



SCALE-OUT FILE SERVICES DESIGN GUIDE

With Microsoft Windows Storage Server 2016
and Pure Storage FlashArray//M20

February 2017



TABLE OF CONTENTS

EXECUTIVE OVERVIEW	3
GOALS AND OBJECTIVES	3
AUDIENCE	3
DESIGN GUIDE PRINCIPLES	3
TOPOLOGY	4
CONFIGURATION & DEPLOYMENT	5
STEP 1 – UCS MANAGER CONFIGURATION	7
ADD BLOCK OF IP ADDRESSES FOR KVM ACCESS	7
EDIT CHASSIS DISCOVERY POLICY	8
ENABLE SERVER, UPLINK, AND FC PORTS	9
ACKNOWLEDGE CISCO UCS CHASSIS	10
CREATE MAC ADDRESS POOLS	11
CREATE WWNN POOLS	12
CREATE WWPN POOLS	14
CREATE UUID SUFFIX POOL	16
CREATE VLANS	18
CREATE VSANS	19
CREATE BOOT POLICIES	20
CREATE SERVICE PROFILE	24
STEP 2 – CONFIGURE BOOT FROM SAN	29
STEP 3 – DEPLOY WINDOWS SERVER	39
STEP 4 – DEPLOYING WINDOWS SERVER FAILOVER CLUSTER	41
DEPLOYING WINDOWS SERVER FAILOVER CLUSTERS	44
STEP 5 – CONFIGURING CONTINUOUSLY AVAILABLE FILE SHARES	46
CONCLUSION	50
APPENDIX 1 – PURE STORAGE COMPONENTS	51
FLASHARRAY//M	51
PURITY OPERATING ENVIRONMENT	52
PURE1® – COMPONENT FEATURES	53

APPENDIX 2 – CISCO COMPONENTS	53
CISCO UCS 5108 CHASSIS	55
CISCO UCS FABRIC INTERCONNECT 6248	55
CISCO NEXUS 5000	55
CISCO UCS B200-M4 SERVERS	56

EXECUTIVE OVERVIEW

Traditional models for IT delivery often cause infrastructure silos, leading to low utilization and high labor costs. Each traditional infrastructure appliance usually has a standalone management interface and requires expensive specialized training. Organizations are constantly challenged to provide a robust, predictable, and dependable environment against a backdrop of ever-increasing volume of data, users, and operations. Given this, the deployment of data center infrastructure can be complex, time consuming, and costly. It is the goal of this design guide to provide a reliable, scalable, easy-to-manage, and high performance platform for consolidated workloads and private cloud environments.

This document describes a reference architecture for deploying a highly-available and scale-out Microsoft® File Server solution with Microsoft Windows® Storage Server 2016. We will showcase the ease of deploying a scalable file services environment with Microsoft Windows Storage Server 2016 native Server Message Block (SMB) support.

Scale-Out File Services is a capability that is designed to provide highly-available scale-out file shares that are continuously available for file-based storage such as home directories and SMB shares. This scenario focuses on how to plan for and deploy Scale-Out File Server.

GOALS AND OBJECTIVES

This document provides a high-level design for customers seeking to implement a green-field deployment or to move/migrate from an existing Network Attached Storage (NAS) solution to a Pure Storage® FlashArray solution with Microsoft Windows Storage Server 2016. We provide guidance on Windows Storage Server 2016 R2 Standard components with Pure Storage for hosting SAN and SMB shares for end users.

AUDIENCE

The target audiences for this document are storage and virtualization administrators, data center architects, field engineers, and server specialists.

DESIGN GUIDE PRINCIPLES

The guiding principles for implementing this reference architecture are:

- **Availability** – Create a design that is resilient and not prone to failure of a single component. For example, we include best practices to enforce multiple paths to storage using MPIO with Windows Server Failover Clustering.
- **Efficient** – Build a solution that leverages the efficiency benefits of an all-flash architecture, including data reduction with in-line deduplication and compression, 100% encryption of data at rest, and an Evergreen system architecture.
- **Simple** – Simplify deployment and ongoing maintenance tasks via automation.

- **Scalable** – Create a design that can start small but easily expand to meet the needs of a growing organization's file system needs.

Please see appendices to this document for a detailed overview of Cisco UCS®, Cisco Nexus®, Pure Storage FlashArray//M, and Microsoft components used in this design guide.

TOPOLOGY

In this section, we will learn more about the high-level topology of the reference architecture. It comprises the following hardware and software components:

HARDWARE

- **Compute** – Cisco UCS Mini [Cisco UCS B200 M4 Blade Server + Cisco UCS 5108 Blade Server Chassis]
- **Network** – Cisco UCS Fabric Interconnect 6324 for external and internal connectivity of IP and FC network.
- **Storage** – Pure Storage FlashArray//M

SOFTWARE

- Cisco UCS Manager
- Cisco UCS Firmware (version 3.1.1E)
- Microsoft Windows Storage Server 2016
- Pure Storage Web Management interface
- Pure Storage PowerShell SDK

Figure 1 shows a detailed topology of the reference architecture configuration. A major goal of the architecture is to build out a highly redundant and resilient infrastructure. We used powerful servers with dual Fibre Channel ports connected redundantly to two SAN switches that were connected to redundant FC target ports on the FlashArray//m. The servers also have redundant network connectivity. It is configured in End-Host Mode.

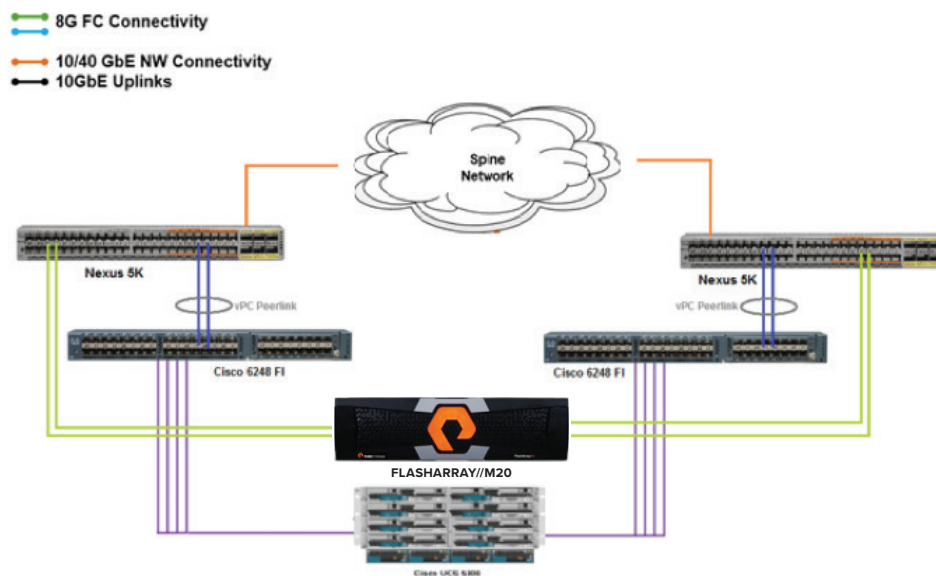


Figure 1. Cisco UCS and Pure Storage FlashArray connectivity

The use case for implementing this reference architecture is:

- **Simplified Provisioning** – Reduce the complexity of provisioning new SAN or SMB shares and reduce time spent on common tasks.
- **Ongoing Maintenance** – Reduce overhead of ongoing maintenance by avoiding the need to tune the environment.
- **Performance** – Enhance infrastructure to improve performance and scale.
- **Availability** – Maintain a highly available application storage and file services environment.
- **Protection** – Integrate data protection to minimize data loss and recover data, as needed.
- **Cost** – Provide a cost-efficient solution that supports the SMB use case with the option to add NFS.

CONFIGURATION & DEPLOYMENT

The configuration of this reference architecture is as follows:

- Pure Storage FlashArray//M20
- Cisco Nexus 5000
- Cisco 6248 Fabric Interconnect
- Purity 4.8.5
- Windows Storage Server 2016

The following Windows Storage Server 2016 components are used:

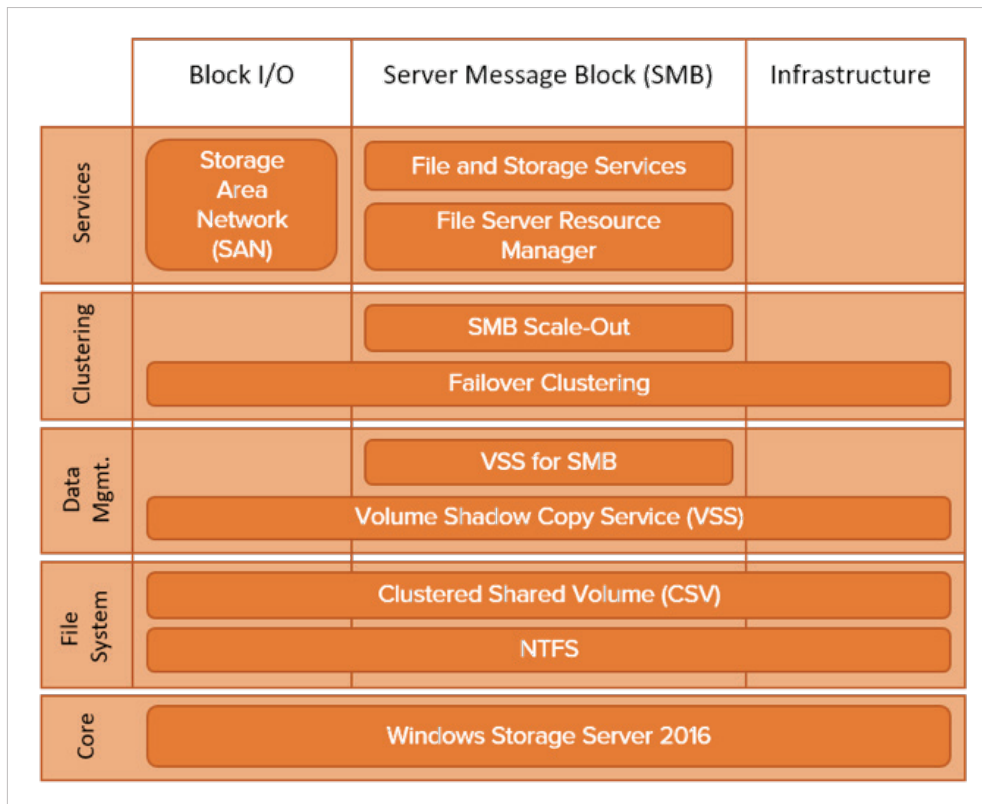


Figure 2. Logical services diagram

The following sections explain the procedure to setup and configure the listed portions of the Cisco UCS environment for SoFS:

1. UCS Manager configuration
2. Configure SAN boot
3. Setup Windows Storage Server 2016
4. Configure Windows Failover Clustering
5. Configure Server Message Block (SMB) for high availability

This deployment guide assumes that the Cisco environment has been fully configured with servers (blades) online, switches, and zoning.

STEP 1 – UCS MANAGER CONFIGURATION

To log in to the Cisco Unified Computing System (UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS 6248UP Fabric Interconnect IP address.
2. Click the **Launch UCS Manager** link to download the Cisco UCS Manager software.
3. If prompted, accept the security certificate.
4. When prompted, enter **admin** as the user name and enter the administrative password.
5. Click **Login** to access the UCS Manager.

ADD BLOCK OF IP ADDRESSES FOR KVM ACCESS

To create a block of IP addresses for server keyboard, video, mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the **LAN** tab.
2. Select **Pools > root > IP Pools > IP Pool ext-mgmt**.
3. In the **Actions** pane, select **Create Block of IPv4 Addresses**.

