



NetApp Verified Architecture

FlexPod Datacenter with SolidFire All-Flash Array Add-On

NVA Deployment

Karthick Radhakrishnan, David Klem, NetApp
April 2017 | NVA-0027-DEPLOY | Version 1.0

TABLE OF CONTENTS

1	Program Summary	3
2	Solution Overview	3
2.1	Solution Technology	3
2.2	Use Case Summary	5
3	Technology Requirements	5
3.1	Hardware Components	6
3.2	Software Components	6
4	Cabling Details for SolidFire Nodes	7
5	Deployment Procedures	8
5.1	Cisco UCS Configuration	8
5.2	Cisco Nexus Switch Configuration	8
5.3	SolidFire Node Configuration	10
	Acknowledgements	31
	References	31

LIST OF TABLES

Table 1)	Hardware components	6
Table 2)	Software components.....	6
Table 3)	VLANs	6

LIST OF FIGURES

Figure 1)	FlexPod Datacenter components	4
Figure 2)	FlexPod Datacenter cabling diagram	5
Figure 3)	Cabling diagram for Cisco Nexus switch and SolidFire nodes	7

1 Program Summary

FlexPod® Datacenter is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), Cisco Nexus family of switches, and NetApp® fabric-attached storage (FAS) systems. FlexPod is an ideal platform for running a variety of virtualization hypervisors and enterprise workloads. FlexPod can be scaled up for greater performance and capacity by adding compute, network, or storage resources individually as needed. It can also be scaled out for both virtualized and nonvirtualized environments that need multiple consistent deployments by rolling out additional FlexPod stacks. FlexPod delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases.

2 Solution Overview

This solution describes the procedure for adding a SolidFire® all-flash storage system into any existing FlexPod Datacenter environment, with an emphasis on multi-tenant workloads demanding minimum performance guarantees. The hardware components included in the design include Cisco compute and networking, NetApp FAS, and SolidFire all-flash block storage system.

2.1 Solution Technology

Figure 1 shows the FlexPod Datacenter with NetApp FAS and SolidFire components and the network connections for a configuration with iSCSI-based storage. This design uses the Cisco Nexus 5000/9000 switches, Cisco UCS C-Series and B-Series servers with the Cisco UCS virtual interface card (VIC), the NetApp FAS family of storage controllers, and SolidFire storage nodes connected in a highly available design by using Cisco virtual port channels (vPCs).

Figure 1 shows the technical components of the solution, and Figure 2 shows the detailed cabling diagram.

Figure 1) FlexPod Datacenter components.

Cisco UCS
C220 M4 C-Series Servers

Cisco UCS
5108 B-Series Blade Chassis
2208XP Chassis FEX Modules
B200 M4 B-Series Blades

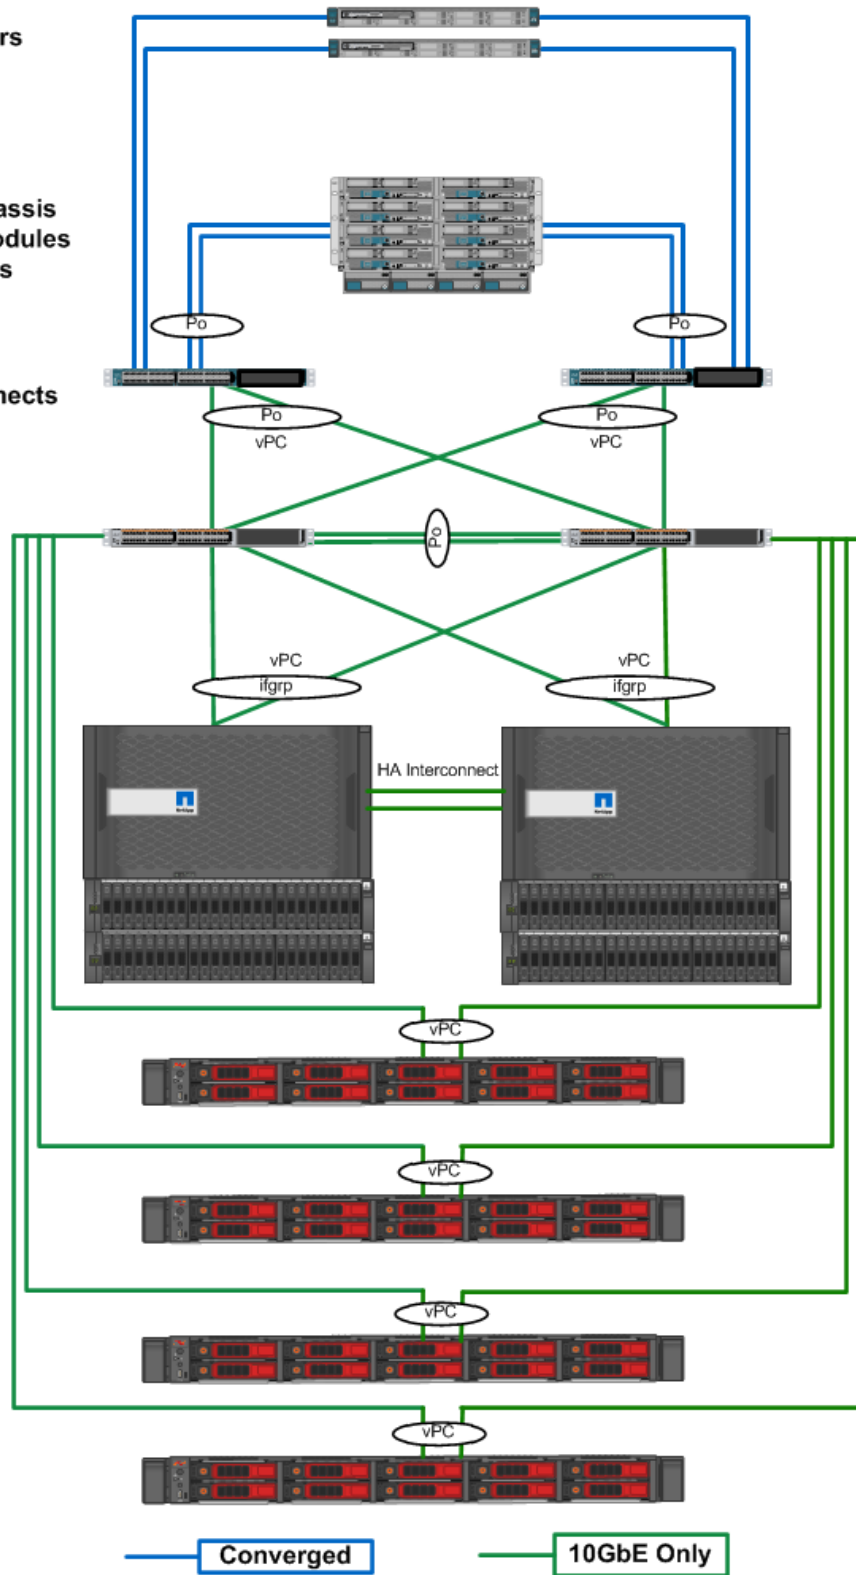
Cisco UCS
6248UP Fabric Interconnects

Cisco Nexus
9372PX or 5548UP
Switches

NetApp FAS8040
Storage Controllers

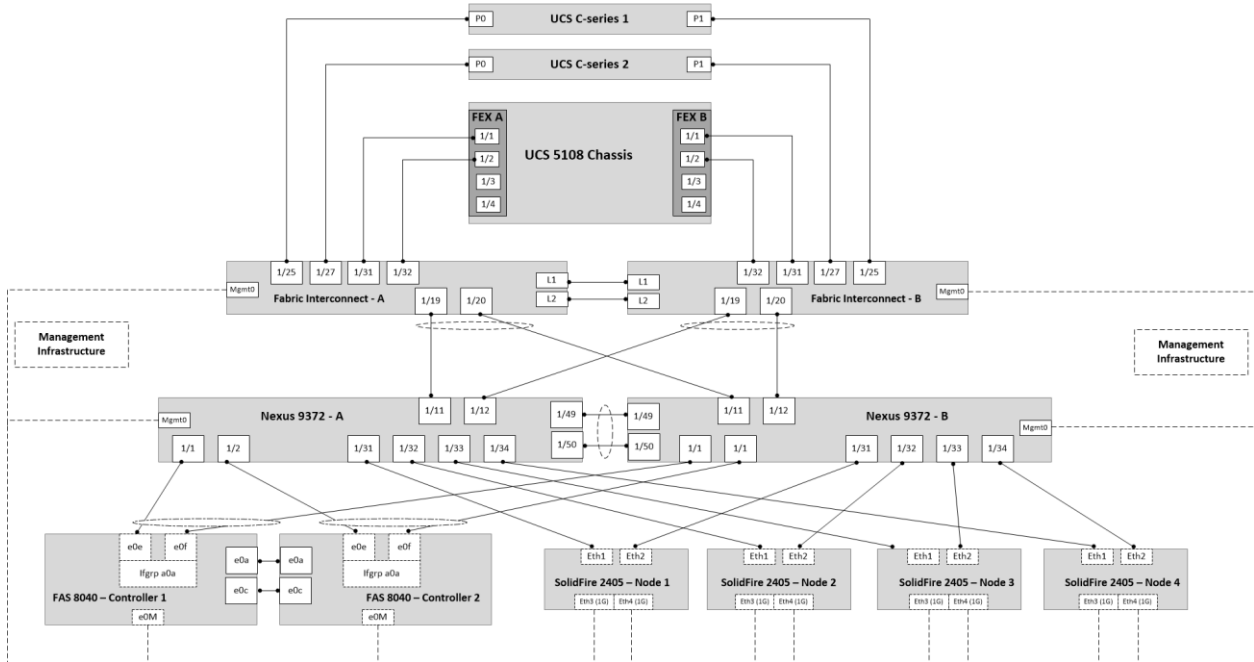
NetApp
DS2246 Disk Shelves

SolidFire 2405



Note: The SF2405 1GbE ports are connected to the management switch.

