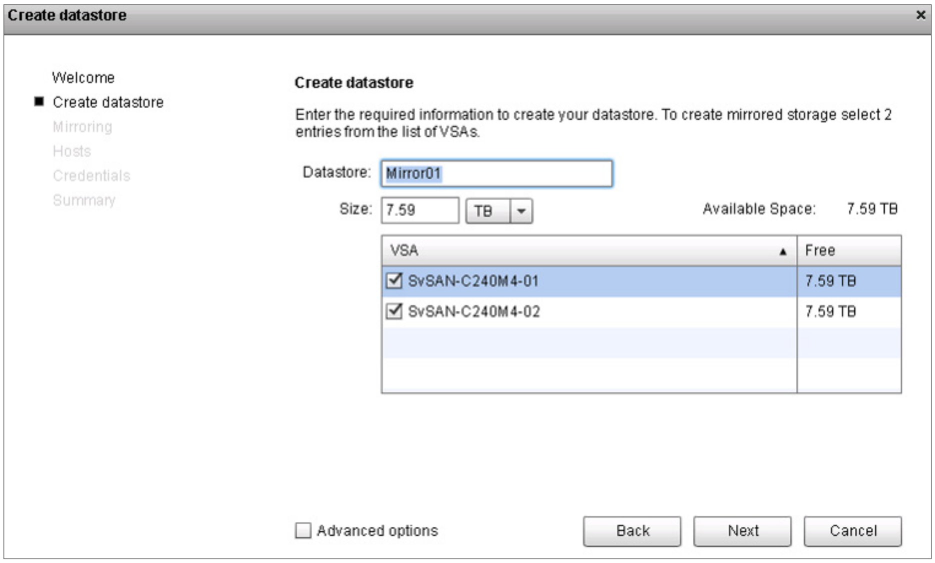

Create datastore

Enter the required information to create your datastore. To create mirrored storage select 2 entries from the list of VSAs.

Datastore: Mirror01

Size: 7.59 TB

Available Space: 7.59 TB

VSA	Free
☑ SvSAN-C240M4-01	7.59 TB
☑ SvSAN-C240M4-02	7.59 TB

☐ Advanced options

7. Click **Next**.
8. Select the neutral storage host for mirroring (Figure 40). The **Advanced** option lets you select the **Up** isolation policy, which does not require an NSH. Click **Next**.

Figure 40. Selecting the NSH for Mirroring

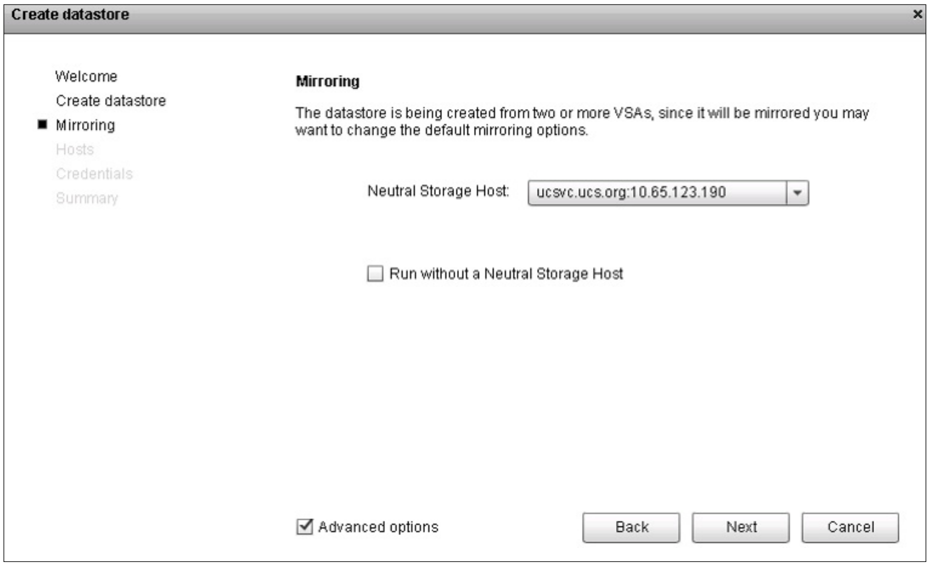

Mirroring

The datastore is being created from two or more VSAs, since it will be mirrored you may want to change the default mirroring options.

Neutral Storage Host: ucsvc.ucs.org:10.65.123.190

□ Run without a Neutral Storage Host

☑ Advanced options

9. Optionally, enable caching on the data store. This option is available only if your license includes caching and if a cache device was enabled when the VSA was deployed. If you enable caching, default cache settings are used; these can be modified at a later time using the WebGUI. Click **Next**.
10. Select the hosts that are to have access to the SvSAN data store (Figure 41). Click **Next.**

Figure 41. Selecting Hosts That Can Access the Data Store

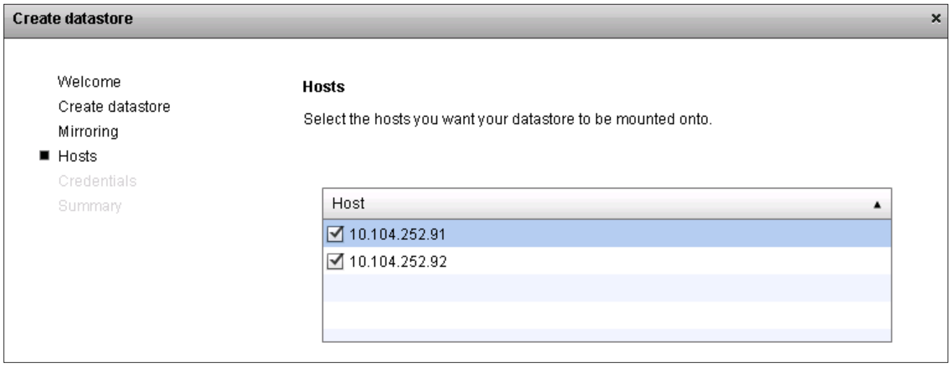

Hosts

Select the hosts you want your datastore to be mounted onto.

Host ▲
☑ 10.104.252.91
☑ 10.104.252.92

11. The first time you create a data store, the VSAs must authenticate with their hosts if they were not already saved during deployment. For each host, provide the ESXi host administrative password (the password that was used when ESXi was installed). Click **Next**.
12. Click **Finish** to start the data-store creation process. You can monitor the data-store creation process in the VMware Recent Tasks window. After this task finishes, the data store is available on the ESXi hosts.
13. When a mirror is created, full synchronization is performed to help ensure that both sides of the mirror are synchronized.

Managing Shared Data Stores

You can manage the shared data stores.

1. To manage shared data stores, in the plug-in, click **Manage Shared Datastores** (Figure 42).

Figure 42. Managing Shared Data Stores

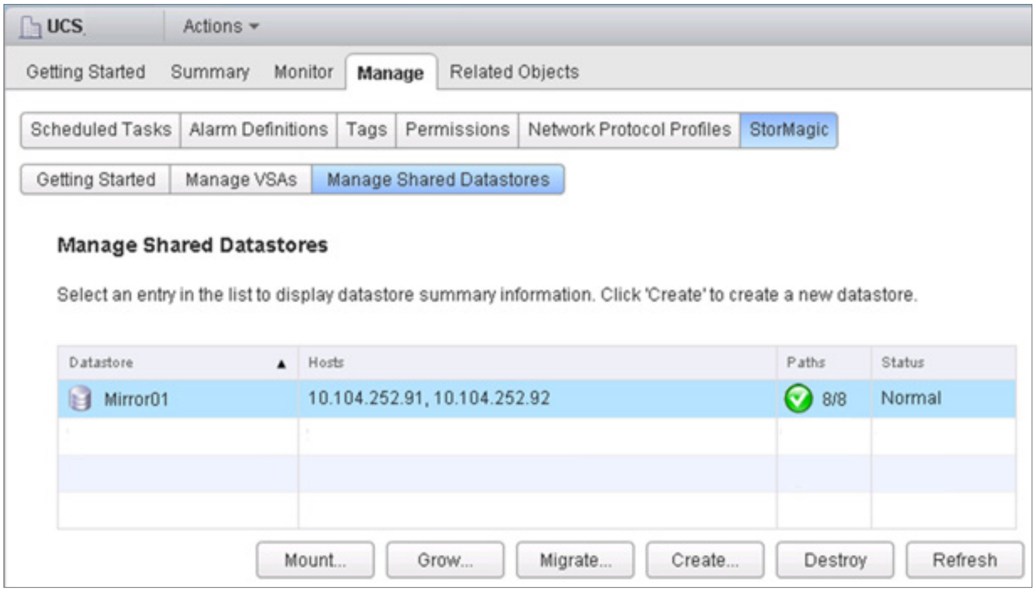

Manage Shared Datastores

Select an entry in the list to display datastore summary information. Click 'Create' to create a new datastore.

Datastore ▲	Hosts	Paths	Status
Mirror01	10.104.252.91, 10.104.252.92	✓ 8/8	Normal

2. The Manage Shared Datastores page displays information about all the VSA-hosted data stores, including the data-store name, the ESXi hosts using the data store, the number of paths to the data store, and the data-store status. Select a data store to view additional information about it:
   - Used and free capacity
   - Target details, including the iSCSI qualified name (IQN) and the 64-bit extended unique identifier (EUI-64); ESXi uses the EUI-64 to generate the name it used for the device
   - Target status, including mirror status if the data store is a mirror
   - Neutral storage host, if the data store is a mirror and an NSH is being used
   - Active paths between the ESXi hosts and the target. Click **Overview** to see a diagram of the paths between an ESXi host and a target
   - Multipathing policy being used

Set a Path Selection Policy

For each storage device, the ESXi host sets the path selection policy on the basis of the claim rules. By default, VMware supports the following path selection policies.