This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

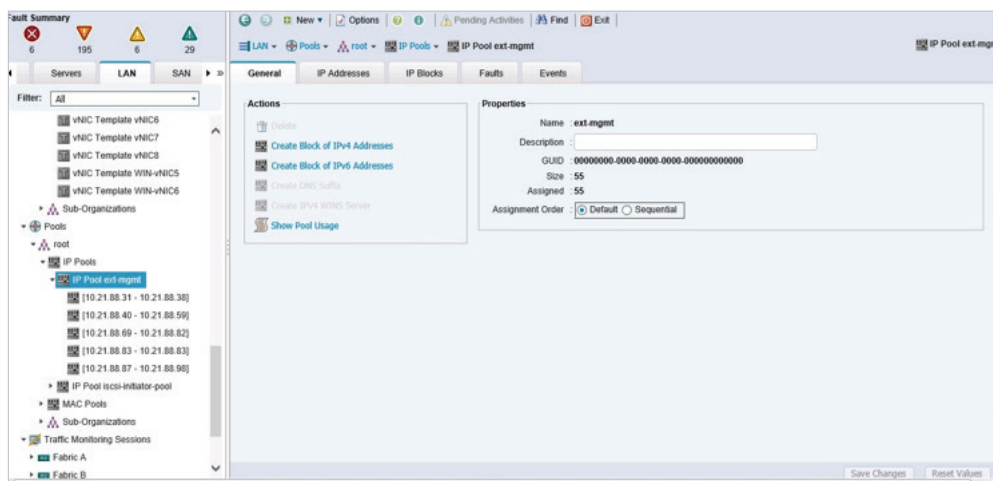


Figure 3. Adding block of IP addresses

4. Enter the following:
 - Starting IP address of the block in the **From** field
 - Number of IP addresses required in the **Size** field
 - Subnet Mask
 - Default Gateway
 - Primary DNS
 - Secondary DNS

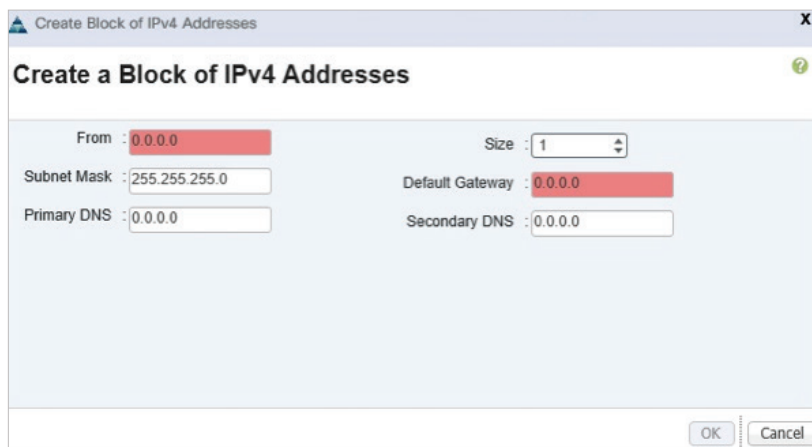


Figure 4. Creating block of IPv4 addresses

5. Click **OK** to create the IP block.
6. Click **OK** in the confirmation message.

EDIT CHASSIS DISCOVERY POLICY

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and C-Series servers. To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, in the left navigation pane, click the **Equipment** tab and select **Equipment** from the list.
2. In the right pane, click the **Policies** tab.
3. Under **Global Policies**, set the **Chassis/FEX Discovery Policy** to 2-link or set it to match the number of uplink ports that are cabled between the chassis and the fabric interconnects.

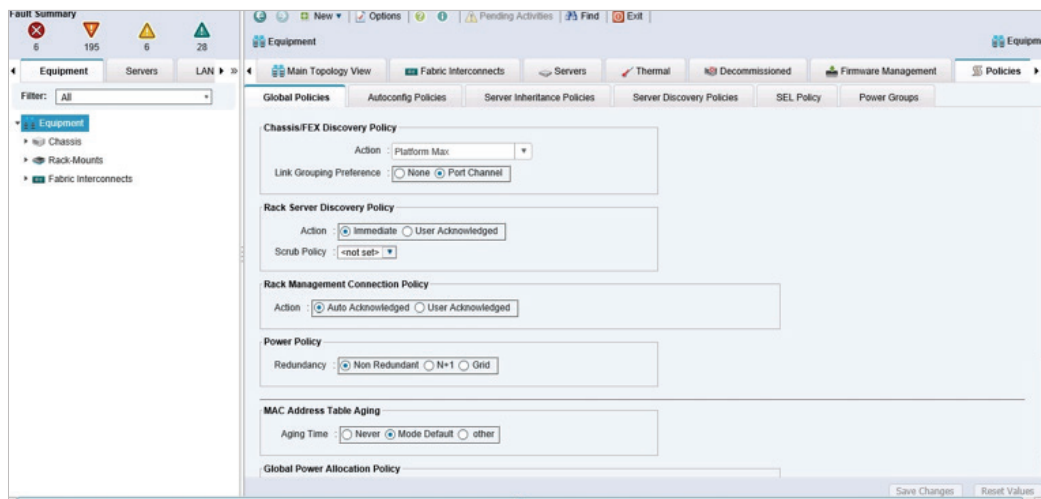


Figure 5. Editing chassis discovery policy

4. Click **Save Changes**.
5. Click **OK**.

ENABLE SERVER, UPLINK, AND FC PORTS

To enable server, uplink, and FC ports, complete the following steps:

1. In Cisco UCS Manager, in the left navigation pane, click the **Equipment** tab.
2. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
3. Expand **Ethernet Ports**.
4. Select the ports that are connected to the chassis or directly connected to C-Series rack servers, right-click them, and select **Configure as Server Port**.
5. Click **Yes** to confirm server ports and click **OK**.
6. Select the ports that are connected to the Cisco Nexus 5000 switches, right-click them, and select **Configure as Uplink Port**.
7. Click **Yes** to confirm uplink ports and click **OK**.
8. Select the ports that are connected to the Cisco Nexus 5000 switches, right-click them, and select **Configure as FC Port**.
9. In the left pane, navigate to **Fabric Interconnect A (primary)**. In the right pane, navigate to the **Physical Ports** tab > **Ethernet Ports** tab. Confirm that ports have been configured correctly in the **If Role** column.

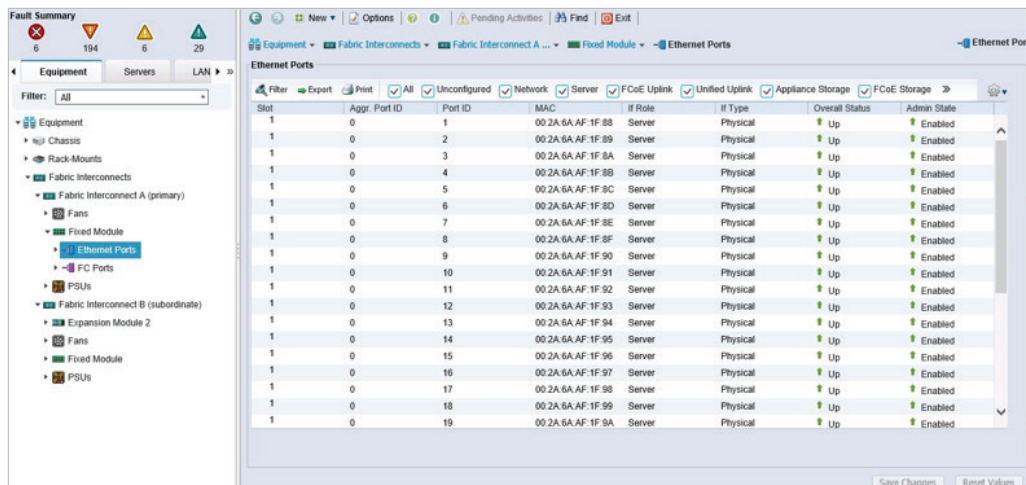


Figure 6. Configuring ports

10. Follow the steps 1 – 9 for **Fabric Interconnect B (subordinate)**.

ACKNOWLEDGE CISCO UCS CHASSIS

To acknowledge all Cisco UCS chassis, complete the following steps:

1. In Cisco UCS Manager, in the left navigation pane, click the **Equipment** tab.
2. Expand **Chassis** and select each chassis that is listed.
3. Right-click each chassis and select **Acknowledge Chassis**.



Figure 7. Acknowledging the chassis

CREATE MAC ADDRESS POOLS

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the left navigation pane.
2. Select **Pools > root**.
3. Right-click **MAC Pools** under the root organization.
4. Select **Create MAC Pool** to create the MAC address pool.
5. Enter the **Name** for MAC pool.
6. **Optional:** Enter a **Description** for the MAC pool.

Keep the Assignment Order at Default.

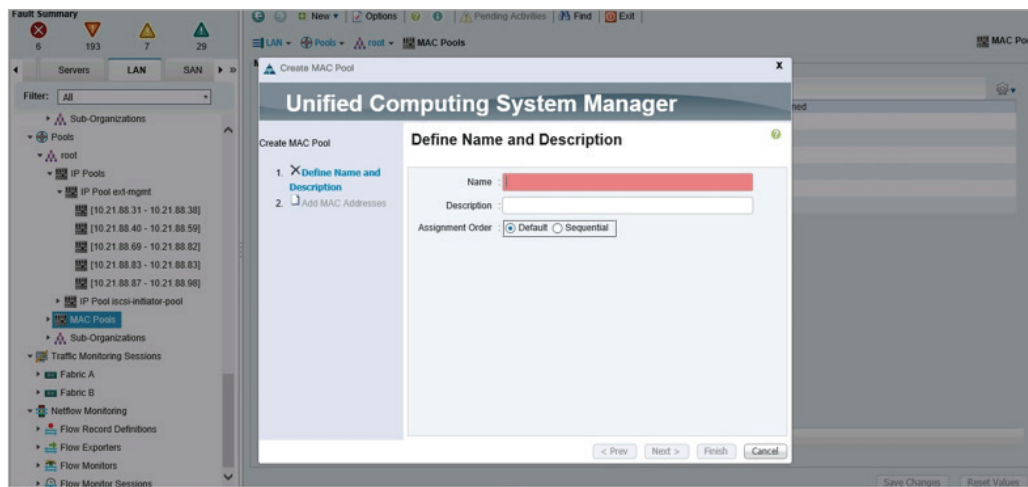


Figure 8. Create MAC addresses pool

7. Click **Next**.

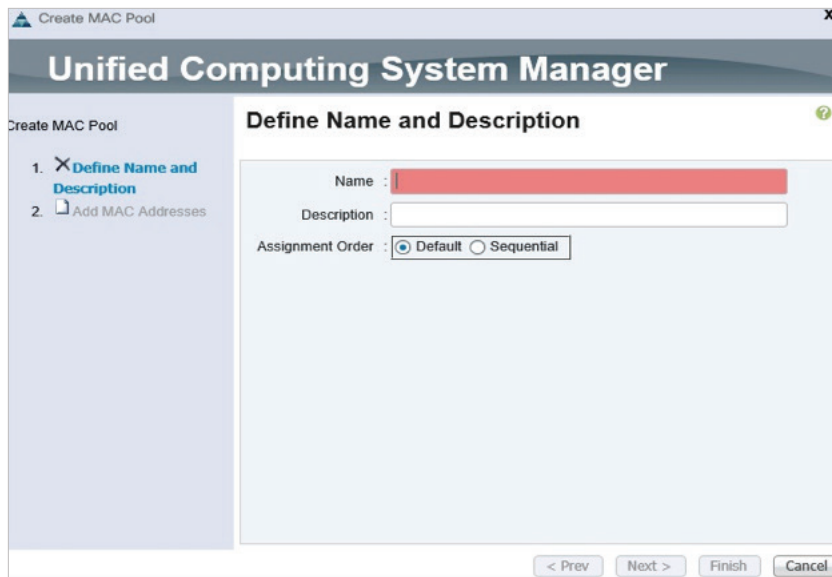


Figure 9. Define name and description of MAC pool

8. Click **Add**.
9. Specify a starting MAC address in the **First MAC address** field.

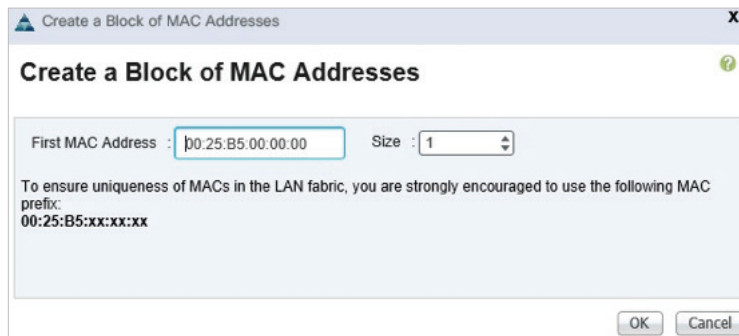


Figure 10. Create block of MAC addresses

10. Click **OK**.
11. Click **Finish**.
12. In the confirmation message window, click **OK**.

CREATE WWNN POOLS

To configure the necessary World Wide Node Name (WWNN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **SAN** tab in the left navigation pane.
2. Select **Pools > root**.

3. Right-click **WWNN Pools**.
4. Select **Create WWNN Pool**.
5. Enter the **Name** for WWNN pool.
6. **Optional:** Add a **Description** for the WWNN pool.
7. Click on **Default** radio button for the **Assignment Order**.

