

StorMagic SvSAN on VMware vSphere 6.0 with Cisco UCS Mini Deployment Guide

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.

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® Mini 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 an edge enterprise to deploy data storage infrastructure that supports its multiple-site environment.

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

Technology Overview

The following hardware and software components are used in the deployment described in this guide:

- Cisco UCS Mini
- Cisco UCS Manager Release 3.0(2c)
- Cisco UCS 6324 Fabric Interconnect
- Cisco UCS B200 M4 Blade Server
- Cisco UCS Virtual Interface Card (VIC) 1340
- VMware vSphere Release 6.0.0b
- StorMagic SvSAN Release 5.3

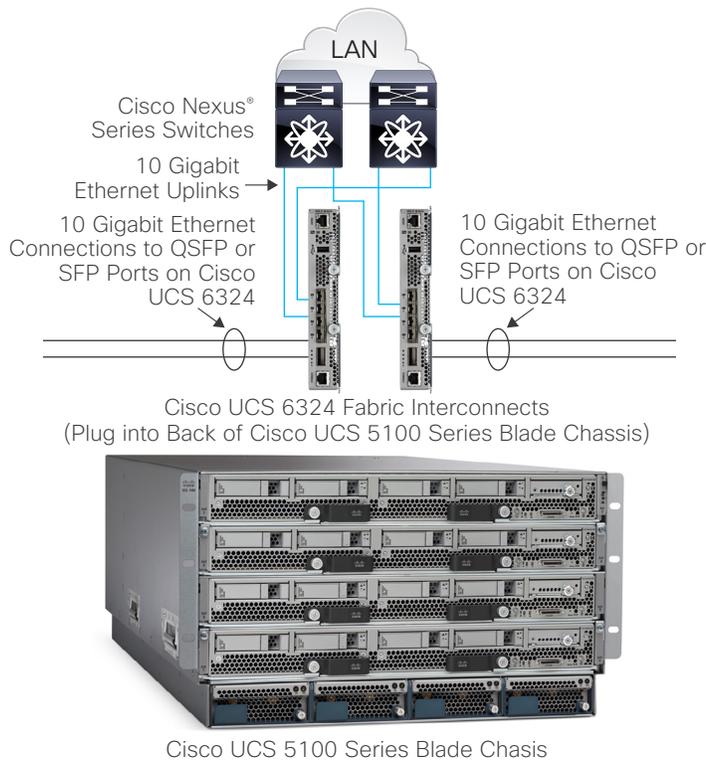
Table of Contents

Introduction	1
Technology Overview	1
Cisco UCS Mini	2
Cisco UCS Manager 3.0	2
Cisco UCS 6324 Fabric Interconnect	3
Cisco UCS B200 M4 Blade Server	3
Cisco UCS VIC 1340	3
VMware vSphere 6.0	3
StorMagic SvSAN	3
Computing	4
Networking	4
Storage	5
Deploying StorMagic SvSAN for VMware vSphere in a Cisco UCS Environment	8
Configuring Cisco UCS	9
Configure Cisco UCS Mini Fabric Interconnects	9
Configure Cisco UCS Manager	9
Synchronize Cisco UCS with NTP	9
Enable Quality of Service in Cisco UCS Fabric	9
Create QoS Policy	9
Enable Uplink Ports	10
Create VLANs	11
Create Local Disk Configuration Policy	12
Create Boot Policy	12
Create BIOS Policy	13
Configure Server Pool and Qualifying Policy	13
Create UUID Suffix Pool	14
Create MAC Address Pool	14
Create an IP Address Pool for KVM Access	15
Create vNIC Templates and LAN Connectivity Policies	15
Create a Service Profile Template	17
Create Service Profiles from Templates	18
Create Virtual Drives	19
Installing VMware ESXi 6.0	23
Prepare VMware ESXi for StorMagic SvSAN Deployment	23
Configure ESXi Host Networking	26
Install StorMagic SvSAN Components on VMware vCenter Server	28
Installing StorMagic SvSAN Software on VMware vCenter Server	30
Installing StorMagic PowerShell Toolkit	31
Deploying StorMagic SvSAN VSAs	32
Recommended StorMagic SvSAN Mirror Configurations	37
Overview	37
Mirrored Targets	37
Cisco UCS Blade Server Expansion and StorMagic SvSAN Target Migration	38
Data Store Spanning Using VMware Extents	41
Creating a Shared Data Store with Mirrored Storage	42
Managing VSAs	44
Configure Jumbo Frames on StorMagic SvSAN Network Interfaces	46
Managing Shared Data Stores	47
Performance Characterization	48
Workload Scenarios	48
Collecting Diagnostic Information for VMware ESXi and StorMagic SvSAN	53
Collect Diagnostic Information for VMware ESXi	53
Collect System Diagnostic Information for StorMagic SvSAN	53
For More Information	54

Cisco UCS Mini

The Cisco Unified Computing System™ (Cisco UCS), originally designed for the data center, is now optimized for branch and remote offices, point-of-sale deployments, and smaller IT environments with Cisco UCS Mini (Figure 1). Cisco UCS Mini is for customers who need fewer servers but still want the robust management capabilities provided by Cisco UCS Manager. This solution delivers servers, storage, and 10-Gbps networking in an easy-to-deploy, compact 6-rack-unit (6RU) form factor. Cisco UCS Mini provides a total computing solution with the proven management simplicity of the award-winning Cisco UCS Manager.

Figure 1. Cisco UCS Mini



Cisco UCS Manager 3.0

Cisco UCS Manager provides unified, embedded management of all Cisco UCS software and hardware components through your choice of an intuitive GUI, a command-line interface (CLI), a Microsoft PowerShell module, or an XML API. Cisco UCS Manager provides a unified management domain with centralized management capabilities and can control multiple chassis and thousands of virtual machines.

The Cisco UCS 6324 Fabric Interconnect hosts and runs Cisco UCS Manager in a highly available configuration, enabling the fabric interconnects to fully manage all Cisco UCS elements. The fabric interconnect supports out-of-band (OOB) management through dedicated 10/100/1000-Mbps Ethernet management ports. Cisco UCS Manager typically is deployed in a clustered active-passive configuration with two Cisco UCS 6324 Fabric Interconnects connected through the cluster interconnect built into the chassis.

Cisco UCS Manager 3.0 supports the Cisco UCS 6324, integrating the fabric interconnect into the Cisco UCS chassis and providing an integrated solution for smaller deployment environments. Cisco UCS Mini simplifies system management and saves costs for smaller-scale deployments. The hardware and software components support Cisco® Unified Fabric, which runs multiple types of data center traffic over a single converged network adapter.

Cisco UCS 6324 Fabric Interconnect

The Cisco UCS 6324 provides management, LAN, and storage connectivity for the Cisco UCS 5108 Blade Server Chassis and direct-connect rack-mount servers. It provides the same full-featured Cisco UCS management capabilities and XML API as the full-scale Cisco UCS solution in addition to integrating with Cisco UCS Central Software and Cisco UCS Director.

The Cisco UCS 6324 uses a cut-through network architecture that supports deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, with switching capacity of up to 500 Gbps independent of packet size and enabled services. Sixteen 10-Gbps links connect to the servers, providing a 20-Gbps link from each fabric interconnect to each server. The product family supports Cisco low-latency lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the interconnect. The solution achieves significant savings in total cost of ownership (TCO) through the use of a Fibre Channel over Ethernet (FCoE)-optimized server design that enables consolidation of network interface cards (NICs), host bus adapters (HBAs), cables, and switches.

The Cisco UCS 6324 Fabric Interconnect is a 10 Gigabit Ethernet, FCoE, and Fibre Channel switch that offers up to 500 Gbps of throughput and up to four unified ports and one scalability port.

Cisco UCS B200 M4 Blade Server

The enterprise-class Cisco UCS B200 M4 Blade Server extends the capabilities of the Cisco UCS portfolio in a half-width blade form factor. The Cisco UCS B200 M4 uses the power of the latest Intel® Xeon® processor E5-2600 v3 series CPUs, with up to 1536 GB of RAM (using 64-GB DIMMs), two solid-state disk (SSD) drives or hard-disk drives (HDDs), and connectivity with throughput of up to 80 Gbps. The Cisco UCS B200 M4 mounts in a Cisco UCS 5100 Series Blade Server Chassis or Cisco UCS Mini blade server chassis. It has a total of 24 slots for ECC registered DIMMs (RDIMMs) or load-reduced DIMMs (LR DIMMs), for up to 1536 GB of total memory capacity (Cisco UCS B200 M4 configured with two CPUs using 64-GB DIMMs). It supports one connector for the Cisco UCS VIC 1340 or 1240 adapter, which provides Ethernet and FCoE connectivity.

Cisco UCS VIC 1340

The Cisco UCS VIC 1340 is a 2-port 40 Gigabit Ethernet or dual 4-port 10 Gigabit Ethernet, FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 supports 2 ports of 40 Gigabit Ethernet.

The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCI Express (PCIe) standards-compliant interfaces to the host that can be dynamically configured as either NICs or HBAs. In addition, the VIC supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.

VMware vSphere 6.0

VMware vSphere 6.0, the industry-leading virtualization platform, empowers users to virtualize any application with confidence, redefines availability, and simplifies the virtual data center. The result is 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.

StorMagic SvSAN

StorMagic SvSAN is a software storage solution that enables enterprises to eliminate downtime for business-critical applications at remote sites, where this disruption directly equates with loss in service and revenue. SvSAN helps ensure high availability through a virtualized shared storage platform, so that these business-critical applications remain operational. This result is achieved by using direct-attached or internal, cost-effective server storage, including SSD storage, and presenting the storage as a virtual SAN (VSAN).

SvSAN supports the industry-leading hypervisors, VMware vSphere and Microsoft Hyper-V. It is installed as a virtual storage appliance (VSA) and requires few server resources to provide the shared storage necessary to enable advanced hypervisor features such as high availability and failover clustering, vMotion and Hyper-V live migration, and VMware Distributed Resource Scheduler (DRS) and Microsoft Dynamic Optimization.

You can deploy SvSAN as a simple 2-node cluster. However, the flexibility of the architecture enables you to scale the virtual infrastructure performance and capacity to meet your changing business needs without affecting service availability. You simply add capacity to existing servers or by increase the size of the SvSAN cluster.

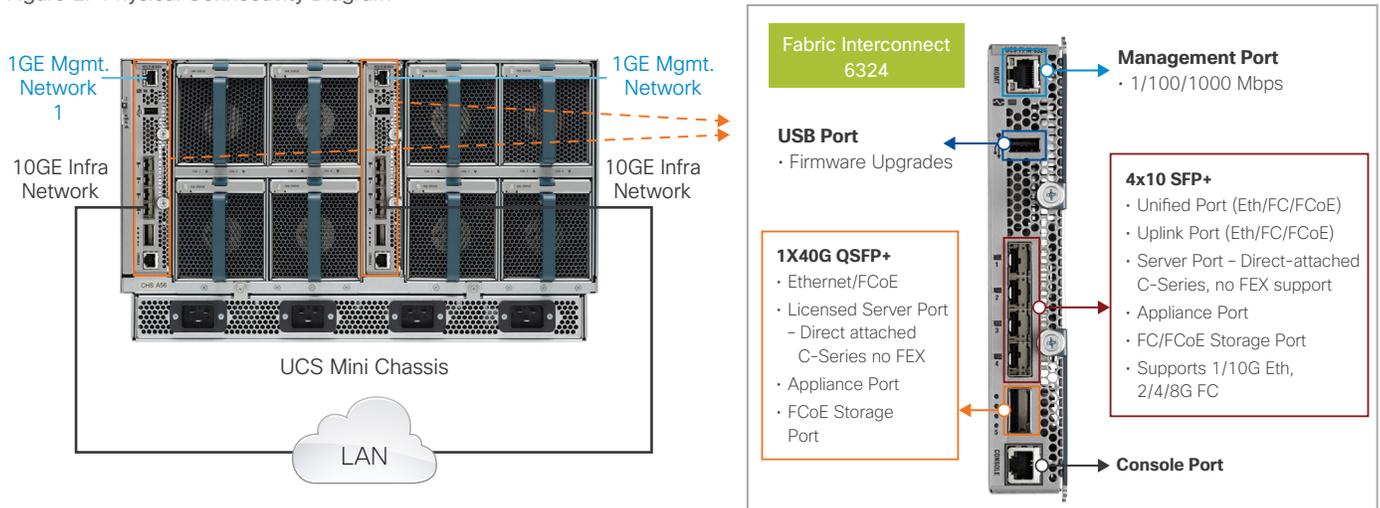
StorMagic SvSAN with Cisco UCS Mini Architecture

The architecture described in this deployment guide is simple. It provides a highly available environment using a combination of features from StorMagic SvSAN, VMware vSphere, and Cisco UCS Mini and using the local storage available on Cisco UCS B200 M4 Blade Servers.

Computing

This deployment uses the Cisco UCS Mini chassis with four Cisco UCS B200 M4 Blade Servers. Each blade server uses the Intel Xeon processor E5-2620 v3, 128 GB of RAM, two 32-GB Cisco FlexFlash Secure Digital (SD) cards, two 1.2-terabyte (TB) 10,000-rpm SAS drives, and the Cisco UCS VIC 1340 (Figure 2).

Figure 2. Physical Connectivity Diagram



Networking

The management ports on both the Cisco UCS 6324 Fabric Interconnects housed inside the Cisco UCS Mini chassis are connected to the management network. In addition, at least one 1- or 10-Gbps port from each side of the fabric is used as an uplink port to connect to the LAN through a pair of upstream switches, as shown in Figure 2. In this example, five virtual NICs (vNICs) are created in the service profile associated with the Cisco UCS B200 M4 servers and will be used to configure the VMware ESXi networking as explained here.

SvSAN uses three types of logical interfaces for different traffic types:

- Management: Used to access the web GUI, plug-in, or CLI; at least one must be present
- Small Computer System Interface over IP (iSCSI): Listens for incoming connection requests from iSCSI initiators; at least one must be present
- Mirror: Used by VSA nodes to communicate data and metadata associated with mirrored volumes

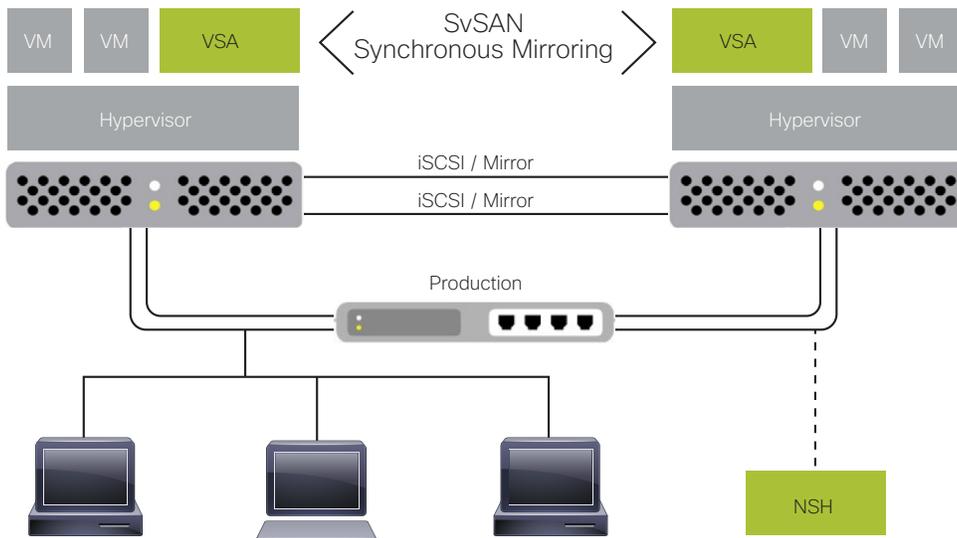
StorMagic recommends use of four or more physical NICs between hosts and SvSAN VSAs to provide the recommended level of redundancy and load balancing.

- vSwitch0 connects directly to the customer's production network, and it is assumed that this is where production virtual machines will reside. The VSA also has an interface on vSwitch0 for management, administration, and monitoring purposes only.
- vSwitch1 is an isolated network, segregated by either VLAN or IP address. This network should not be used for any other traffic apart from VSA iSCSI and mirror traffic, prevent contention and provide the best storage throughput.

- vSwitch2 is an isolated network, segregated by either VLAN or IP address. This network should not be used for any other traffic apart from VSA iSCSI and mirror traffic, to prevent contention and provide the best storage throughput.
- vSwitch3 is an isolated network, segregated by either VLAN or IP address. This network is dedicated to vMotion traffic.

Figure 3 shows network logical connectivity.

Figure 3. Logical Network Connectivity Diagram



Storage

In this deployment, each blade has two 1.2-TB 6-Gbps SAS disks and two 32-GB FlexFlash SD cards (Figure 4). The two 1.2-TB SAS drives are configured as RAID 0 striped and presented as two virtual drives (Figure 5). The first virtual drive is fixed for 2 TB, and the other is configured as a second virtual drive.

Figure 4. Disk Drives in Cisco UCS B200 M4 Server

Motherboard CIMC CPUs Memory Adapters HBAs NICs iSCSI vNICs Storage								
Controller LUNs Disks								
+ - Filter Export Print								
Name	Size (MB)	PID	Serial	Operability	Drive State	Presence	Device Type	
Controller SAS 1								
Disk 1	1143455	HUC101212CSS600	LOGXB10G	Operable	Online	Equipped	HDD	
Disk 2	1143455	HUC101212CSS600	LOGXVLBG	Operable	Online	Equipped	HDD	

Figure 5. RAID 0 Group with Two Virtual Drives in Cisco UCS B200 M4 Server

Motherboard CIMC CPUs Memory Adapters HBAs NICs iSCSI vNICs Storage					
Controller LUNs Disks					
Filter Export Print					
Name	Size (MB)	Raid Type	Operability	Presence	Bootable
Controller SAS 1					
Virtual Drive SvSAN-RDM	2097152	RAID 0 Striped	Operable	Equipped	True
Virtual Drive SvSAN_OS	189758	RAID 0 Striped	Operable	Equipped	False

The two 32-GB FlexFlash SD cards are configured for mirroring (RAID 1) in the Cisco UCS service profile's local disk configuration policy.

The ESXi OS is installed on the FlexFlash SD. The VSA deployment uses the 185-GB virtual disk and must be configured as persistent storage on the ESXi host. The second virtual drive is allocated explicitly to the VSA in a raw disk mapping. The VSA sees the virtual disk as a pool of storage from which iSCSI targets can be carved. These iSCSI targets can be mirrored with other VSAs to create highly available storage.

This deployment uses four mirror volumes: two mirror volumes between each pair of SvSAN VSAs, as shown in Figure 6. All iSCSI mirror volumes are presented to all ESXi hosts, but mirroring occurs only between a pair of SvSAN VSAs.

A mirrored target is a target that is mirrored between a pair of VSAs. Therefore, there are two identical copies of the target data when the system is running normally: one on each VSA. Any data sent to one copy, or plex, is automatically copied to the other. This copying is performed synchronously, so that if an initiator sends data to one plex, the VSA does not acknowledge the request until the data is present on both plexes. A synchronous mirror can be used to provide highly available storage, because an initiator can access any plex at any time, and if one plex fails, the initiator can continue without interruption by using the other plex.