Figure 2) FlexPod Datacenter cabling diagram.



2.2 Use Case Summary

The primary use case for this solution is to provide a scale-out storage option for block-based iSCSI workloads within the FlexPod Datacenter environment using SolidFire all-flash array. This use case also provides multitenant workloads with guaranteed, minimum performance service-level agreements (SLAs) that are required through SolidFire quality of service (QoS).

This document assumes the FlexPod Datacenter environment is configured as per any of our [CVD](#) best practices and describes only the deployment procedures and best practices to add a SolidFire all-flash block storage system in an existing FlexPod Datacenter environment. The server operating system is VMware vSphere ESXi, and a VMware vCenter Server is installed to manage the ESXi instances. The document leverages any existing FlexPod Datacenter environment for boot and existing workloads.

3 Technology Requirements

Cisco, NetApp, and VMware have interoperability matrixes that must be referenced to determine support for any specific implementation of FlexPod. The “FlexPod Datacenter Technical Specifications” document details the hardware and configuration requirements for FlexPod.

For more information, see the following links:

- [NetApp Interoperability Matrix Tool \(IMT\)](#)
- [Cisco UCS Hardware and Software Interoperability Tool](#)
- [TR-4036: FlexPod Datacenter Technical Specifications](#)

3.1 Hardware Components

Table 1 lists the hardware components used for this solution validation. However, any supported FlexPod hardware component or SolidFire node component can be used in this solution.

Table 1) Hardware components.

Layer	Hardware	Quantity
Compute	Cisco UCS 5108 chassis	1
	Cisco UCS B200 M4 blades with VIC 1240	2
	Cisco UCS C220 M4 rack-mount servers	2
Network	Cisco Nexus 9372PX	2
Storage	FAS8040	HA pair
	Disk shelf: DS4246 with 24x900GB disks	2
	SolidFire 2405	4

3.2 Software Components

Table 2 lists the software components used for this solution validation. However, any supported software component can be used in this solution.

Table 2) Software components.

Layer	Software	Version
Compute	Cisco UCS infrastructure software bundle	2.2(6)
	Cisco UCS server bundle	2.2(6)
Network	Cisco Nexus switch software (system and kick start)	NX-OS 7.0(3)I1(3)
Storage	NetApp clustered Data ONTAP®	8.3.2
	SolidFire Element OS	8.4
Hypervisor	VMware vSphere ESXi	6.0
	VMware vCenter	6.0
	Enic and fnic drivers	2.1.2.42 (enic) 1.6.0.5 (fnic)

Table 3 lists the VLANs used for this solution validation.

Table 3) VLANs.

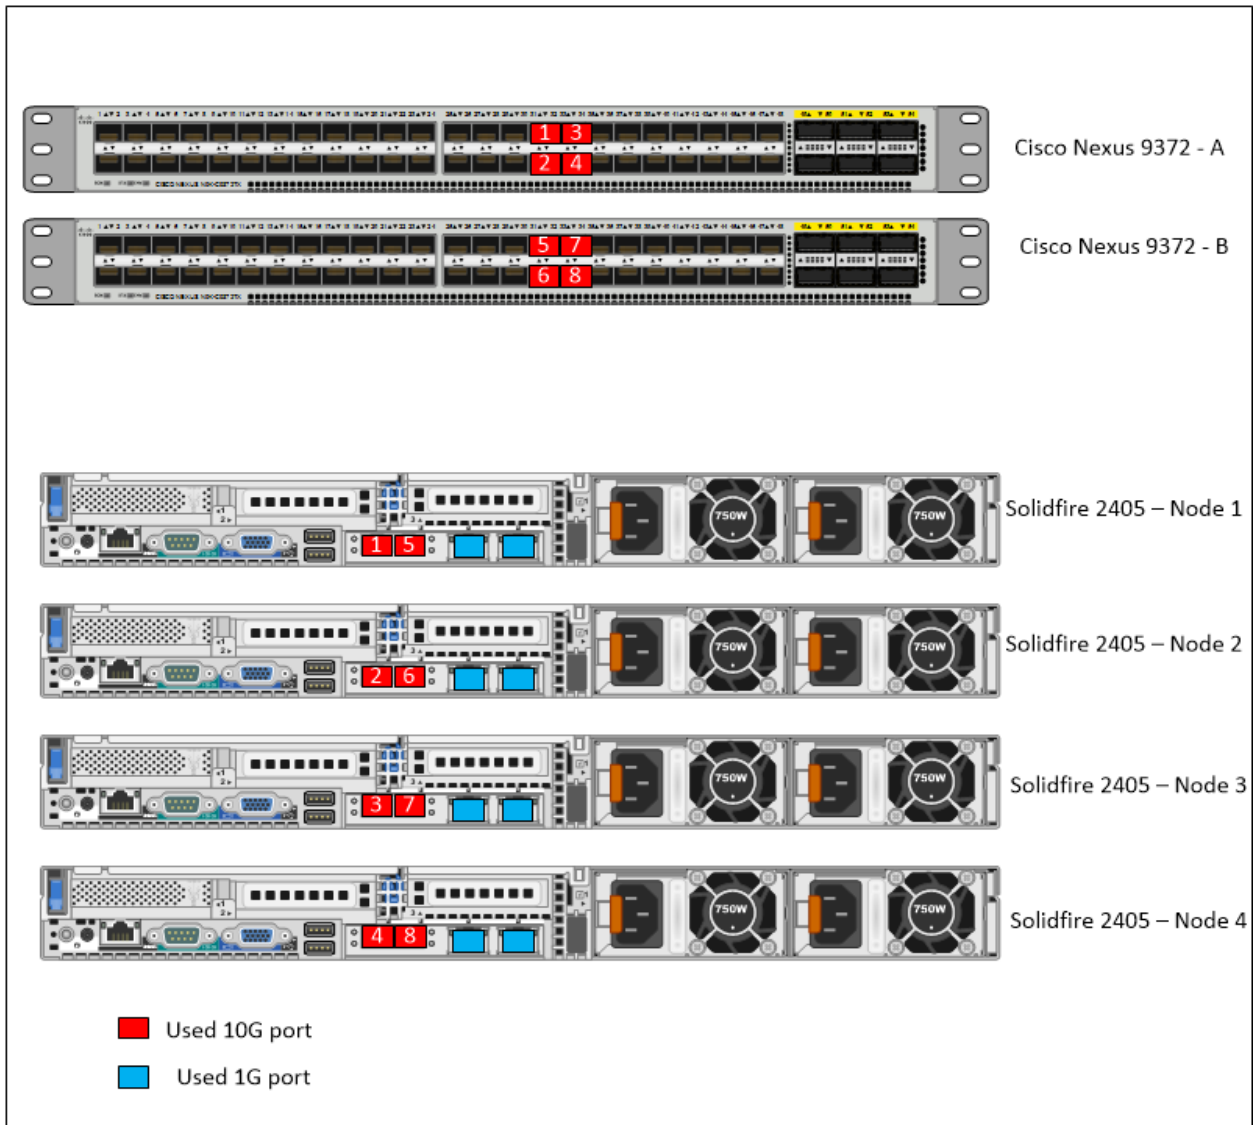
VLAN Name	VLAN
NATIVE-VLAN	2
IB-MGMT-VLAN	3317
iSCSI-STORAGE-VLAN	3318

VLAN Name	VLAN
iSCSI-VLAN-ID_TENANT_1	3342

4 Cabling Details for SolidFire Nodes

Figure 3 shows the cabling diagram for Cisco Nexus switches and SolidFire nodes.