- Fixed
- Most recently used
- Round-robin

Configure a Fixed Local iSCSI Path Policy

For best performance in the SvSAN and Cisco UCS C-Series configuration, you should use a fixed path policy for any StorMagic iSCSI mirrors. The preferred path should be configured to use either of the network paths pointing the host's local VSA. Figure 43 shows a sample configuration.

Figure 43. Sample Network Configuration.

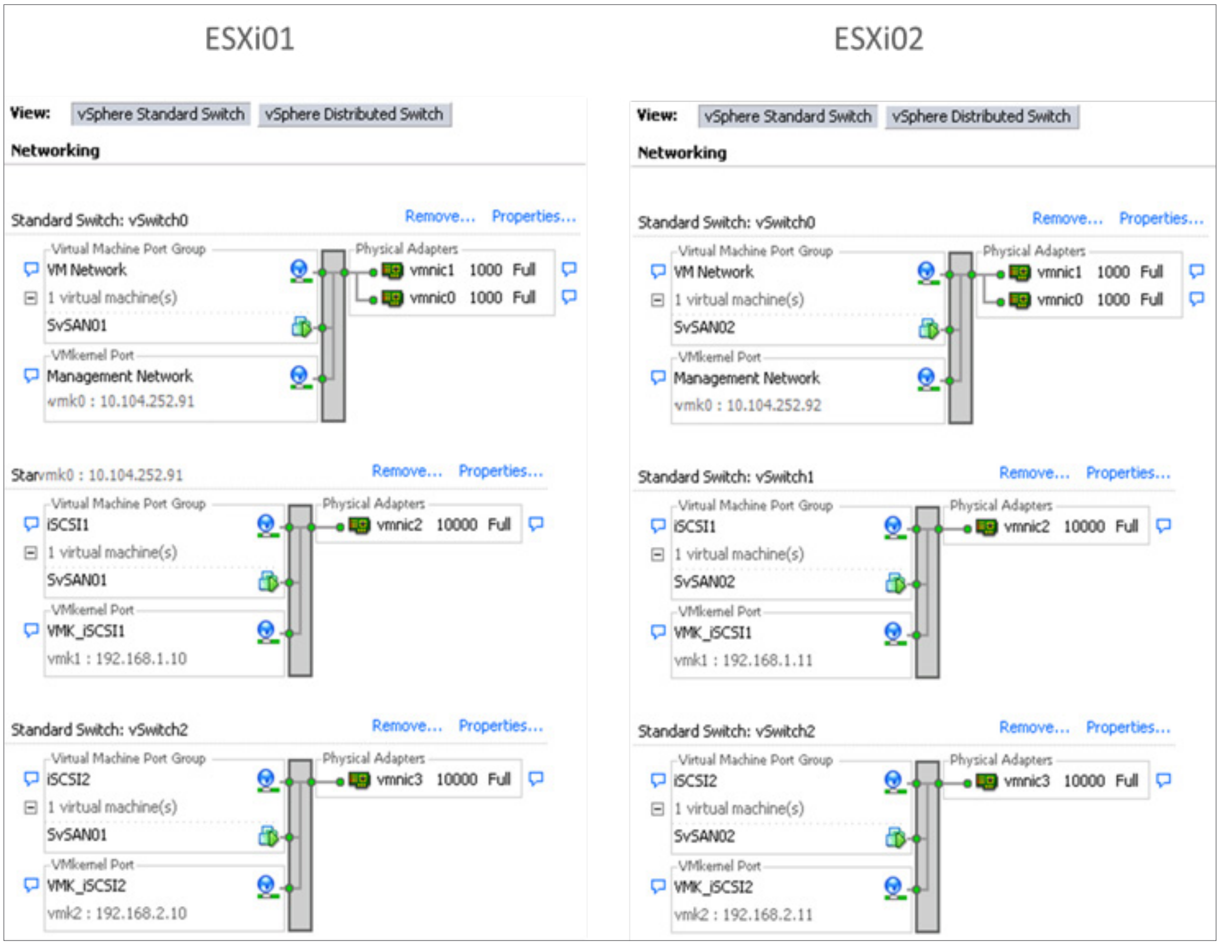

ESXi01

View: vSphere Standard Switch | vSphere Distributed Switch

Networking

Standard Switch: vSwitch0 — Remove... Properties...

Virtual Machine Port Group
VM Network
1 virtual machine(s)
SvSAN01

VMkernel Port
Management Network
vmk0 : 10.104.252.91

Physical Adapters
vmnic1 1000 Full
vmnic0 1000 Full

Starvmk0 : 10.104.252.91 — Remove... Properties...

Virtual Machine Port Group
iSCSI1
1 virtual machine(s)
SvSAN01

VMkernel Port
VMK_iSCSI1
vmk1 : 192.168.1.10

Physical Adapters
vmnic2 10000 Full

Standard Switch: vSwitch2 — Remove... Properties...

Virtual Machine Port Group
iSCSI2
1 virtual machine(s)
SvSAN01

VMkernel Port
VMK_iSCSI2
vmk2 : 192.168.2.10

Physical Adapters
vmnic3 10000 Full

ESXi02

View: vSphere Standard Switch | vSphere Distributed Switch

Networking

Standard Switch: vSwitch0 — Remove... Properties...

Virtual Machine Port Group
VM Network
1 virtual machine(s)
SvSAN02

VMkernel Port
Management Network
vmk0 : 10.104.252.92

Physical Adapters
vmnic1 1000 Full
vmnic0 1000 Full

Standard Switch: vSwitch1 — Remove... Properties...

Virtual Machine Port Group
iSCSI1
1 virtual machine(s)
SvSAN02

VMkernel Port
VMK_iSCSI1
vmk1 : 192.168.1.11

Physical Adapters
vmnic2 10000 Full

Standard Switch: vSwitch2 — Remove... Properties...

Virtual Machine Port Group
iSCSI2
1 virtual machine(s)
SvSAN02

VMkernel Port
VMK_iSCSI2
vmk2 : 192.168.2.11

Physical Adapters
vmnic3 10000 Full

Table 6 shows the VSA iSCSI IP addresses for the sample configuration

Table 6. VSA iSCSI IP Addresses for the Example in Figure 43

	vSwitch1	vSwitch2
VSA1 IP addresses	192.168.1.50	192.168.2.50
VSA2 IP addresses	192.168.1.51	192.168.2.51

VSA1's IP addresses are local to ESXi01. Therefore, any read or write I/O from the software iSCSI initiator will remain within the host's virtual network when the appropriate preferred path is configured. VSA1 will service all read requests locally and will send write requests to VSA2 to help ensure data consistency across the mirror. The same is true for ESXi02 and VSA2, with VSA2 helping ensure that write requests are sent to VSA1 for data consistency.

The VSAs will automatically aggregate the two physical network links to provide greater throughput for write I/O between the VSAs, reducing the potential for the network to become a bottleneck. Virtual machines on each host will benefit from having a dedicated local array to service read I/O without any contention from the virtual machines on the second host (Figure 44).

Figure 44. Local iSCSI Path

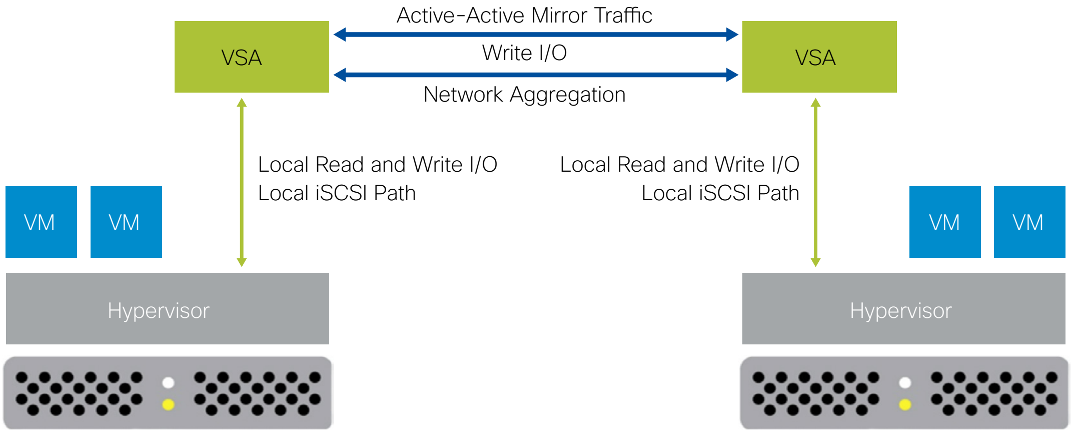

The fixed-path policy must be configured for each iSCSI volume on each host. From the vSphere Web Client, follow these steps:

1. Select the first host in the vSphere inventory.
2. Choose **Manage > Storage**.
3. Select the StorMagic iSCSI disk
4. In the lower window in the **Properties** pane, click **Edit Multipathing** (Figure 45)