Figure 6. StorMagic SvSAN Mirror Target with Four Cisco UCS B200 M4 Servers

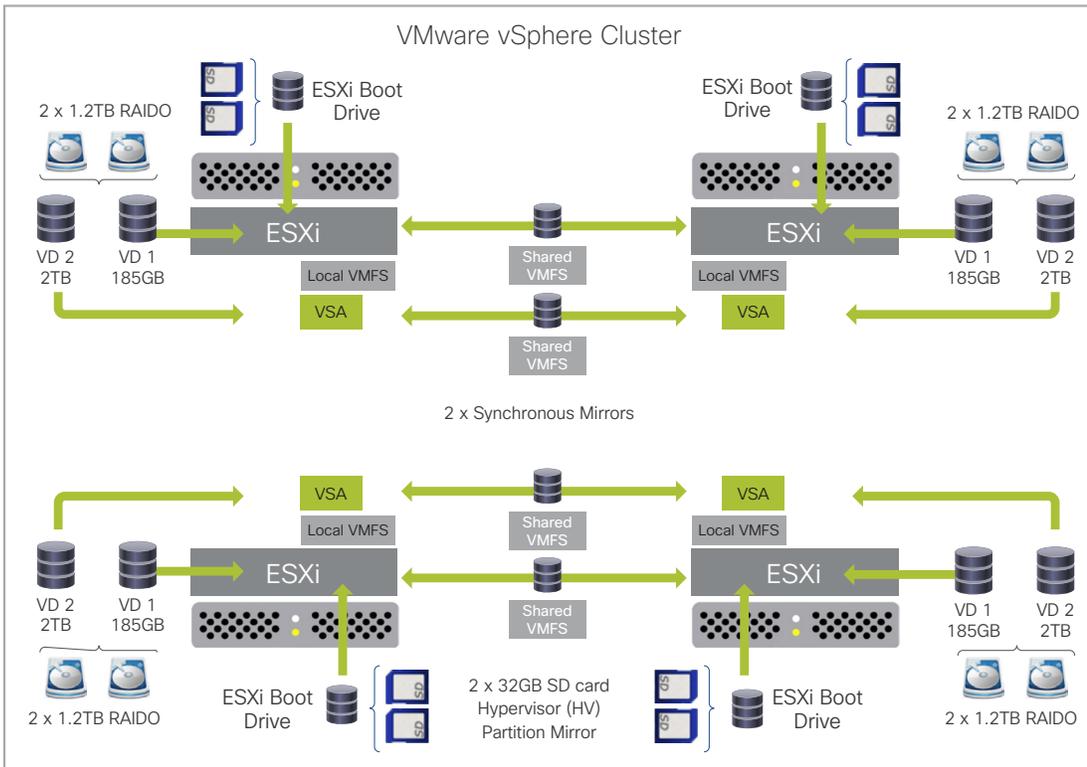


Table 1 lists the hardware and software components used in this deployment.

Table 1. Hardware Software Components

Component	Description
Cisco UCS	<ul style="list-style-type: none"> • 1 Cisco UCS Mini chassis • 2 Cisco UCS 6324 Fabric Interconnects with Cisco UCS Manager 3.0(2c) • 4 Cisco UCS B200 M4 Blade Servers <ul style="list-style-type: none"> - 2 x Intel Xeon processor E5-2620 v3 CPUs per blade server - 8 x 16-GB DDR4 2133-MHz RDIMMs, PC4-17000, dual-rank, x4, and 1.2V per blade server - 2 x 1.2-TB SAS 10,000-rpm SFF HDDs per blade server - 1 UCSB-MRAID12G controller per blade server - 1 Cisco UCS VIC 1340 (UCSB-MLOM-40G-03) per blade server - 2 x 32-GB Cisco Flexible Flash cards (UCS-SD-32G-S) per blade server
Cisco UCS firmware	<ul style="list-style-type: none"> • Cisco UCS Manager Release 3.0(2c) • Infrastructure firmware Release 3.0(2c)A
VMware vSphere	<ul style="list-style-type: none"> • VMware vCenter Release 6.0.0b • VMware ESXi Release 6.0.0b
StorMagic SvSAN	Release 5.3

Note: New Software and Firmware releases should be installed as listed in Cisco UCS HW and SW Interoperability Matrix: <http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html> and SvSAN certified and supported version appears on the VMware HCL: <http://www.vmware.com/resources/compatibility/search.php>

Table 2 shows the virtual machine sizing assumptions used for this Cisco UCS with StorMagic SvSAN deployment.

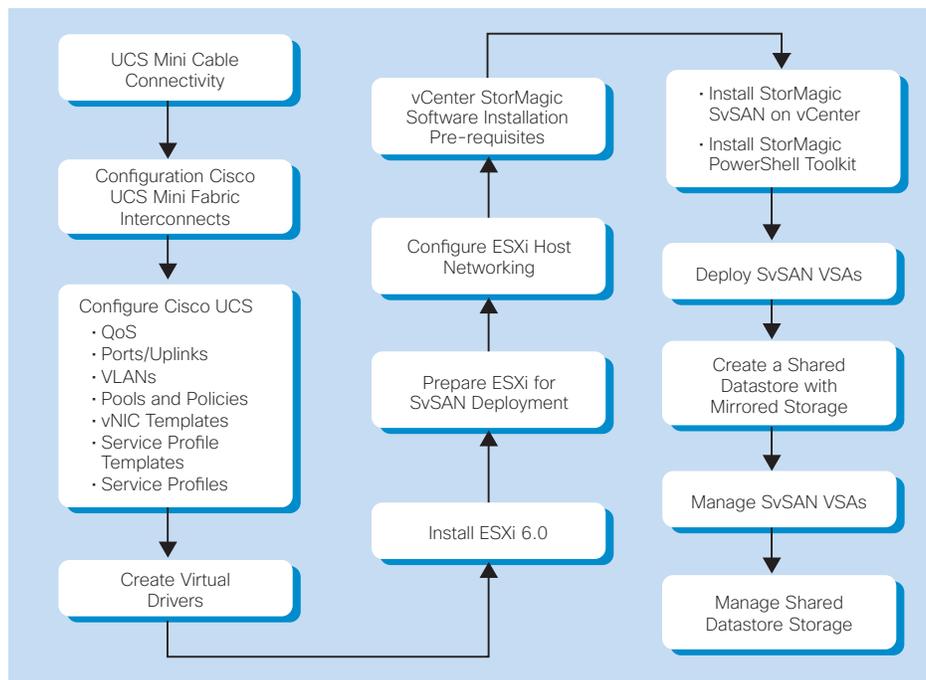
Table 2. Virtual Machine Configuration

Item	Description
Virtual machine instance size	4 virtual CPUs (vCPUs) with 6 GB of RAM 50-GB OS LUN; 100-GB raw LUN for I/O test

Deploying StorMagic SvSAN for VMware vSphere in a Cisco UCS Environment

The flow chart in Figure 7 shows the high-level steps for deploying StorMagic SvSAN for VMware vSphere in a Cisco UCS environment.

Figure 7. High-Level Steps for Deploying StorMagic SvSAN for VMware vSphere on Cisco UCS Mini with Cisco UCS B200 M4 Blade Servers



Configuring Cisco UCS

This section presents a high-level procedure for configuring the Cisco UCS Mini environment to deploy StorMagic SvSAN for VMware vSphere.

Configure Cisco UCS Mini Fabric Interconnects

Perform the initial setup of the Cisco UCS 6324 Fabric Interconnects in a cluster configuration. Refer to the following URL for step-by-step instructions to perform the initial system setup for a cluster configuration: http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/3-0/b_UCSM_GUI_User_Guide_3_0/b_UCSM_GUI_User_Guide_3_0_chapter_0101.html.

Configure Cisco UCS Manager

Log in to Cisco UCS Manager and configure the pools, policies, and service profiles as presented in the sections that follow.

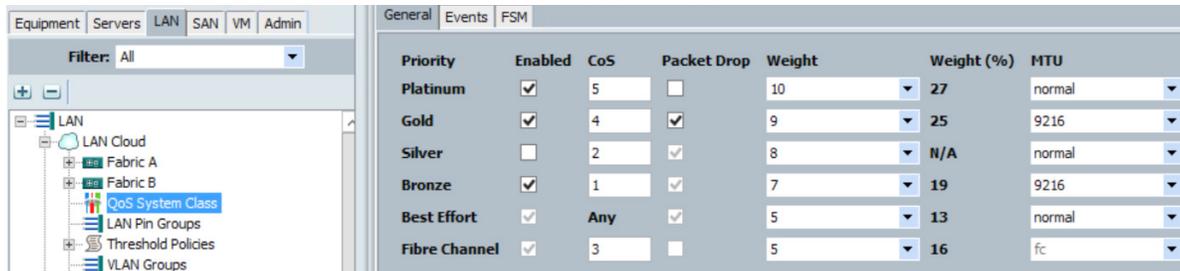
Synchronize Cisco UCS with NTP

In Cisco UCS Manager, select the Admin tab and navigate to Time Zone Management. Select a time zone and add an NTP server address to synchronize Cisco UCS with an NTP server.

Enable Quality of Service in Cisco UCS Fabric

In Cisco UCS Manager, select the LAN tab and navigate to LAN Cloud > QoS System Class to enable the quality-of-service (QoS) priorities. Enter the maximum transmission unit (MTU) size (Figure 8).

Figure 8. Cisco UCS Manager: Enable QoS System Class



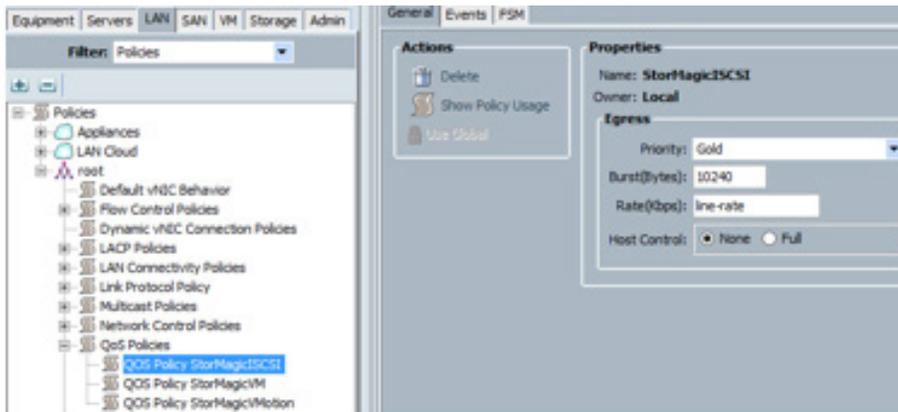
Create QoS Policy

Select the LAN tab and navigate root > Policies > QoS Policies. Create three policies with the information provided in Table 3 and Figure 9.

Table 3. QoS Policies for Different Traffic

QoS Policy Name	Priority Selection	Other Parameters
StorMagicISCSI	Gold	Default
StorMagicVM	Platinum	Default
StorMagicVMotion	Bronze	Default

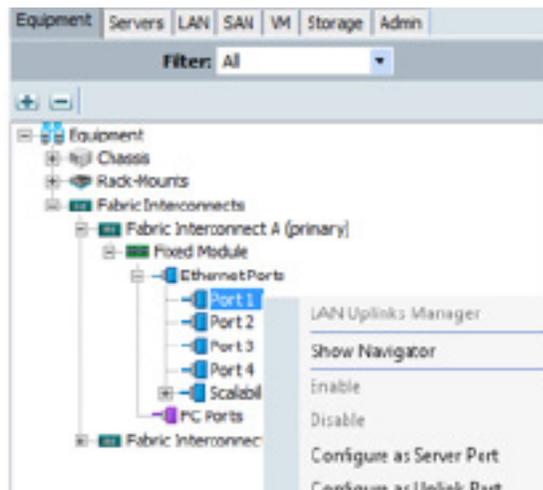
Figure 9. Cisco UCS Manager: Create QoS Policy



Enable Uplink Ports

In Cisco UCS Manager, select the **Equipment** tab and navigate to **Fabric Interconnects > Fabric Interconnects A > Fixed Module > Ethernet Ports**. Configure the ports connected to upstream switches as uplink ports (Figure 10). Repeat the same process on Fabric Interconnect B.

Figure 10. Cisco UCS Manager: Enable Uplink Ports



Create VLANs

Select the LAN tab and navigate to **LAN > LAN Cloud**. Create one VLAN in each Fabric A (StorMagicISCSI-A VLAN) and Fabric B (StorMagicISCSI-B VLAN) for the iSCSI storage traffic, as shown in Figure 11.

Figure 11. Cisco UCS Manager: Create VLANs for iSCSI Storage Traffic

The screenshot displays two configuration panels for VLANs in Cisco UCS Manager. The top panel shows the configuration for 'StorMagicISCSI-A' with a VLAN ID of 253. The bottom panel shows the configuration for 'StorMagicISCSI-B' with a VLAN ID of 254. Both configurations include a 'Fault Summary' section with four status icons (red X, orange triangle, yellow triangle, green triangle) and a 'Properties' section with the following details:

- Name: StorMagicISCSI-A (top) / StorMagicISCSI-B (bottom)
- Native VLAN: No
- Network Type: Lan
- Locale: External
- Owner: Local
- Multicast Policy Name: <not set>
- Multicast Policy Instance: org-root/mc-policy-default
- Sharing Type: None Primary Isolated Community
- VLAN ID: 253 (top) / 254 (bottom)
- Fabric ID: A (top) / B (bottom)
- If Type: Virtual
- Transport Type: Ether

Select the **LAN** tab and navigate to **LAN > VLANs**. Create the global VLANs for your environment, as shown in Figure 12. For this deployment, two VLANs were created, one for management traffic and one for vMotion traffic, as shown in Figure 12.

Figure 12. Create Global VLANs for Management and vMotion Traffic

The screenshot displays the 'VLANs' configuration page in Cisco UCS Manager. It includes a table listing the created VLANs and a detailed view of the 'Mgmt' VLAN configuration.

Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN Name	Multicast Po
VLAN Mgmt (604)	604	Lan	Ether	No	None		
VLAN Mgmt2 (121)	121	Lan	Ether	No	None		
VLAN Mgmt3 (602)	602	Lan	Ether	No	None		
VLAN StorMagicVmotion (255)	255	Lan	Ether	No	None		
VLAN default (1)	1	Lan	Ether	Yes	None		

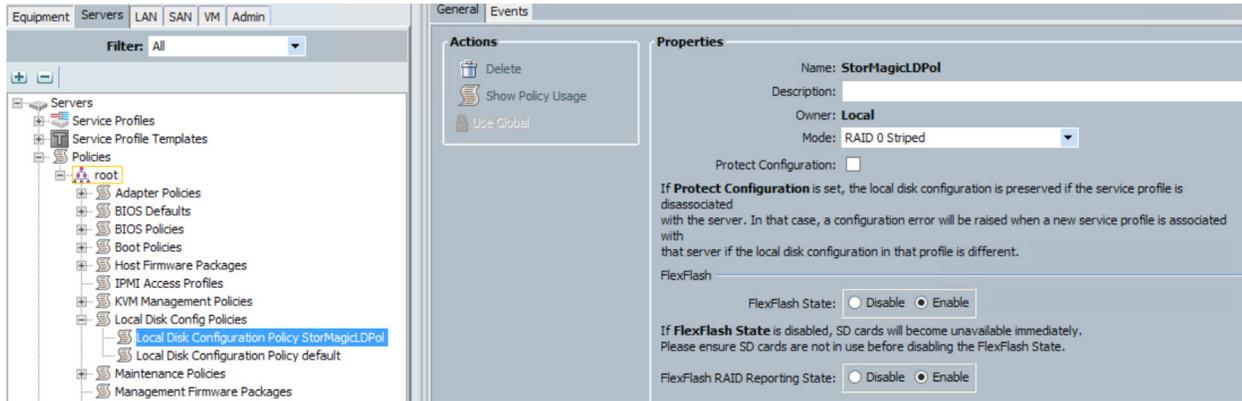
The 'Details' section for the 'Mgmt' VLAN (VLAN ID: 604) includes the following configuration:

- Name: Mgmt
- Native VLAN: No
- Network Type: Lan
- Locale: External
- Owner: Local
- Multicast Policy Name: <not set>
- Multicast Policy Instance: org-root/mc-policy-default
- Sharing Type: None Primary Isolated Community
- VLAN ID: 604
- Fabric ID: Dual
- If Type: Virtual
- Transport Type: Ether

Create Local Disk Configuration Policy

Select the **Servers** tab and navigate to **Policies > root > Local Disk Configuration Policy**. Create a local disk configuration policy. For Mode, choose RAID 0 Striped, and enable both FlexFlash State and FlexFlash RAID Reporting State, as shown in Figure 13.

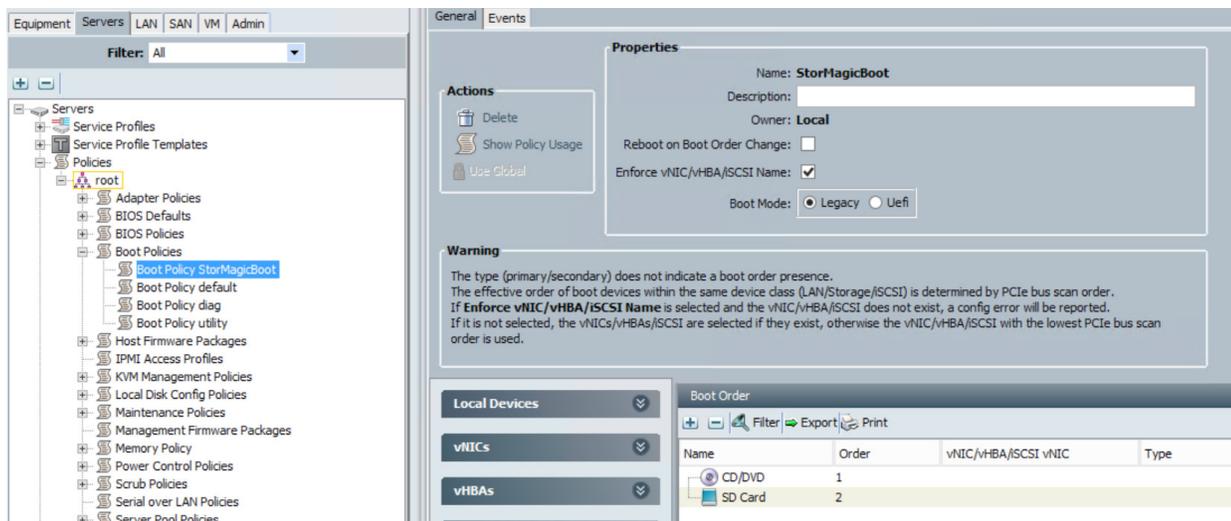
Figure 13. Create Local Disk Configuration Policy



Create Boot Policy

Select the **Servers** tab and navigate to **Policies > root > Boot Policies**. Create a boot policy by adding CD/DVD and SD Card in the Boot Order pane, as shown in Figure 14.