Figure 3) Cabling diagram for Cisco Nexus switch and SolidFire nodes.



5 Deployment Procedures

This document assumes that the FlexPod Datacenter environment is already configured. This document provides detailed steps for attaching SolidFire nodes into an existing FlexPod Datacenter environment.

Deploying the solution involves the following tasks:

- Cisco UCS configuration
- Cisco Nexus switch configuration
- SolidFire node configuration

Note: NetApp recommends configuring Bond1G ports to a 1G management switch for SolidFire node management. This procedure is not covered in this document.

5.1 Cisco UCS Configuration

Follow any existing FlexPod Datacenter documentation to create service profiles and other Cisco UCS configurations. No additional configurations are required to add SolidFire nodes to an existing FlexPod Datacenter environment except the addition of the iSCSI tenant VLAN in vNIC A and vNIC B of a service profile.

5.2 Cisco Nexus Switch Configuration

The procedures in this section describe how to extend the switch configuration to add SolidFire nodes in a FlexPod Datacenter environment. This document assumes that the global configuration, license feature, and vPC are already configured. For more details, see any existing FlexPod Datacenter [documentation](#).

Create iSCSI VLANs

This document covers only the iSCSI-related configuration. If your iSCSI VLANs are already configured, skip this step. To create iSCSI VLANs, complete the following step:

1. From the global configuration mode, run the following commands:

```
vlan <<iSCSI-VLAN-ID_TENANT_1>>  
name iSCSI-VLAN-TENANT-1
```

Add Individual Port Descriptions

To add individual port descriptions for troubleshooting activity, complete the following steps:

Cisco Nexus 9000 A

1. From the global configuration mode, run the following commands:

```
interface Eth1/31  
description SF1:eth1  
exit  
interface Eth1/32  
description SF2:eth1  
exit  
interface Eth1/33  
description SF3:eth1  
exit  
interface Eth1/34  
description SF4:eth1  
exit
```


Cisco Nexus 9000 B

1. From the global configuration mode, run the following commands:

```
interface Eth1/31
description SF1:eth2
exit
interface Eth1/32
description SF2:eth2
exit
interface Eth1/33
description SF3:eth2
exit
interface Eth1/34
description SF4:eth2
exit
```

Create Port Channels

To create the necessary port channels between devices, complete the following step on both switches:

Cisco Nexus 9000 A and Cisco Nexus 9000 B

1. From the global configuration mode, run the following commands:

```
interface Po1
description SF1:Bond10G
exit
interface Eth1/31
channel-group 1 mode active
exit
interface Po2
description SF2:Bond10G
exit
interface Eth1/32
channel-group 2 mode active
exit
interface Po3
description SF3:Bond10G
exit
interface Eth1/33
channel-group 3 mode active
exit
interface Po4
description SF4:Bond10G
exit
interface Eth1/34
channel-group 4 mode active
exit
```

Configure Port Channel Parameters

To configure port channel parameters, complete the following step on both switches:

Cisco Nexus 9000 A and Cisco Nexus 9000 B

1. From the global configuration mode, run the following commands:

```
int Po1
switchport mode trunk
switchport trunk native vlan <<NATIVE-VLAN>>
switchport trunk allowed vlan <<iSCSI-STORAGE-VLAN>>, <<iSCSI-VLAN-ID_TENANT_1>>,
spanning-tree port type edge trunk
mtu 9216
vpc 1

int Po2
switchport mode trunk
```

```

switchport trunk native vlan <<NATIVE-VLAN>>
switchport trunk allowed vlan <<iSCSI-STORAGE-VLAN>>, <<iSCSI-VLAN-ID_TENANT_1>>
spanning-tree port type edge trunk
mtu 9216
vpc 2

int Po3
switchport mode trunk
switchport trunk native vlan <<NATIVE-VLAN>>
switchport trunk allowed vlan <<iSCSI-STORAGE-VLAN>>, <<iSCSI-VLAN-ID_TENANT_1>>
spanning-tree port type edge trunk
mtu 9216
vpc 3

int Po4
switchport mode trunk
switchport trunk native vlan <<NATIVE-VLAN>>
switchport trunk allowed vlan <<iSCSI-STORAGE-VLAN>>, <<iSCSI-VLAN-ID_TENANT_1>>
spanning-tree port type edge trunk
mtu 9216
vpc4

```

Note: This document assumes that the VPCs are already configured in the FlexPod Datacenter environment.

Note: If new tenant is created, make sure to modify the allowed VLAN in the previous configuration.

5.3 SolidFire Node Configuration

This document assumes that your SolidFire hardware is racked, cabled, and powered on. The SolidFire cluster hardware must be appropriately installed and cabled so that network communications and configuration management communications can be established. Instructions for setting up the SolidFire hardware are provided in the hardware box in which it was shipped. For more cabling information, see Figure 3.

Configure SolidFire Bond1G and Bond10G Network Using Terminal User Interface

To configure the SolidFire nodes using the terminal user interface (TUI), complete the following steps:

1. Using the USB and VGA ports on the back side of the SolidFire node, attach the keyboard and monitor to the node.
2. Power on the node.
3. The TUI displays on the tty1 terminal with the Network Settings tab. Make sure that the static IP address is configured for the SolidFire nodes.

Note: A node with a DHCP-assigned IP address cannot be added to a cluster.

4. Set the 1G interface settings as follows:
 - a. Enter <<var_solidfire_node01_mgmt_ip>> in the IP Address field.
 - b. Enter <<var_solidfire_mgmt_mask>> in the Subnet Mask field.
 - c. Enter <<var_solidfire_mgmt_gateway>> in the Gateway Address field.
 - d. Leave the other settings at the default.
 - e. Press S to save the settings. Enter Y to confirm.

```
Bond1G
Method      : static

IP Address  : 172.21.161.64
--->

Subnet Mask : 255.255.255.0
--->

Gateway Address : 172.21.161.1
--->

MTU         : 1500
--->

DNS Servers : 10.61.184.251, 10.61.184.252
--->

Search Domains : cie.netapp.com
--->

Bond Mode   : ActivePassive [ActivePassive, ALB, LACP]
--->

Status      : UpAndRunning [Down, Up, UpAndRunning]
--->

Virtual Network Tag : 0
--->

Routes      : Number of routes: 0.
--->
```

Note: To enter text in each field, press the Enter key to open the edit mode. Upon completion, press the Enter key again to close the edit mode. Use the arrow keys to navigate the fields.

5. Set the 10G interface settings as follows:

- a. Enter <<var_solidfire_node01_storage_ip>> in the IP Address field.
- b. Enter <<var_solidfire_storage_mask>> in the Subnet Mask field.
- c. Enter <<var_solidfire_storage_gateway>> in the Gateway Address field.
- d. Enter 9000 in the MTU field.
- e. Enter LACP in the Bond Mode field.
- f. Enter <<var_iscsi_default_vlan_id>> in the Virtual Network Tag field and press S to save.
- g. Leave the other settings at the default.

Note: The gateway address is optional in a basic configuration of the 10G interfaces. Virtual Network Tag is optional and is only required if it is the primary network for SolidFire.

```
Bond10G
Method          : static
IP Address      : 172.21.162.64
---->
Subnet Mask     : 255.255.255.0
---->
Gateway Address : 172.21.162.1
---->
MTU             : 9000
---->
Bond Mode       : LACP [ActivePassive, ALB, LACP]
---->
LACP Rate       : Fast [Fast, Slow]
---->
Status          : UpAndRunning [Down, Up, UpAndRunning]
---->
Virtual Network Tag : 0
---->
Routes          : Number of routes: 0.
---->
```

- h. Press S to save the settings and enter Y to confirm.
6. Repeat steps 1 to 5 for all the SolidFire nodes.

Configure SolidFire Cluster

To configure the SolidFire clusters complete the following steps:

1. From the TUI, press C to navigate to Cluster Settings.
2. Enter <<var_solidfire_node01>> in the Hostname field.
3. Enter <<var_solidfire_cluster>> in the Cluster field.

Note: Use the same cluster name on all the SolidFire nodes.

4. Leave the other fields at the default.
5. Press S to save the settings and then press Y to accept and save the settings. This operation may take a few minutes to complete.

```
cluster
  Role          : Storage
  Hostname      : SF-AEZD
                ---> SF3010-node01
  Cluster       :
                ---> SF3010-Cluster
  Cluster Membership : Available
  Version       : 9.0.0.1549
  Cluster Interface : Bond10G
  Management Interface : Bond1G [Bond10G, Bond1G]
                    --->
  Storage Interface  : Bond10G
```

6. Perform steps 1-5 on all SolidFire nodes using their respective values.

Note: SolidFire requires a minimum of four nodes.