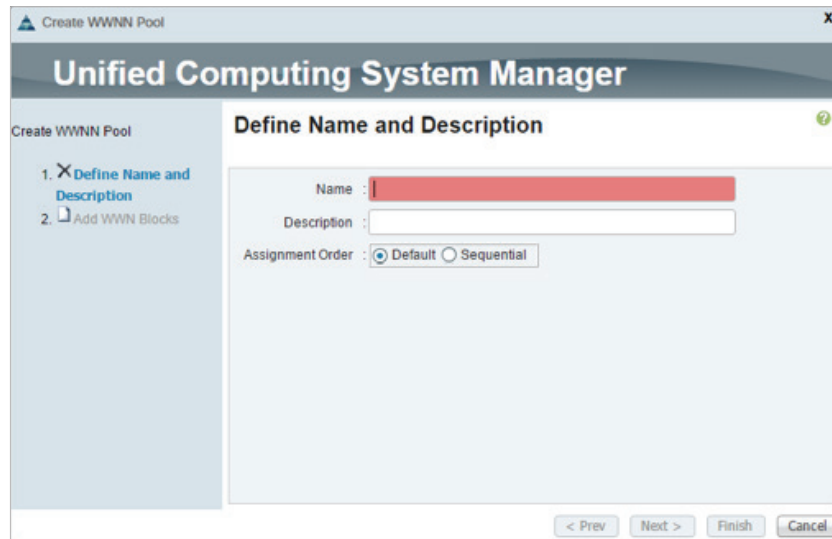


Figure 11. Define name and description of WWNN pool

8. Click **Next**.



Figure 12. Add WWNN block

9. Click **Add** to add a block of WWNNs.
10. Either retain the default block of WWNNs, or specify a base WWNN.

11. Specify a **Size** for the WWNN block.

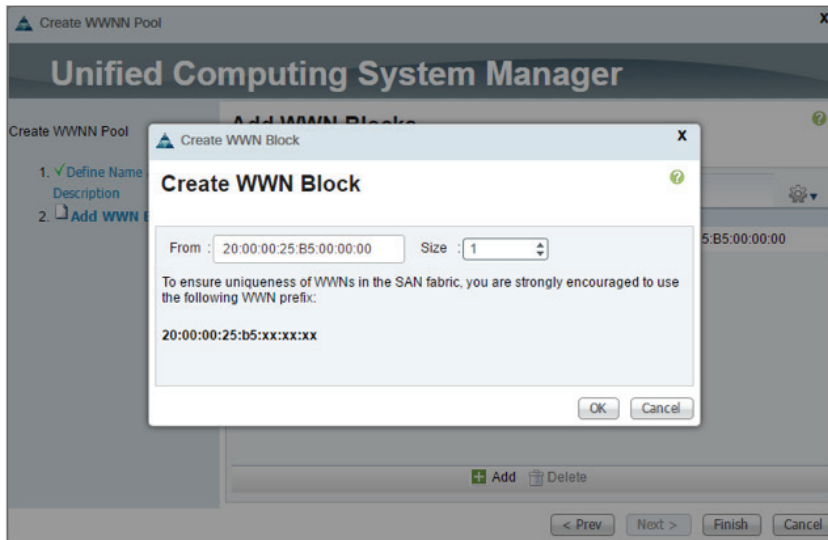


Figure 13. Create WWN block

12. Click **OK**.
13. Click **Finish**.
14. Click **OK**.

CREATE WWPN POOLS

To configure the necessary World Wide Port Name (WWPN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **SAN** tab in the left navigation pane.
2. Select **Pools > root**.
3. Right-click **WWPN Pools**.
4. Select **Create WWPN Pool**.
5. Enter the **Name** for WWPN pool.
6. **Optional:** Enter a **Description** for this WWPN pool.
7. Click the **Default** radio button for **Assignment Order**.
8. Click **Next**.

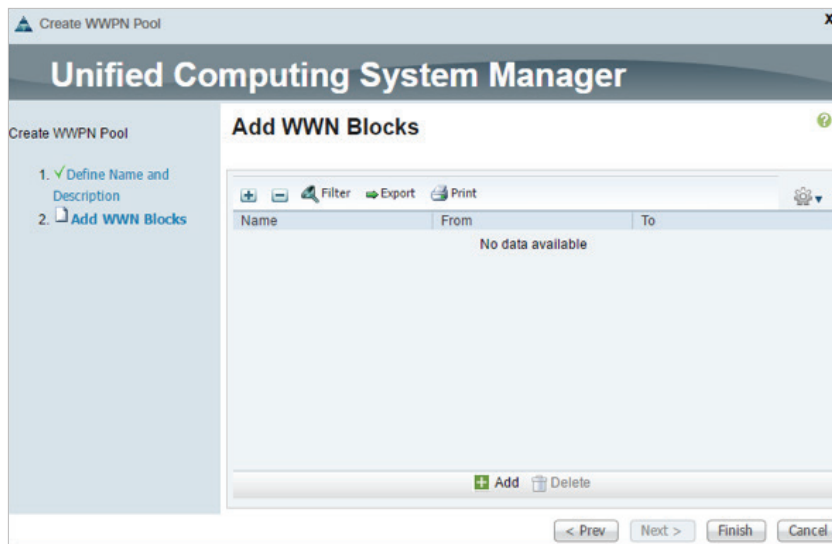


Figure 14. Add WWN blocks

9. Click **Add** to add a block of WWPNS.
10. Specify the starting WWPNS in the block.
11. Specify a **Size** for the WWPNS block.

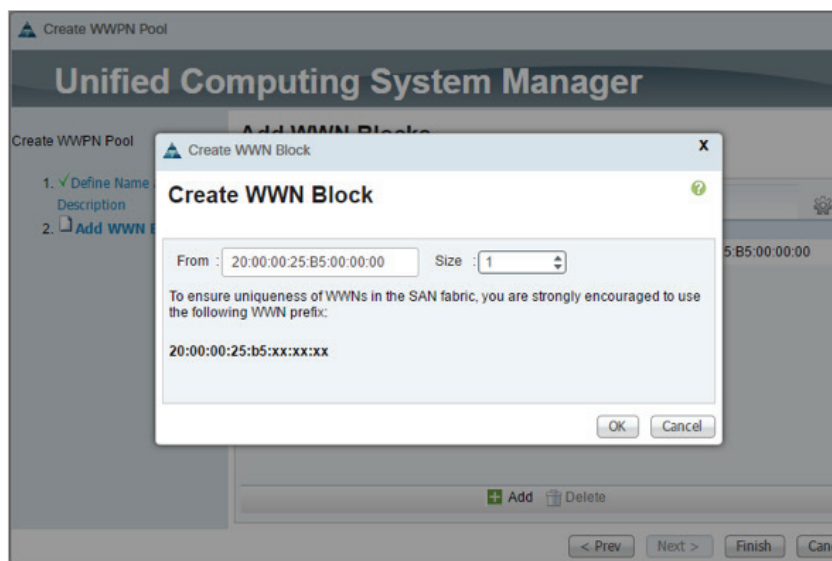


Figure 15. Create WWPNS block

12. Click **OK**.

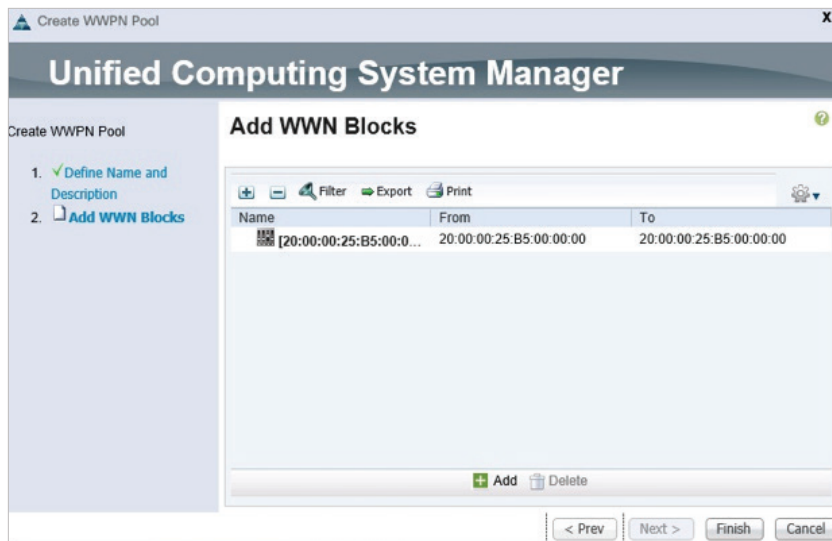


Figure 16. Complete the creation of WWPN pool

13. Click **Finish** to create the WWPN pool.
14. Click **OK**.

CREATE UUID SUFFIX POOL

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the left navigation pane.
2. Select **Pools > root**.
3. Right-click **UUID Suffix Pools**.
4. Select **Create UUID Suffix Pool**.
5. Enter the **Name** for UUID suffix pool.
6. **Optional:** Enter a **Description** for UUID suffix pool.
7. Click the **Derived** radio button for **Prefix**.
8. Click the **Default** radio button for the **Assignment Order**.

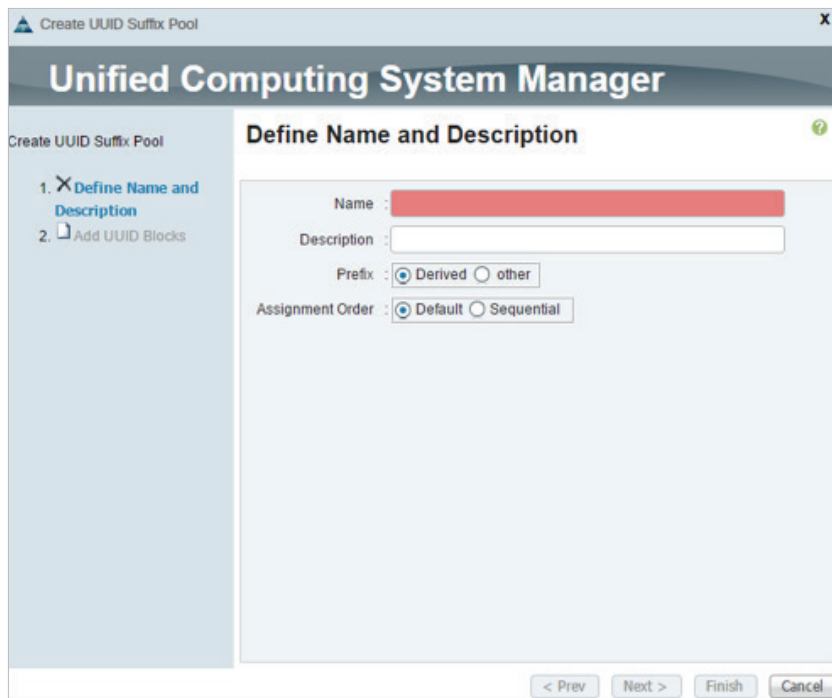


Figure 17. Create UUID suffix pool

9. Click **Next**.

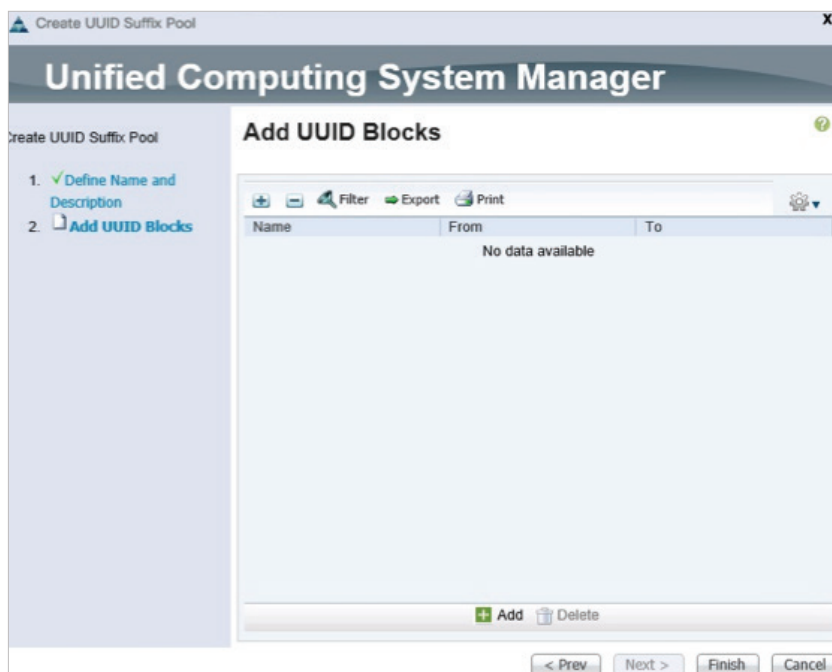


Figure 18. Add UUID block

10. Click **Add** to add a block of UUIDs.

11. Select from option as the default setting.
12. Specify a **Size** for the UUID block.

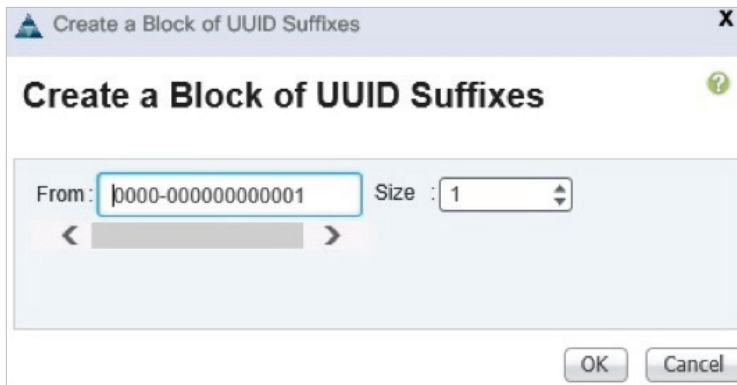


Figure 19. Create block of suffixes

13. Click **OK**.
14. Click **Finish**.
15. Click **OK**.

CREATE VLANS

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select **LAN > LAN Cloud**.
3. Right-click **VLANs**.
4. Select **Create VLANs**.
5. Enter the **Name** for VLAN to be used for management traffic.
6. Click the **Common/Global** radio button for the scope of the VLAN.
7. Enter **VLAN IDs** as the ID of the management VLAN.
8. Click the **None** radio button for **Sharing Type**.