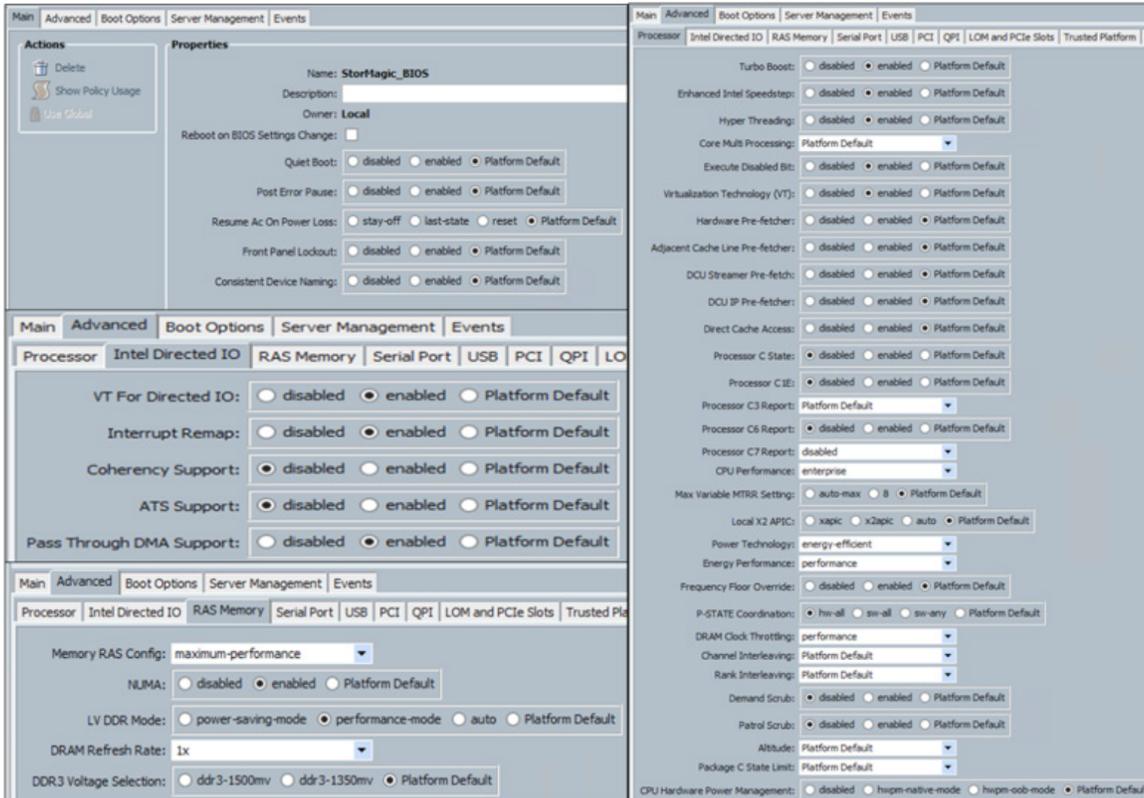
Figure 14. Create Boot Policy



Create BIOS Policy

Select the **Servers** tab and navigate to **Policies > root > BIOS Policies**. Create a BIOS policy with settings as shown in Figure 15.

Figure 15. Create BIOS Policy



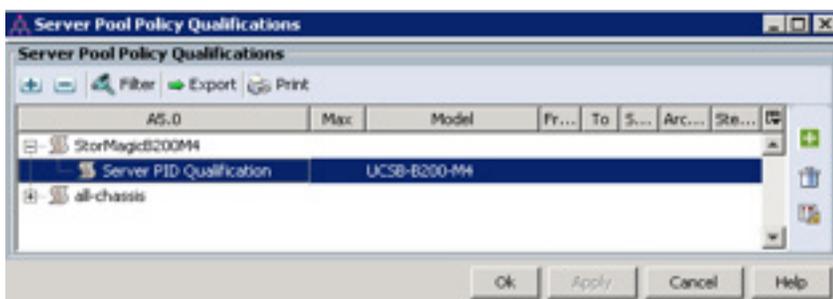
Configure Server Pool and Qualifying Policy

In this section, you create server pool, server pool policy qualifications, and server pool policy to auto populate servers depending on the configuration.

Select the **Servers** tab and navigate to **Pools > root > Server Pool**. Right-click to create a server pool. Provide a name and click Finish without adding any servers to the pool.

Select the **Servers** tab and navigate to **Policies > root > Server Pool Policy Qualifications**. Right-click to create server pool policy qualifications with Server PID Qualifications as the criteria. Enter **UCSB-B200-M4** for the model, as shown in Figure 16.

Figure 16. Create Server Pool Policy Qualifications



Select the **Servers** tab and navigate to **Policies > root > Server Pool Policies**. Right-click to create a server pool policy. Provide a name and, from the drop-down list, select the target pool and qualification created in the previous section (Figure 17).

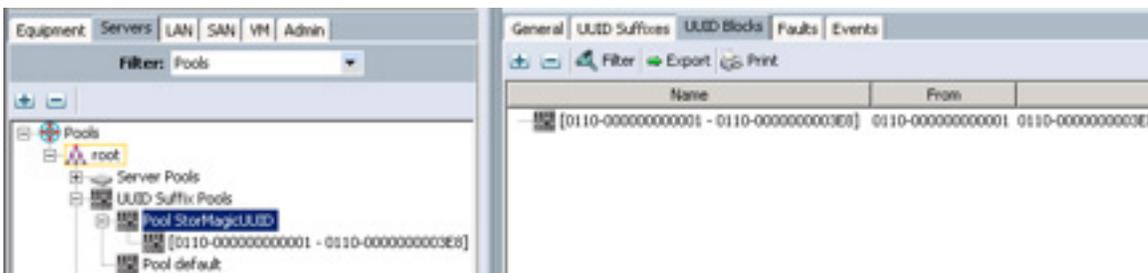
Figure 17. Create Server Pool Policy



Create UUID Suffix Pool

Select the **Servers** tab and navigate to **Pools > root > UUID Suffix Pools**. Right-click to create a unique user ID (UUID) suffix pool and then add a block of UUIDs to the pool (Figure 18).

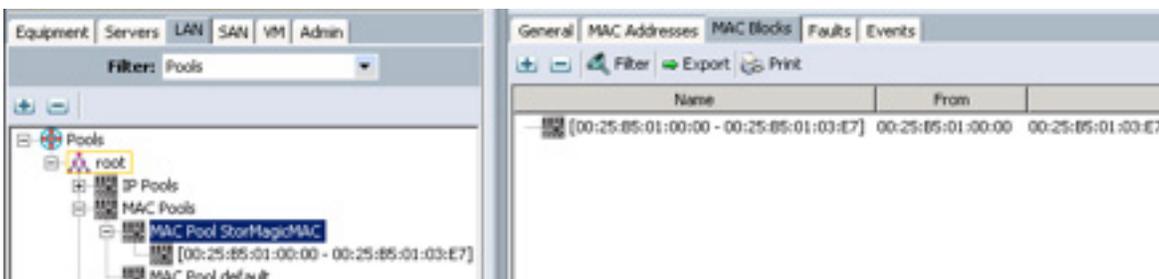
Figure 18. Create UUID Suffix Pool



Create MAC Address Pool

Select the **LAN** tab and navigate to **Pools > root > MAC Pools**. Right-click to create a MAC address pool and then add a block of MAC addresses to the pool (Figure 19).

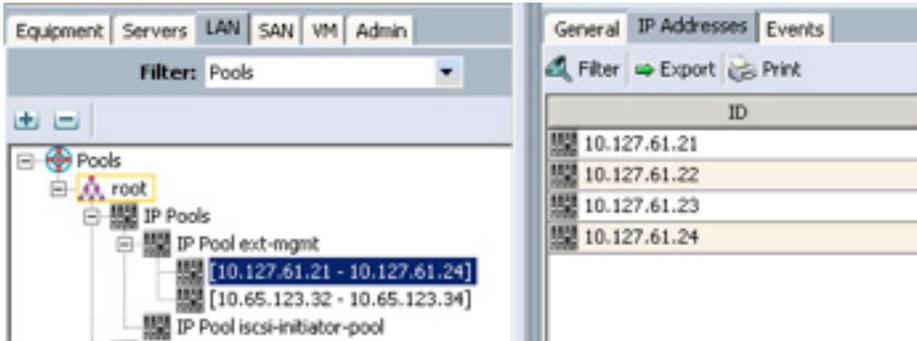
Figure 19. Create MAC Address Pool



Create an IP Address Pool for KVM Access

Select the **LAN** tab and navigate to **Pools > root > IP Pools > IP Pool ext-mgmt**. Right-click to create a block of IPv4 addresses for Kernel-based Virtual Machine (KVM) access (Figure 20).

Figure 20. Create IP Pool for KVM Access



Create vNIC Templates and LAN Connectivity Policies

Select the **LAN** tab and navigate to **Policies > root > vNIC Templates**. Right-click to create vNIC templates using Table 4 and Figure 21.

Table 4. vNIC Templates

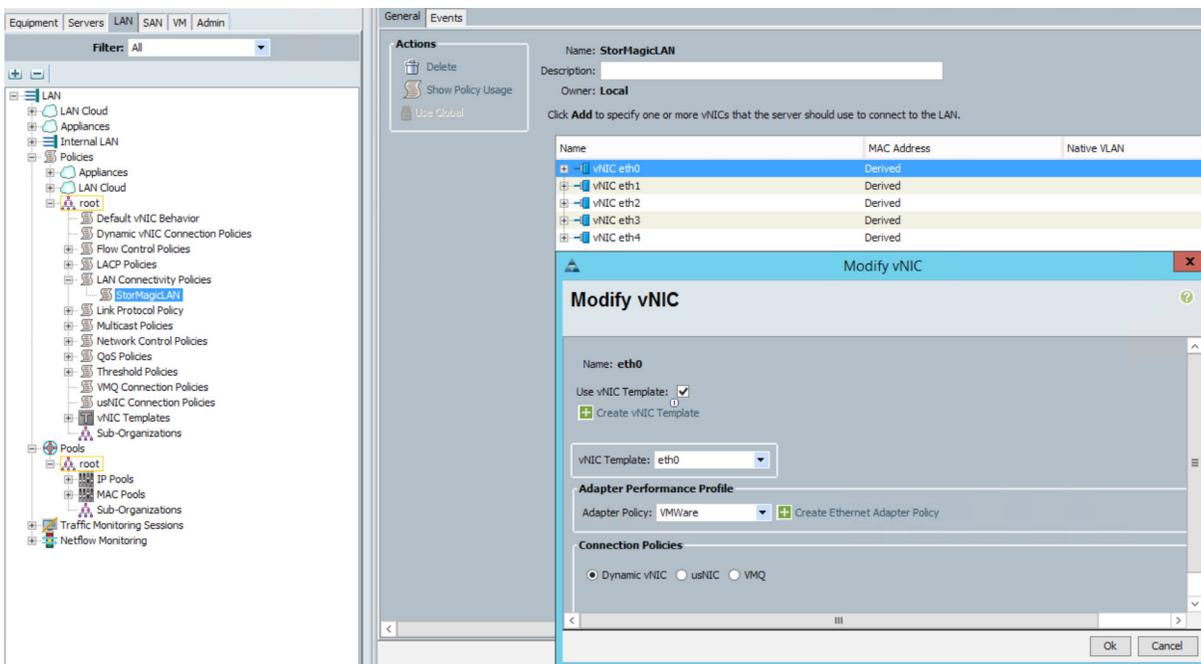
vNIC Template Name	Fabric ID	Enable Failover (Yes or No)	Template Type	VLANs Allowed	Native VLAN	MTU	MAC Address Pool	QoS Policy
eth0	Fabric A	No	Updating	Mgmt	Mgmt	1500	StorMagicMAC	StorMagicVM
eth1	Fabric B	No	Updating	Mgmt	Mgmt	1500	StorMagicMAC	StorMagicVM
eth2	Fabric A	No	Updating	StorMagicISCSI-A	StorMagicISCSI-A	9000	StorMagicMAC	StorMagicISCSI
eth3	Fabric B	No	Updating	StorMagicISCSI-B	StorMagicISCSI-B	9000	StorMagicMAC	StorMagicISCSI
eth4	Fabric B	Yes	Updating	StorMagicVMotion	StorMagicVMotion	9000	StorMagicMAC	StorMagicVMotion

Figure 21. Create vNIC Template



Create a LAN connectivity policy using the vNIC templates you just created. Select VMware as the adapter policy for the adapter performance profile (Figure 22).

Figure 22. Create LAN Connectivity Policy



Create a Service Profile Template

From the Servers tab, create a service profile template using all the pools and policies that you created in the preceding sections (Figure 23).

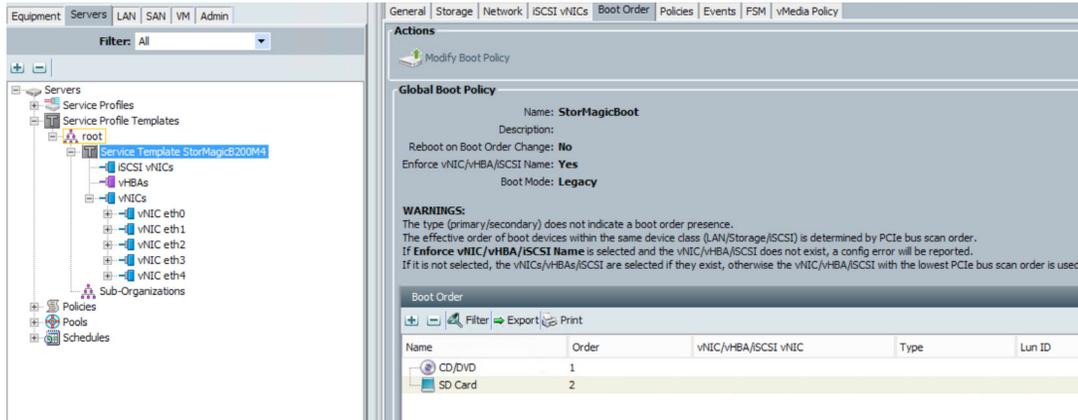
1. In the navigation pane, click the **Servers** tab.
2. On the Servers tab, expand **Servers > Service Profiles**.
3. Expand the node for the organization for which you want to create the service profile. If the system does not include multitenancy, expand the root node.
4. Right-click the organization and select **Create Service Profile (expert)**.
5. In the **Identify Service Profile** panel, specify the service profile name and UUID; then click **Next**. You can also provide an optional description for this service profile. If the UUID is not available, you can also create a UUID suffix pool from this panel.

Note: To create a service profile quickly, you can click **Finish** after you specify the name. Cisco UCS Manager creates a new service profile with the specified name and all system default values.

6. (Optional) In the **Networking** panel, specify the required information in the **Dynamic vNIC Connection Policy** and **LAN Connectivity** sections; then click **Next**.
You can create a dynamic vNIC connection policy and LAN connectivity policy from this panel
7. (Optional) In the **Storage** panel, specify the SAN configuration information, such as local storage policy, SAN connectivity, World Wide Name (WWNN), and VSAN; then click **Next**. You can also create a local disk configuration policy and SAN connectivity policy from this panel.
8. (Optional) In the **Zoning** panel, specify the required zoning information; then click **Next**. You can also create virtual HBA (vHBA) initiator groups from this panel.
9. (Optional) In the **vNIC/vHBA Placement** panel, specify the placement method and PCI order; then click **Next**. You can also create a placement policy from this panel.
10. (Optional) In the **Server Boot Order** panel, specify the **Boot Policy** from the drop-down list; then click **Next**. You can also create a boot policy from this panel.
11. (Optional) In the **Maintenance Policy** panel, specify the maintenance policy; then click **Next**. You can also create a new maintenance policy and specify a maintenance schedule from this panel.
12. (Optional) In the **Server Assignment** panel, specify the server assignment from the drop-down list and the power state to apply on assignment; then click **Next**. You can also create a server pool or a host firmware package from this panel.
13. (Optional) In the **Operational Policies** panel, specify the system operating information, such as the BIOS configuration, external IPMI management configuration, management IP address, monitoring configuration (thresholds), power control policy configuration, and scrub policy; then click **Finish**.

Note: To set up an out-of-band IPv4 address or an in-band IPv4 or IPv6 address, click the respective tabs and complete the required fields.

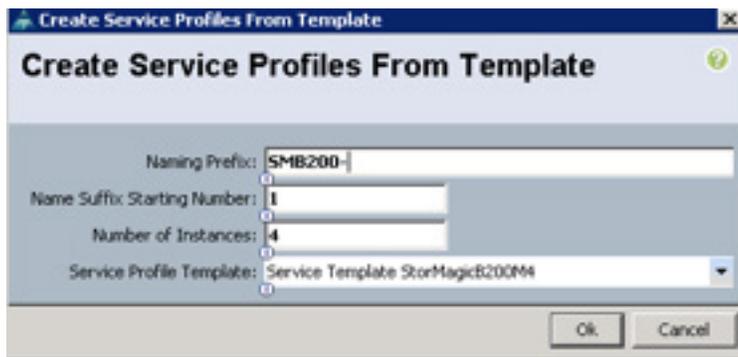
Figure 23. Create Service Profile Template



Create Service Profiles from Templates

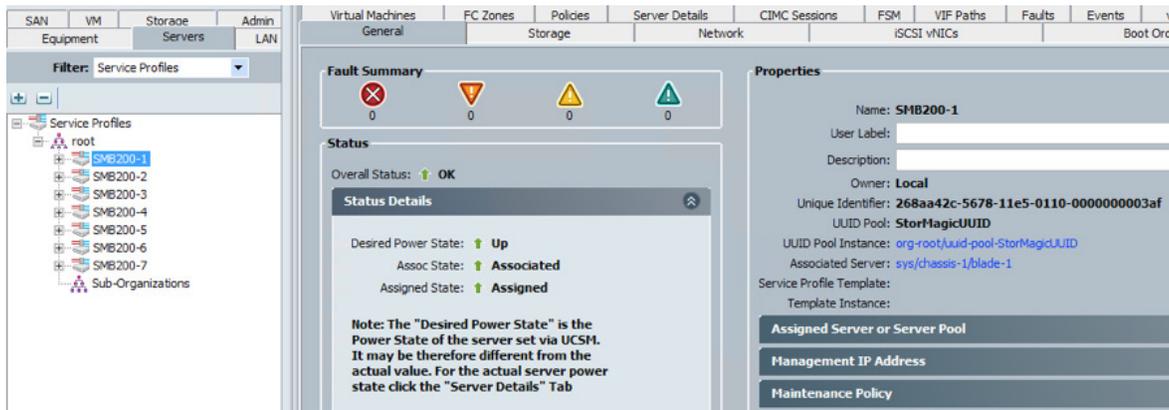
From the **Servers** tab, create service profiles from a template. Provide a naming prefix, a starting number, and the number of profile instances you want to create, and select the service profile template created in the preceding step (Figure 24).

Figure 24. Create Service Profiles from Template



This configuration creates four service profiles and automatically associates them with the blade servers that match the selection criteria chosen in the server pool policy (Figure 25).

Figure 25. Service Profiles Associated with Blade Servers



Create Virtual Drives

Select a service profile, click Boot Server, and launch the KVM console. Enter the Cisco FlexStorage Controller BIOS Configuration Utility (Ctrl+R) to configure the RAID 0 group with the two 1.2T-B drives and create two virtual drives. Complete this RAID configuration on all the four Cisco UCS B200 M4 servers, as shown in Figure 26.