Create SolidFire Cluster by Using Web UI

You can create a cluster from any node. Creating a new cluster initializes a node as the communications owner for a cluster and establishes network communications for each node in the cluster.

To create a cluster by using the web UI, complete the following steps:

1. In a browser window, enter any node management IP (MIP) address. The Create a New Cluster page appears automatically.
2. All of the nodes are automatically displayed in the Nodes pane.

Create a New Cluster

Node: SF3010-node01 **Status:** Searching for cluster SF3010-Cluster

Management VIP :

ISCSI (Storage) VIP :

Data Protection :

Create Username :

Create Password :

Repeat Password :

EULA


By creating this cluster you signify you have read and agree to the terms of the End User License Agreement ("EULA"), available at solidfire.com/eula, and you represent that you have the authority to enter into this agreement personally, or if you have named a company as customer, on behalf of that customer and bind the customer to the terms of this agreement.

I Agree

Nodes

IP Address	Version	Include
172.21.162.64	9.0.0.1549	<input checked="" type="checkbox"/>
172.21.162.65	9.0.0.1554	<input checked="" type="checkbox"/>
172.21.162.66	9.0.0.1554	<input checked="" type="checkbox"/>
172.21.162.67	9.0.0.1549	<input checked="" type="checkbox"/>

3. Configure the following fields:
 - Management VIP: <<MVIP address>>
 - ISCSI (Storage) VIP: <<SVIP address>>
 - Create User Name: <<username>>
 - Create Password: <<password>>
 - Repeat Password: <<password>>
4. Select the I Agree checkbox. Click Create Cluster.
5. Type the <<var_solidfire_cluster_mgmt_ip>> address in a web browser and enter the authentication credentials.
6. When prompted to add your available drives, click Add Drives.



7. Select Cluster and click the Nodes tab to verify that all four nodes are active.

SolidFire Reporting Management Data Protection

Settings SNMP LDAP Drives Nodes FC Ports Network

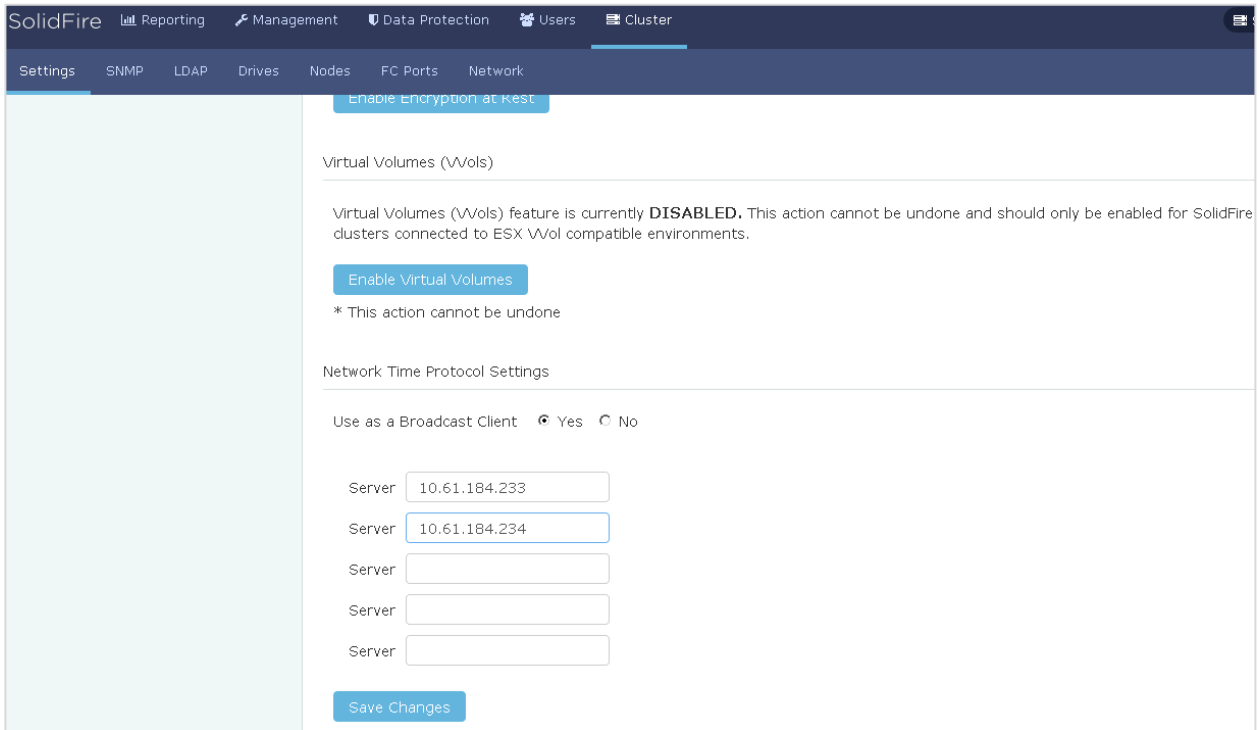
Active Pending PendingActive Filter

<input type="checkbox"/>	Node ID	Node Name	Available 4k IOPS	Node Role
<input type="checkbox"/>	4	SF3010-node04	50,000	Cluster Master
<input type="checkbox"/>	3	SF3010-node03	50,000	Ensemble Node
<input type="checkbox"/>	2	SF3010-node02	50,000	Ensemble Node
<input type="checkbox"/>	1	SF3010-node01	50,000	Ensemble Node

Configure Cluster Full Settings and NTP

To configure the cluster full settings and NTP, complete the following steps:

1. Open a web browser and navigate to the cluster MVIP address.
2. Navigate to Cluster > Settings.
3. Click Cluster.
4. In the Cluster Full Settings section, enter 3 and click Save Changes.
5. Click Back to Settings.
6. In the Network Time Protocol Settings section, click the Broadcast Client option.
7. In the Server field, enter the desired NTP address.
8. Click Save Changes.



Create Tenant Account

Tenant accounts are billable accounts that have access to the storage resources on a SolidFire storage cluster. These accounts enable access to volumes on the cluster through an iSCSI connection and require a Challenge-Handshake Authentication Protocol (CHAP) identification and authorization before a connection can be made.

To create a new tenant account, complete the following steps:

1. Open a web browser and navigate to the cluster MVIP address.
2. Navigate to Management > Accounts.



3. Click Create Account.
4. Enter a new user name.
5. In the CHAP Settings section, enter the initiator secret and target secret passwords for CHAP node session authentication.

Note: Leave the fields blank to autogenerate the passwords. Although Volume Access Groups do not use CHAP authentication, the volume creation still requires an account to be assigned.
6. Click Create Account.

Create a New Account ✕

Account Details

Username

CHAP Settings

Initiator Secret

Target Secret

Create Volume

To create a new volume, complete the following steps:

1. Open a web browser and navigate to the cluster MVIP address.
2. Navigate to Management > Volumes.
3. Click Create Volume.
4. Enter the volume name.

Note: In the Volume Name field, you can enter letters, digits, or dashes (-).

5. Click the Account drop-down list and select the tenant account that is to have access to the volume.
6. Enter the total size of the volume.
7. Select whether or not to enable the 512k block emulation.

Note: This option is necessary to support operating systems that do not recognize native 4k drives, such as VMware ESX. By default, this option is selected.

8. Set the Quality of Service Settings values or accept the default values.
9. Click Create Volume.

Create a New Volume
✕

Volume Details

Volume Name

Volume Size Block Size 512e 4k

Account
 [Create Account?](#)

Quality of Service

IO Size	Min IOPS	Max IOPS	Burst IOPS
4 KB	<input style="width: 50px;" type="text" value="50"/>	<input style="width: 50px;" type="text" value="15000"/>	<input style="width: 50px;" type="text" value="15000"/>
8 KB	31 IOPS	9375 IOPS	9375 IOPS
16 KB	19 IOPS	5556 IOPS	5556 IOPS
262 KB	1 IOPS	385 IOPS	385 IOPS

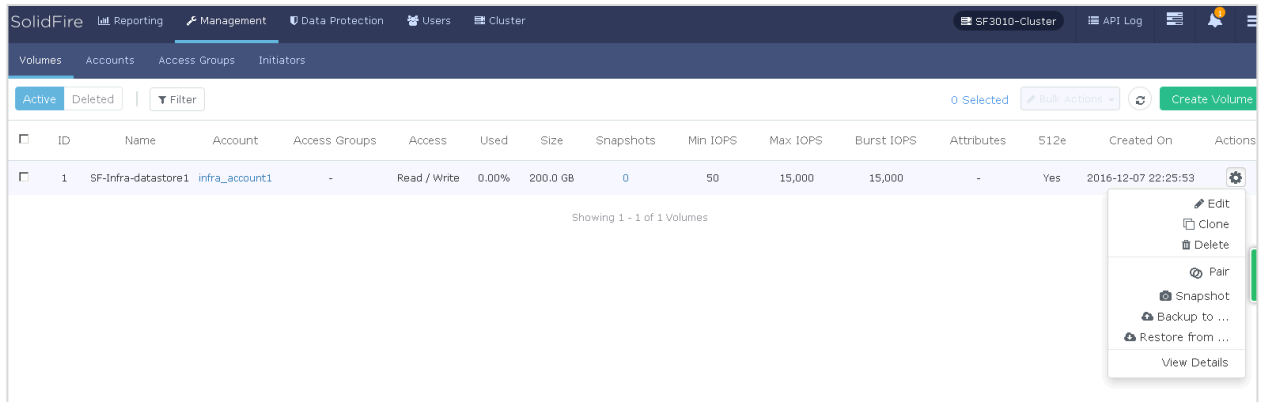
Max Bandwidth 104.86 MB/sec 104.86 MB/sec

Create Volume
Cancel

Create Volume Snapshot Copy

To create new volume, complete the following steps:

1. Open a web browser and navigate to the cluster MVIP address.
2. Navigate to Management > Volumes.
3. Click Active Volumes.
4. Under Actions, click Settings and then select Snapshot.