Figure 20. Create VLAN

9. Click **OK**, and then click **OK** again.

CREATE VSANS

To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Expand the **SAN > SAN Cloud Tree**.
3. Right-click **VSANs**.
4. Select **Create VSAN**.
5. Enter the **Name** for the VSAN for Fabric A.
6. Click the **Disabled** radio button for **FC Zoning**.
7. Enter **VSAN ID**.

Figure 21. Create VSAN

8. Click **OK**, and then click **OK** again to create the VSAN.
9. Right-click **VSANs**.
10. Select **Create VSAN**.
11. Enter the **Name** for the VSAN for Fabric B.
12. Click the **Disabled** radio button for **FC Zoning**.
13. Click the **Fabric B** radio button.
14. Enter **VSAN ID** for Fabric B.

CREATE BOOT POLICIES

To create boot policies for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, in the left navigation pane, click the **Servers** tab.
2. Select **Policies > Root**.
3. Right-click **Boot Policies**.
4. Select **Create Boot Policy**.
5. Enter the **Name** for the boot policy.
6. **Optional:** Enter a **Description** for the boot policy.
7. Uncheck the **Reboot on Boot Order Change** checkbox.

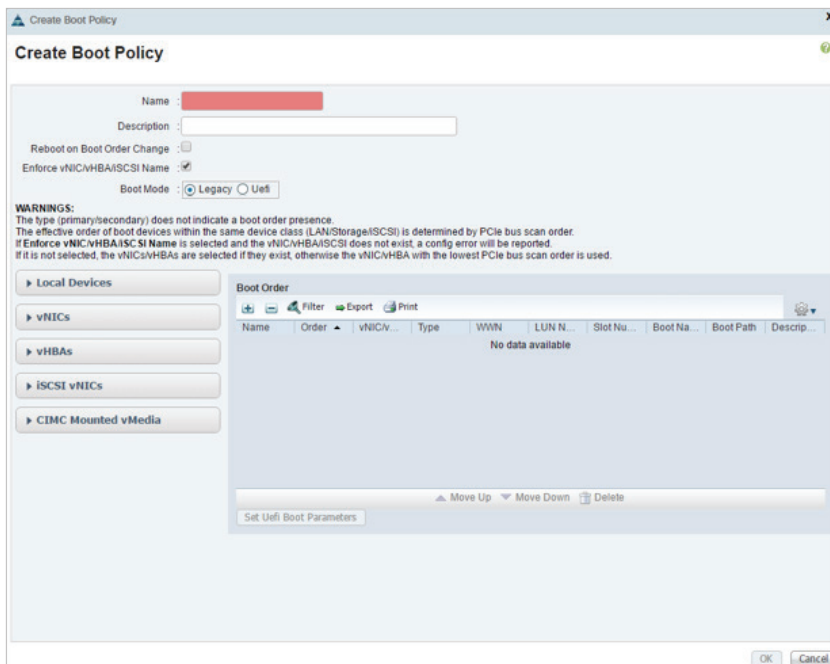


Figure 22. Create boot policy.

8. Expand the vHBAs drop-down menu and select **Add SAN Boot**.
9. In the **Add SAN Boot** dialog box, enter **Fabric-A** as value for vHBA field.
10. Click the **Primary** radio button for the **Type** option.

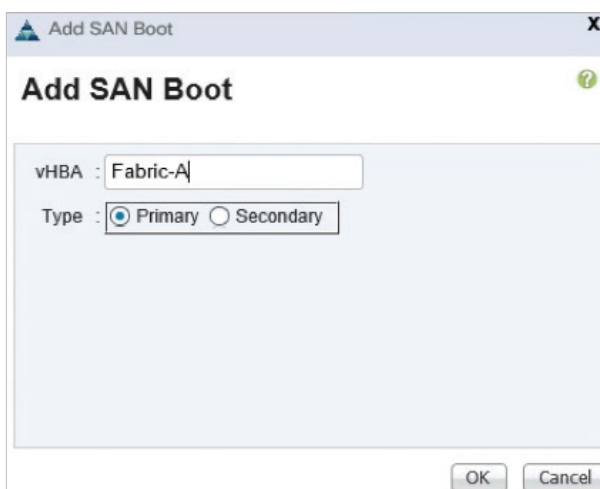


Figure 23. Add SAN boot on Fabric A

11. Click **OK** to add the SAN boot initiator.
12. From the vHBA drop-down menu, select **Add SAN Boot Target**.
13. Enter **1** as the value for **Boot Target LUN**.

14. Enter the WWPN for CT0.FC0.
15. Click the **Primary** radio button for the **Type** option.

*To obtain this information, log in to pure storage, and under the **system tab** in the left panel, click the host connection for the target Ports.*

16. Click the **Primary** radio button for the SAN boot target type.
17. Click **OK** to add the SAN boot target.
18. From the vHBA drop-down menu, select **Add SAN Boot Target**.
19. Enter **1** as the value for Boot Target LUN.
20. Enter the WWPN for CT1.FC0.
21. Click the **Secondary** radio button for the **Type** option.

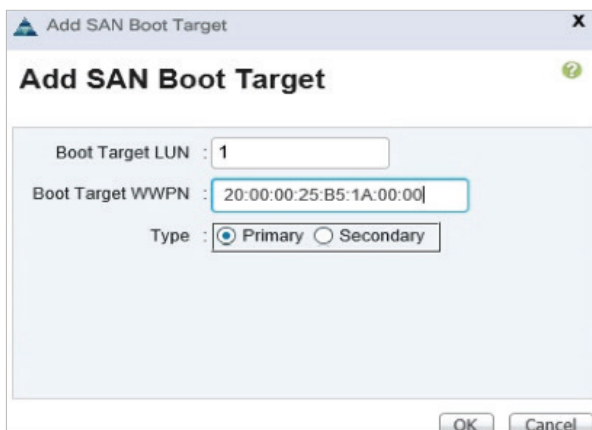


Figure 24. Add SAN boot target for Fabric A

22. Click **OK** to add the SAN boot target.
23. From the vHBA drop-down menu, select **Add SAN Boot**.
24. In the **Add SAN Boot** dialog box, enter **Fabric-B** as the value in the **vHBA** field.
25. By default, the SAN boot Type will be set to **Secondary** and the options will be disabled.

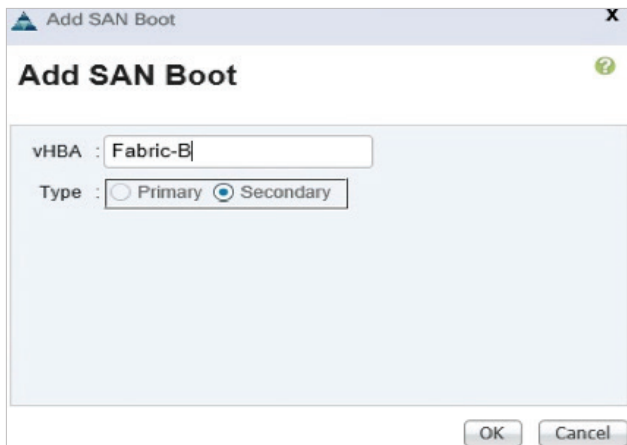


Figure 25. Add SAN boot on Fabric B

26. Click **OK** to add the SAN boot initiator.
27. From the vHBA drop-down menu, select **Add SAN Boot Target**.
28. Enter **1** as the value for **Boot Target LUN**.
29. Enter the WWPN for CT0.FC1.
30. Click the **Primary** radio button for the SAN boot target type.
31. Click **OK** to add the SAN boot target.
32. From the vHBA drop-down menu, select **Add SAN Boot Target**.
33. Enter **1** as the value for **Boot Target LUN**.
34. Enter the WWPN for CT1.FC1

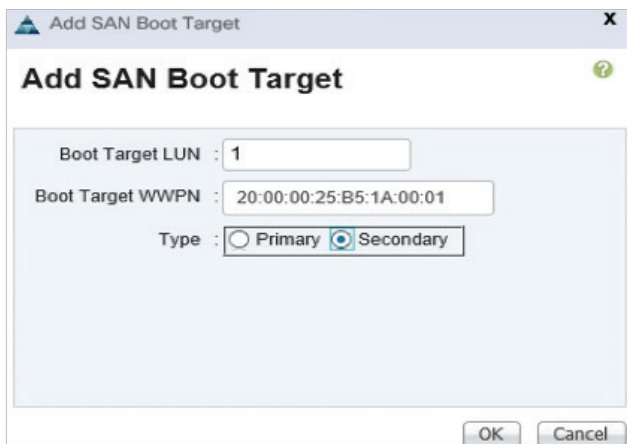
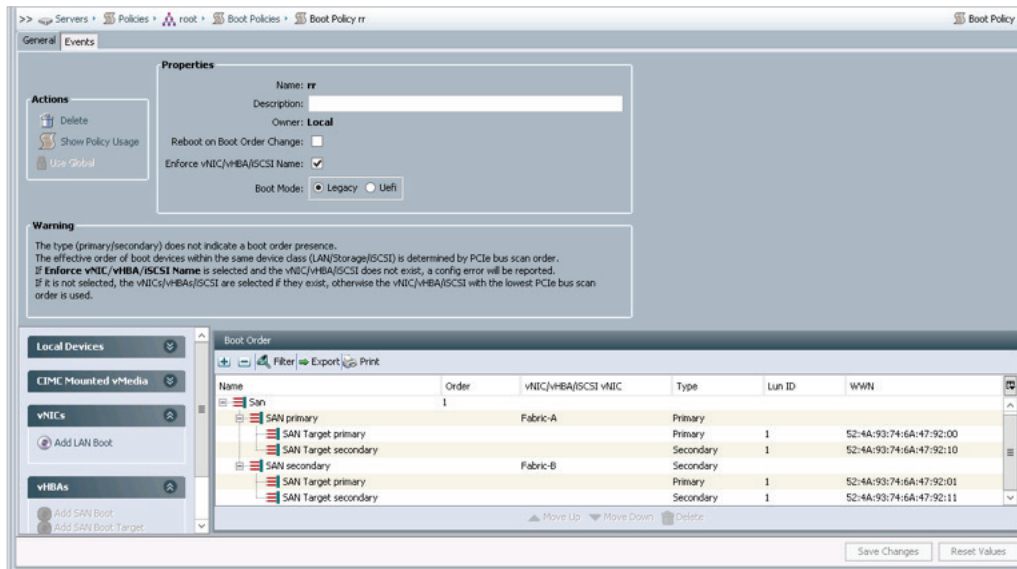


Figure 26. Add SAN boot target for Fabric B

35. Click the **Secondary** radio button for the **Type** option.
36. Click **OK** to add the SAN boot target.



37. Click **OK**, then click **OK** again to create the boot policy.

CREATE SERVICE PROFILE

To create service profile, complete the following steps:

1. In Cisco UCS Manager, in the left navigation pane, click the **Servers** tab.
2. Select **Service Profile > Root**.
3. Right-click root.
4. To open the Create Service Profile Export wizard, select **Create Service Profile Export**.
5. Identify the Service Profile:
 - a. Enter the **Name** for the service profile.
 - b. Under UUID, select the **UUID_Pool**.
 - c. Click **Next**.



Figure 28. Identify service profile

6. Configure the Networking options:
 - a. Retain the default setting for **Dynamic vNIC Connection Policy**.
 - b. Click the **Expert** radio button to configure the LAN connectivity.



Figure 29. Configure networking

- c. Click the upper **Add** button to add a vNIC1.
 - d. In the **Create vNIC** dialog box, enter the **Name** for vNIC.
 - e. Under MAC, Select the **Mac_Pool**.
 - f. In the Fabric ID, select the **Enable Failover**.
 - g. Select the **Management Vlan**.
 - h. Under Pin Group, Select the **Pin Group**.
 - i. From the **Adapter Policy** list, select **Windows**.