1. To enter the LSI BIOS settings, press **Ctrl+R** while the system is booting (Figure 26).

Figure 26. Open the BIOS Settings



```
LSI MegaRAID SAS-MFI BIOS
Version 6.19.05.0 (Build May 07, 2014)
Copyright(c) 2014 LSI Corporation

HA -0 (Bus 1 Dev 0) Cisco FlexStorage 12G SAS RAID Controller
Battery Status: Not present
PCI Slot Number: 11

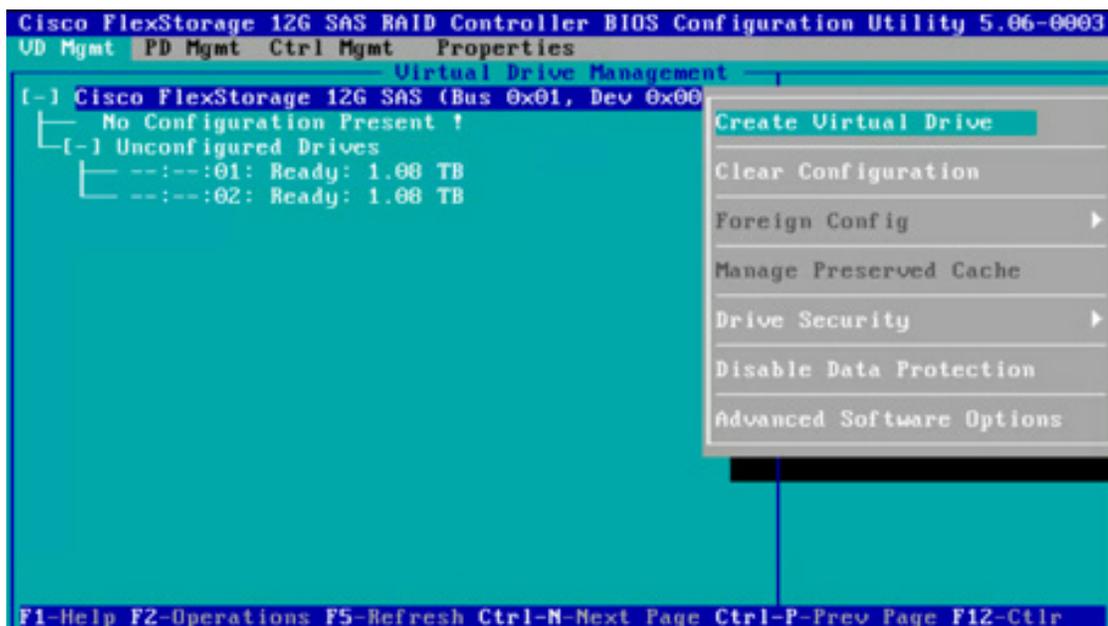
ID LUN VENDOR    PRODUCT                REVISION    CAPACITY
--- --  -
6  0  SEAGATE  ST1200MM0007          0002        1144641MB
7  0  SEAGATE  ST1200MM0007          0002        1144641MB
  0  LSI     Virtual Drive          RAID0        2097152MB
  1  LSI     Virtual Drive          RAID0        189758MB

0 JBOD(s) found on the host adapter
0 JBOD(s) handled by BIOS

2 Virtual Drive(s) found on the host adapter.
2 Virtual Drive(s) handled by BIOS
Press <Ctrl><R> to Run MegaRAID Configuration Utility
```

2. On the LSI BIOS page, select the adapter, and press **F2** for operations, and select **Create New Virtual Drive** (Figure 27)

Figure 27. Create a New Virtual Drive



```
Cisco FlexStorage 12G SAS RAID Controller BIOS Configuration Utility 5.06-0003
UD Mgmt PD Mgmt Ctrl Mgmt Properties
Virtual Drive Management
[-] Cisco FlexStorage 12G SAS (Bus 0x01, Dev 0x00)
  [-] No Configuration Present !
    [-] Unconfigured Drives
      --:--:01: Ready: 1.08 TB
      --:--:02: Ready: 1.08 TB
      Create Virtual Drive
      Clear Configuration
      Foreign Config
      Manage Preserved Cache
      Drive Security
      Disable Data Protection
      Advanced Software Options

F1-Help F2-Operations F5-Refresh Ctrl-N-Next Page Ctrl-P-Prev Page F12-Ctrl
```

3. For the RAID level, select RAID 0. Two HDDs should be available for configuring the new virtual drive. Enter the sizes and names of the virtual drive.
 - a. Create one virtual drive with 2 TB of space. This drive will be used by SvSAN (Figure 28).
 - b. Select the Advanced option and configure virtual drive properties. Then initialize the virtual drive (Figure 29).
 - c. Create the second virtual drive with the remaining space. This drive will be used for SvSAN virtual machines (Figure 30).
 - d. Select the Advanced option and configure virtual drive properties. Then initialize the virtual drive (Figure 31).

Figure 28. Create the First Virtual Drive

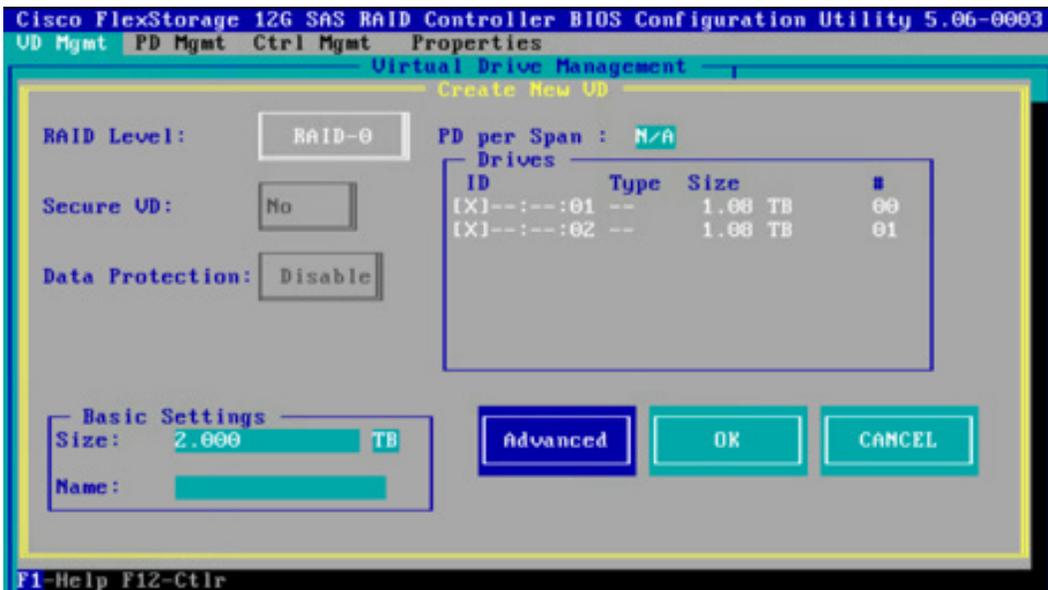


Figure 29. Select Advanced Properties for the First Virtual Drive

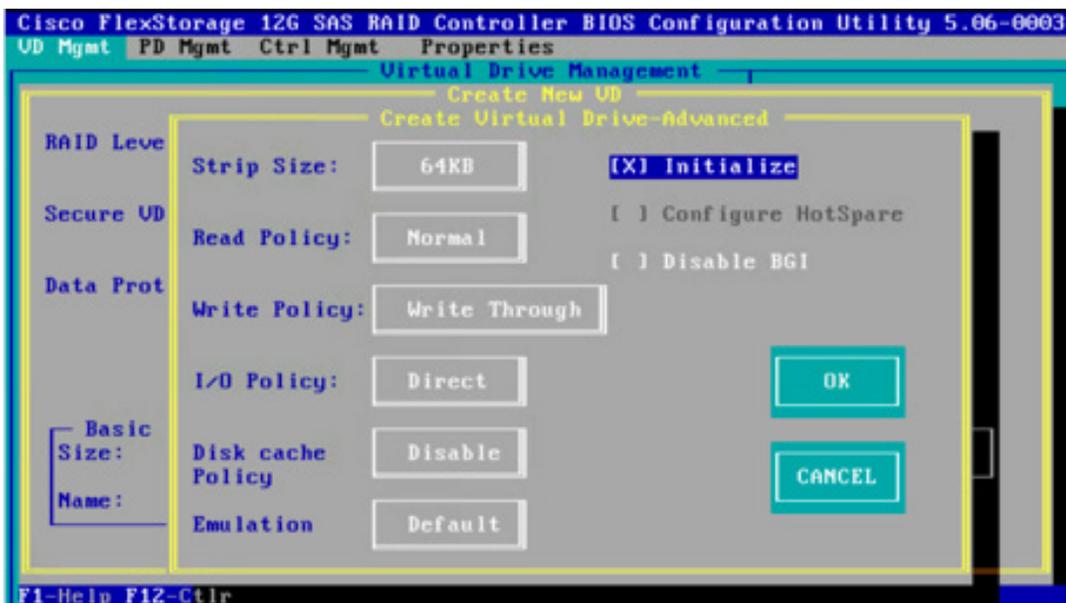


Figure 30. Create a Second Virtual Drive

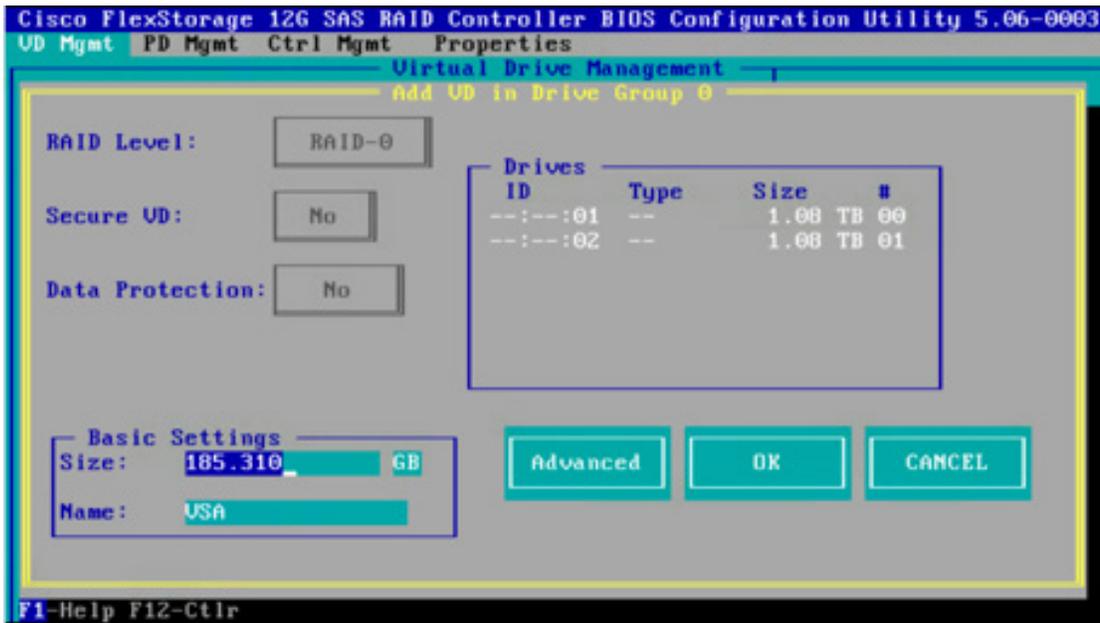
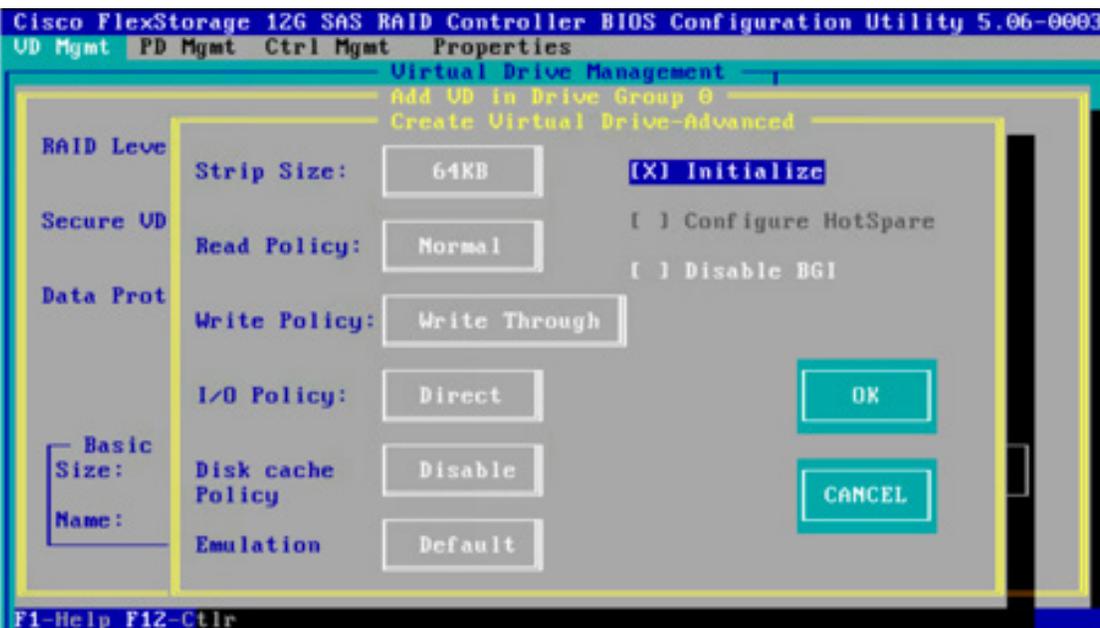
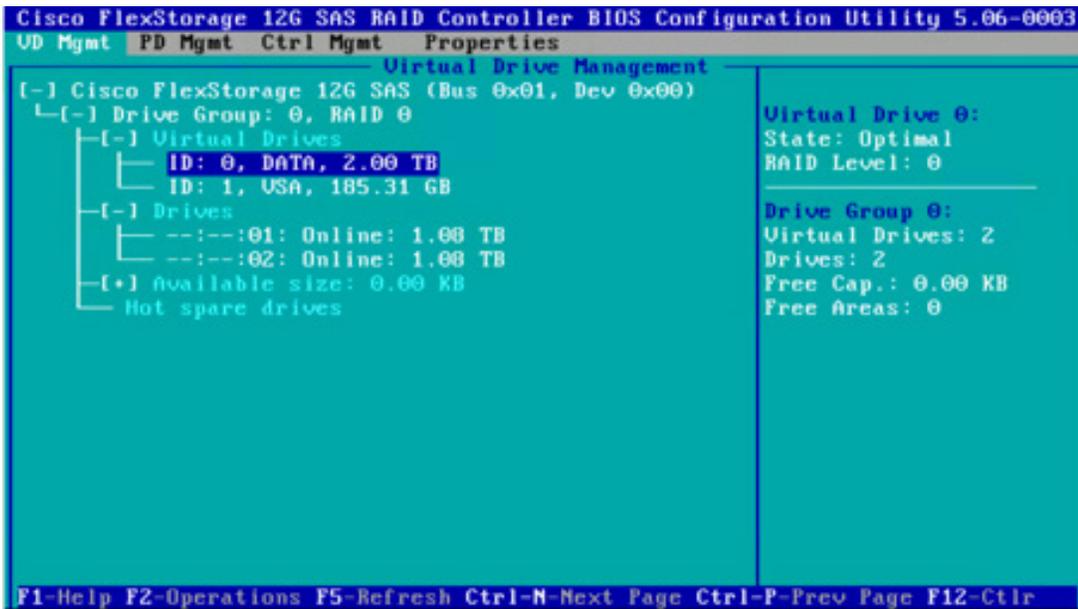


Figure 31. Select Advanced Properties for the Second Virtual Drive



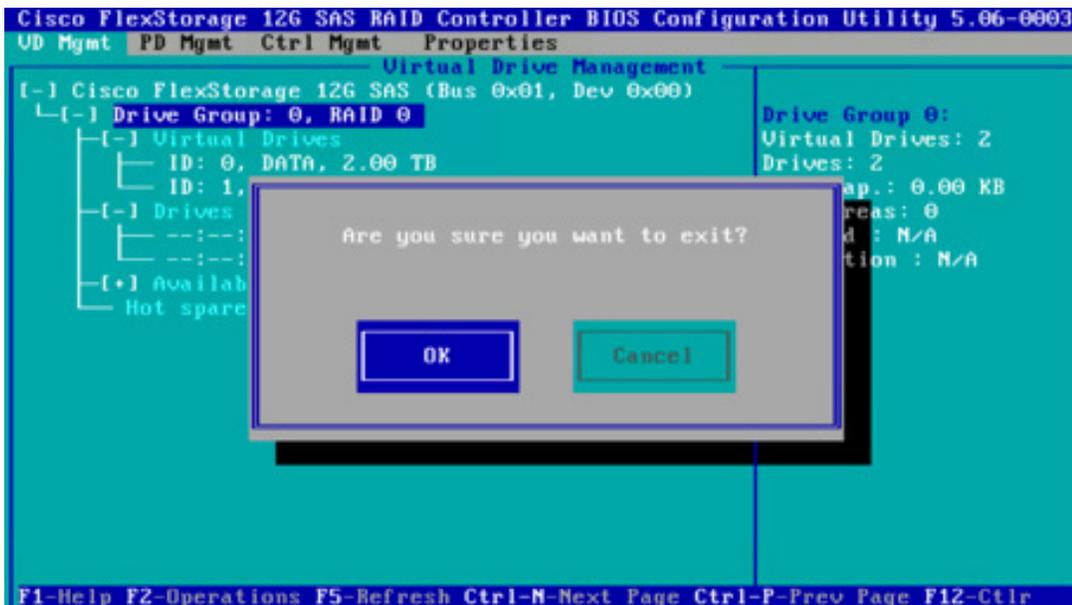
4. After the virtual drives have been created, the details should be available on main page of the LSI settings (Figure 32).

Figure 32. View the Virtual Drive Details



5. Press **Esc** to exit from this menu (Figure 33). Then press **Ctrl+Alt+Del** to reboot the server.

Figure 33. Exit the Menu

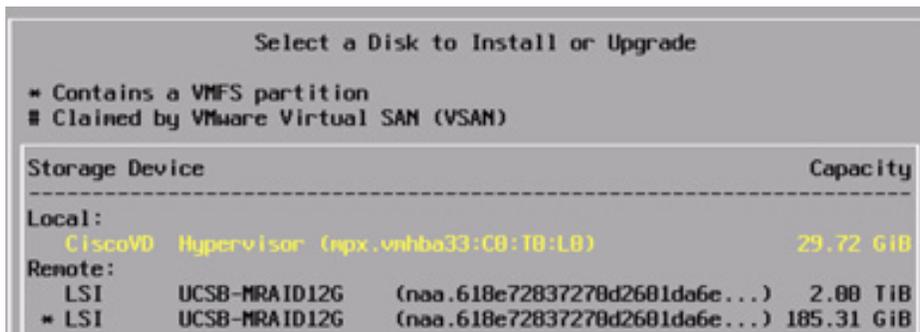


Installing VMware ESXi 6.0

Perform the procedure in this section on all Cisco UCS B200 M4 Blade Servers.

1. Download the Cisco Custom image for VMware ESXi 6.0 from the following URL: https://my.vmware.com/web/vmware/details?productId=490&downloadGroup=OEM-ESXI60GA-CISCO#product_downloads.
2. In Cisco UCS Manager, select a blade server and launch a KVM console session. From the console's **Virtual Media** tab, select **Map CD/DVD** and mount the Cisco custom ISO image that you downloaded in the previous step.
3. Select the CiscoVD Hypervisor partition, as shown in the Figure 34.