Figure 45. Editing Path Properties

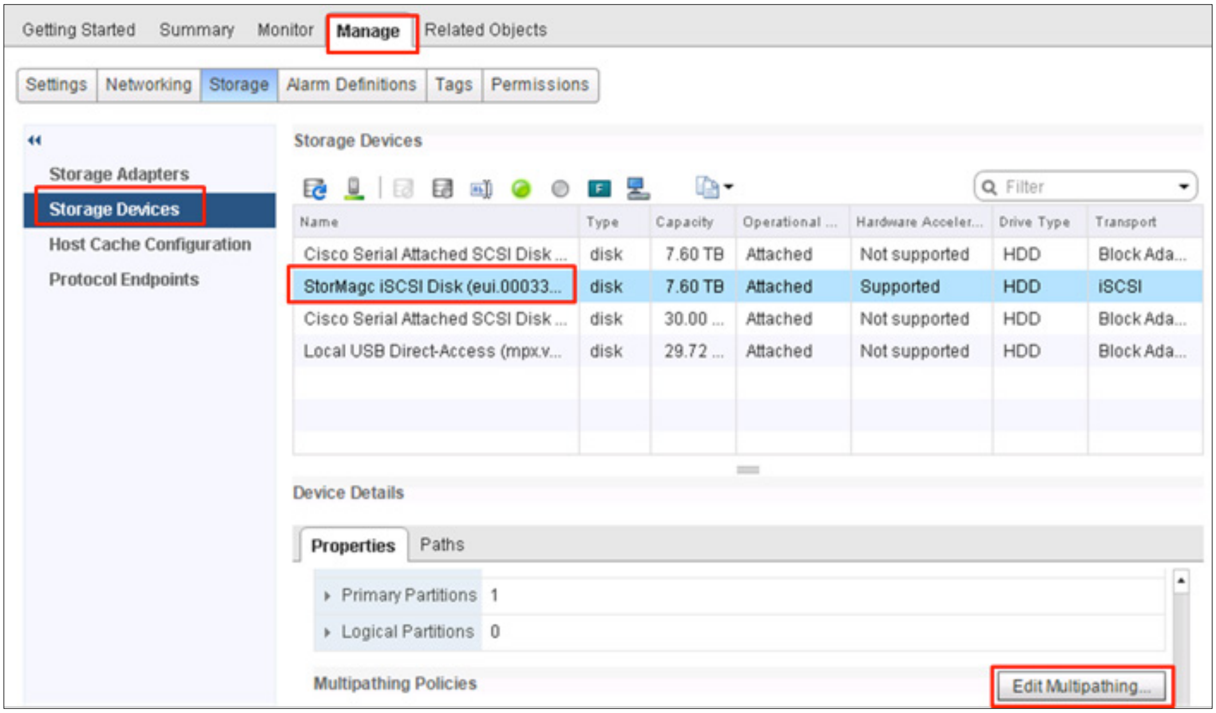

Storage Devices

Name	Type	Capacity	Operational ...	Hardware Acceler...	Drive Type	Transport
Cisco Serial Attached SCSI Disk ...	disk	7.60 TB	Attached	Not supported	HDD	Block Ada...
StorMagc iSCSI Disk (eui.00033...	disk	7.60 TB	Attached	Supported	HDD	iSCSI
Cisco Serial Attached SCSI Disk ...	disk	30.00 ...	Attached	Not supported	HDD	Block Ada...
Local USB Direct-Access (mpx.v...	disk	29.72 ...	Attached	Not supported	HDD	Block Ada...

Device Details

Primary Partitions	1
Logical Partitions	0

Multipathing Policies

5. In the drop down list, select **Fixed (VMware)**.
6. Select the path that points to the local VSA of the host (Figures 46 and 47) and click **OK**.

Figure 46. Selecting the Path to the Local VSA (ESXi01)

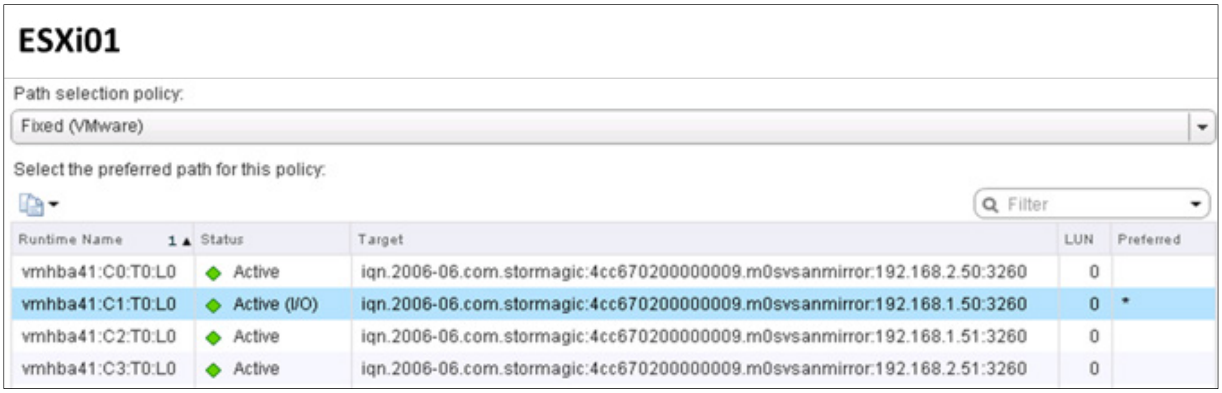

ESXi01

Path selection policy:

Fixed (VMware)

Select the preferred path for this policy:

Runtime Name 1 ▲	Status	Target	LUN	Preferred
vmhba41:C0:T0:L0	◆ Active	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.2.50:3260	0	
vmhba41:C1:T0:L0	◆ Active (I/O)	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.1.50:3260	0	*
vmhba41:C2:T0:L0	◆ Active	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.1.51:3260	0	
vmhba41:C3:T0:L0	◆ Active	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.2.51:3260	0	

Figure 47. Selecting the Path to the Local VSA (ESXi02)

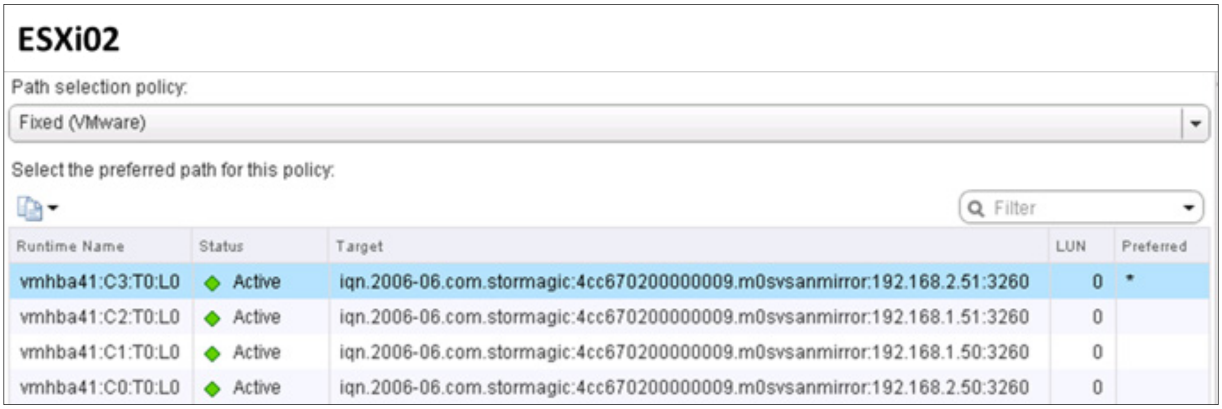

ESXi02

Path selection policy:

Fixed (VMware)

Select the preferred path for this policy:

Runtime Name	Status	Target	LUN	Preferred
vmhba41:C3:T0:L0	Active	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.2.51:3260	0	*
vmhba41:C2:T0:L0	Active	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.1.51:3260	0	
vmhba41:C1:T0:L0	Active	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.1.50:3260	0	
vmhba41:C0:T0:L0	Active	iqn.2006-06.com.stormagic:4cc670200000009.m0svsanmirror:192.168.2.50:3260	0	

In the sample IP addresses, ESXi01 can be set as either 192.168.1.50 or 192.168.2.50, where these IP addresses are the iSCSI and mirror interfaces on the local VSA (VSA1). If you do not know the IP addresses of the local VSA, select the VSA virtual machine from the vSphere inventory and choose **Summary > View All IP Addresses**.

7. Repeat the preceding steps for any other StorMagic iSCSI disks and other ESXi hosts to help ensure that the path points to the local VSA.

Configuring Cisco UCS C240 M4 Server and StorMagic SvSAN Expansion

The Cisco UCS C240 M4 server is designed for both performance and expandability over a wide range of storage-intensive infrastructure workloads, from big data to collaboration.

The Cisco UCS C240 M4 supports up to 24 SFF drives. The SvSAN and Cisco UCS C240 M4 architecture supports two disk expansion packs. Each disk expansion pack has eight 1.2-TB SAS HDDs and is configured with RAID 5. Figure 48 shows the SvSAN and Cisco UCS C240 M4 Server storage architecture with a single disk expansion pack.

Figure 48. StorMagic SvSAN and Cisco UCS C240 M4 Server Storage Architecture with One Disk Expansion Pack

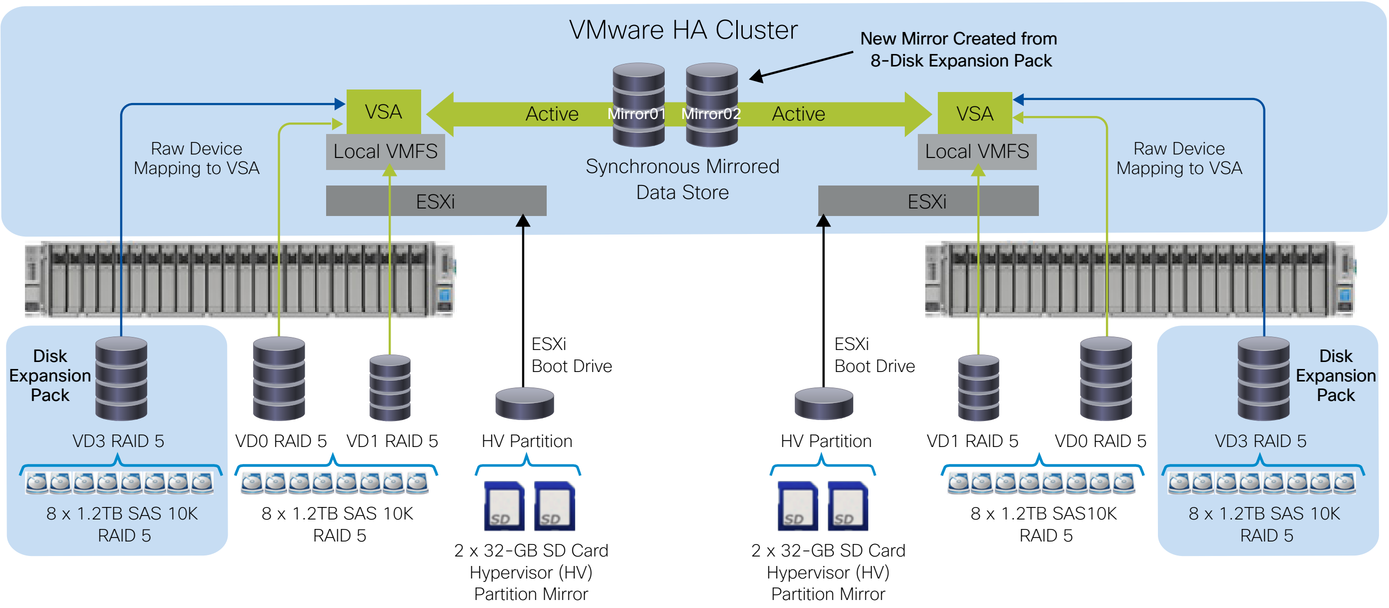

Configure RAID 5

Create a single RAID 5 drive group with one virtual drive is created using eight 1.2-TB SAS drives.