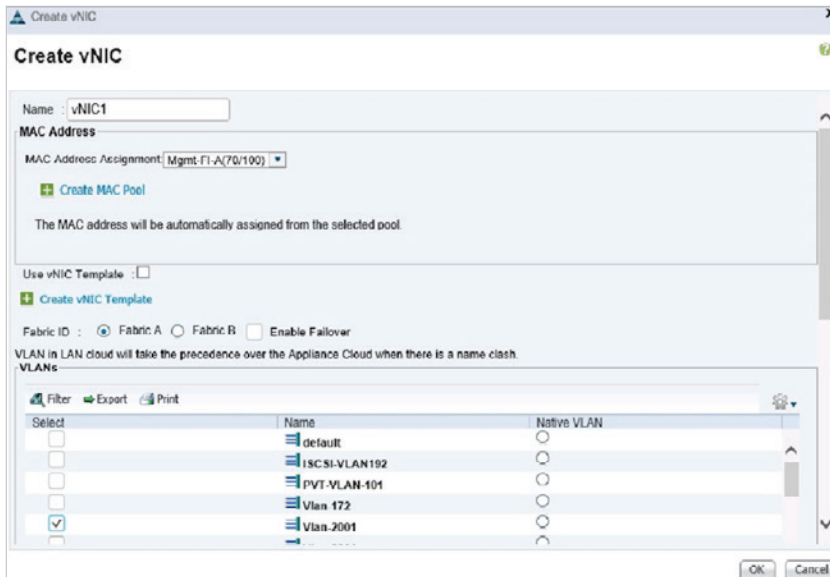


Figure 30. Creating virtual network interface (vNIC)

- j. Click **OK**, and then click the upper **Add** button to add a vNIC2 (above Same Steps).

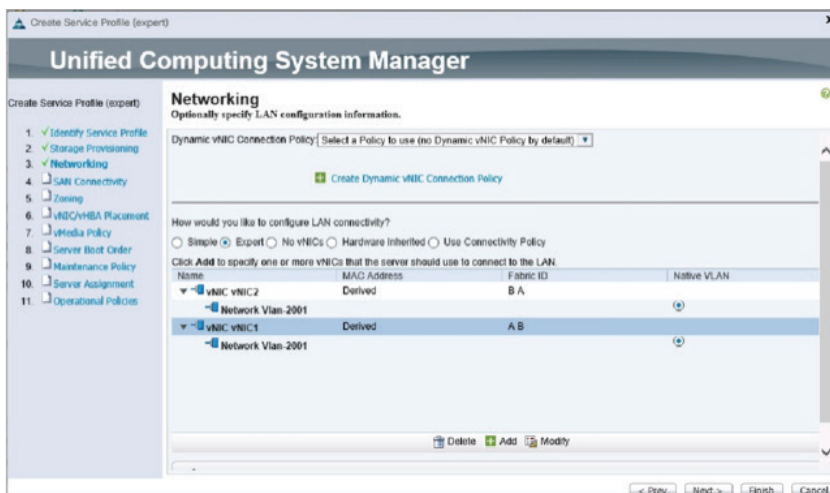


Figure 31. vNIC added

- k. Click **Next**.
7. Configure the Storage options:
 - a. Retain the default setting for **Local Disk Configuration Policy**.
 - b. Click the **Expert** radio button to configure the SAN connectivity.
 - c. From the **WWNN Assignment** list, select **WWNN_Pool**.
 - d. Click the **Add** button to add a vHBA.

- e. In the **Create vHBA** dialog box, enter the **Name** for vHBA1.
- f. Under WWPN, select the **WWPN_Pool**.
- g. In the Fabric ID, select the **Fabric A**.
- h. Under VSAN, **Select the VSAN for Fabric A**.
- i. Under Pin Group, **Select the Pin Group**.
- j. From the **Adapter Policy** list, select **Windows**.

Figure 32. Create virtual host bus adapter (vHBA)

- k. Click **OK**, and then click the upper **Add** button to add a vHBA2 (follow same steps).

Figure 33. SAN connectivity

- I. Click **Next**.
8. Set no Zoning options and click **Next**.
9. Set the vNIC/vHBA placement options:
 - a. Assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
 - vHBA1
 - vHBA2
 - vNIC1
 - vNIC2
 - b. Review the table to verify that all of the vNICs and vHBAs were assigned to the policy in the appropriate order.

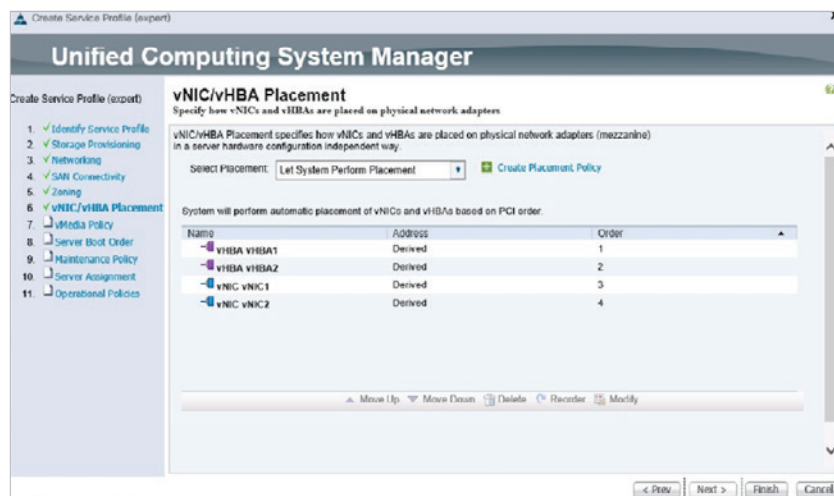


Figure 34. Placement of vNIC and vHBAs

10. Set no vMedia Policy and click **Next**.
11. Set the Server Boot Order:
 - a. From the Boot Policy list, select **Boot-Policy from SAN**.
 - b. Review the table to verify that all the boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.

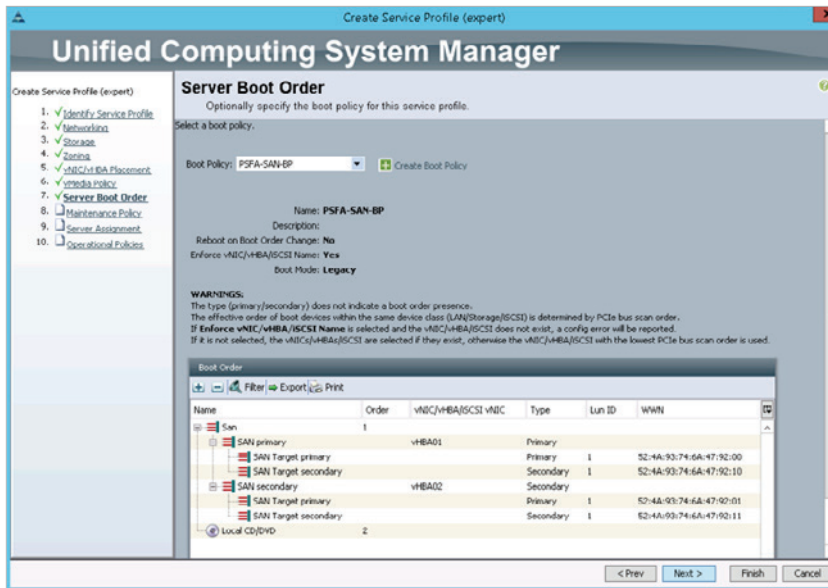


Figure 35. Server boot order view

- c. Click **Next**.
12. Add a Maintenance Policy:
 - a. Confirm that maintenance policy is set to default.
13. Specify the Server Assignment:
 - a. Assign later, Click **Next**.
14. Operational Policies.
 - a. Default BIOS settings.
15. Click **Finish** to create the service profile template.
16. Click **OK** in the confirmation message.

STEP 2 – CONFIGURE BOOT FROM SAN

The first step prior to installing and configuring the Windows Storage Server 2016 hosts is to provision the storage required for SAN boot. Before the storage can be provisioned, you must obtain the connectivity information (Fibre Channel WWNs) for each host so they can be configured on the FlashArray. The quickest way to do this is to use UCS Manager.



Figure 36. Cisco UCS Manager

1. Launch the UCS Manager by navigating to Fully-Qualified Domain Name (FQDN) or IP.
2. After clicking **Launch UCS Manager**, a download prompt for the Java Network Launch Protocol file (.jnlp) will appear. Download and run this file – note that Java Runtime Environment (JRE) 2 must be installed to run the UCS Manager.

It is recommended to use the HTML version of the UCS Manager.

3. Login to the UCS Manager with proper credentials (default user name is “admin”).

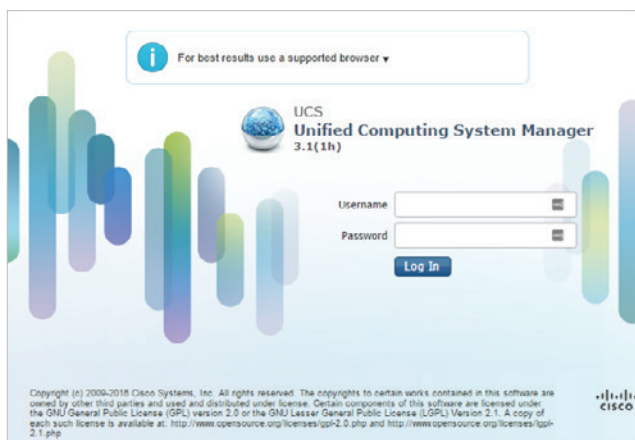


Figure 37. UCS Manager login

- In the **Equipment** tab, select the **Chassis** object and then expand one of the server objects down to the **HBA** object: **Chassis > Servers > Server # > Adapters > HBAs > HBA #**. Select each of the “HBA” objects to record their WWNs.

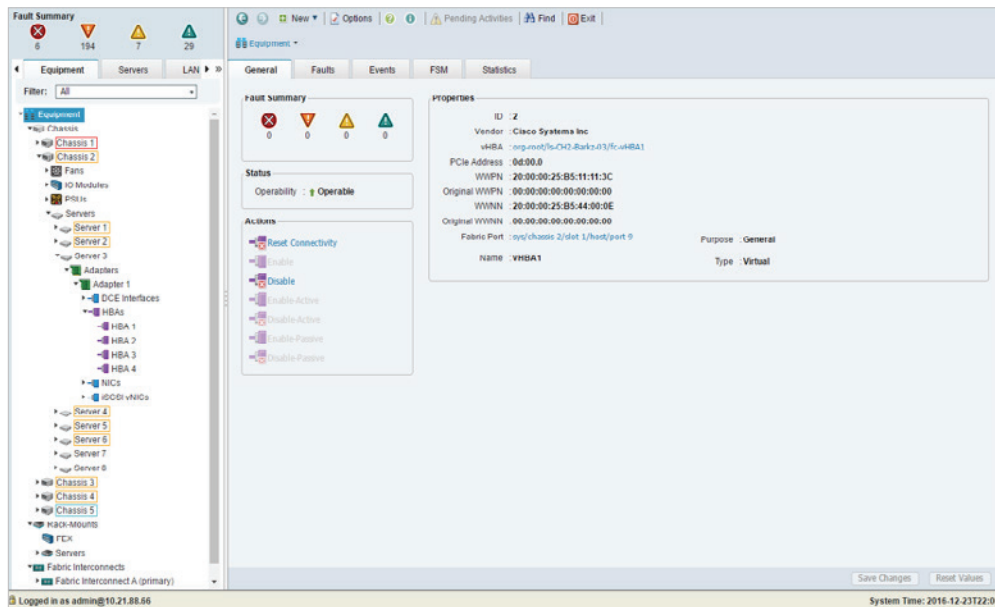


Figure 38. Server 3 showing HBA adapter information

- The WWN is listed on the right panel after selecting an HBA under the **General** tab in the **Properties** box. The correct WWN will be listed as the WWPN (worldwide port name). Gather the WWN of each HBA for each server and note them down. For example, HBA 2 WWPN shown below.

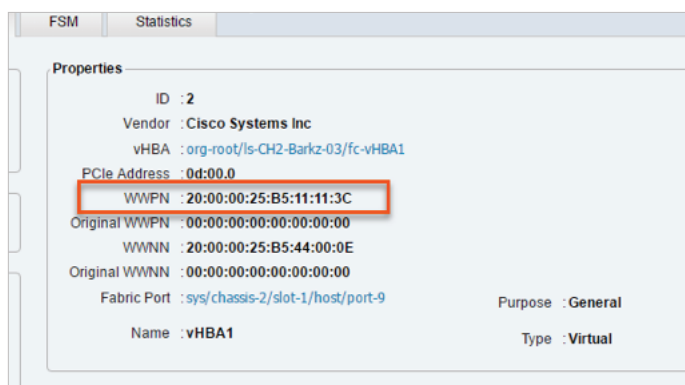



Figure 39. HBA 2 WWPN

- Once all the WWNs have been gathered for the respective hosts, login to the FlashArray GUI via its virtual IP address with a supported internet browser.