5. In the Create Snapshot of Volume page:
 - a. Enter a name for the Snapshot copy.
 - b. In the Retention section, select your desired option.
 - c. In the Schedule section, select Create Snapshot Schedule.
 - d. Enter the schedule name and select Schedule Type.
6. Click Create Schedule.

Create Snapshot of Volume
✕

Volume Details

ID: 1 Name: SF-Infra-datastore1

Account: infra_account1

Slice Count: 1 512e: Yes

IQN: iqn.2010-01.com.solidfire:u051.sf-infra-datastore1.1

General

New Snapshot Name

Include Snapshot in Replication When Paired

Retention

Keep Forever

Set Retention Period

Schedule

Take Snapshot Now

Create Snapshot Schedule

New Schedule Name

Schedule Type

Recurring Schedule

Create Snapshot every

SUN

MON

TUE

WED

THU

FRI

SAT

Time of Day (UTC) ⓘ

:

Create Schedule

Cancel

Create Volume Access Group and Attach Volumes

To create a new volume access group, complete the following steps:

1. Open a web browser and navigate to the cluster MVIP address.
2. Navigate to Management > Access Groups.
3. Click Create Access Group.



4. In the Create a New Access Group page:
 - a. Enter a name for the volume access group.
 - b. Click the Create Initiator link.
 - c. Enter an IQN in the Initiators text box and click Create.
 - d. After creating the initiator, select the initiator and click Add Initiator.

Note: To gather the vNIC iSCSI qualified name (IQN) information, launch the Cisco UCS Manager GUI. In the navigation pane, select the Servers tab. Expand Servers > Service Profiles > root. Click each service profile and then click the iSCSI vNICs tab on the right. The initiator name is displayed at the top of the page under Service Profile Initiator Name.
 - e. Click Add Initiator.
 - f. In the Attach Volumes section, select the volume from the Volumes drop-down list and click Attach Volume.
5. Click Create Access Group.

Create a New Access Group
✕

Volume Access Group Details

Name

Add Initiators

Initiators

Create Initiator?

Initiators			1 ▼
ID	Name	Alias	
2	iqn.1992-08.com.cisco:ucs-host:1	-	✕

Attach Volumes

Volumes

Attached Volumes			1 ▼
ID	Name		
1	SF-Infra-datastore1		✕

Create New VLAN

To create a new VLAN, complete the following steps:

1. Open a web browser and navigate to the cluster MVIP address.
2. Navigate Cluster > Network.
3. Click Create New VLAN.
4. Configure the following fields:
 - VLAN Name: ESX-iSCSI-VLAN-Tenant-1
 - VLAN Tag: <<iSCSI-VLAN-ID_TENANT_1>>
 - SVIP: <<IP_address>>
 - Netmask: <<IP_netmask>>
5. In the IP Address blocks section, enter the starting IP address and size.
6. Click Create VLAN.

Create a New VLAN
✕

VLAN Name

VLAN Tag SVIP

Netmask

Enable VRF

Description

IP Address Blocks

Starting IP

Size

Connect VMware vSphere to SolidFire

To connect the VMware vSphere environment to SolidFire, complete the following steps:

Create VMkernel Adapters

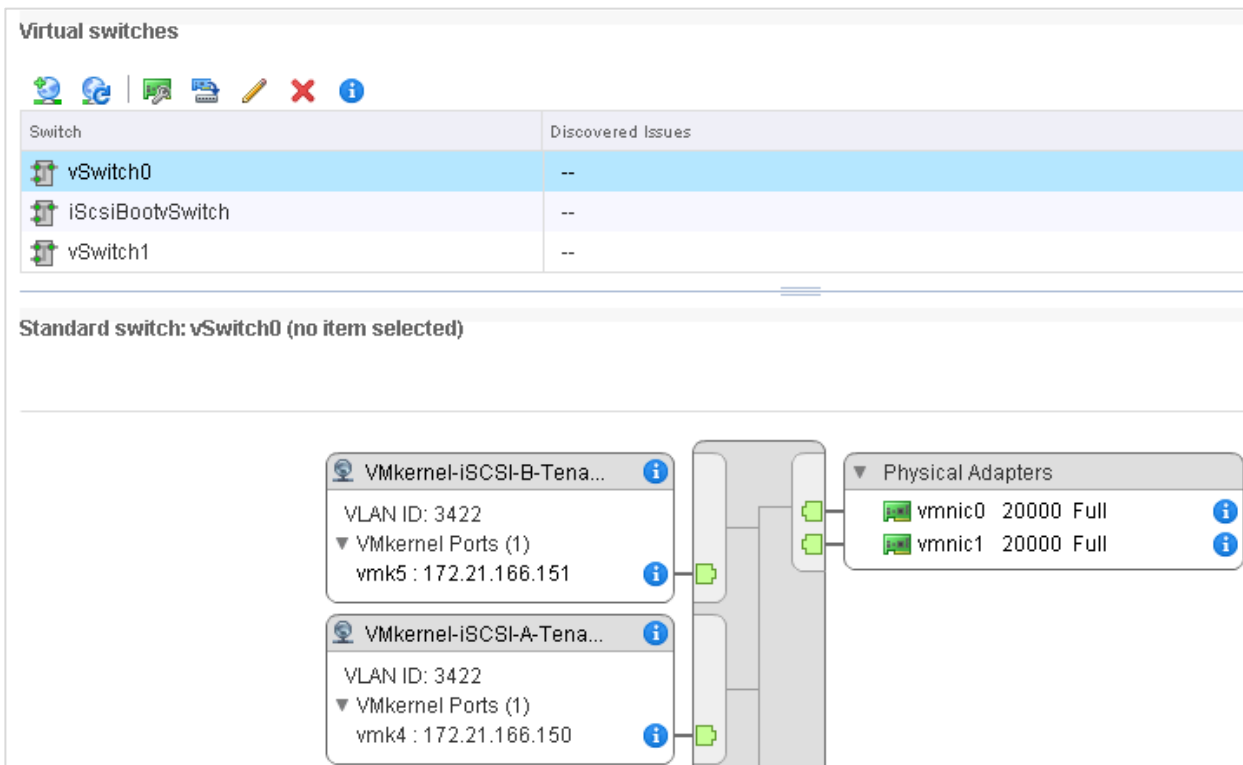
Make sure that the following prerequisites are met:

- One vSphere vSwitch or vSphere distributed switch with at least two physical network uplinks.
- One or more network connections between the ESXi host and SolidFire storage.

To create the VMkernel adapters, complete the following steps:


1. Log in to VMware vCenter using the VMware vSphere Client.
2. From the home page, navigate to Hosts and Clusters.
3. On the right pane, click Manage > Networking > VMkernel adapters.