Figure 34. VMware ESXi Installation Storage Device Selection



4. After the ESXi installation is complete, reboot the system and configure the management IP address.

Prepare VMware ESXi for StorMagic SvSAN Deployment

1. Log in to the ESXi host using the vSphere client, navigate to **Configuration > Storage**, and click **Add Storage** to add persistent storage to the ESXi host for SvSAN VSA deployment.
2. Create a storage array type plug-in (SATP) rule to help ensure that ESXi uses the round-robin path policy and reduces the path I/O operation per second (IOPS) count to 1 IOPS for all StorMagic iSCSI volumes (Figure 35).

```
esxcli storage nmp satp rule add --satp "VMW_SATP_ALUA" --psp "VMW_PSP_RR" -O iops=1 --vendor="StorMagic" --model="iSCSI Volume"
```

Figure 35. Create SATP Rule

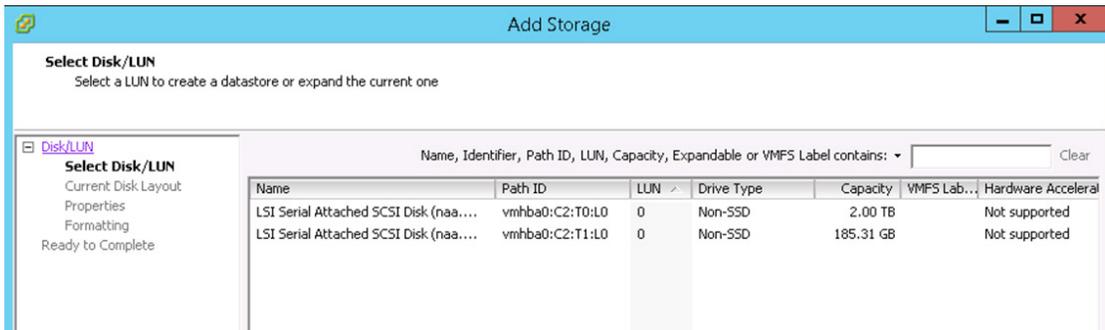


This step helps ensure the even distribution of I/O across mirrored volumes and is a best practice for StorMagic on Cisco UCS Mini architecture. You can use the following command to verify that the setting has taken effect:

```
esxcli storage nmp device list
```

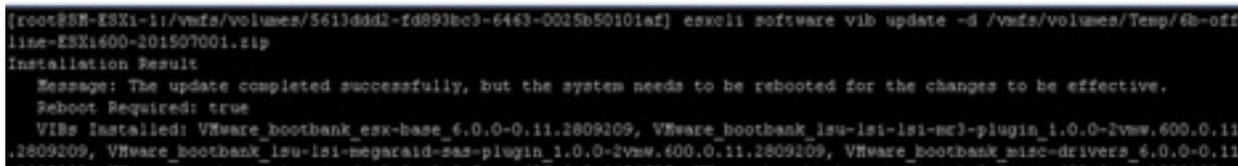
- For the storage type, choose **Select Disk/LUN** and create a data store using the 185-GB LUN, as shown in Figure 36.

Figure 36. Add Persistent Storage to VMware ESXi Host



- Put the ESXi host into maintenance mode:
`vim-cmd hostsvc/maintenance_mode_enter`
- From the ESXi shell, install the latest ESXi version from the following link, as shown in Figure 37:
<https://my.vmware.com/group/vmware/patch-search>.
`esxcli software vib update -d /vmfs/volumes/Temp/<filename>.zip`

Figure 37. Upgrade VMware ESXi



- Download the Cisco UCS B200 M4 device drivers image `ucs-bxxx-drivers.3.0.2a.iso` from the following URL and upgrade the drivers for the NIC, storage, and other components, as shown in Figure 38: [https://software.cisco.com/download/release.html?mdfid=283853163&flowid=25821&softwareid=283853158&release=2.2\(6\)&reind=AVAILABLE&rellifecycle=&reltype=latest](https://software.cisco.com/download/release.html?mdfid=283853163&flowid=25821&softwareid=283853158&release=2.2(6)&reind=AVAILABLE&rellifecycle=&reltype=latest)
 Command: `esxcli software vib install --no-sig-check -v /vmfs/volumes/Temp/<filename>.vib`

Figure 38. Upgrade Device Driver



Note: ESX drivers can be updated using ZIP file or VIB file. For Zip use `-d` flag and for VIB file, use `-v` flag.

- Verify the path policy (Figure 39)

Figure 39. Verify Path Policy Setting

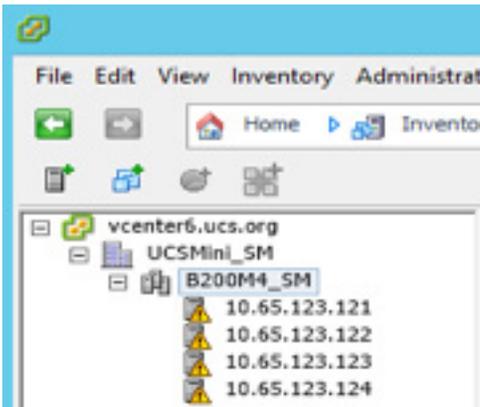
```

vcenter6.ucs.org
Device Display Name: MirrorBA
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: (implicit_support=on,explicit_support=off,explicit_allow=on,allow_followover=on;action_on_retry_error=off;(TPO_up=1,TPO_state=AO))
TPO_up=1,TPO_state=AO))
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config: (policy=rr,iops=1,bytes=10485760,useABO=0;lastPathIndex=0;numIOsPending=0,numBytesPending=0)
Path Selection Policy Device Custom Config:
Working Paths: vmwaa3:C1:T1:L0, vmwaa3:C1:T1:L0, vmwaa3:C1:T1:L0, vmwaa3:C1:T1:L0
Is CSB: false
  
```

Figure 39 shows MirrorBA. The two fields that identify the changes are:

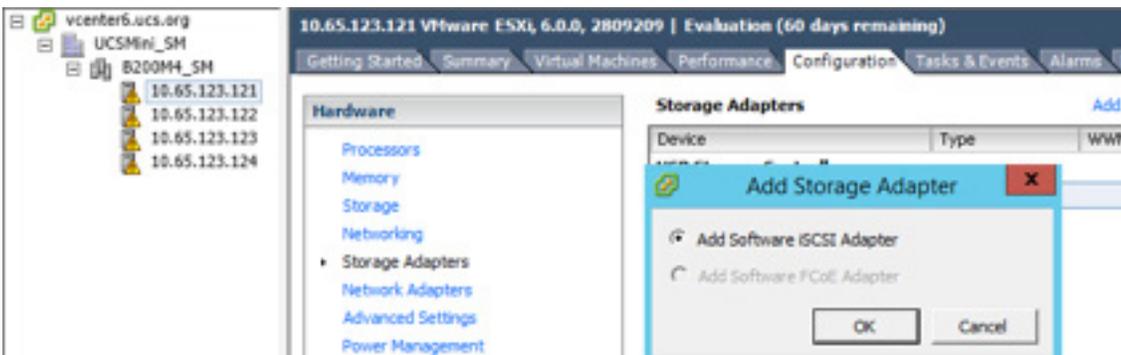
- Path Select Policy: VMW_PSP_RR
 - Path Selection Policy Device Config: (policy=rr,iops=1....)
8. Log in to vCenter and add all the ESXi hosts to the managing vCenter Server on which the StorMagic software is installed.
 9. Under **Datacenter**, create a cluster and add all the ESXi hosts on which the StorMagic SvSAN will be deployed, as shown in the Figure 40.

Figure 40. vCenter Hosts and Clusters



10. Select a host and select the **Configuration** tab. Choose **Storage Adapters** and click Add to add the software iSCSI adapter (Figure 41).

Figure 41. Add Software iSCSI Adapter



Note: SAN connectivity is configured for the ESXi host iSCSI software adapter:

- The iSCSI port (TCP port 3260) must be enabled in the VMware firewall.
- The iSCSI software initiator must be enabled on ESXi hosts.
- Secure Shell (SSH) must be enabled on the ESXi hosts to allow VSAs to be deployed.

SvSAN uses the network ports shown in Table 5.

Table 5. Network Ports

Component or Service	Protocol or Port
Discovery	UDP 4174
XMLRPC server	TCP 8000
Inter-VSA communications	TCP 4174
SMD XML RPC	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

For additional port numbers used by SvSAN, see the following URL: http://www.stormagic.com/manual/SvSAN_5-2/en/Content/port-numbers-used-by-SvSAN.htm.

Configure ESXi Host Networking

1. Create virtual standard switches in all the ESXi hosts as described here:
 - vSwitch0
 - VMkernel Port: Management network
 - Virtual Machine Port Group: Production virtual machines and SvSAN management
 - vSwitch1
 - VMkernel Port: iSCSI traffic
 - Virtual Machine Port Group: SvSAN iSCSI and mirror traffic
 - vSwitch2
 - VMkernel Port: iSCSI traffic
 - Virtual Machine Port Group: SvSAN iSCSI and mirror traffic
 - vSwitch3
 - VMkernel Port: vMotion
2. Edit all the vSwitches and VMkernel ports used for iSCSI, mirror, and vMotion traffic and set the MTU size to 9000, as shown in Figures 42 and 43.

Figure 42. MTU Settings for vSwitch

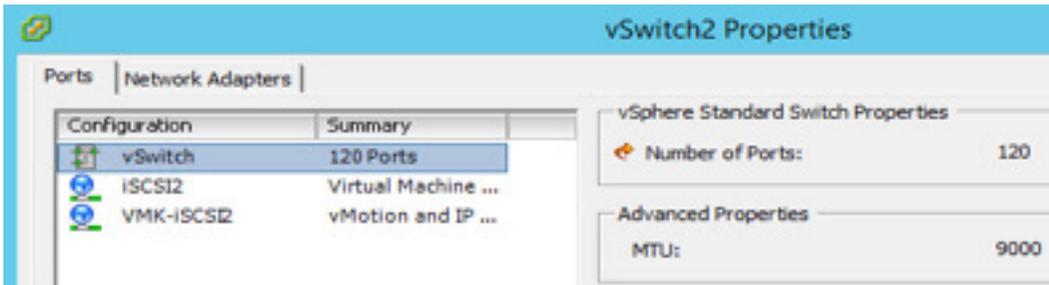


Figure 43. MTU Settings for VMkernel Port

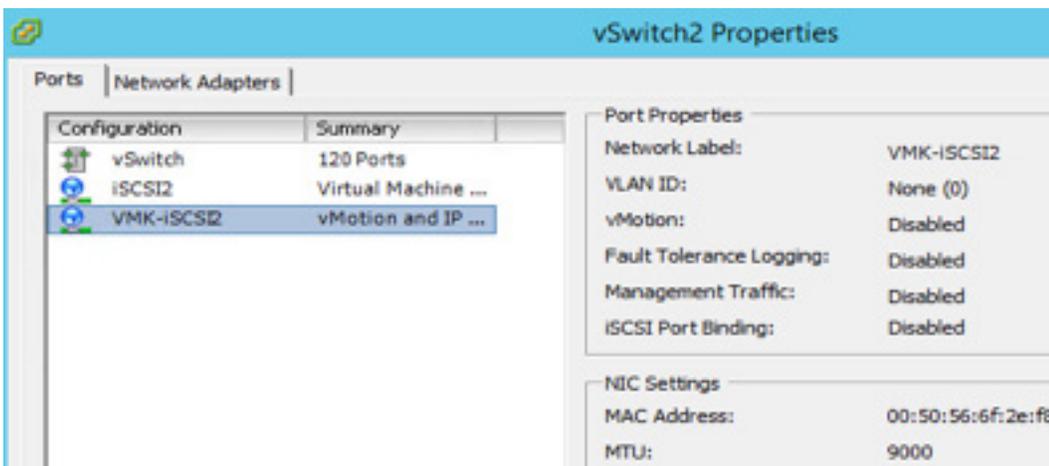
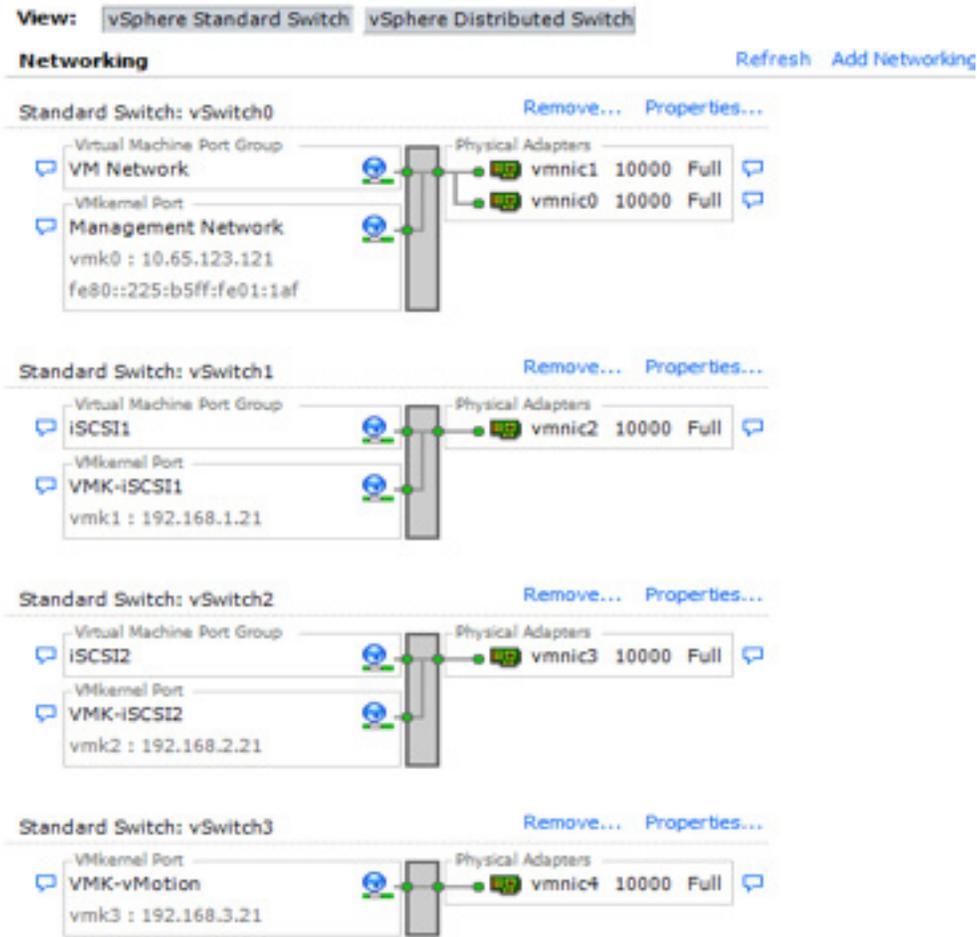


Figure 44 shows all the vSwitches and the VMkernel and virtual machine port groups created on all the ESXi hosts.

Figure 44. VMware ESXi Networking Configuration



Install StorMagic SvSAN Components on VMware vCenter Server

The SvSAN management components are installed directly on the vCenter Server. This package includes a numbers of services that are required to orchestrate VSA deployment and storage provisioning. A Cisco UCS Director module is available that includes workflows for automated VSA deployment and storage provisioning. In addition, a PowerShell module is available to allow administrators to write scripts for deployment, provisioning, and management tasks.

Table 6 lists the prerequisites for vCenter and StorMagic software Installation

Table 6. VMware vCenter and StorMagic SvSAN Version Compatibility

VMware vCenter Release	StorMagic SvSAN 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	Compatible

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

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

Installing StorMagic SvSAN Software on VMware vCenter Server

Follow the steps here to install the SvSAN software on a vCenter Server.

1. Select **Run as Administrator** and run the setup.exe file provided in the downloaded zip file.
2. Click **Next**. The end-user license agreement opens.

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

3. To accept the agreement, select the **Accept** check box and then click **Next**.
4. Click **Typical**. This option will install the Neutral Storage Service on vCenter Server, together with the plug-in components.
5. After the installation has begun, you must provide a valid vCenter username and password. If a user credentials prompt does not appear, right-click setup.exe and then select **Run as Administrator**. If the installation appears to hang, check that the pop-up box is not hidden behind another window. Finish the setup process after the installation process is complete.

Installing StorMagic PowerShell Toolkit

There are no strict requirements for installing the StorMagic PowerShell toolkit. You can install it on the vCenter server or on a scripting server.

1. To install the StorMagic PowerShell Toolkit, run the file setupStorMagicPowerShellToolkit.exe. The StorMagic PowerShell Toolkit Setup wizard opens.
2. Complete the wizard. There is only one component to install.

Note: You can uninstall the StorMagic PowerShell Toolkit from Windows by navigating to the **Control Panel** and choosing **Programs and Features** and proceeding in the usual way.

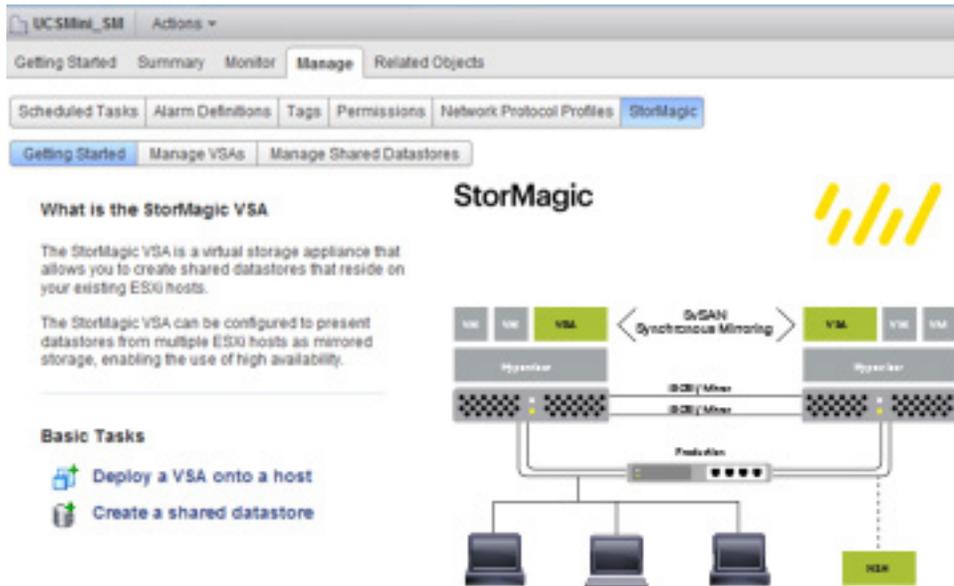
In the vSphere Web Client, a StorMagic management tab will be visible when you choose **vCenter > Hosts and Clusters**, and a data center object will appear in the inventory tree.

Deploying StorMagic SvSAN VSAs

Repeat the following steps for each VSA to be deployed.

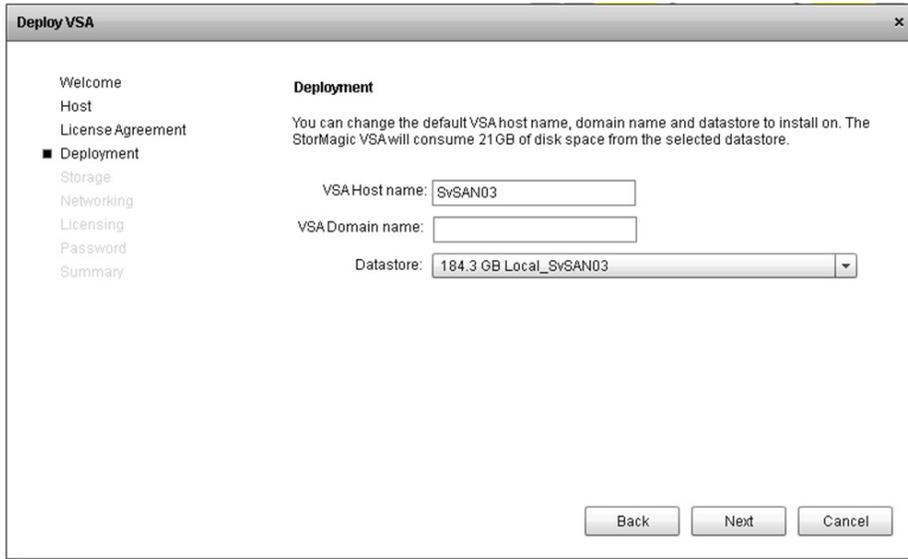
1. Open the plug-in. In vSphere Web Client, select your data center and then click **Manage > StorMagic** (Figure 45).