- Click on the **Storage** tab in the **Pure Storage FlashArray GUI** and click the plus sign  next to the **Hosts** panel on the left. Click **Create Host Group**; a dialog appears. Provide a name for the host group. For this deployment guide the name **SoFS-HGroup** is given.




Figure 40. Create Host Group





Figure 41. Name new host group

- Now that the host group is created hosts can be added. This will allow for private boot volumes to be provisioned to each host. To do so, select the newly created host group under the **Hosts** panel on the left. Click on the gear icon or as shown the 'hamburger' menu  in the lower right panel that appears in the **Hosts** tab. Select the **Add Hosts** option.

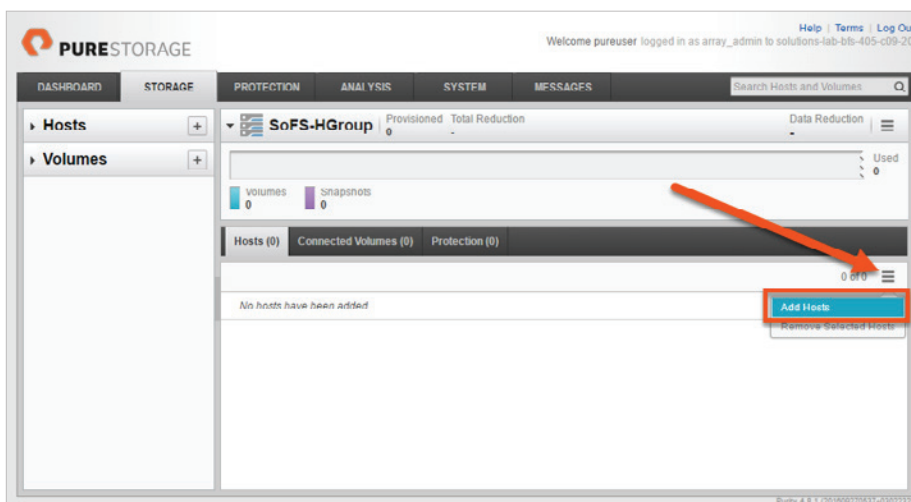


Figure 42. Create host using Add Hosts menu item

- Before the hosts can be added to the host group, they need to be created on the FlashArray. Creation of a host means assigning a friendly name to one or more initiators (WWNs). In the window that popped up, click the **Create New Host** button.

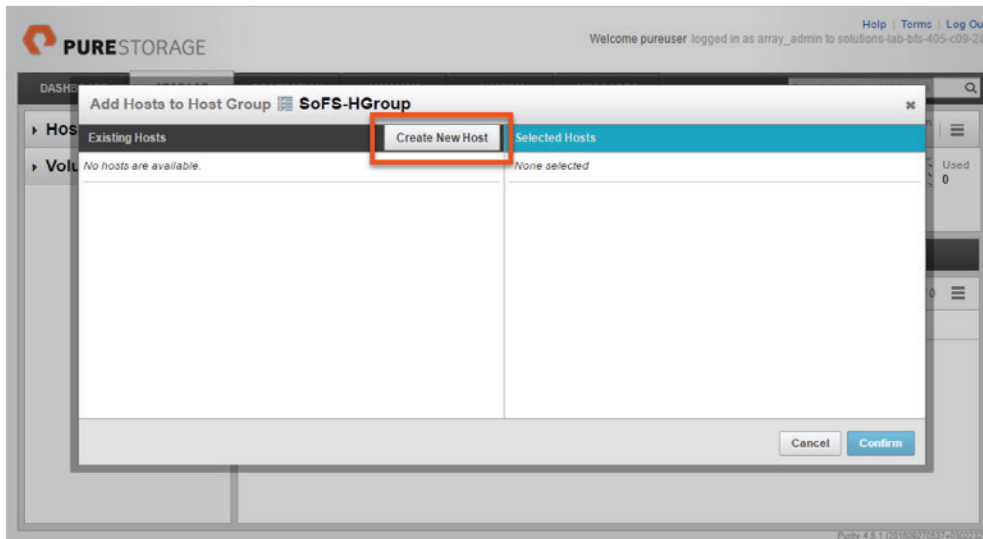


Figure 43. Create new host

10. Click the **Create Multiple** button at the bottom of the dialog. This will allow for quick creation of all hosts. Enter a name for the hosts followed directly (no space) with a # sign (and optionally a suffix). The # sign will be replaced with a number incremented by one for each additional host. Also indicate a starting number, a count (number of hosts to create), and how many digits (a value of 1 would make the first number “1”, a value of 2 would make it “01” etc.).

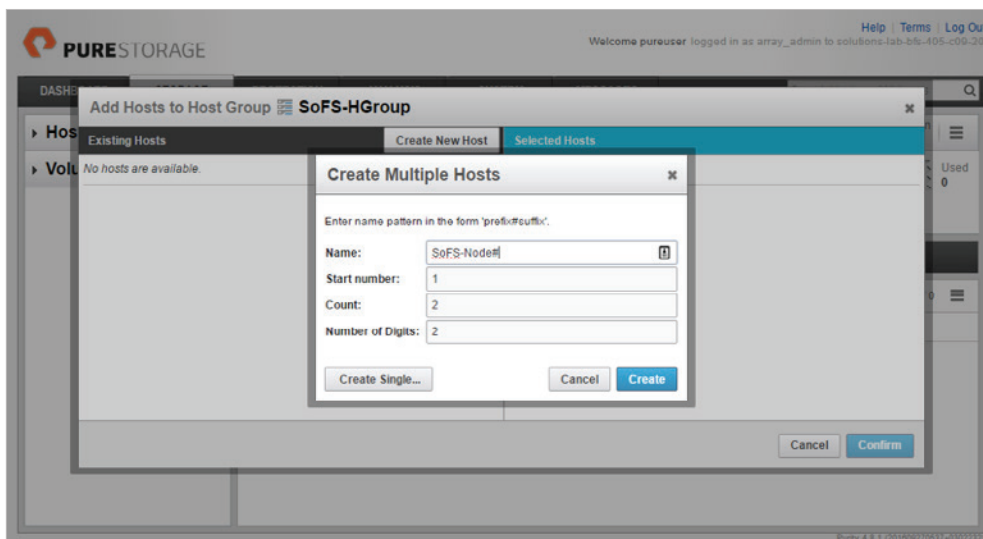


Figure 44. Create multiple hosts

11. Now that all the hosts are added to the host group, the proper WWNs need to be associated with each host, so that the volumes can be connected. Click **Confirm**. The hosts are now populated in the host group's **Hosts** panel. Select the first host by clicking on the name.




Figure 45. Hosts added to host group

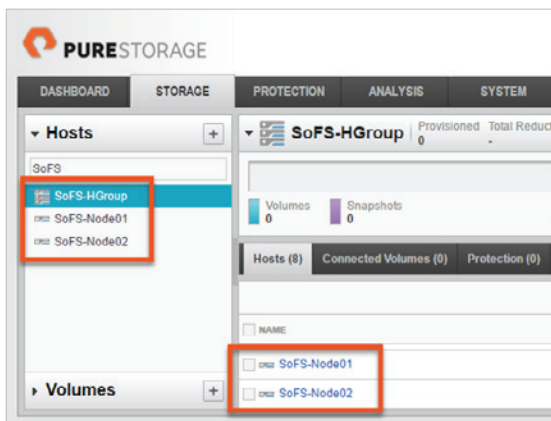


Figure 46. Host group and hosts

12. In the new panel that appears, click on the **Host Ports** tab, then click on the gear icon on the right-hand side and select **Configure Fibre Channel WWNs**.

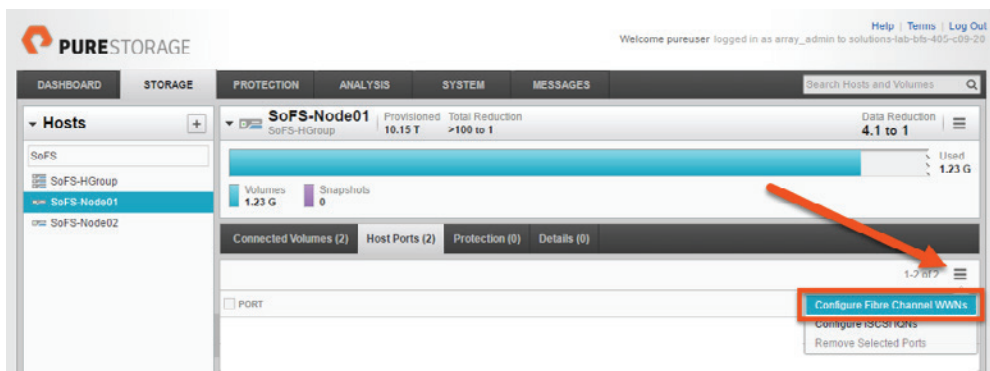


Figure 47. Configure Fibre Channel WWNs

13. In the window that appears, a series of WWNs will appear in the **Existing WWNs** table on the left. Refer to the information gathered earlier and click on the proper WWNs for the selected host to add them. Clicking them will push them from the **Existing WWNs** table on the left to the **Selected WWNs** table on the right. Doing so and clicking the **Confirm** button adds them to the host. Repeat this process for all the hosts.

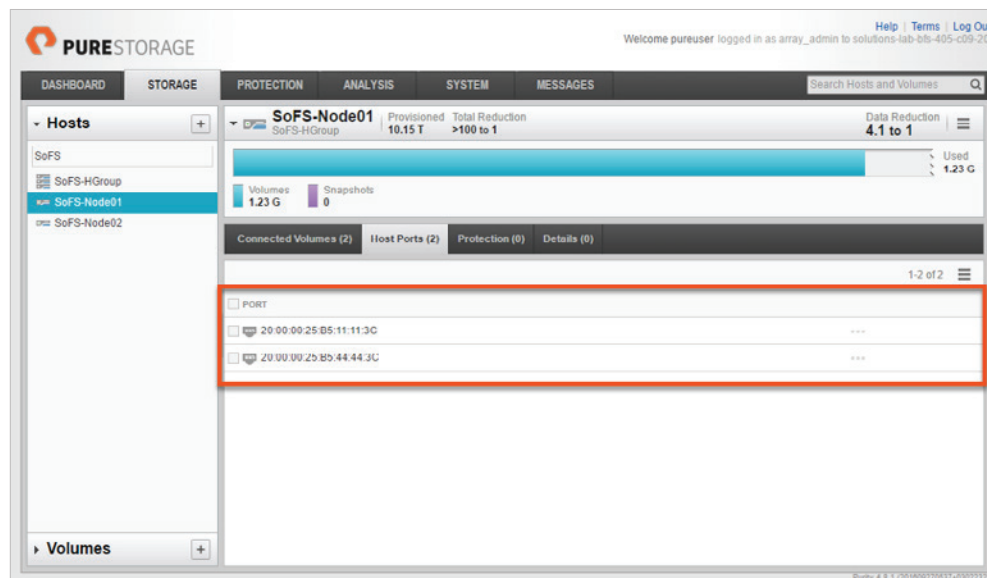


Figure 48. SoFS-Node01 showing two added WWPNS

Once all the hosts have been configured, the boot volumes can be configured. But first, it is important to verify proper and redundant connectivity from each host to the FlashArray. The FlashArray GUI has a very simple way to check connectivity – click on the **System** tab and then **Host Connections**. In the right-hand pane, a list of all the hosts will appear with their WWNs and indicate whether or not that host has redundant connections (connected to both controllers). If one or more hosts do not have the green “Redundant Connections” label, re-check zoning and cabling until they do.