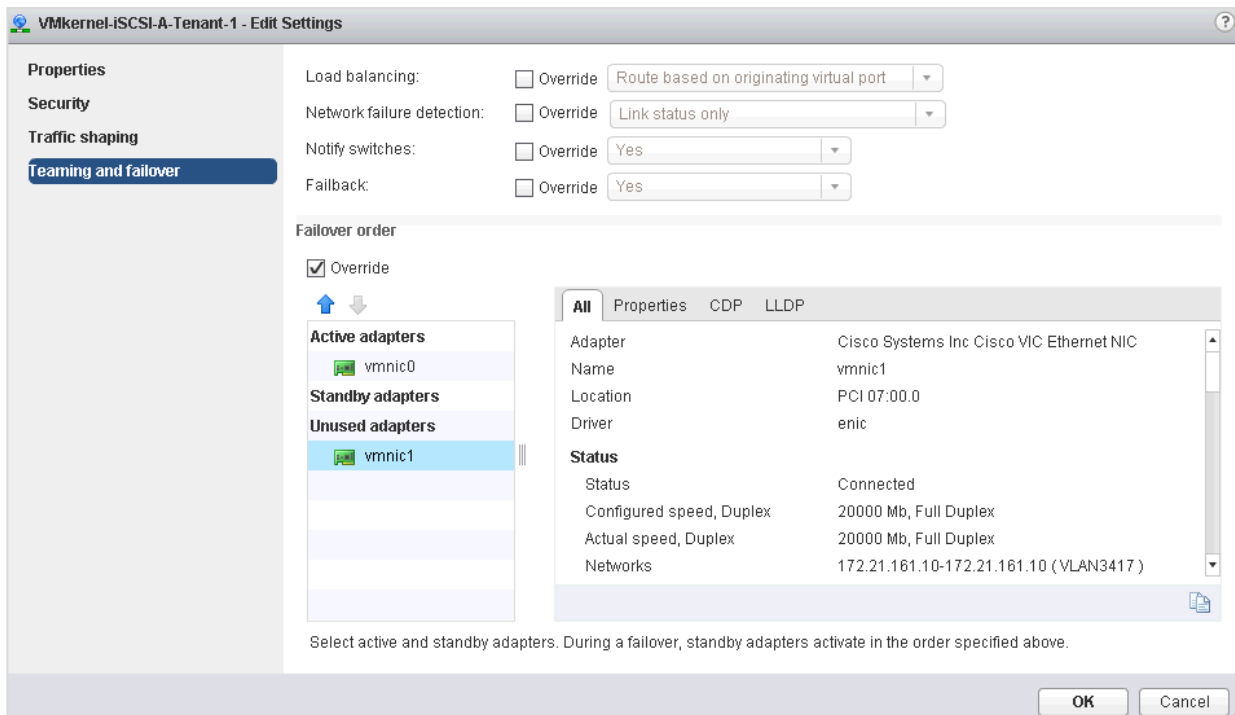
4. Click the add () button.
5. In the Add Networking wizard, select the following options:
 - a. Select VMkernel Network Adapter for the Connection type.
 - b. Select the target device as vSwitch or Virtual Distributed Switch.
 - c. In Port Properties, enter the network label and VLAN ID.
 - d. In IPV4 settings, enter the static IP and subnet mask.
6. Click Finish.
7. (Optional) Create another VMkernel adapter in the same subnet.



Configure iSCSI Multipathing

A VMkernel interface should be configured for each physical network interface to be included in the multipathing configuration. By default, all uplinks are active for each port group. Configure each port group used for iSCSI and its VMkernel interface to override the vSwitch physical interface failover order to configure a single active uplink per iSCSI port group.

1. Log in to VMware vCenter using the VMware vSphere Web Client.
2. From the home page, navigate to Hosts and Clusters.
3. On the right pane, click Manage > Networking > Virtual Switches.
4. Select the desired vSwitch and VMkernel interfaces.
5. Click the Edit settings () button.
6. In the VMkernel Edit settings pane, select Teaming and failover.
7. Select Override under Failover Order.
8. Move one of the VMNICs to Unused adapters and click OK.

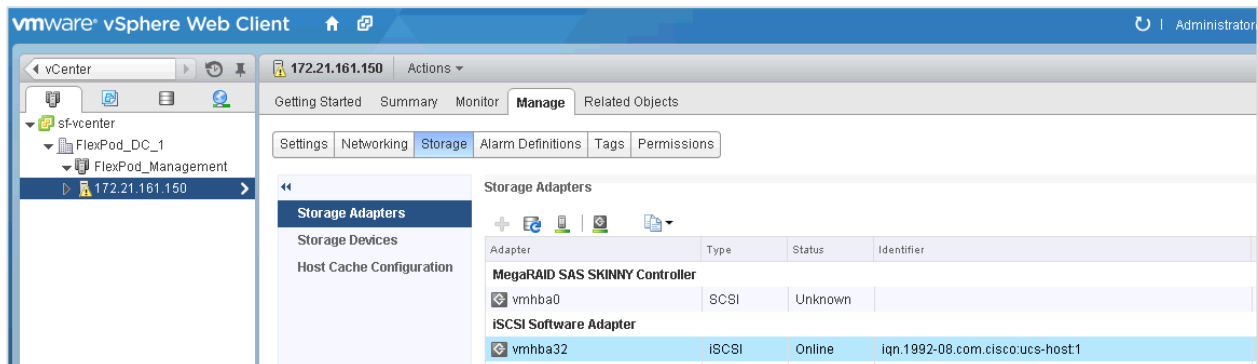


9. Repeat steps 1-8 for the other VMkernel adapter.

Note: vmnic1 is unused adapter for VMkernel-iSCSI-A-Tenant-1 and vmnic0 is unused adapter for VMkernel-iSCSI-B-Tenant-1.

Binding VMkernel Interfaces to the iSCSI Adapter

1. Log in to VMware vCenter by using the VMware vSphere Web Client.
2. From the home page, navigate to Hosts and Clusters.
3. Select the ESXi host to which you want to add the SolidFire iSCSI datastore.
4. Select the Manage tab.
5. Select the Storage tab.



6. In the left pane, click Storage Adapters.
7. Select the adapter under iSCSI Software Adapter.
8. Under Adapter Details, select the Targets tab.
9. Click Static Discovery.

10. Click Add.

The screenshot shows the 'Storage Adapters' configuration page. The left sidebar contains 'Storage Adapters', 'Storage Devices', and 'Host Cache Configuration'. The main content area has tabs for 'Settings', 'Networking', 'Storage', 'Alarm Definitions', 'Tags', and 'Permissions'. The 'Storage Adapters' section includes a table with the following data:

Adapter	Type	Status	Identifier
MegaRAID SAS SKINNY Controller			
vmhba0	SCSI	Unknown	
iSCSI Software Adapter			
vmhba32	iSCSI	Online	iqn.1992-08.com.cisco:ucs-host:1

Below the table, the 'Adapter Details' section is shown with tabs for 'Properties', 'Devices', 'Paths', 'Targets', 'Network Port Binding', and 'Advanced Options'. The 'Targets' tab is selected, showing a table of iSCSI servers and target names:

iSCSI server	Target Name
172.21.163.52:3260	iqn.1992-08.com.netapp:sn.87231f39bd2c11e6891400a09864ecbd:vs.3
172.21.163.51:3260	iqn.1992-08.com.netapp:sn.87231f39bd2c11e6891400a09864ecbd:vs.3
172.21.164.52:3260	iqn.1992-08.com.netapp:sn.87231f39bd2c11e6891400a09864ecbd:vs.3
172.21.164.51:3260	iqn.1992-08.com.netapp:sn.87231f39bd2c11e6891400a09864ecbd:vs.3

An 'Add...' button is located at the bottom right of the Targets table.

11. In the iSCSI Server field, enter the SVIP.

12. In the iSCSI Target Name field, enter the volume IQN.

Note: The IQN value can be retrieved from the SolidFire Element UI by selecting the Modify Volume option.

13. If you are using a one-way CHAP, complete the following steps:

- Select Use Unidirectional CHAP if required by target from the Authentication Method drop-down list.
- In the CHAP Name field, enter the CHAP user name for the SolidFire tenant account that you previously created.
- In the CHAP Secret field, enter the SolidFire initiator secret that you previously created.
- Click OK.

vmhba32 - Add Static Target Server

iSCSI Server:

Port:

iSCSI Target Name:

Authentication Settings

Inherit settings from parent

Authentication Method:

Outgoing CHAP Credentials (target authenticates the initiator)

Name: Use initiator name


Secret:

Incoming CHAP Credentials (initiator authenticates the target)

Name: Use initiator name




Secret:











OK **Cancel**

14. Click the Network Port Binding tab and add the VMkernel interfaces using the  button.

Adapter Details

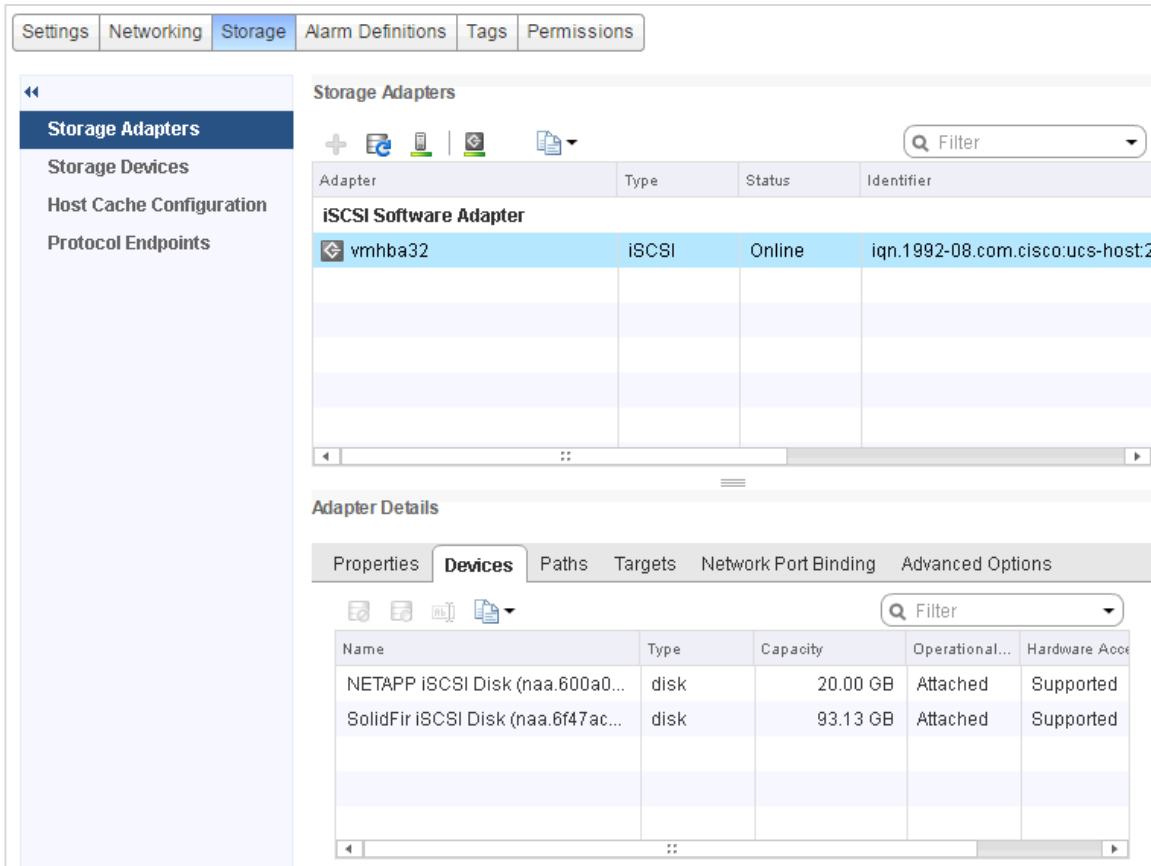
Properties Devices Paths Targets **Network Port Binding** Advanced Options

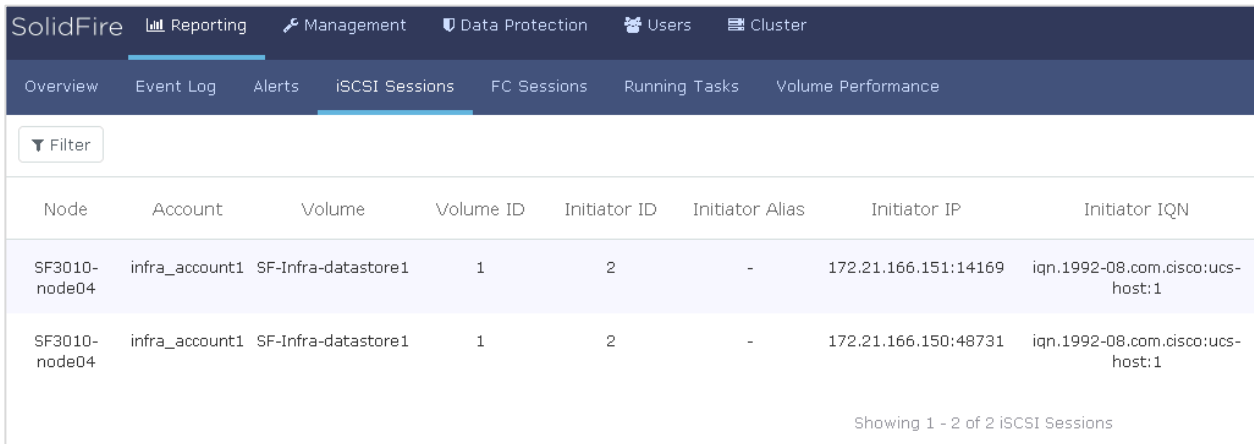
Port Group	VMkernel Ad...	Port Group Policy	Path Status	Physical Network Adapter
 VMkernel-iSCSI-A-...	 vmk4	 Compliant	 Not used	 vmnic0 (20 Gbit/s, Full)
 VMkernel-iSCSI-B-...	 vmk5	 Compliant	 Not used	 vmnic1 (20 Gbit/s, Full)

15. Rescan the storage adapter.

16. In the left pane, select Storage Devices and check for the new storage device.



17. From SolidFire Cluster user interface, make sure the iSCSI sessions are created for both VMkernel interfaces.




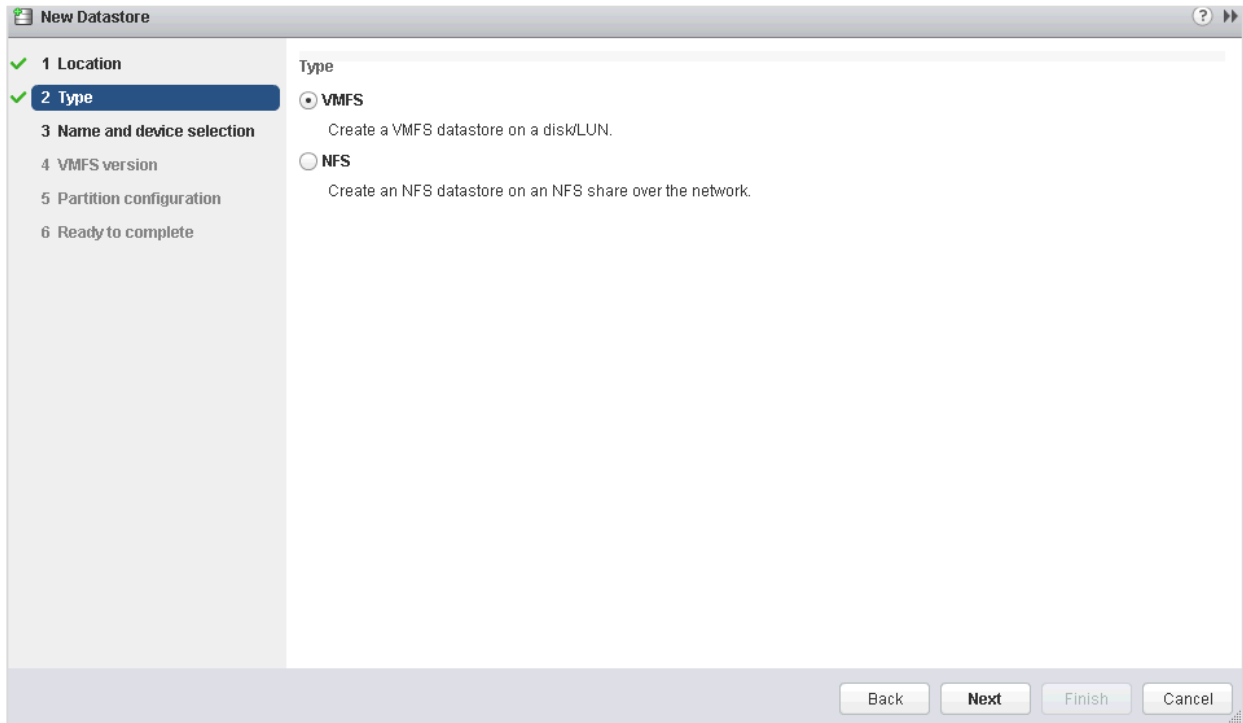
18. Repeat steps 1 to 17 for all of the ESXi hosts in the cluster.