Figure 45. VMware vSphere Web Client with StorMagic



2. Click **Deploy a VSA onto a host**. The deployment wizard opens. Click **Next**.
3. Select a host on which to deploy a VSA. Click **Next**.
4. Read and accept the terms and conditions. Click **Next**.
5. Enter a host name for the VSA (unique on your network), a domain name, and the data store on which the VSA virtual machine will reside. The data store should be on internal or direct-attached server storage. Two virtual machine disks (VMDKs) are created on this data store: a 512-MB boot disk and a 20-GB journal disk (Figure 46).

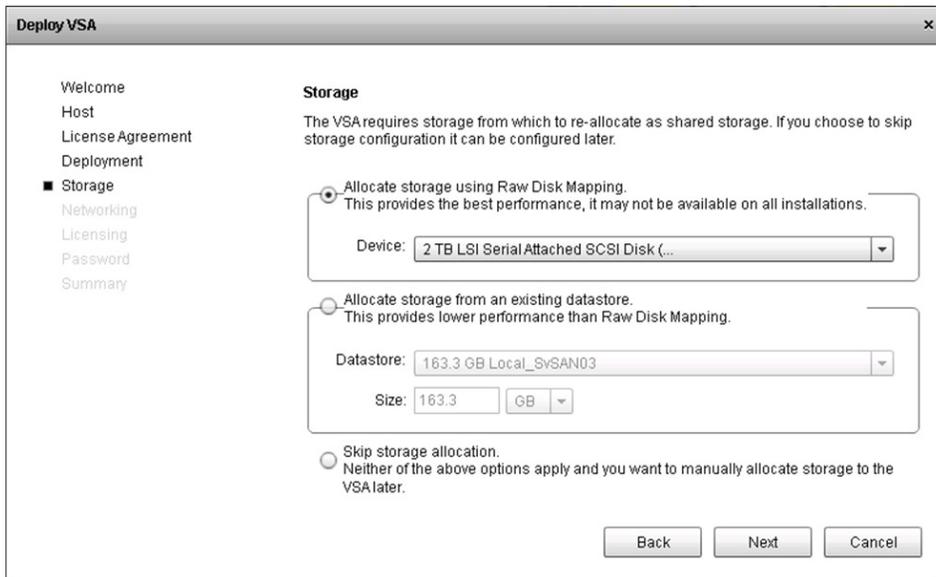
Figure 46. Deploy SvSAN VSA



- Choose the desired storage allocation technique (Figure 47). Raw device mapping (RDM) devices 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).

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

Figure 47. Deploy VSA: Storage



- If you want to allocate multiple RDM devices as a pool, select the **Advanced** options check box. If the pool you want is listed, select it. Otherwise, click **Add**. The Create Storage Pool window opens.
- Click **Next** and select the **Skip cache allocation** radio button in the **Caching** configuration window.

9. The **Networking** page can be left to acquire IP addresses using Dynamic Host Configuration Protocol (DHCP), or the addresses can be set statically. Select the Network interface and click **Configure** to select the interface traffic types. In the setup here, the VM Network interface is configured for Management, and the iSCSI1 and iSCSI2 interfaces are configured for iSCSI and mirroring traffic, as shown in the Figures 48 and 49.

Figure 48. Deploy SvSAN VSA: Networking

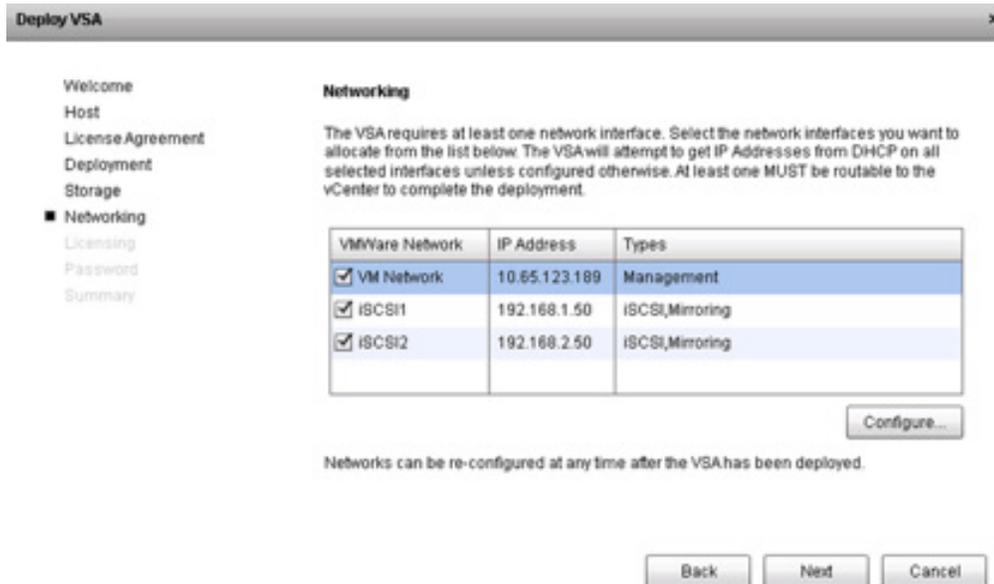
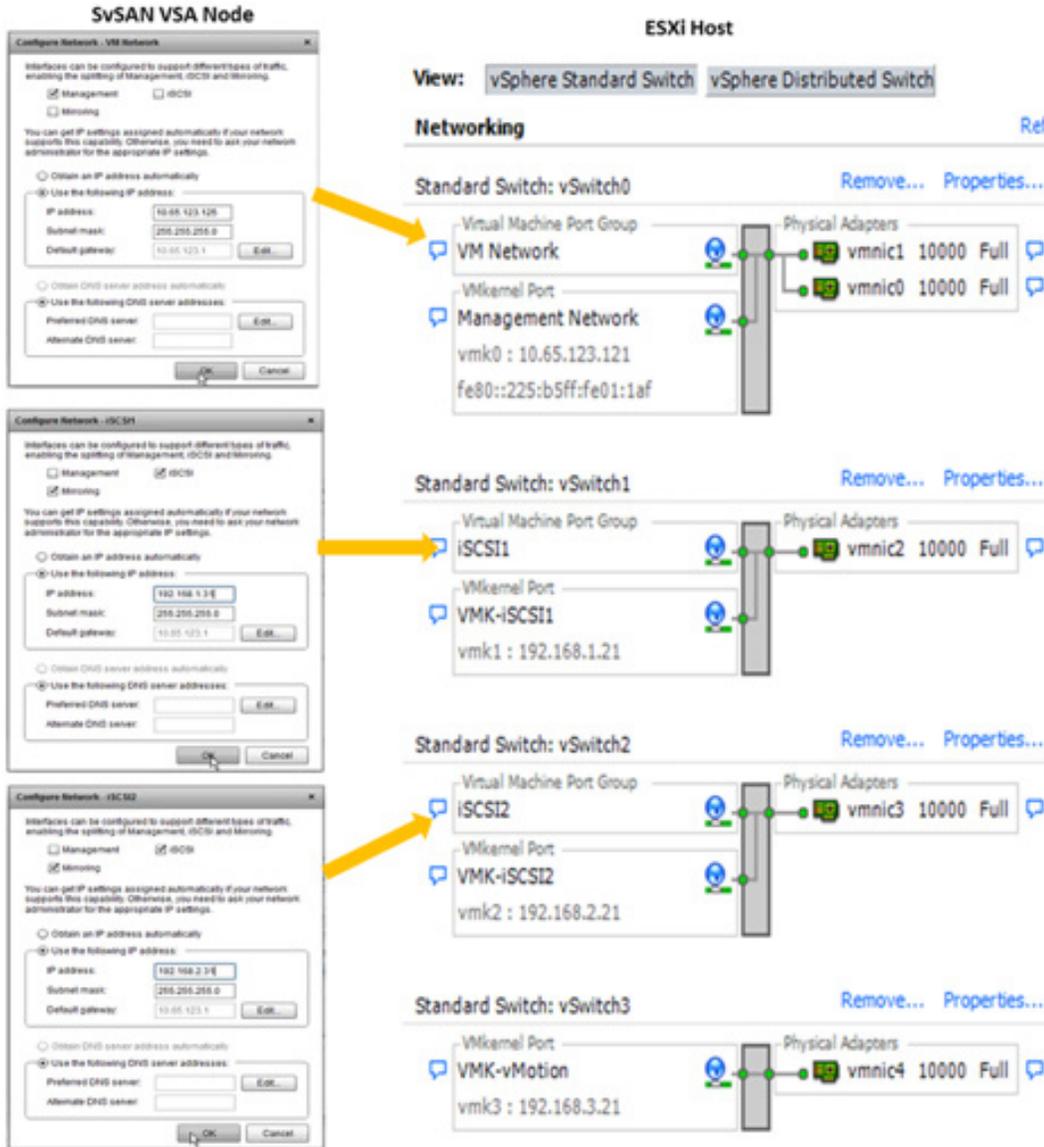


Figure 49. Deploy VSA: Networking Configuration



Note: Multiple networks are shown 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 the 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 the VSA needs to go through 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. For more information about licensing click http://www.stormagic.com/manual/SvSAN_5-2/en/index.htm#system.htm

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. The StorMagic deploy OVF task provides an overall percentage of the deployment. When this task has completed, the VSA will have booted and be operational.
 14. After the VSA has been deployed, it is available for the creation of data stores.

Note: The VSA requires at least one network interface to be routable to vCenter. Otherwise, the VSA deployment will fail. This requirement helps ensure that after the VSA virtual machine has booted, it can communicate back to the vCenter server on which the SvSAN deployment task manager resides.

Recommended StorMagic SvSAN Mirror Configurations

Overview

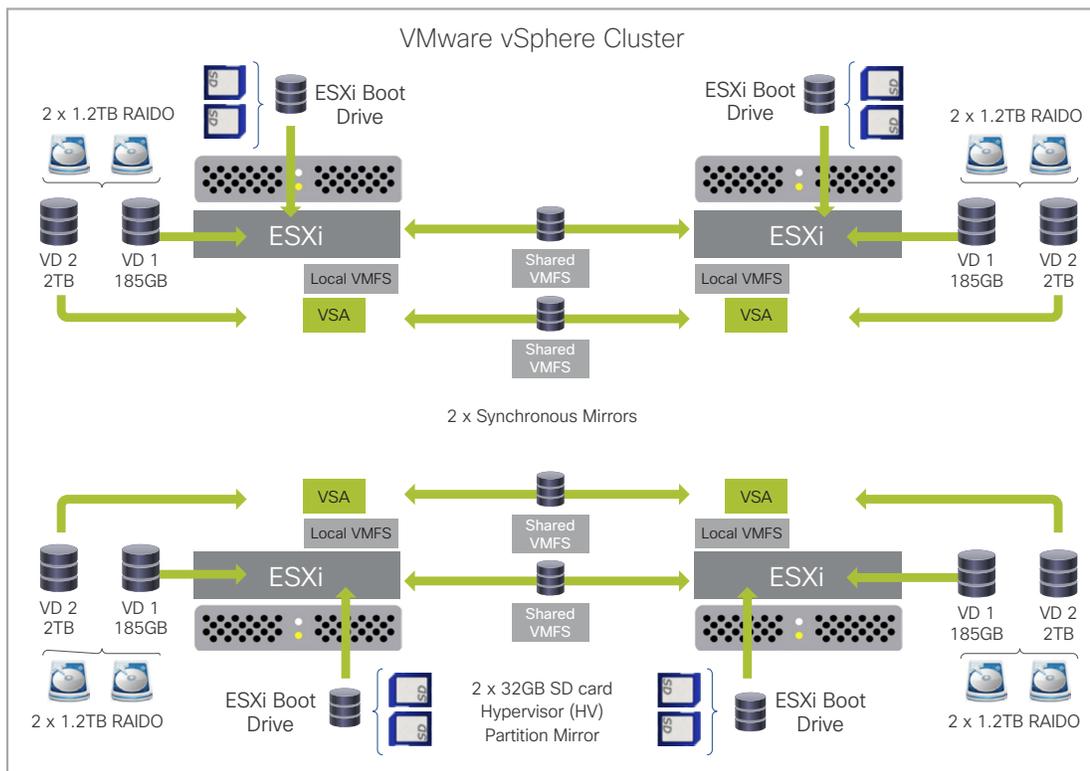
To allow the VSA storage to dynamically scale with odd numbers of Cisco UCS blades, you should configure each VSA to divide its pooled storage into two iSCSI mirrors. This configuration also adheres to VMware's default storage heartbeat policy, which requires a minimum of two data stores when deploying a two-blade configuration.

Mirrored Targets

A mirrored target is a target that is mirrored between a pair of VSAs, with each VSA hosting one side of the mirror (called a mirror plex). Therefore, there are two identical copies of the target data when the system is running normally: one on each VSA. Any data sent to one plex is automatically copied to the other. This copy operation is performed synchronously, meaning that if an initiator sends data to one plex, the VSA does not acknowledge the request until the data is present on both plexes. A synchronous mirror can be used to provide highly available storage. An initiator can access any plex at any time, and if one plex fails. The initiator can continue without interruption by using the other plex.

As previously mentioned, the pool of storage on each VSA should be divided into two iSCSI mirrors (Figure 50). This approach allows dynamic storage migration when you add blades and VSAs to the cluster. This configuration enhances the reliability of the solution in the event of multiple blade failures, and reduces the possibility of any storage contention, by spreading multiple virtual machines across mirrored targets.

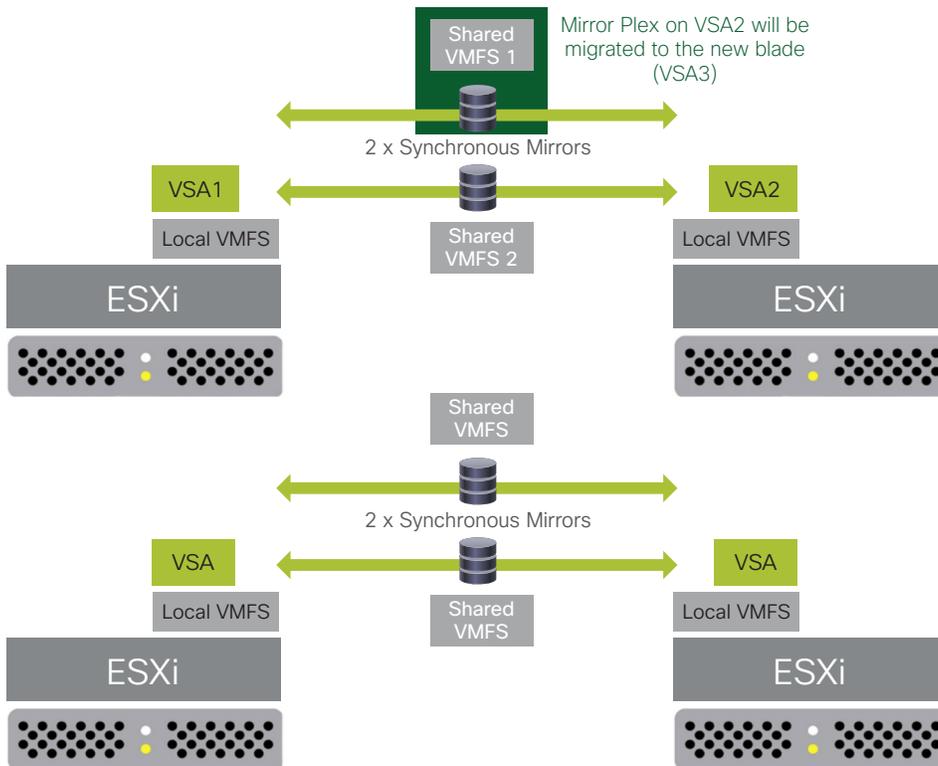
Figure 50. Four-Blade Sample Mirror Configuration



Cisco UCS Blade Server Expansion and StorMagic SvSAN Target Migration

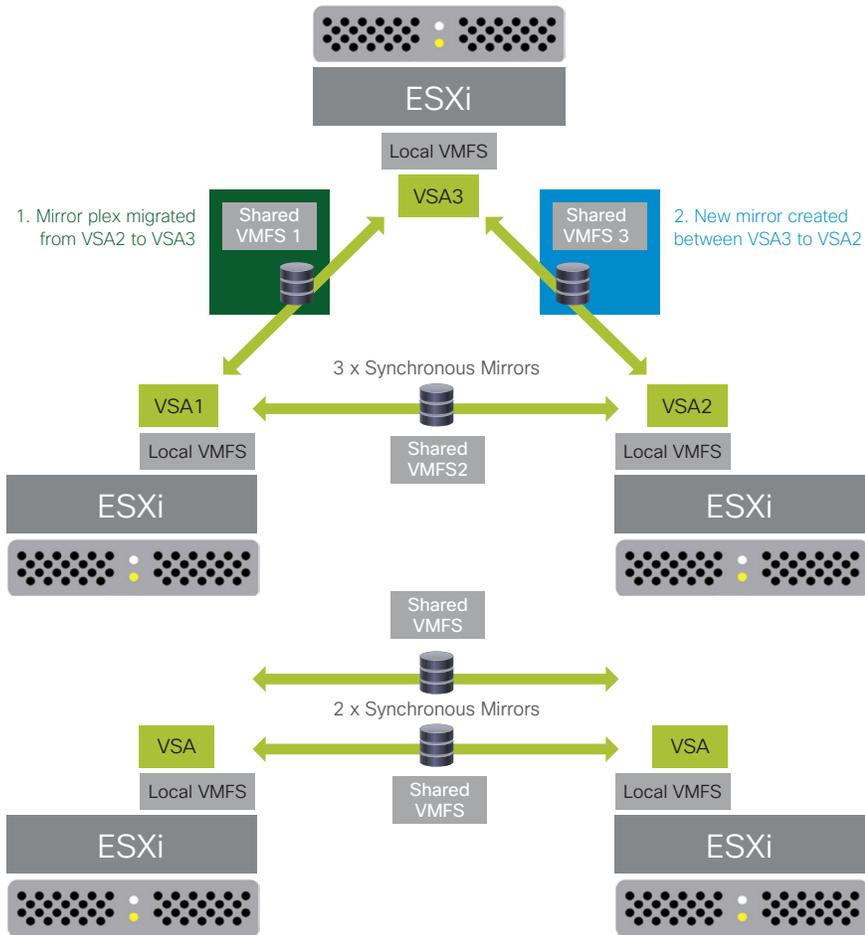
Implementing two mirrors across blade servers allows dynamic target migration when you add blades: for example, when you add a fifth blade to the four-blade configuration discussed previously. You can use SvSAN management tools to automate this process, migrating plexes from two blades to other blades. This approach helps ensure that storage I/O is spread evenly across all the Cisco UCS blade servers (Figure 51).