1. In the Integrated Management Controller, navigate to **Storage** and select **Cisco 12G SAS Modular Raid Controller**.
2. Select **Create Virtual Drives from Unused Physical Drives** and create a virtual disk using all the available space in the RAID 5 drive group (Figure 49).

Figure 49. Creating a Virtual Disk

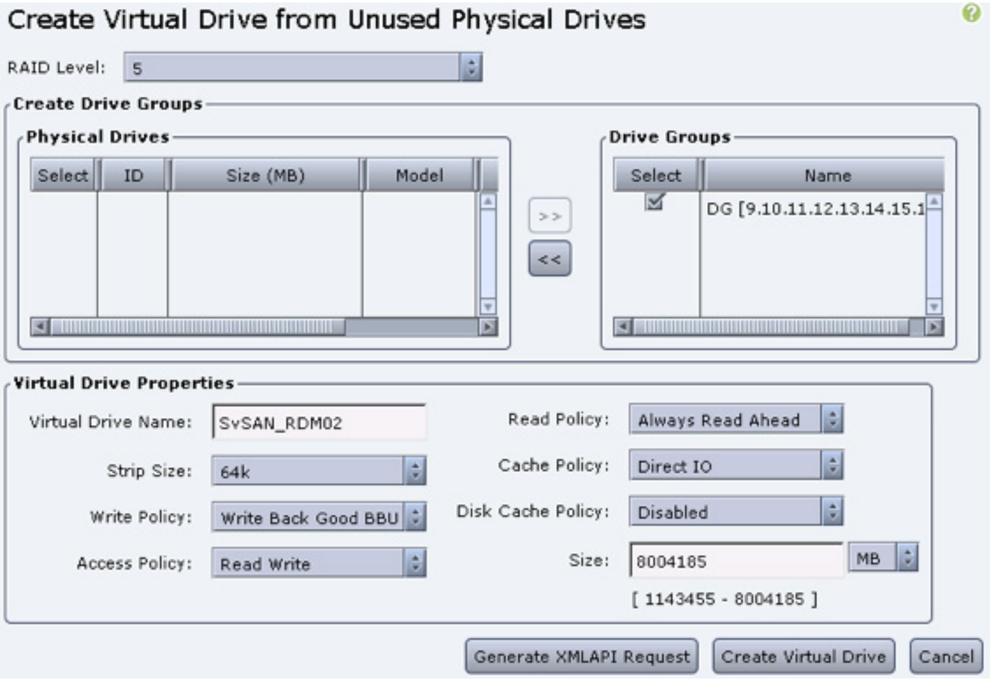

Create Virtual Drive from Unused Physical Drives

RAID Level: 5

Create Drive Groups

Physical Drives

Select	ID	Size (MB)	Model

Drive Groups

Select	Name
☑	DG [9.10.11.12.13.14.15.1

Virtual Drive Properties

Virtual Drive Name: SvSAN_RDM02

Strip Size: 64k

Write Policy: Write Back Good BBU

Access Policy: Read Write

Read Policy: Always Read Ahead

Cache Policy: Direct IO

Disk Cache Policy: Disabled

Size: 8004185 MB

[1143455 - 8004185]

3. Repeat the preceding steps if you have Disk Expansion Pack 2.

Allocate Additional Storage to a VSA for Mirroring

Raw device mapping, or RDM, provides a mechanism that gives a virtual machine direct access to a LUN on the physical storage subsystem.

When you map a LUN to a VMFS volume, vCenter Server creates an RDM file that points to the raw LUN. Encapsulating disk information in a file allows vCenter Server to lock the LUN so that only one virtual machine can write to it at a time.

If you use a RAID controller, a RAID array with one or more RAID LUNs or volumes will be created, using the management utility supplied with the RAID controller. Each RAID LUN can then be assigned to a local VSA using a raw device mapping. The RDM gives the VSA direct access to the device, without the overhead of additional mapping entailed when you use a VMDK. In some cases, you may be able to create the RDM using vSphere Web Client; in other cases, you may need to use the ESXi console commands to create the RDM, which can then be assigned to the SvSAN as an existing disk drive.

For information about managing the RAID controller, consult the documentation supplied with it. For more information about RDMs, consult the vSphere documentation at http://www.vmware.com/pdf/vmfs-best-practices-wp.pdf.

The Cisco UCS C240 M4 supports up to 24 SFF drives. The SvSAN and Cisco UCS C240 M4 architecture supports two disk expansion packs, and each disk expansion pack has eight 1.2-TB SAS HDDs and is configured with RAID 5.

Identify Physical Devices (RAID LUNs) on the System

To identify physical devices on the system, use either the ESXi host console or an SSH client to connect to the host. SSH is not enabled by default in ESXi. You can enable SSH on each host by clicking the **Configuration** tab and choosing **Security Profile**.

To use the ESXi console to list and identify the physical devices (RAID LUNs), enter the **ls** command at the command line:

```
{root@ESXi01 ~]# ls –lh /vmfs/devices/disks/
```

Figure 50 shows an example of the output.

Figure 50. Sample Output

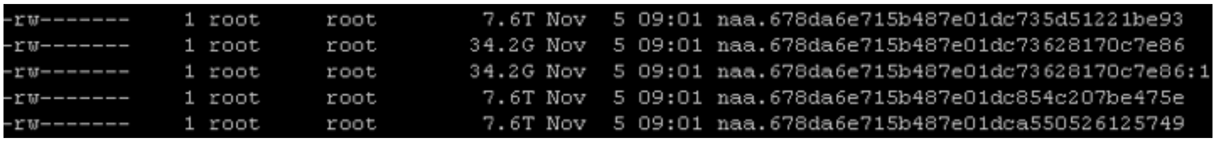

```
-rw-------    1 root     root         7.6T Nov  5 09:01 naa.678da6e715b487e01dc735d51221be93
-rw-------    1 root     root        34.2G Nov  5 09:01 naa.678da6e715b487e01dc73628170c7e86
-rw-------    1 root     root        34.2G Nov  5 09:01 naa.678da6e715b487e01dc73628170c7e86:1
-rw-------    1 root     root         7.6T Nov  5 09:01 naa.678da6e715b487e01dc854c207be475e
-rw-------    1 root     root         7.6T Nov  5 09:01 naa.678da6e715b487e01dca550526125749
```

To use vSphere Web Client to list and identify the physical devices (RAID LUNs), follow these steps:

1. Select the ESXi host and click **Manage**.
2. Select **Storage** and click **Storage Devices**. You will see the device names (Figure 51).

Figure 51. Using VMware vSphere Web Client to List Physical Devices

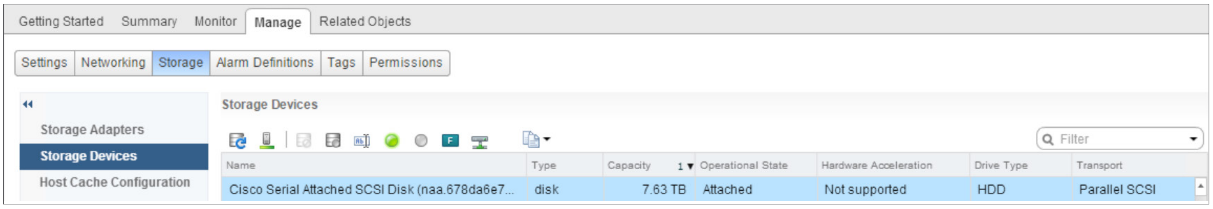

Name	Type	Capacity	Operational State	Hardware Acceleration	Drive Type	Transport
Cisco Serial Attached SCSI Disk (naa.678da6e7...	disk	7.63 TB	Attached	Not supported	HDD	Parallel SCSI

Add an RDM Disk to a StorMagic VSA

Create a mapping file that resides on a VMFS data store and points to the LUN. Although the mapping file has the same .vmdk extension as a regular virtual disk file, the mapping file contains only mapping information. The virtual disk data is stored directly on the LUN. The mapping file to an RDM is created by using the **vmkfstools** command:

```
{root@esxi1~]# vmkfstools -z /vmfs/devices/disks/naa.device_id /vmfs/volumes/SvSAN_LocalDisk/SvSAN-
C240M4-01/rawdisk-naa.device_id.vmdk
```

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click **Add** and select **Hard Disk**; then click **Next**.
3. In the **Select a Disk** window, select **Use an Existing Virtual Disk**.
4. In the **Select Existing Disk** window, click **Browse** and select the RDM that was created using the **vmkfstools** command (Figure 52).