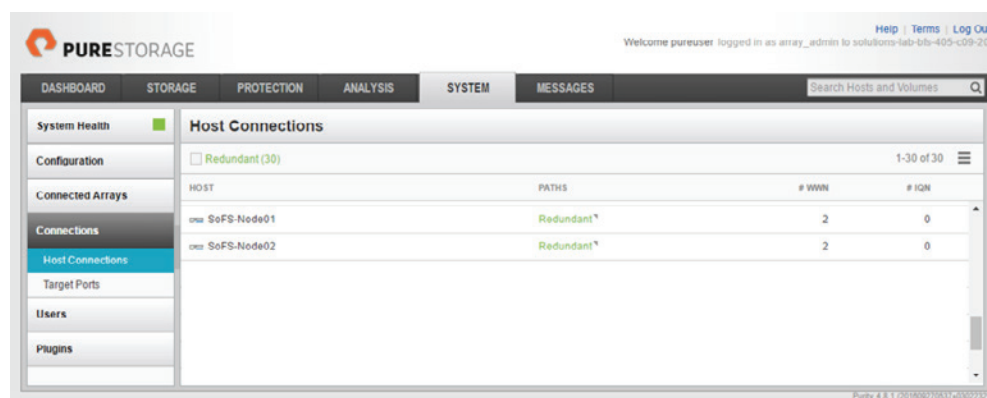


Figure 49. Host Connections

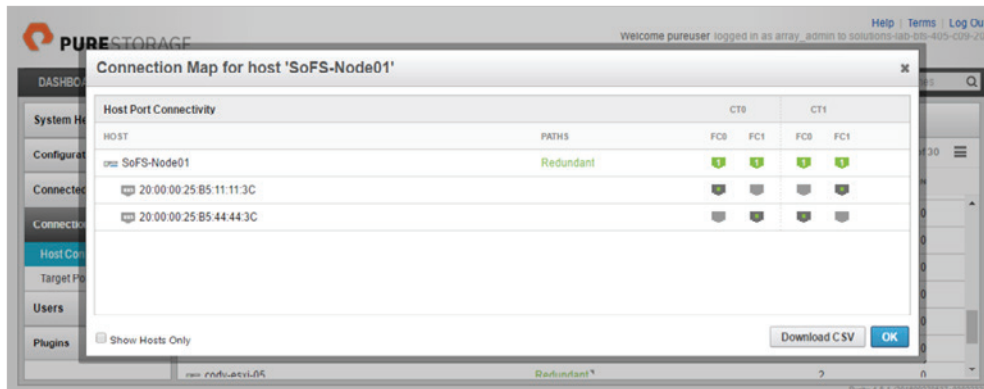


Figure 50. SoFS-Node01 showing CT0 and CT1 FC ports

If all the hosts are redundantly connected to the FlashArray, it is time to create and provision boot volumes to each host. All boot volumes can be created at once, in a similar fashion to how the hosts were created.

1. Click on the **Storage** tab in the FlashArray GUI and then click the plus **+** sign next to the **Volumes** panel on the left. In the dialog that appears, click the **Create Multiple...** button.

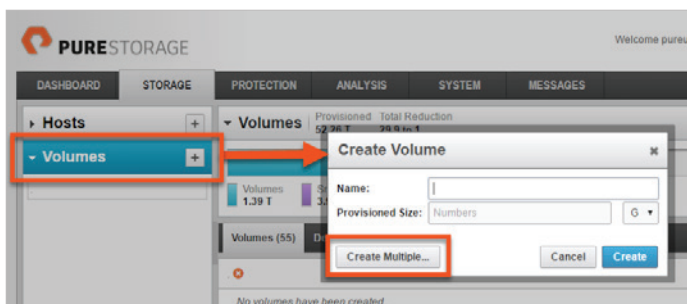


Figure 51. Create multiple volumes

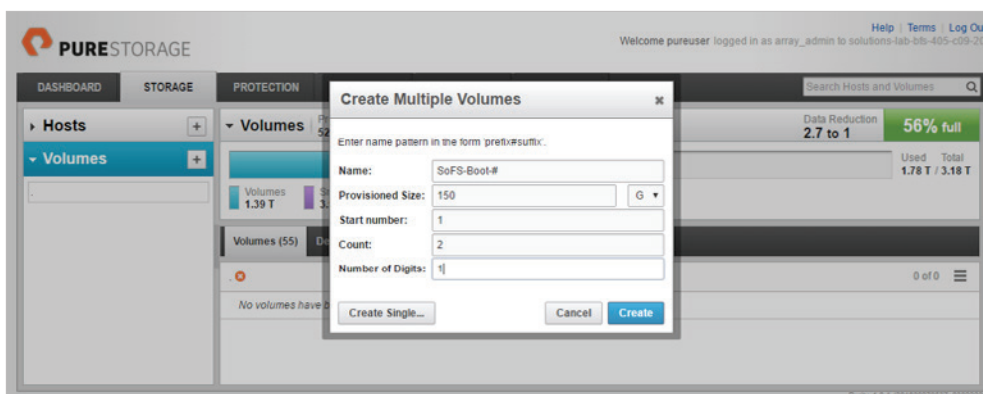


Figure 52. Example of creating multiple SoFS-Boot volumes

2. Enter a name for the volumes, followed directly (no space) with a # sign (and optionally a suffix). The # sign will be replaced with a number incremented by one for each additional volume. Also, indicate a size, starting number, a count (number of volumes to create) and how many digits (a value of 1 would make the first number “1”, a value of 2 would make it “01” etc.).
3. Once the volumes have been created, they can be connected to the hosts. It is important to remember that these volumes should be connected individually, one per each host. Boot volumes should not be connected to the host group object which would effectively share them with every host in the host group. Instead, they should be accessed exclusively by the booting host.
4. Now that the host group is created hosts can be added. This will allow for private boot volumes to be provisioned to each host. To do so, select the newly created host group under the **Hosts** panel on the left. Click on the **Connected Volumes** tab on the right-hand panel that appears. Then click on the ‘hamburger’ icon ☰ in the lower right panel that appears in the Hosts tab. Select the **Add Hosts** option.

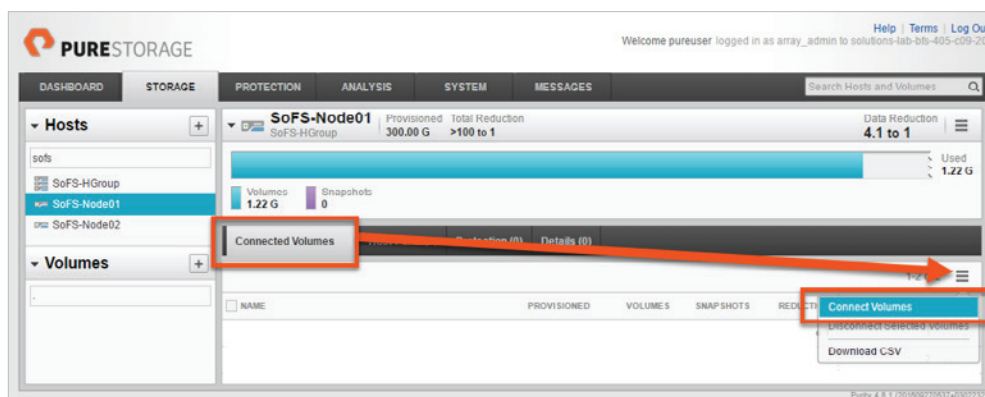


Figure 53. Connect volumes

5. Find the proper boot volume for the selected host and click on the volume in the **Existing Volumes** table on the left to move it to the **Selected Volumes** table on the right. Click the **Confirm** button to complete the operation and to provision the volume. Repeat this process for each host, assigning one unique volume directly to each host. Windows Storage Server 2016 can now be installed.

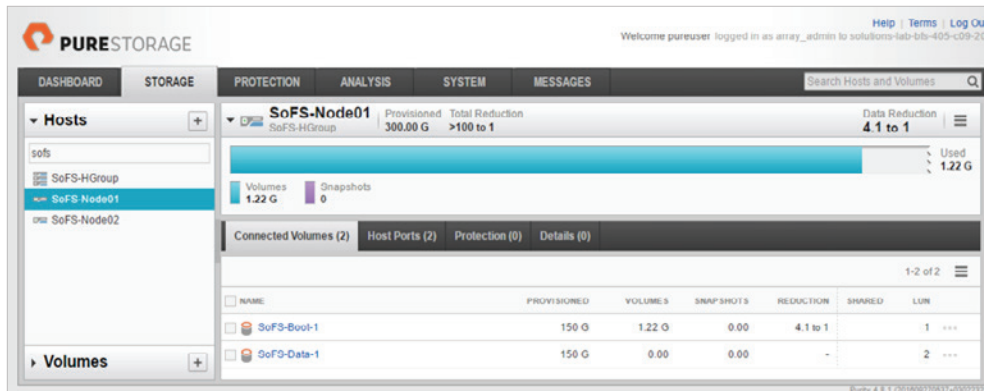


Figure 54. SoFS-Node01 with connected boot and data volumes

This process can be automated using the Pure Storage PowerShell SDK. The automated workflow for Steps 1 – 5 is as follows:

```
$FlashArray = New-PfaArray -EndPoint solutions-lab-bfs-405-c09-20.lab.purestorage.com`
```

```
-Credentials (Get-Credential) -IgnoreCertificateError
```

```
ForEach ($i in 1..2) {
```

```
    $bootVol = New-PfaVolume -Array $FlashArray -VolumeName "SoFS-Boot-$i" `
```

```
    -Unit G -Size 150
```

```
    $dataVol = New-PfaVolume -Array $FlashArray -VolumeName "SoFS-Data-$i" `
```

```
    -Unit G -Size 150
```

```
    New-PfaHostVolumeConnection -Array $FlashArray -VolumeName $bootVol.name `
```

```
    -HostName "SoFS-Node0$i" -LUN 1
```

```
    New-PfaHostVolumeConnection -Array $FlashArray -VolumeName $dataVol.name `
```

```
    -HostName "SoFS-Node0$i" -LUN 2
```

```
}
```

```
Get-PfaVolumes -Array $FlashArray | ?{ $_.name -like '*SoFS-Boot*' }
```

```
Get-PfaHostVolumeConnections -Array $FlashArray -Name 'SoFS-Node01'
```

```
Get-PfaHostVolumeConnections -Array $FlashArray -Name 'SoFS-Node02'
```

STEP 3 – DEPLOY WINDOWS SERVER

After you create Cisco UCS service profiles, you must associate available servers to the created service profiles.

1. Navigate to the **Server** tab in UCS Manager.
2. Locate the Service Profile created in **Step 1 – UCS Manager Configuration**.
3. From the **root** node, right click the <<service profile name>> (Eg. CH2-Barkz-03) and click **Associate with Server Pool**.

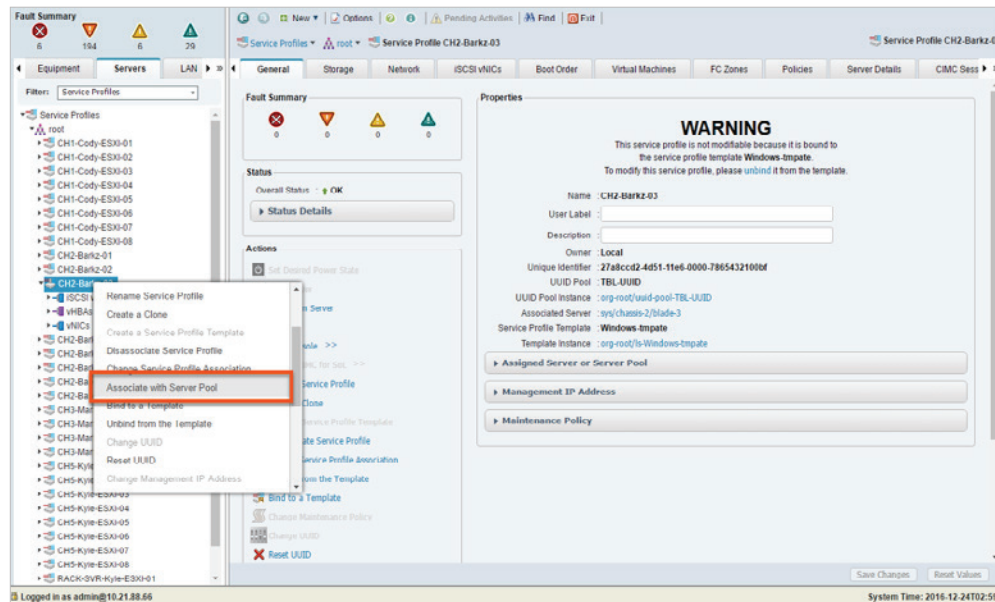


Figure 55. Associate with Server Pool

4. From the **Equipment** console, right click the required server, and select **KVM Console**.

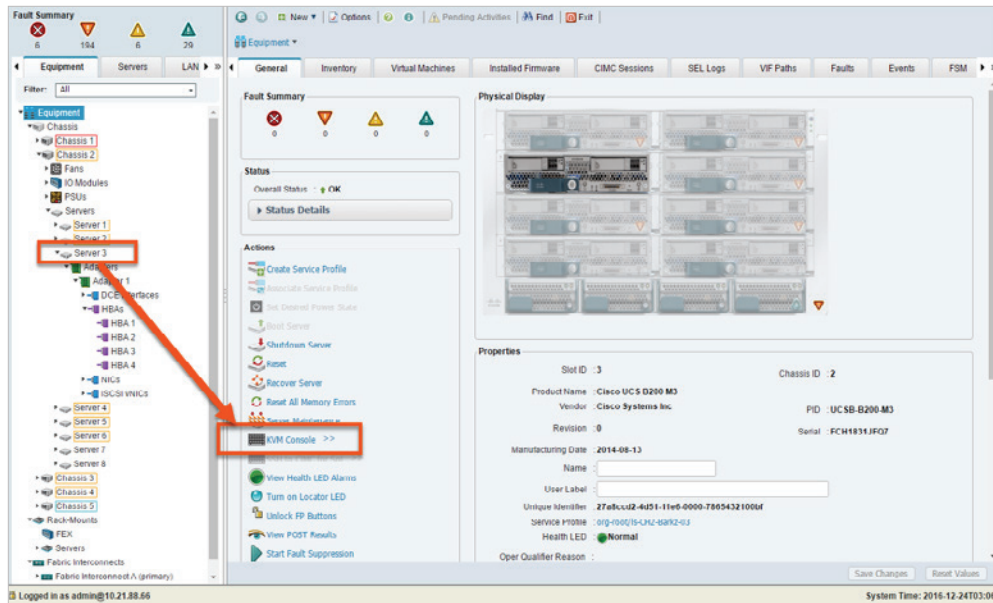


Figure 56. Access the KVM console

5. Ensure that the SAN boot LUNs are discovered and displayed on the BIOS screen.




Figure 57. Pure Storage boot LUNs

6. Configure the **Virtual Media** by **Activating Virtual Devices** and then map the Windows Storage Server 2016 ISO image to automatically start installation after a **Reset**.
7. After starting the installation, select the Cisco VIC driver pack from the appropriate ISO image using the Virtual Device as in Step 6. Windows Storage Server 2016 will continue setup until successful completion.