Figure 51. Blade Server Expansion and StorMagic SvSAN Target Migration



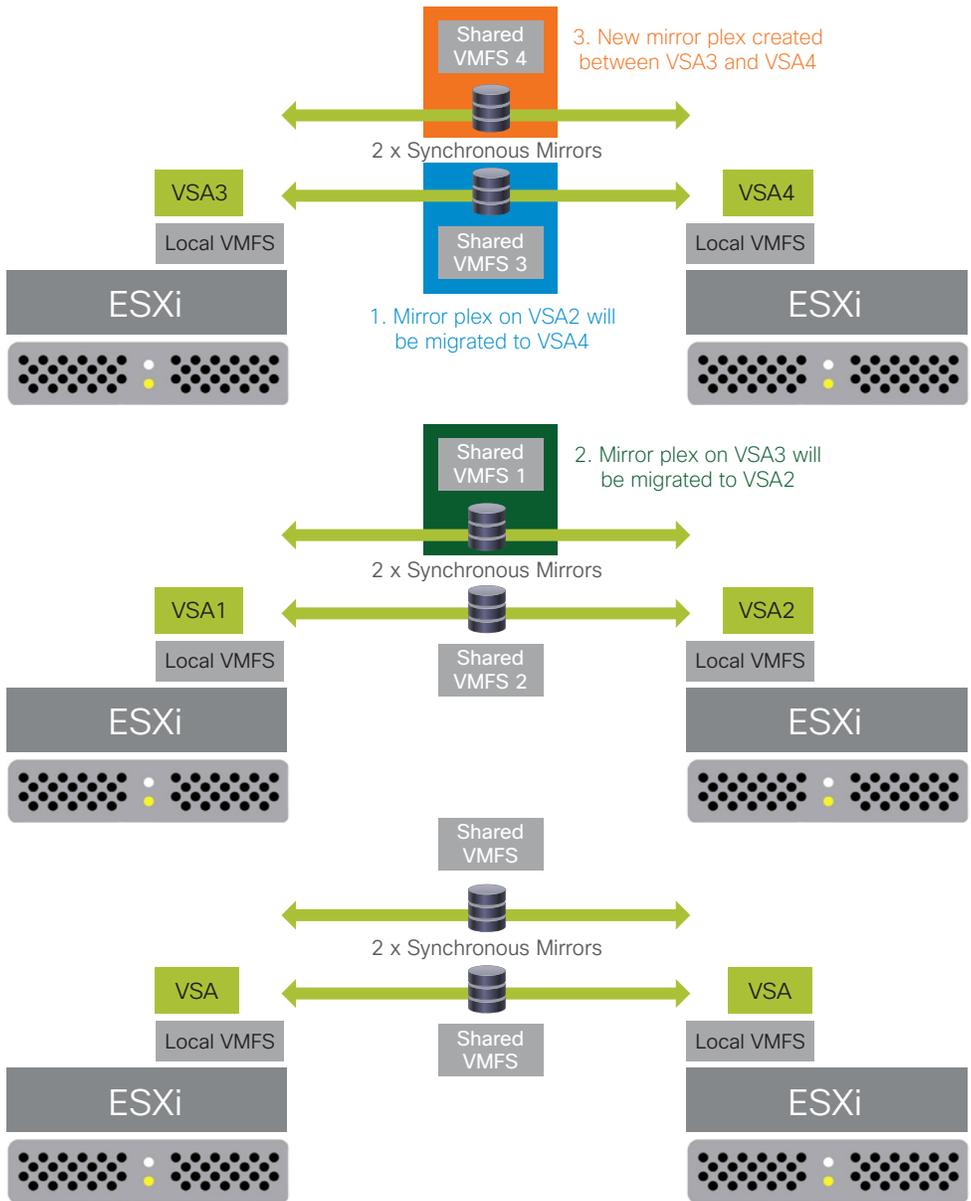
The mirror plex on VSA2 that houses VMFS1 will be migrated to the newly introduced blade VSA (Figure 52). The storage will remain online throughout this operation, but additional I/O will be observed as the plex data is written to the new VSA.

Figure 52. Five Blade Servers: Sample Mirror Configuration After Migration



After the plex migration is complete, a new mirror will be created between the new VSA and the VSA from which the plex was migrated, loading the storage evenly across the three blades. The other two blades in the environment remain unchanged (Figure 53).

Figure 53. Six Blade Servers: Sample Mirror Configuration



When you add a sixth blade to the solution, you can apply the same procedure without taking the storage offline. You can apply this process to N number of blades introduced to the solution.

Data Store Spanning Using VMware Extents

Organizations can aggregate multiple iSCSI mirrored volumes in a single data store by using VMware extents. Using StorMagic with the Cisco UCS Director module, vSphere plug-in, and PowerShell scripting module, the aggregation process can be automated. Aggregation allows the administrator to see all mirrored iSCSI volumes in all the blades as a single data store, which may be required if the virtual machines exceed the capacity of a single mirror. Note, however, that this configuration may introduce storage contention. As a performance best practice, data stores should be provisioned on single mirrors, not by using extents.

For more information, see http://www.stormagic.com/manual/SvSAN_5-3/en/Content/datastore-create-vs.htm#extend-spanned-mirrored-datastore-plugin.

Creating a Shared Data Store with Mirrored Storage

1. In vSphere Web Client, select your data center and navigate to **Manage > StorMagic**. The plug-in opens.
2. Click **Create a shared datastore**. The Create Datastore wizard opens. Click **Next**.
3. Name the data store.
4. Set the provisioned size. To set the size so that all the available space is used, select **Use all**. To divide this storage equally across the two mirrors, set the size to **1023.99 GB**. This setting will divide the available 1.99 TB equally across two mirrors.
5. To create mirrored storage, select two VSAs.
6. By default, the pool to be used on each VSA is selected automatically (Figure 54). To specify the pool yourself, select the **Advanced options** check box.

Figure 54. Create Shared Data Store with Mirrored Storage

Create datastore

Welcome

- Create datastore
- Mirroring
- Hosts
- Credentials
- Summary

Create datastore

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

Datastore:

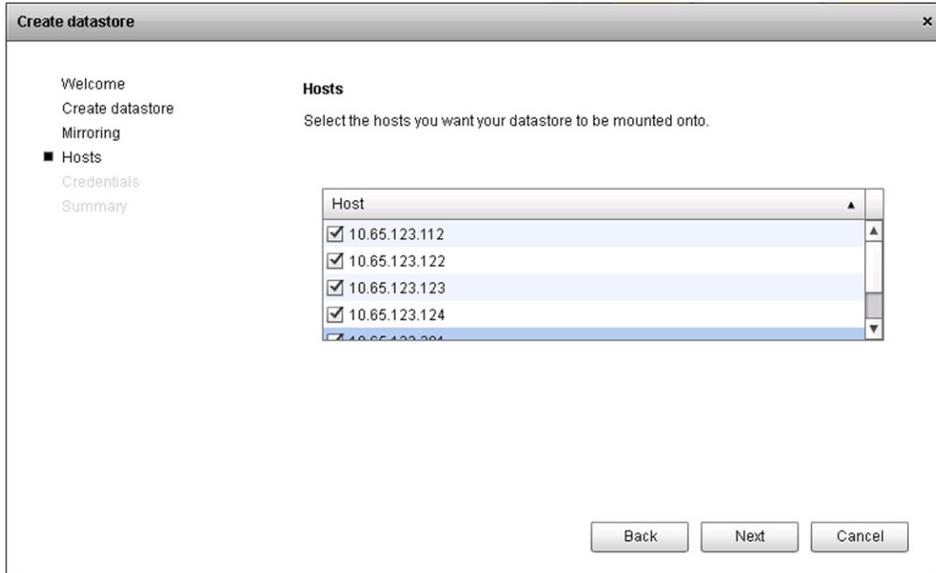
Size: Available Space: 1.99 TB

VSA	Free
<input checked="" type="checkbox"/> SvSAN01	1.99 TB
<input checked="" type="checkbox"/> SvSAN02	1.99 TB
<input type="checkbox"/> SvSAN03	1.99 TB
<input type="checkbox"/> SvSAN04	1.99 TB

Advanced options

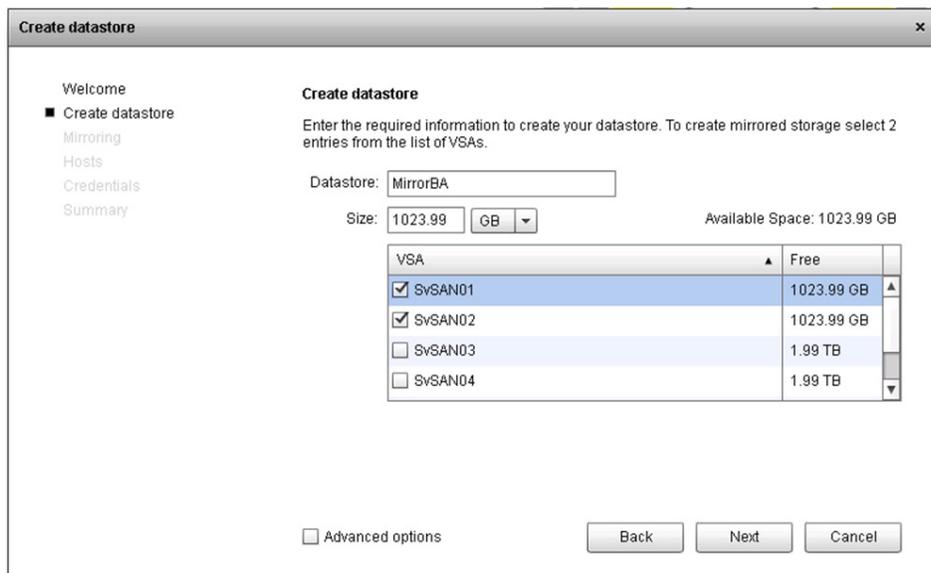
7. Click **Next**.
8. Select the neutral storage host for mirroring. Select one of the other VSAs running on another blade. The Advanced option lets you select the Up isolation policy, which does not require an NSH but is not a best practice. See http://www.stormagic.com/manual/SvSAN_5-2/en/index.htm#mirror-isolation-policy.htm for more information about isolation policies. Click **Next**.
9. Optionally, enable caching on the data store. This option is available only if your license includes caching and a cache device was enabled when the VSA was deployed. If you enable caching, default cache settings are used; you can modify these later using the web GUI. Click **Next**.
10. Select the hosts that are to have access to the SvSAN data store (Figure 55). Click **Next**.

Figure 55. Select Hosts to Mount Data Store



11. The first time you create a data store, the VSAs must authenticate with their hosts. 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 creation of the data store. 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.
13. When a mirror is created, a full synchronization is performed to help ensure that both sides of the mirror are synchronized.
14. Repeat the steps in this section to create a second data store on the VSAs or another data store using the other two SvSAN nodes, as shown in Figure 56.

Figure 56. Create Mirrored Storage



Managing VSAs

To check the status of your VSAs, in the plug-in, click **Manage VSAs** (Figure 57).

Figure 57. Manage VSAs

The screenshot displays the 'Manage VSAs' interface. At the top, there are navigation tabs: 'Getting Started', 'Summary', 'Monitor', 'Manage', and 'Related Objects'. Under 'Manage', there are sub-tabs: 'Scheduled Tasks', 'Alarm Definitions', 'Tags', 'Permissions', 'Network Protocol Profiles', and 'StorMagic'. Below these, there are further sub-tabs: 'Getting Started', 'Manage VSAs', and 'Manage Shared Datastores'. The main content area is titled 'Manage VSAs' and includes a brief instruction: '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
SvSAN01	10.65.123.125	Normal
SvSAN02	10.65.123.126	Normal
SvSAN03	10.65.123.127	Normal
SvSAN04	10.65.123.128	Normal

Buttons: Restore... Log In... Refresh

Host details for SvSAN01:

- Hostname: SvSAN01
- Management URL: <https://10.65.123.125/>
- Status: Normal
- SAN Name: Unconfigured
- System Discovery ID: 2620064CD826
- Serial Number: C9DAC89222E8
- Firmware Version: 5.2.3439
- License Key: Licensed

Capacity and Usage:

- Capacity: 1.99 TB
- Used: 1.92 TB
- Free: 73.98 GB

This action displays a list of 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 system firmware version.

You can monitor the mirror synchronization status by managing the VSAs directly. From the StorMagic plug-in, click **Manage VSAs**, select the VSA you want to manage, and click the management URL. This process launches a new tab directly connected to the VSA management interface. The default username is **admin**, and the password is the password supplied at VSA deployment.

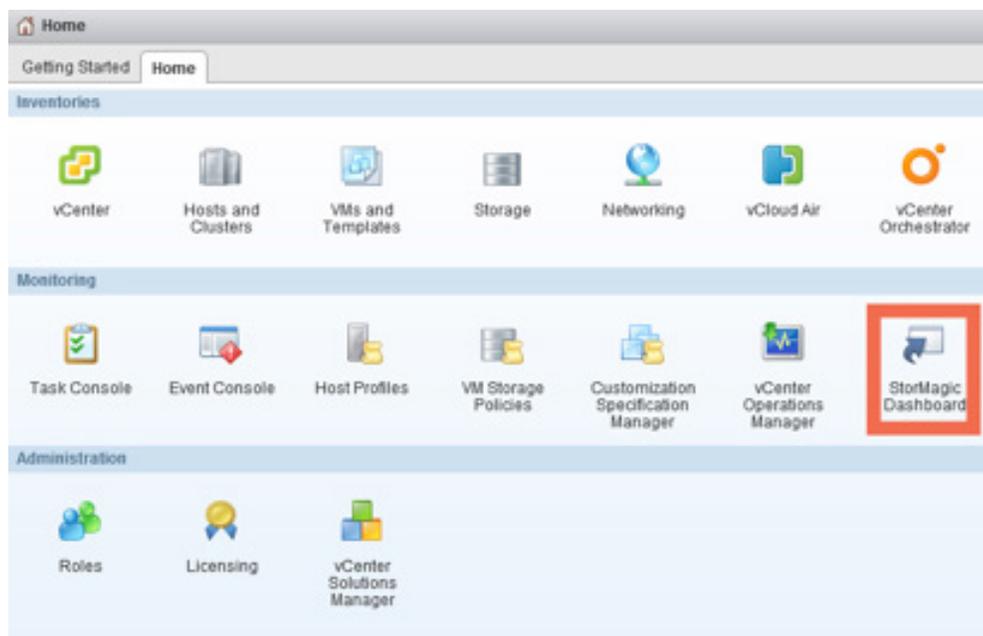
You can also use the VSA dashboard to monitor VSAs in real time. The VSA state reflects the current system status of the VSA. For example, if a mirror is not synchronized, the VSA system status will be listed as Warning because the mirror is currently degraded. When the mirror synchronizes, the VSA system status will change to Healthy, showing that the VSA is in an optimal running state (Figure 58).

Figure 58. Manage VSAs

State	Name	IPs	Serial	Host	Last Seen
Healthy	SvSAN01	192.168.2.31,192.168.1.31,10.65.123.125,10.65.123.126	2620064CD826	10.65.123.111	Fri Oct 16 19:54:21 IST 20
Healthy	SvSAN02	10.65.123.126,192.168.1.32,10.65.123.127,10.65.123.128	9D4E4DEA9DBE	10.65.123.122	Fri Oct 16 19:55:21 IST 20
Healthy	SvSAN03	10.65.123.127,192.168.1.33,10.65.123.128,192.168.1.34	C8E75F924642	10.65.123.123	Fri Oct 16 19:53:21 IST 20
Healthy	SvSAN04	10.65.123.128,192.168.1.34,10.65.123.129,10.65.123.130	8D9B7BAF8A39	10.65.123.124	Fri Oct 16 19:58:21 IST 20

You can access the dashboard from the vSphere Web Client homepage (Figure 59).

Figure 59. StorMagic Dashboard in vSphere Web Client



Configure Jumbo Frames on StorMagic SvSAN Network Interfaces

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

1. Log in to the vSphere Web Client and select **Hosts and Clusters**.
2. Select the data center in which the VSAs are deployed. Choose **Manage > StorMagic > Manage VSAs**.
3. Select the VSA. A management URL will appear. Click this URL.
4. A new tab opens 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, click **Edit**.
7. Edit the MTU size to the desired size to match the rest of the environment network configuration. Then click **Apply** (Figure 60).

Figure 60. Setting the MTU

Edit Network Device

MAC	00:50:56:A9:68:C9 (vSwitch0:VM Network)
Device Name	eth0
Driver	vmxnet3
MTU	<input type="text" value="1500"/>
TXQLEN	<input type="text" value="1000"/>
LRO	<input checked="" type="checkbox"/>
TSO	<input checked="" type="checkbox"/>

Network Devices

MAC	Carrier	Speed	Driver	MTU	TXQLEN	LRO	TSO
00:50:56:A9:3C:00 (vSwitch0:VM Network)	✓	10000 Mb/s	vmxnet3	1500	1000	✓	✓
00:50:56:A9:1B:19 (vSwitch1:iSCSI1)	✓	10000 Mb/s	vmxnet3	9000	1000	✓	✓
00:50:56:A9:5F:74 (vSwitch2:iSCSI2)	✓	10000 Mb/s	vmxnet3	9000	1000	✓	✓

Managing Shared Data Stores

To manage shared data stores, in the plug-in, click **Manage Shared Datastores** (Figure 61).

Figure 61. Manage Shared Datastores

Manage Shared Datastores

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

Datastore	Hosts
MirrorAB	10.65.123.123, 10.65.123.124, 10.65.123.121, 10.65.123.122
MirrorBA	10.65.123.123, 10.65.123.124, 10.65.123.121, 10.65.123.122
MirrorCD	10.65.123.123, 10.65.123.124, 10.65.123.121, 10.65.123.122
MirrorDC	10.65.123.123, 10.65.123.124, 10.65.123.121, 10.65.123.122

Mount... Grow... Migrate... Create... Destroy Refresh

ESXi Host: 10.65.123.123 Capacity: 1023.75 GB
Device: eui.00033989cb390002 Used: 877.19 GB
Free: 146.56 GB
Multipath Policy: Fixed (VMware)

Path	VSA	State
iqn.2006-06.com.stormagic:89cb390200000014.m0mirrorab:192.168.1.31:3260	SvSAN01	Active (I/O)
iqn.2006-06.com.stormagic:89cb390200000014.m0mirrorab:192.168.1.32:3260	Unknown	Active
iqn.2006-06.com.stormagic:89cb390200000014.m0mirrorab:192.168.2.32:3260	Unknown	Active
iqn.2006-06.com.stormagic:89cb390200000014.m0mirrorab:192.168.2.31:3260	SvSAN01	Active

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.

Performance Characterization

The Cisco UCS Mini with Cisco UCS B200 M4 Blade Servers plus StorMagic software hyperconverged solution is designed to handle branch-office and remote-office workloads that are more computation or network focused than I/O focused.

This section describes the types of workloads most suitable for this solution and performance test results from a four-blade configuration (starter kit) and a six-blade configuration (starter kit plus a two-blade expansion pack).

Table 7 lists the workloads that are suitable to run on a Cisco UCS Mini and Cisco UCS B200 M4 plus StorMagic solution. The application profiles and I/O patterns listed here do not represent a complete list of all workloads that can be made to run; however, the majority of the workloads are expected to be similar.