Figure 52. Selecting the RDM

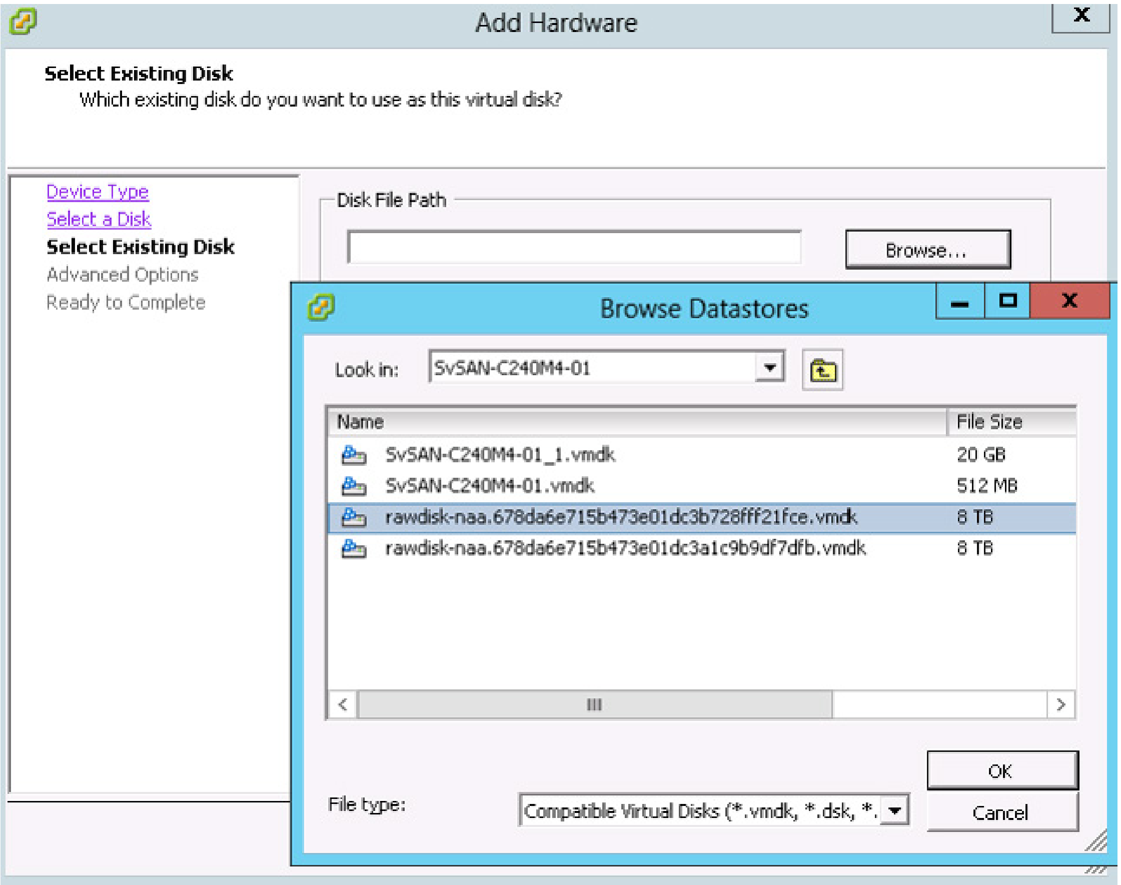

Select Existing Disk

Which existing disk do you want to use as this virtual disk?

Look in: SvSAN-C240M4-01

Name	File Size
SvSAN-C240M4-01_1.vmdk	20 GB
SvSAN-C240M4-01.vmdk	512 MB
rawdisk-naa.678da6e715b473e01dc3b728fff21fce.vmdk	8 TB
rawdisk-naa.678da6e715b473e01dc3a1c9b9df7dfb.vmdk	8 TB

File type: Compatible Virtual Disks (*.vmdk, *.dsk, *.

5. In the **Advanced Options** window, for **Virtual Device Node**, choose **SCSI1:0**. In the **Mode** section, select **Independent** and **Persistent** (Figure 53).

Figure 53. Setting Advanced Options

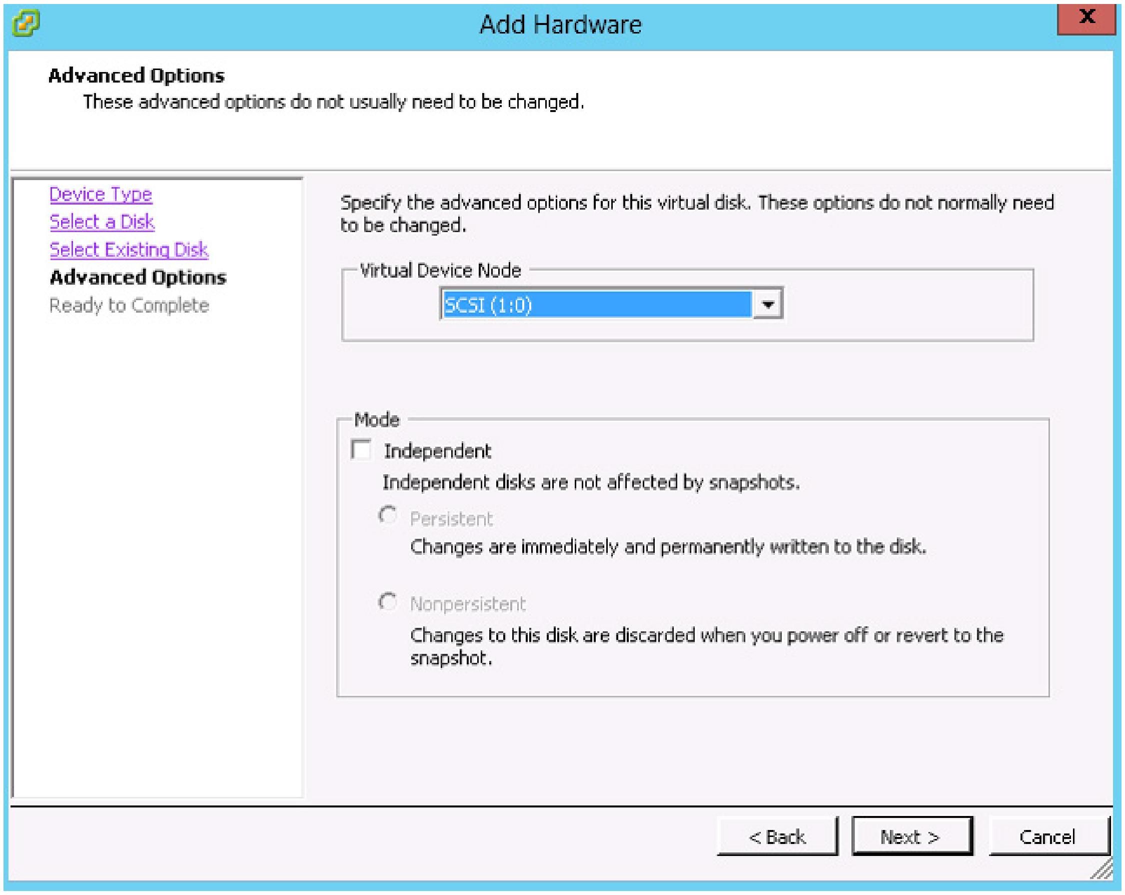

Add Hardware

Advanced Options

These advanced options do not usually need to be changed.

Specify the advanced options for this virtual disk. These options do not normally need to be changed.

Virtual Device Node: SCSI (1:0)

Mode

□ Independent

Independent disks are not affected by snapshots.

○ Persistent

Changes are immediately and permanently written to the disk.

○ Nonpersistent

Changes to this disk are discarded when you power off or revert to the snapshot.

6. Complete the wizard to create the virtual disk.
7. Repeat the preceding steps if you have multiple RDM disks.

Create Pools and Shared Data Stores with Mirrored Storage

You can add pools and created shared data stores with mirrored storage.

Creating a Pool

Normally a VSA has a single pool, created automatically when the VSA is deployed, but you can add pools at any time.

You can add pools without restarting the VSA (for ESXi hypervisors only). To do this, add the storage to the VSA virtual machine, use the Rescan Subsystems action on the Storage Pools page, and then use the Create Pool action to create a pool using the new storage.

To create a pool, follow these steps:

1. From vSphere Client, select the StorMagic VSA, or log in to the VSA from a web browser and click the **StorMagic** tab.
2. Select **Pools** and click **Create Pool**.
3. In the **Create Storage Pool** window, enter the pool name.
4. Select the subsystem to use. Unless you are using a RAID controller, the only subsystem is **Direct Device Mapping**.
5. For the RAID level, select **JBOD** for the pool.
6. Select the physical disk devices to use for the pool.
7. Optionally, specify the performance level. This setting indicates to you or other users of the system the performance (for example, the speed) of the storage device. Select a number from 1 to 10. A higher number indicates better performance. Pools created with versions of SvSAN prior to Release 5.2 are given the value 0 (unclassified).
8. Optionally, specify any other information about the device in the **Category** field. Then click **Create** (Figure 54).