Mount Datastore

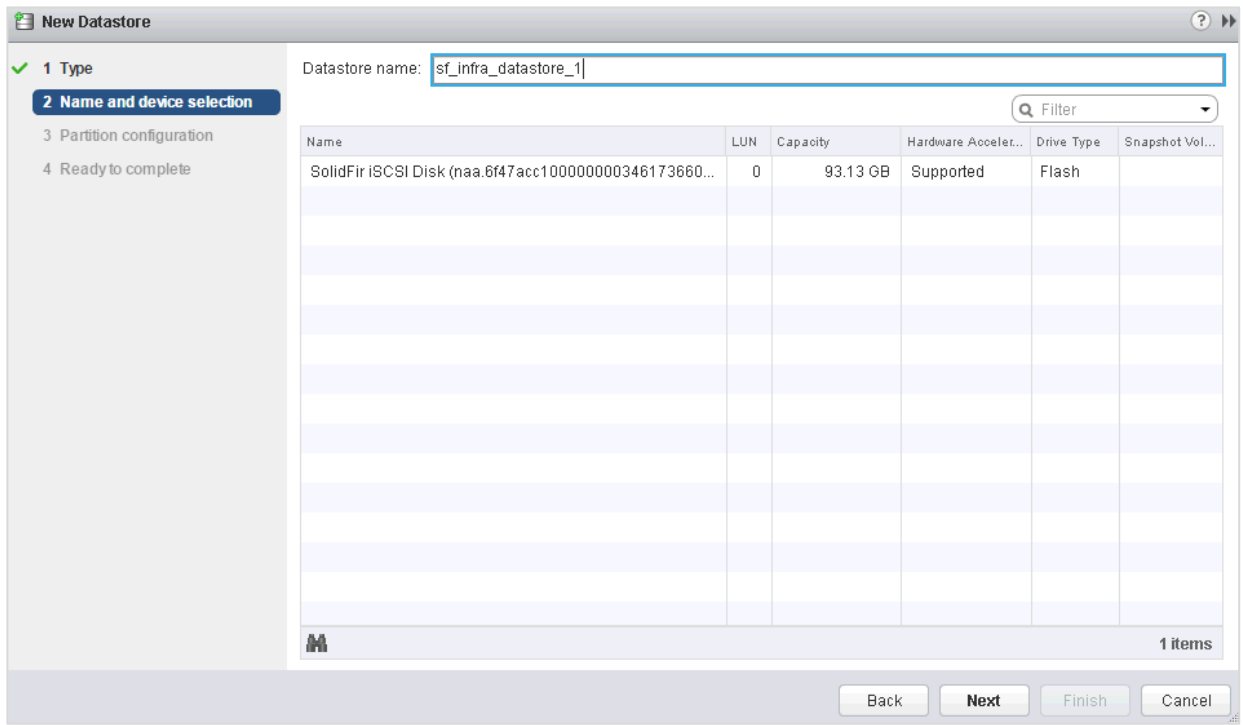
To mount the datastore, complete the following steps:

1. In the VMware vSphere Web Client, open the ESXi host.
2. Select the Related Objects tab.
3. Select the Datastore tab.

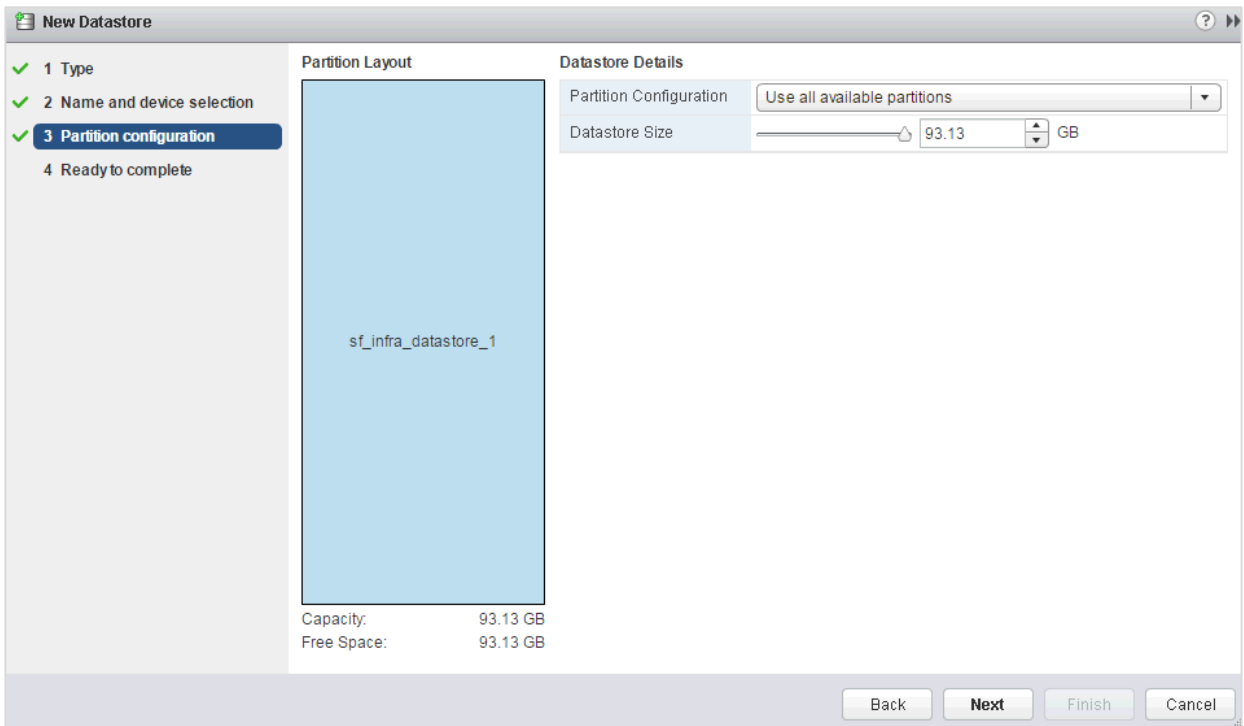
4. Click the Add Datastore () button.
5. Select the VMFS type and click Next.



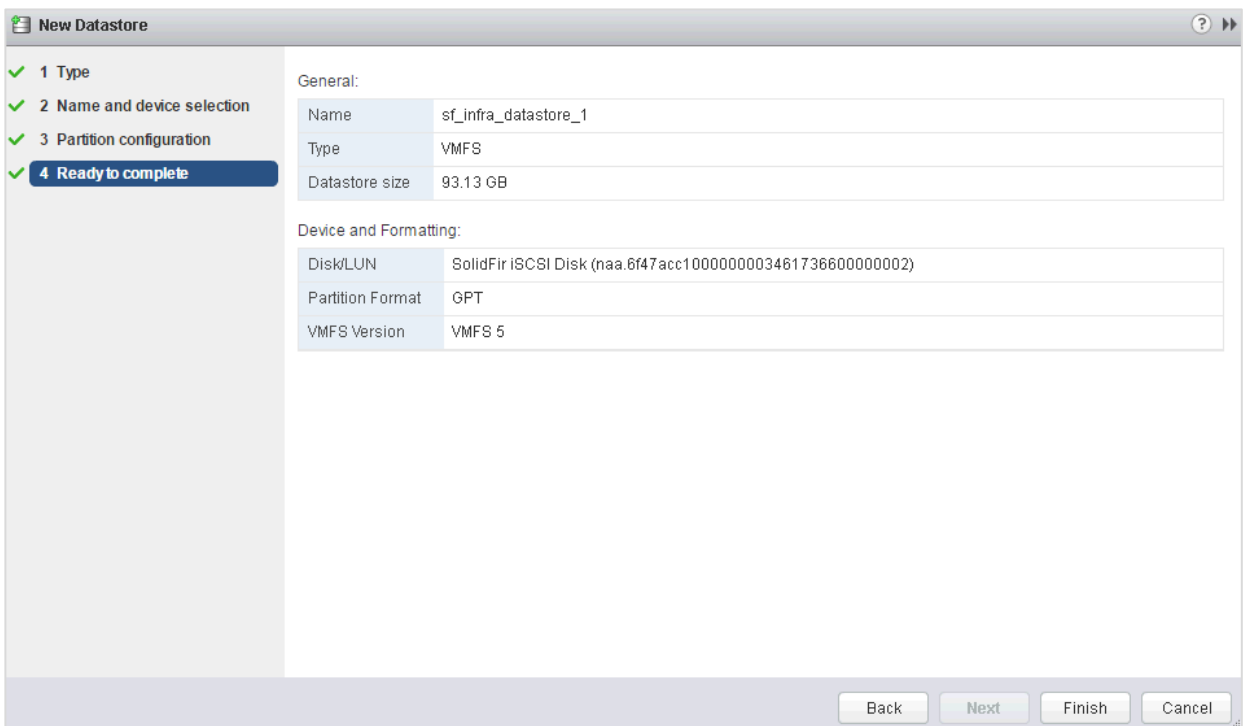
6. In the Datastore Name field, enter a datastore name, select the storage device from the list, and click Next.



7. From the Partition Configuration drop-down list, select your partition layout and click Next.



8. Review the datastore information and click Finish.



Acknowledgements

The authors of this document would like to thank the following people for their support and contribution to the design, validation, lab support, and creation of this NetApp Verified Architecture (NVA):

- Dave Derry, NetApp
- Bhavin Shah, NetApp
- Chad Smith, NetApp

References

This report references the following documents and resources:

- FlexPod Datacenter with VMware vSphere 6.0:
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi60_n9k.html
- VMware vSphere and vSphere with Operations Management:
<http://www.vmware.com/in/products/vsphere>
- SolidFire Active Support:
<http://www.solidfire.com/platform/support>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY NETAPP OR ITS PARTNERS.

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NVA-0027-DEPLOY-0417