Table 7. Application I/O Profile: Cisco UCS Mini and Cisco UCS B200 M4 plus StorMagic Solution

Application Profile	Access Mode	Read:Write Ratio	Block Size	Performance Metric
Web front end	Random	100:0	8 KB	IOPS and response time (milliseconds [ms])
Online transaction processing (OLTP)	Random	80:20	8 KB	IOPS and response time (ms)
Decision support system, business intelligence, and video on demand (VoD)	Sequential	100:0 and 0:100	256/512 KB	Bandwidth or transfer rate (MBps)

Each disk provides approximately 140 to 150 IOPS with a response time of less than 20 milliseconds for 100 percent random-read workloads when the virtual machine volumes are distributed across the entire disk. However, greater variations are noticed in IOPS and response time when the virtual machine volumes are of smaller size, as a result of the mechanical head movement between the inner and outer cylinders of the hard disk.

For the purposes of this testing, 100 GB of hard-disk space was allotted to each virtual machine (which is typical in remote-office and branch-office deployments, with four vCPUs, 6 GB of RAM, and 100 GB of storage on each mirrored volume).

Workload Scenarios

The following two scenarios were tested and validated for this solution:

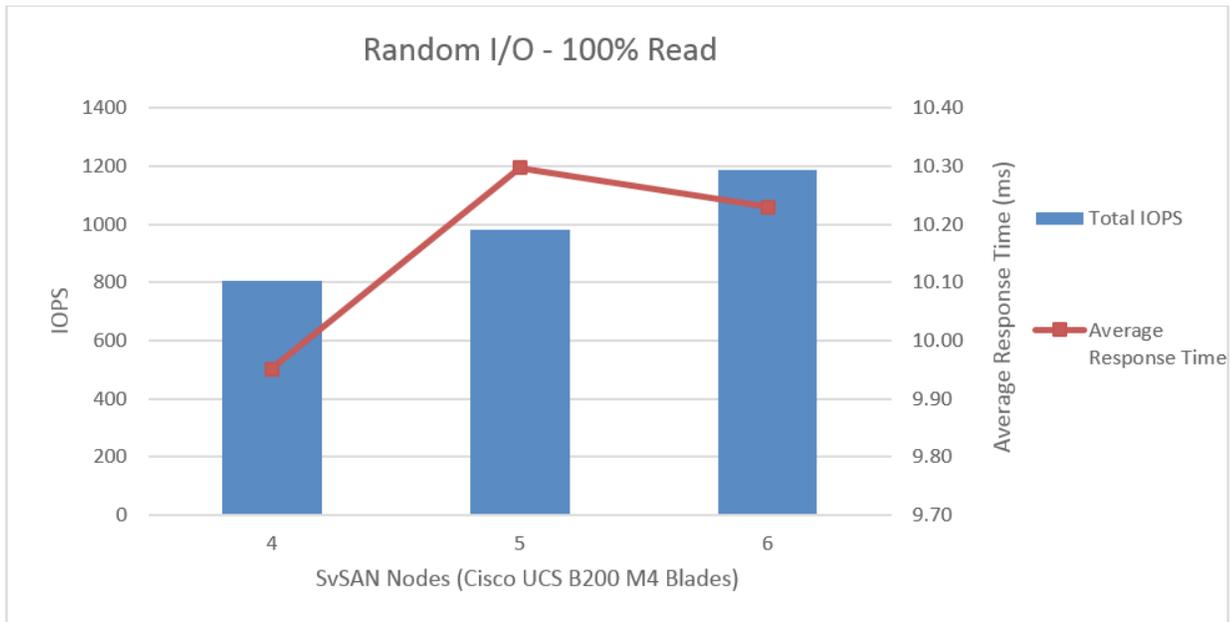
- Scenario A, with one virtual machine per host (total of four virtual machines [starter kit] or six virtual machines [starter kit plus 2-blade expansion pack]): This scenario sets the baseline for I/O performance, because no resource contention occurs at the computing and network levels. The workload generated is purely I/O intensive and helps identify the maximum achievable throughput and response time for the system.
- Scenario B, with three virtual machines per host (total of 12 virtual machines [starter kit] or 18 virtual machines [starter kit plus 2-blade expansion pack]): This scenario helps estimate the achievable throughput and response time under load. It uses three virtual machines per mirror to illustrate scalability.

In both scenarios, the access profiles listed in Table 7 were configured in the Iometer tool with a queue depth of 2. The tests used an 8-KB block size because most applications use an 8-KB block size. A 4-KB block size was also tested, and the results matched those for the 8-KB block size; hence, the graphs show the 8-KB block size.

Scenario A: One Virtual Machine per Blade

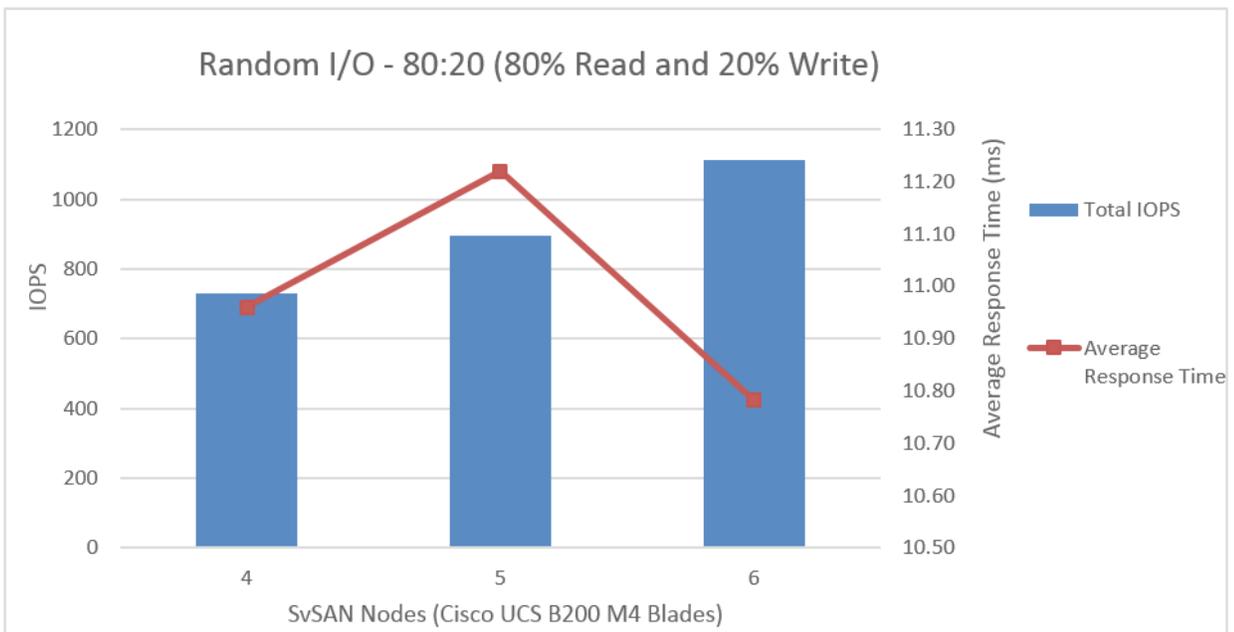
Figure 62 shows that the system is capable of delivering up to about 1200 IOPS (random read) with a response time of less than 20 milliseconds, which aligns with the expected disk performance

Figure 62. 8-KB Random I/O on One Virtual Machine per Blade: 100 Percent Random Read



For a random read and write workload, IOPS scales linearly with additional nodes with acceptable response times (less than 20 milliseconds). A six-node system can deliver up to 1113 IOPS with a response time of 10.78 milliseconds (Figure 63).

Figure 63. 8-KB Random I/O on One Virtual Machine per Blade: 80 Percent Read and 20 Percent Write



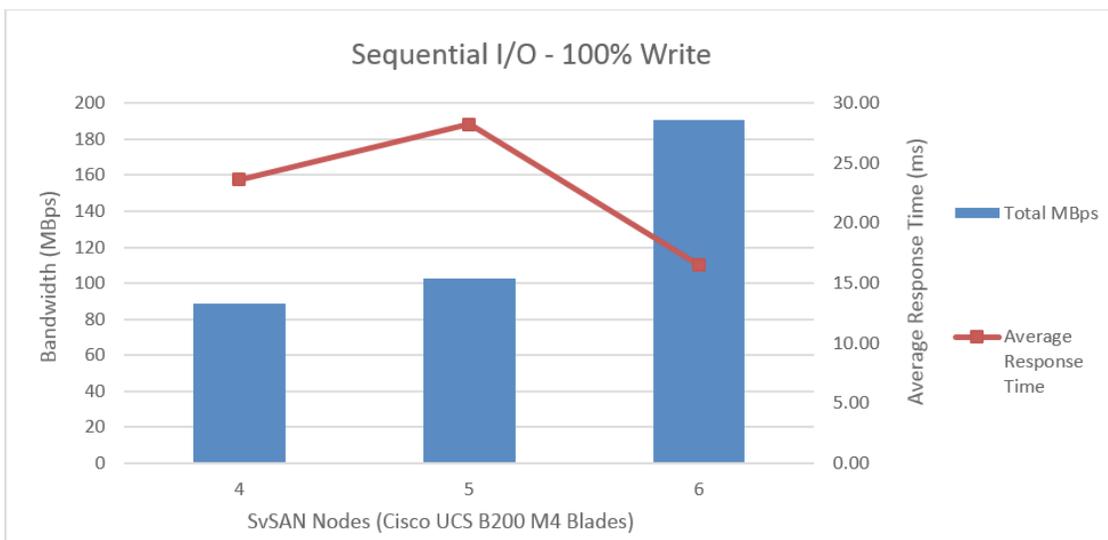
Sequential read I/O also demonstrates that the bandwidth scales linearly with additional nodes with acceptable response times (less than 20 milliseconds). In a six-node setup, a peak bandwidth of 605 MBps was reached without exceeding acceptable response times (Figure 64).

Figure 64. 256-KB Sequential I/O on One Virtual Machine per Blade: 100 Percent Read



Sequential write I/O also demonstrates that the bandwidth scales linearly with additional nodes with acceptable response times. Sequential write bandwidth reached 190 MBps (Figure 65).

Figure 65. 256-KB Sequential I/O on One Virtual Machine per Blade: 100 Percent Write

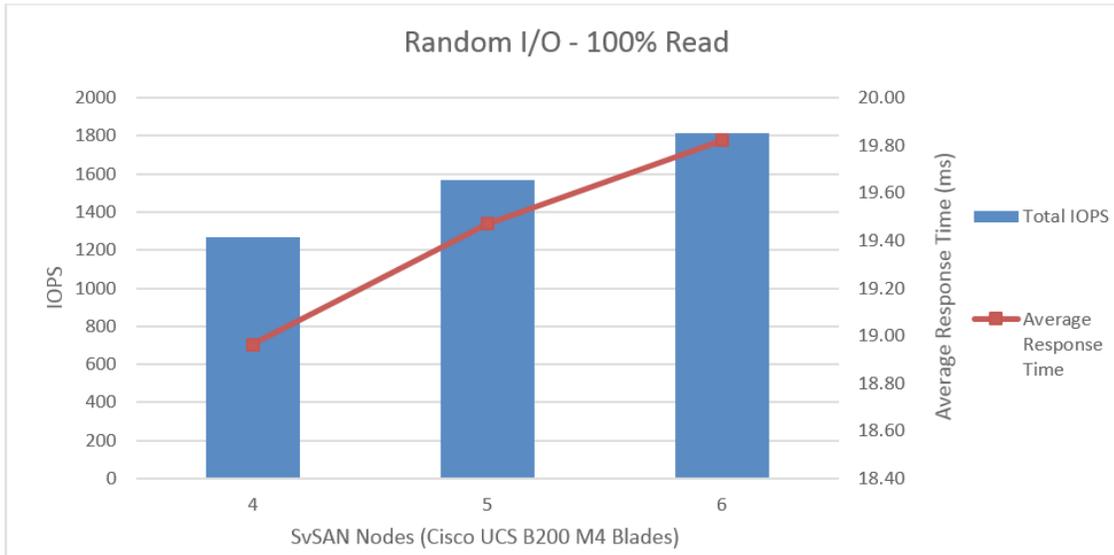


The response time variations (the fourth and fifth nodes have longer response times than the sixth node) are the result of the inherent random access levels at the disk, although write operations are sent sequentially at the application level, so caution must be exercised when running 100 percent sequential write workloads.

Scenario B: Three Virtual Machines per Blade

With three virtual machines loaded and 8-KB random read I/O, the system is capable of delivering up to 1816 IOPS within 20 milliseconds, which is an acceptable response time (Figure 66).

Figure 66. 8-KB Random I/O on Three Virtual Machines per Blade: 100 Percent Read



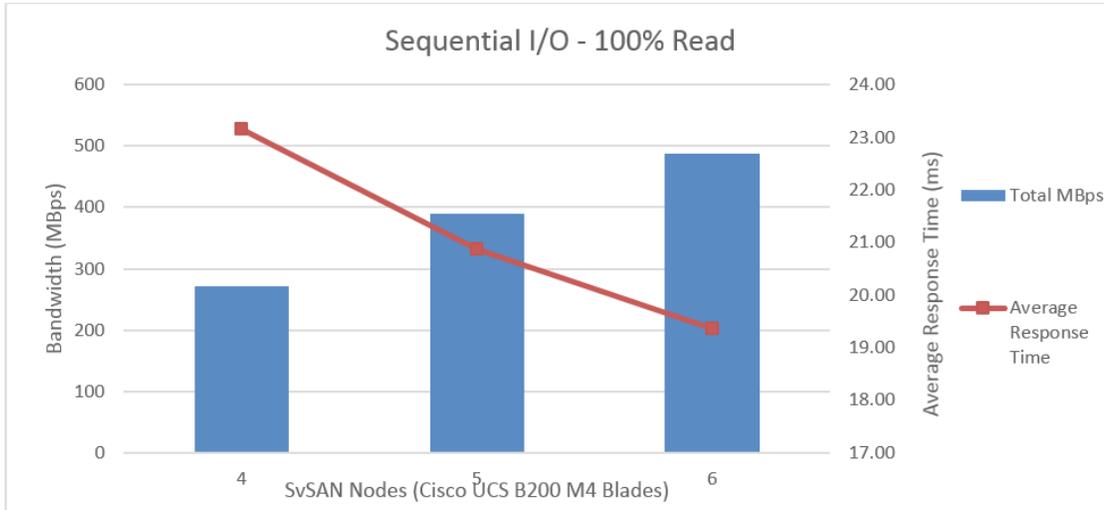
With three virtual machines loaded and 8-KB random read and write I/O (80 percent read and 20 percent write), the system is capable of delivering up to about 1614 IOPS with a response time of 22 milliseconds. Because this workload had 20 percent write activity, the response time showed a slight increase over the generally accepted response time limit (Figure 67).

Figure 67. 8-KB Random I/O on Three Virtual Machines per Blade: 80 Percent Read and 20 Percent Write



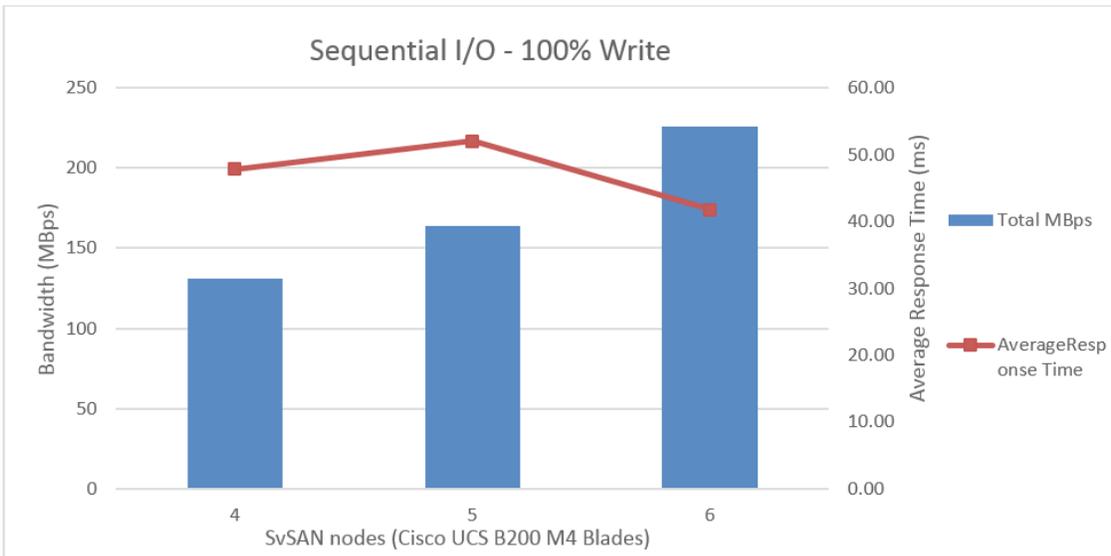
With three virtual machines loaded system, the 256-KB sequential read access profile reached a peak bandwidth of 487 MBps (Figure 68).

Figure 68. 256-KB Sequential I/O on Three Virtual Machines per Blade: 100 Percent Read



With three virtual machines loaded, the 256-KB sequential write access profile reached a peak bandwidth of 226 MBps (Figure 69). The increase in response times occurred because all the virtual machines perform 100 percent write operations.

Figure 69. 256-KB Sequential I/O on Three Virtual Machines per Blade: 100 Percent Write



This data is similar to what is seen in the case with a single virtual machine per blade (Figure 65). The response time variations (the fourth and fifth nodes have longer response times than the sixth node) are aggravated further because this scenario uses more virtual machines. Hence, caution must be exercised while running such high sequential write-intensive workloads.

Collecting Diagnostic Information for VMware ESXi and StorMagic SvSAN

Collect Diagnostic Information for VMware ESXi

VMware technical support technicians routinely request diagnostic information from you to respond to a support request. 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.

For the procedures here to obtain 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. Use the `vm-support -h` command for a list of options available on a given version of ESXi or ESX.

A compressed bundle of logs is produced and stored in a file with a `.tgz` extension in one of these locations:

- `/var/tmp/`
 - `/var/log/`
 - The current working directory
3. 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 Client. See the VMware knowledgebase article at <http://kb.vmware.com/kb/653>.

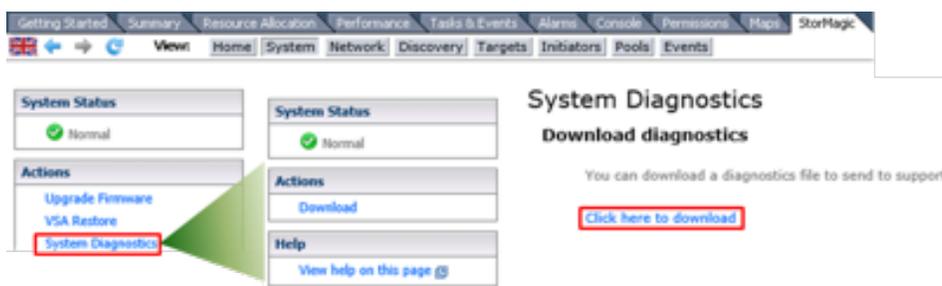
Collect System Diagnostic Information for StorMagic SvSAN

The 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, follow these steps:

1. Log in to SvSAN and click **System**.
2. Select **System Diagnostics** and click **Click here to download** (Figure 70).

Figure 70. Collect System Diagnostic Information



For More Information

- <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-mini/index.html>
- http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/3-0/b_UCSM_GUI_User_Guide_3_0/b_UCSM_GUI_User_Guide_3_0_chapter_010.html
- <http://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>
- <https://www.vmware.com/files/pdf/vsphere/VMW-WP-vSPHR-Whats-New-6-0.pdf>
- http://www.stormagic.com/manual/SvSAN_5-3/en/



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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

C11-736761-01 4/16