Figure 54. Creating a Storage Pool

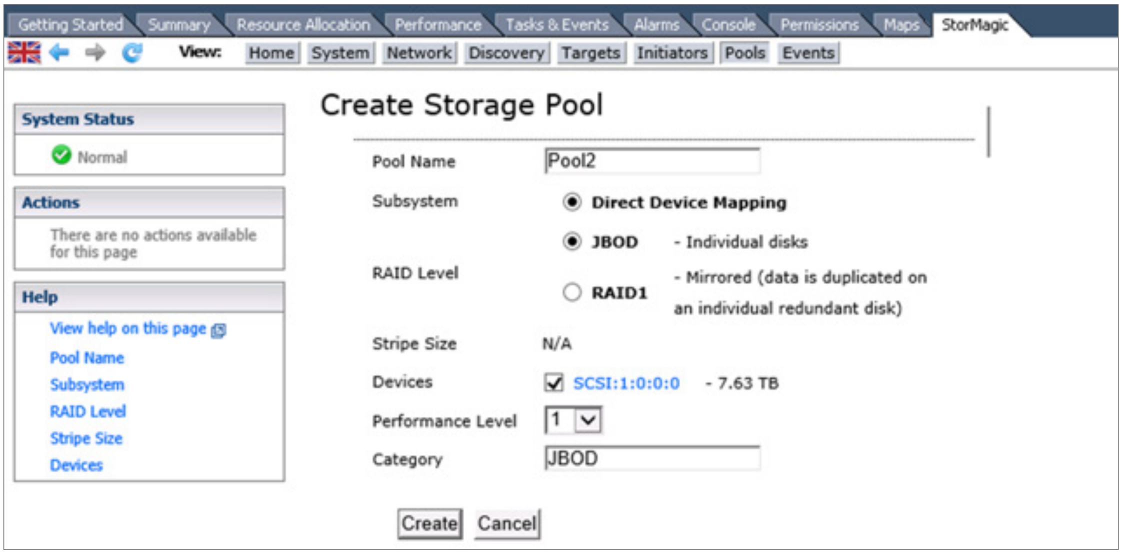

Create Storage Pool

Pool Name: Pool2

Subsystem: ● Direct Device Mapping

RAID Level: ● JBOD - Individual disks
○ RAID1 - Mirrored (data is duplicated on an individual redundant disk)

Stripe Size: N/A

Devices: ☑ SCSI:1:0:0:0 - 7.63 TB

Performance Level: 1

Category: JBOD

Create Cancel

9. If you have multiple RDM disks, repeat the preceding steps to create a pool.
10. Repeat the preceding steps to create the pool on the second StorMagic VSA.

Creating a Shared Data Store with Mirrored Storage

1. In vSphere Web Client, select your data center and then choose **Manage > StorMagic**. The plug-in opens (Figure 55).

Figure 55. Opening the Plug-in

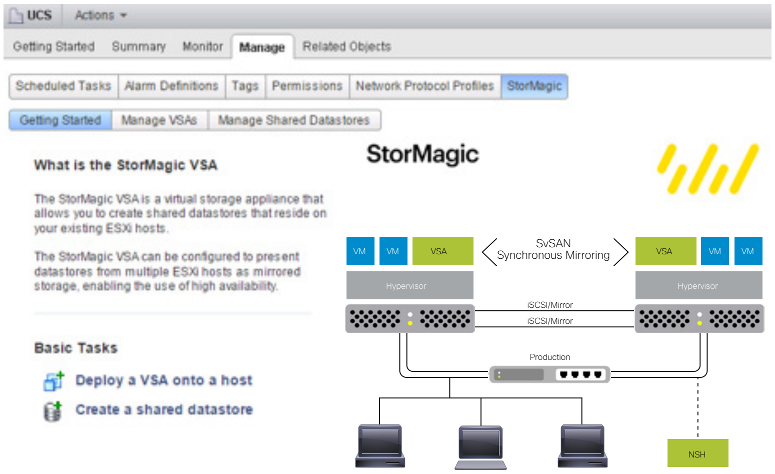

What is the StorMagic VSA

The StorMagic VSA is a virtual storage appliance that allows you to create shared datastores that reside on your existing ESXi hosts.

The StorMagic VSA can be configured to present datastores from multiple ESXi hosts as mirrored storage, enabling the use of high availability.

Basic Tasks

Deploy a VSA onto a host

Create a shared datastore

2. Click **Create a shared datastore**. The Create Datastore wizard opens. Click **Next**.
3. Enter the data-store name.
4. Set the provisioned size. To set the size so that all the available space is used, select **Use all**.
5. To create mirrored storage, select two VSAs.
6. By default, the pool to be used on each VSA is selected automatically (Figure 56). To specify the pool, select the **Advanced options** checkbox.

Figure 56. Selecting the Pool Automatically

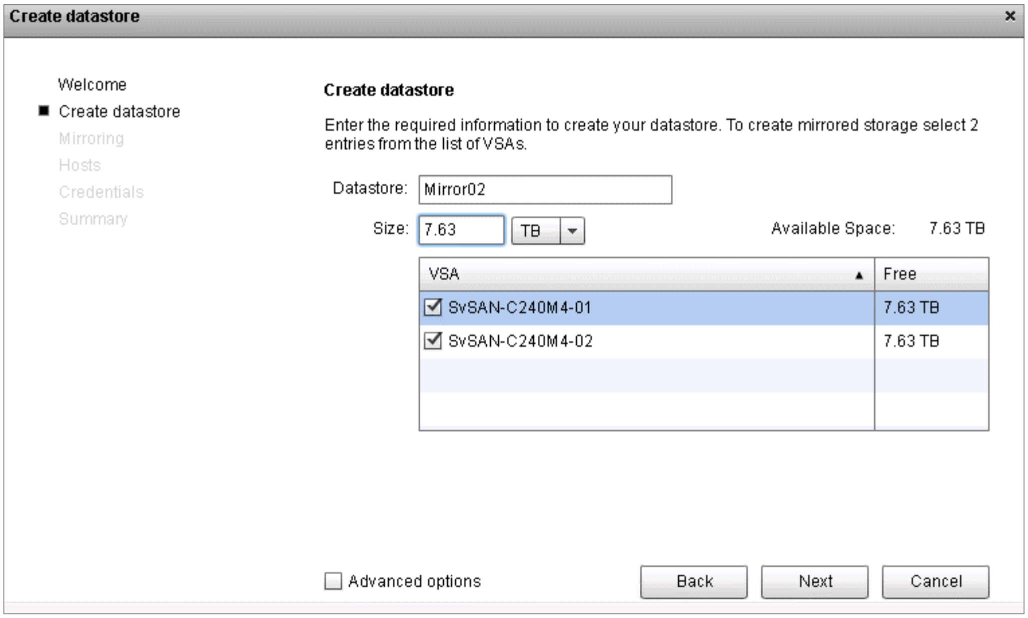

Create datastore

Enter the required information to create your datastore. To create mirrored storage select 2 entries from the list of VSAs.

Datastore: Mirror02

Size: 7.63 TB

Available Space: 7.63 TB

VSA ▲	Free
☑ SvSAN-C240M4-01	7.63 TB
☑ SvSAN-C240M4-02	7.63 TB

☐ Advanced options

7. Click **Next**.
8. Select the neutral storage host for mirroring (Figure 57). The **Advanced** option lets you select the **Up** isolation policy, which does not require an NSH. Click **Next**.

Figure 57. Selecting the NSH for Mirroring

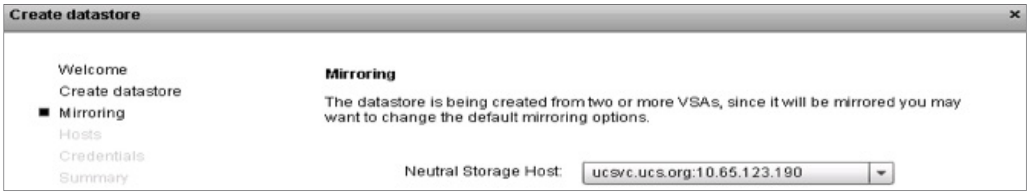

Mirroring

The datastore is being created from two or more VSAs, since it will be mirrored you may want to change the default mirroring options.

Neutral Storage Host: ucsvc.ucs.org:10.65.123.190

9. Optionally, enable caching on the data store. This option is available only if your license includes caching and if a cache device was enabled when the VSA was deployed. If you enable caching, default cache settings are used; these can be modified at a later time using the WebGUI. Click **Next**.
10. Select the hosts that are to have access to the SvSAN data store (Figure 58). Click **Next**.

Figure 58. Selecting the Hosts That Can Access the Data Store

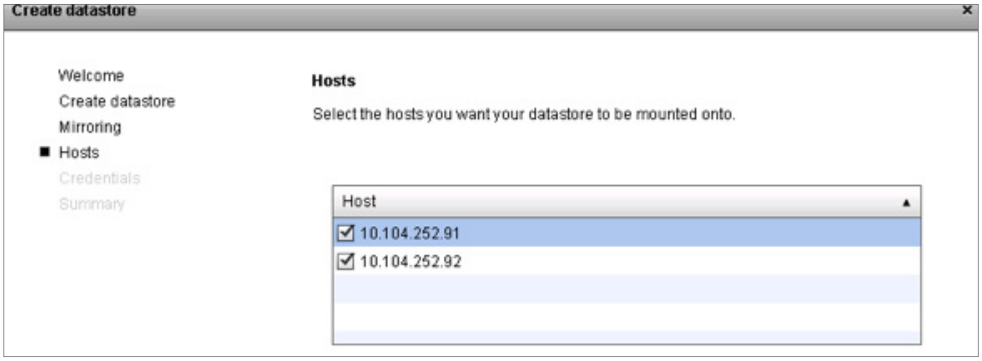

Hosts

Select the hosts you want your datastore to be mounted onto.

Host
☑ 10.104.252.91
☑ 10.104.252.92

11. The first time you create a data store, the VSAs must authenticate with their hosts if they were not already saved during deployment. For each host, provide the ESXi host administrative password (the password that was used when ESXi was installed). Click **Next**.
12. Click **Finish** to start the data-store creation process. You can monitor the data-store creation process in the VMware Recent Tasks window. After this task has completed, the data store is available on the ESXi hosts (Figure 59).
13. When a mirror is created, full synchronization is performed to help ensure that both sides of the mirror are synchronized.
14. Repeat the preceding steps if you have multiple storage pools.

Figure 59. Shared Data Stores Ready for Use

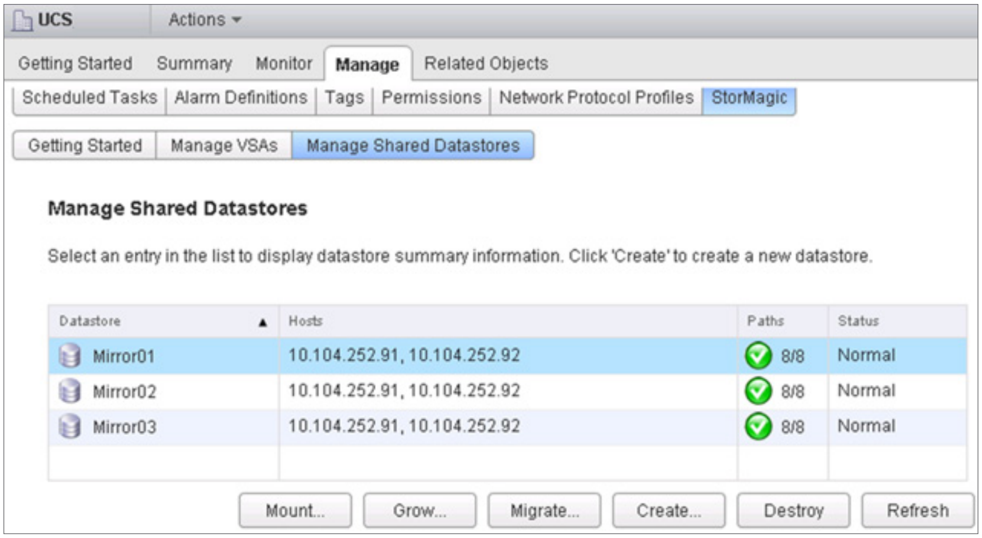

Manage Shared Datastores

Select an entry in the list to display datastore summary information. Click 'Create' to create a new datastore.

Datastore ▲	Hosts	Paths	Status
Mirror01	10.104.252.91, 10.104.252.92	✓ 8/8	Normal
Mirror02	10.104.252.91, 10.104.252.92	✓ 8/8	Normal
Mirror03	10.104.252.91, 10.104.252.92	✓ 8/8	Normal

Note: Use of a single mirrored volume per LUN and pool managed by the VSA is recommended.

Performance Characterization

Table 7 shows the workloads suited to run on the Cisco UCS C240 M4 and StorMagic solution. Each virtual machine running on the system can exercise the I/O patterns shown in the table. The table does not list all workloads that this solution can run, but the workload characterization provides a baseline for further estimation and sizing.

Table 7. Application I/O Profile: Cisco UCS C240 M4 and StorMagic Solution

Application Profile	RAID Type	Access Mode	Read:Write Ratio	Block Size	Metric
Online transaction processing (OLTP)	5	Random	80:20	8 KB	I/O operations per second (IOPS) and response time (milliseconds [ms]) Decision support system, business intelligence, and video on demand (VoD)
Decision support system, business intelligence, and video on demand (VoD)	5	Sequential	100:0 and 0:100	256/512 KB	Transfer rate (MBps)

For the Cisco UCS C240 M4 and StorMagic solution, each server is configured with twenty-four 1.2-TB 10,000-rpm disks. Each disk provides about 140 to 150 IOPS, with a response time of less than 20 milliseconds. The large disk size results in greater variation in the response times, because of the mechanical head movement between the inner and outer cylinders. Performance varies based on the size of the volume that is in use. For the test validation purposes here, three mirrored volumes are created, each with eight 1.2-TB disks. Each virtual machine is allocated with 100 GB of disk space on the mirrored volume.

Workload Scenarios

The following three scenarios were tested and validated on the Cisco UCS C240 M4 and StorMagic solution:

- Scenario A, with 10 virtual machines on mirrored volume 1: This scenario helps estimate the achievable throughput and response time under an average system load. Ten virtual machines are running on two servers, with each host carrying five virtual machines loaded on one mirrored volume. This scenario mimics a typical deployment use case for a Cisco UCS C240 M4 and StorMagic solution with eight disks.
- Scenario B, with 10 virtual machines on mirrored volume 1, and 10 virtual machines on mirrored volume 2: This scenario helps estimate the achievable throughput and response time under a system with a medium-size load. Twenty virtual machines are running on two servers, with each host carrying 10 virtual machines loaded on two mirror volumes. This scenario mimics a typical deployment use case for a Cisco UCS C240 M4 and StorMagic solution with 16 disks.
- Scenario C, with 10 virtual machines on mirrored volume 1, 10 virtual machines on mirrored volume 2, and 10 virtual machines on mirrored volume 3: This scenario helps estimate the achievable throughput and response time on a system with a heavy load. Thirty virtual machines are running on two servers, with each host carrying 15 virtual machines loaded on three mirrored volumes. This scenario mimics a typical deployment use case for a Cisco UCS C240 M4 and StorMagic solution with 24 disks.

All the scenarios used the I/O profiles shown in Table 7. The results are discussed in the following sections. A 4-KB block size was tested. The results matched those for an 8-KB block size; hence, the graphs show an 8-KB block size.

Scenario: Ten Virtual Machines per Mirrored Volume (30 Virtual Machines on 3 Mirrored Volumes and 15 Virtual Machines on Each Host)

Note: All the graphs in this section show the results for 10, 20, and 30 virtual machines.

With 8-KB 100 percent random read I/O and a queue depth of 4 on 10, 20, and 30 virtual machines, the system is capable of delivering up to about 9900 IOPS within a response time of 13 milliseconds. IOPS shows a linear increase as the number of virtual machines running I/O processing is increased (Figure 60).

Figure 60. 8-KB Random I/O: 100 Percent Read and Queue Depth 4

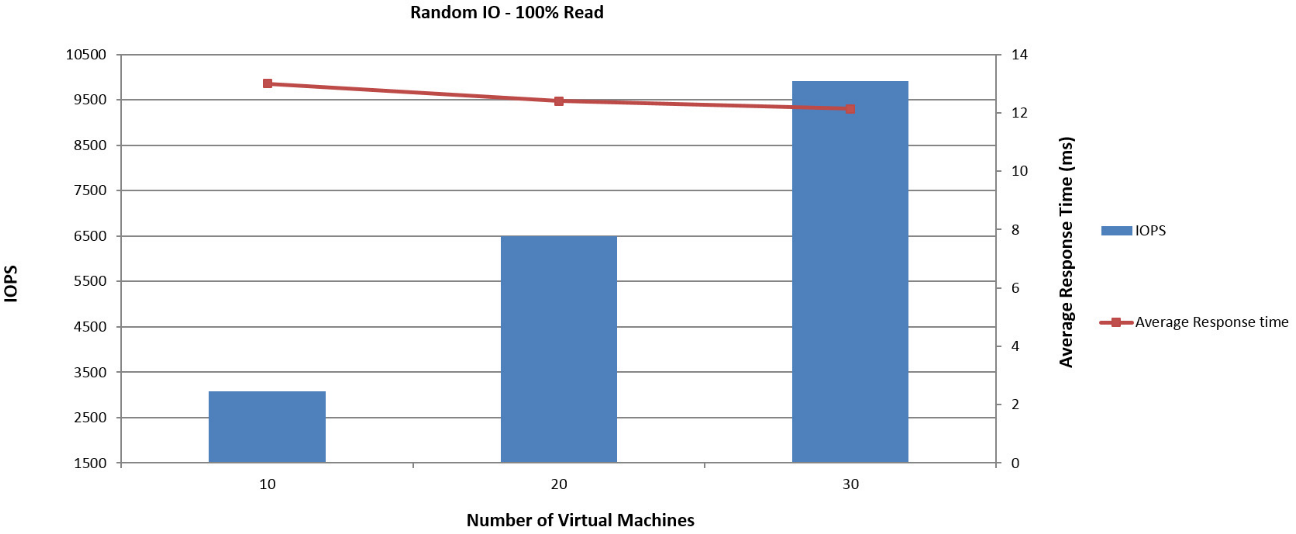

With 8-KB random read and write I/O (80 percent read and 20 percent write) and a queue depth of 4 on 10, 20, and 30 virtual machines, the system is capable of delivering up to about 5500 IOPS with a response time of 22 milliseconds. Because the workload is 20 percent write operations, the response time increases slightly over that for a workload with 100 percent read operations (Figure 61).

Figure 61. 8-KB Random I/O: 80 Percent Read and 20 Percent Write and Queue Depth 4

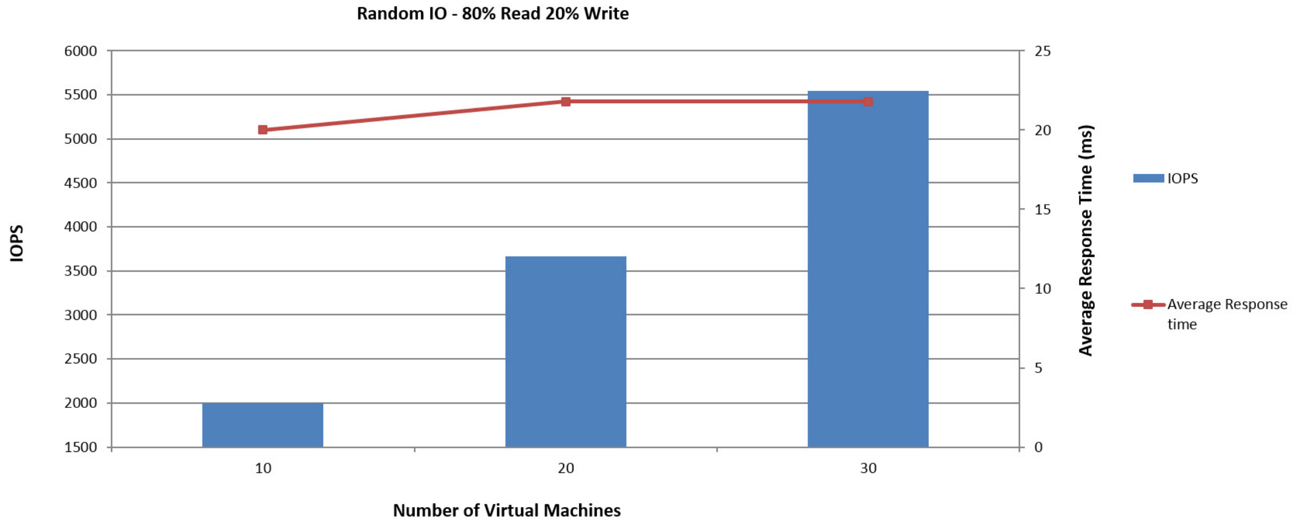

The sequential read and write workloads stay sequential when only one virtual machine is present in the mirror. However, their I/O changes to random at the disk level when more virtual machines are present, although sequential I/O occurs at the individual virtual machine level. Hence, the recommended approach is to use time scheduling among the virtual machine workloads when heavy sequential I/O is expected.

With 256-KB 100 percent sequential read I/O on 10, 20, and 30 virtual machines, the system is capable of delivering bandwidth of up to about 3100 MBps within a response time of 2.5 milliseconds. The bandwidth increases linearly as the number of virtual machines running I/O is increased (Figure 62).

Figure 62. 256-KB Sequential I/O: 100 Percent Read and Queue Depth 1

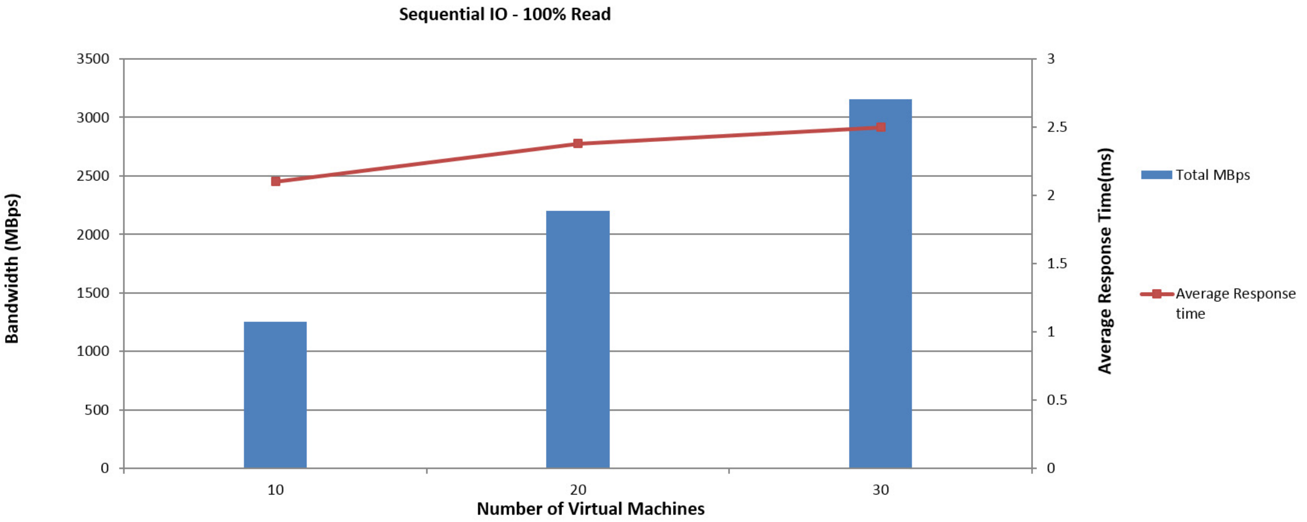

With 256-KB 100 percent sequential write I/O on 1 and 10 virtual machines, the system is capable of delivering bandwidth of up to about 816 MBps within a response time of 3.5 milliseconds (Figure 63). This test was run only on 1 and 10 virtual machines because it uses the available RAID cache to the maximum.

Figure 63. 256-KB Sequential I/O: 100 Percent Write and Queue Depth 1

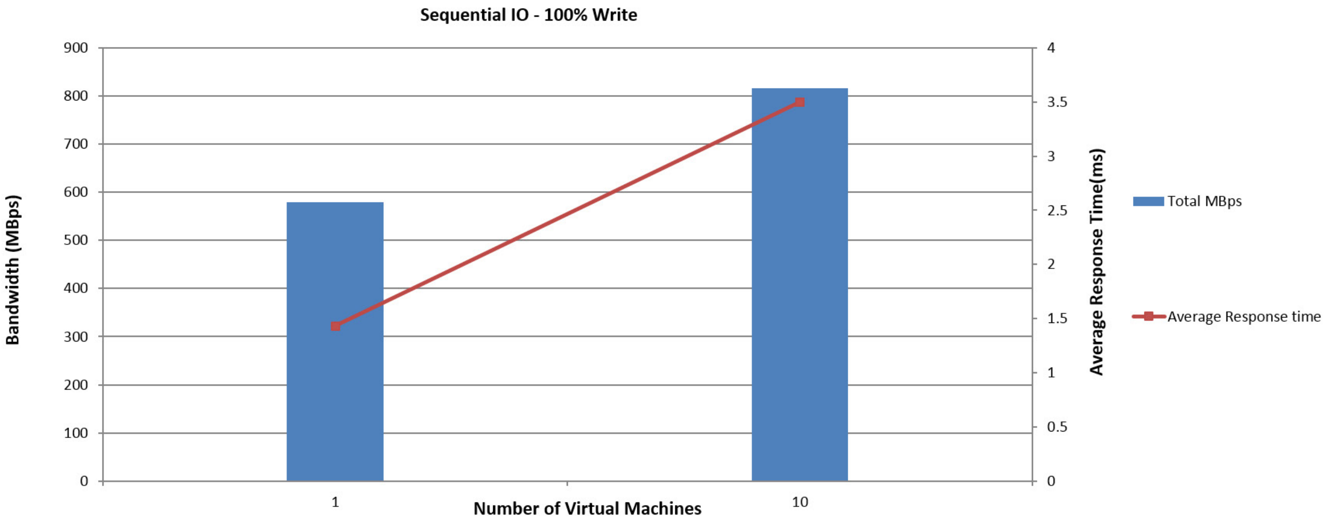

Collecting Diagnostic Information for VMware ESXi and StorMagic SvSAN

Collect Diagnostic Information for VMware ESXi

VMware technical support staff routinely requests diagnostic information from you when a support request is addressed. This diagnostic information contains product-specific logs and configuration files from the host on which the product is running. This information is gathered using a specific script or tool within the product.

This section provides procedures for obtaining diagnostic information for an ESXi or ESX host using the vm-support command-line utility.

1. Open a console to the ESXi host.
2. Run the command **vm-support**.

Note: You can specify additional options to customize the log bundle collection. Run the **vm-support -h** command to see a list of the options available for a specific version of ESXi or ESX.

3. A compressed bundle of logs is produced and stored in a file with a .tgz extension. The file is stored in one of these locations:
   - /var/tmp/
   - /var/log/
   - The current working directory
4. After the log bundle has been collected and downloaded to a client, upload the logs to the SFTP or FTP site. For more information, see Uploading Diagnostic Information to VMware (1008525).

   Note: You can obtain diagnostic information from ESX and ESXi hosts using the vSphere or VMware Infrastructure (VI) Client. See the VMware knowledgebase article at http://kb.vmware.com/kb/653

Collect System Diagnostic Information for StorMagic SvSAN

StorMagic support staff may request a system diagnostics dump. The dump contains various logs that can be used to help diagnose the cause of a problem.

To obtain a system diagnostics dump, use these steps:

1. Log in to SvSAN and click **System**.
2. Select **System Diagnostics** and click **Click here to download** (Figure 64).

Figure 64. Downloading System Diagnostics Information

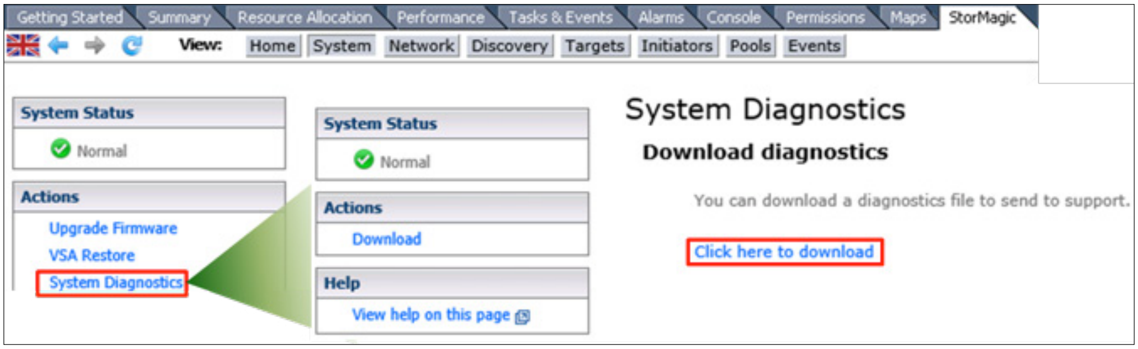

System Diagnostics

Download diagnostics

You can download a diagnostics file to send to support.

For More Information

- http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c240-m4-rack-server/index.html
- http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1227/index.html
- https://www.vmware.com/files/pdf/vsphere/VMW-WP-vSPHR-Whats-New-6-0.pdf
- http://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/unified-computing/whitepaper_C11-718938.pdf
- http://www.stormagic.com/manual/SvSAN_5-3/en/
- https://software.cisco.com/download/type.html?mdfid=286281345&flowid=71442

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**