See the *UCS Driver Installation for Common Operating Systems for VIC Drivers*, <http://www.cisco.com/c/en/us/support/docs/servers-unified-computing/ucs-manager/116349-technote-product-00.html>.




Figure 58. UCS Boot Drivers

STEP 4 – DEPLOYING WINDOWS SERVER FAILOVER CLUSTER

1. After successful completion of Windows Storage Server 2016 installation and reboots from the previous section, start **Server Manager** and select the **Local Server** within **Server Manager > Add Roles and Features Wizard**. From the Features section, enable **Multipath I/O** and **Failover Clustering**.

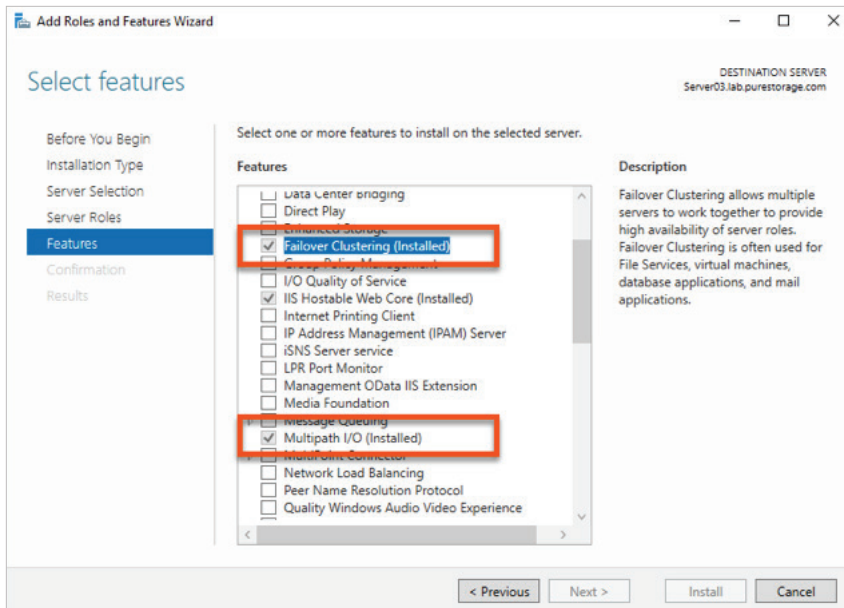


Figure 59. Add Roles and Features Wizard

2. Click **Next** and complete the installation process. This will not require a reboot.
3. Open a Windows PowerShell console and enter:
New-MSDSMSupportedHW -VendorId 'PURE' -ProductId 'FlashArray'.
4. From **Server Manager**, click **Tools > MPIO**. The MPIO Properties screen appears.
 From the **MPIO Devices** tab, ensure that the Pure Storage FlashArray is discovered.

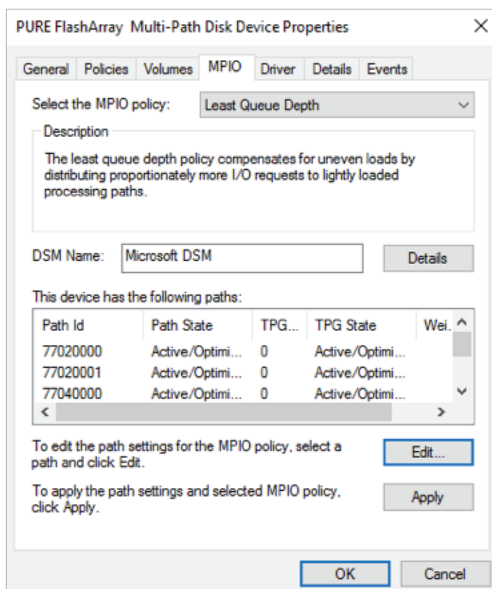


Figure 60. Windows Storage Server 2016 MPIO Properties

- ```
mpclaim -s -d
```

```

Administrator: Command Prompt
Microsoft Windows [Version 10.0.14393]
(c) 2016 Microsoft Corporation. All rights reserved.

C:\Windows\system32>mpclaim -s -d

For more information about a particular disk, use 'mpclaim -s -d #' where # is the MPIO disk number.

MPIO Disk System Disk LB Policy DSM Name

MPIO Disk5 Disk 5 RR Microsoft DSM
MPIO Disk4 Disk 4 RR Microsoft DSM
MPIO Disk3 Disk 3 RR Microsoft DSM
MPIO Disk2 Disk 2 RR Microsoft DSM
MPIO Disk1 Disk 1 RR Microsoft DSM
MPIO Disk0 Disk 0 RR Microsoft DSM

C:\Windows\system32>

```

Figure 60 shows that all MPIO Disks are using the Load Balancing policy of Least Queue Depth (LQD). This is documented as part of the Pure Storage Windows Server Best Practices located at <https://support.purestorage.com>.

- 
- The screenshot shows the 'Connection Map for host 'SoFS-Node01'' window in the Pure Storage Connect console. The window displays the host port connectivity for the host 'SoFS-Node01'. The connectivity is shown as redundant, with green status indicators for all paths. The window also includes a 'Download CSV' button and a 'Show Hosts Only' checkbox.
- | HOST                    | PATHS     | CT0 |     | CT1 |     |
|-------------------------|-----------|-----|-----|-----|-----|
|                         |           | FC0 | FC1 | FC0 | FC1 |
| one SoFS-Node01         | Redundant |     |     |     |     |
| 20:00:00:25:85:11:11:3C |           |     |     |     |     |
| 20:00:00:25:85:44:44:3C |           |     |     |     |     |



**PURE**STORAGE®

## DEPLOYING WINDOWS SERVER FAILOVER CLUSTERS

### CREATING FAILOVER CLUSTERING FEATURE

1. Start Server Manager.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or Feature-based** installation, and then click **Next**.
5. On the **Select destination server** page, click the server where you want to install the feature, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, select the **Failover Clustering** check box.
8. To install the failover cluster management tools, click **Add Features**, and then click **Next**.
9. On the **Confirm installation selections** page, click **Install**.
10. When the installation is completed, click **Close**.
11. Repeat this procedure on every server that you want to add as a failover cluster node.

### VALIDATE THE CLUSTER CONFIGURATION

1. Before you create a failover cluster, you must validate the nodes have been configured as required.
2. To do this on a server, start Server Manager, and then on the **Tools** menu, click **Failover Cluster Manager > Create Cluster Wizard**.

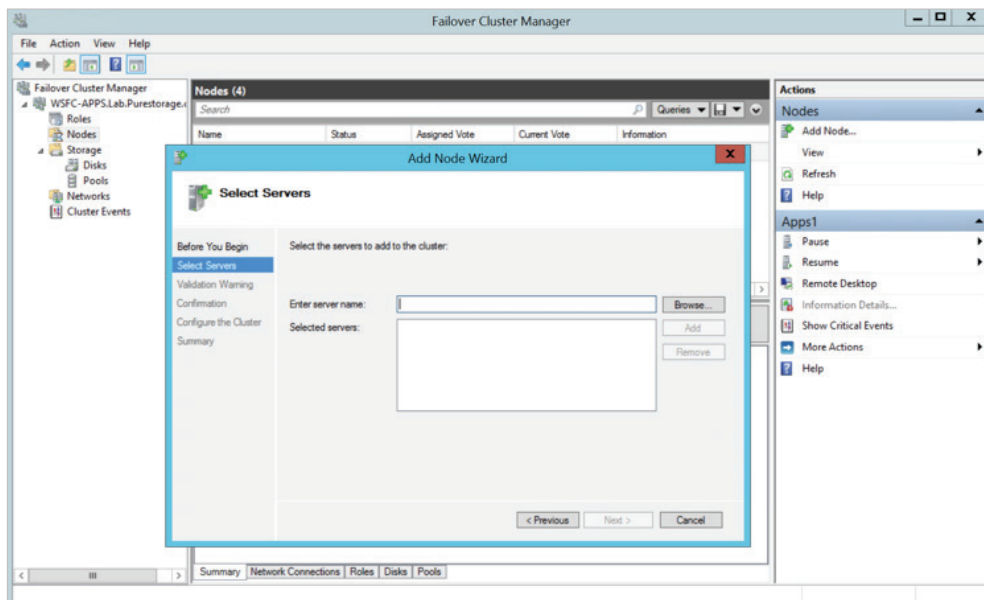


Figure 63. Create Cluster Wizard

3. If the **Select Servers** page appears, in the **Enter server name** textbox, enter the NetBIOS name or the fully qualified domain name (FQDN) of a server to participate in the failover cluster, then click **Add**.
4. From the **Validate Warning** page, click **Run all tests (recommended)**, and then click **Next**.

*Do not select the **Reboot on Boot Order Change** checkbox.*

5. On the **Access Point for Administering the Cluster** page, In the **Cluster Name** box, enter a name for the cluster.
6. Click **Next** and **Finish**.

Once you've completed creating the clusters, verify that the storage, nodes, disk, and network are up and running:

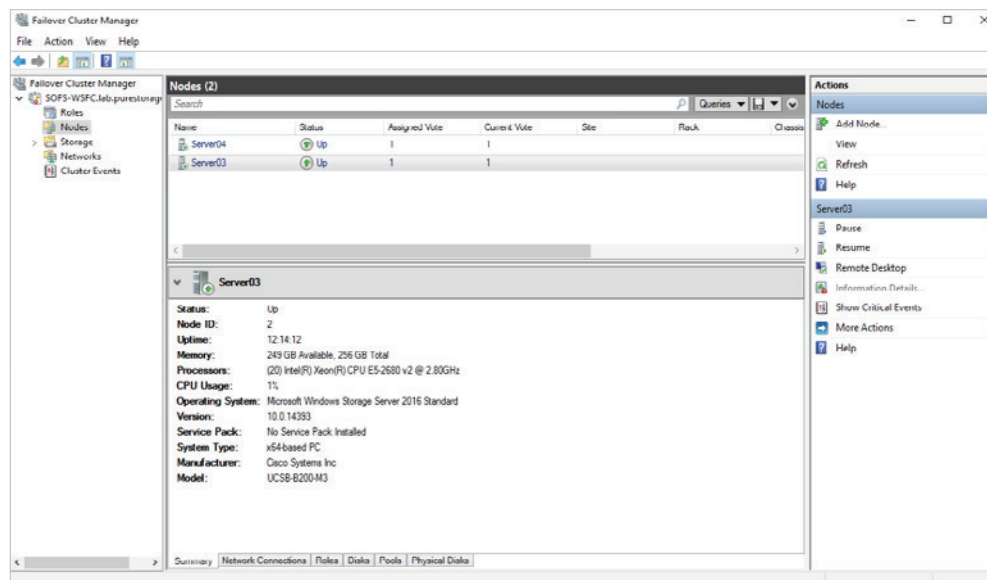


Figure 64. Windows Server Failover Cluster

## STEP 5 – CONFIGURING CONTINUOUSLY AVAILABLE FILE SHARES

Before you begin creating Continuously Available File Shares (CAFS), you must ensure that the following pre-requisites are met:

- Operating System is Windows Storage Server 2016 and at least a two-node Windows Server Failover Cluster.
- The File and Storage Services and File Server role must be installed on all cluster nodes.
- SMB client computers must be running Windows 8 or higher and/or Windows Server 2012 or higher to take advantage of the new SMB Transparent Failover capability.
- The clustered file server must be configured with one or more file shares that use the new continuously available setting.

*We recommend that you design your network to contain multiple pathways between nodes to prevent the network from becoming a single point of failure. You can also consider using network adapter teaming and multiple switches and/or redundant routers to add resiliency to your network configuration.*

7. From any of the nodes in the cluster, open **Failover Cluster Manager**, and then right-click the **Roles** node in the left navigation pane. This opens the **High Availability Wizard**.

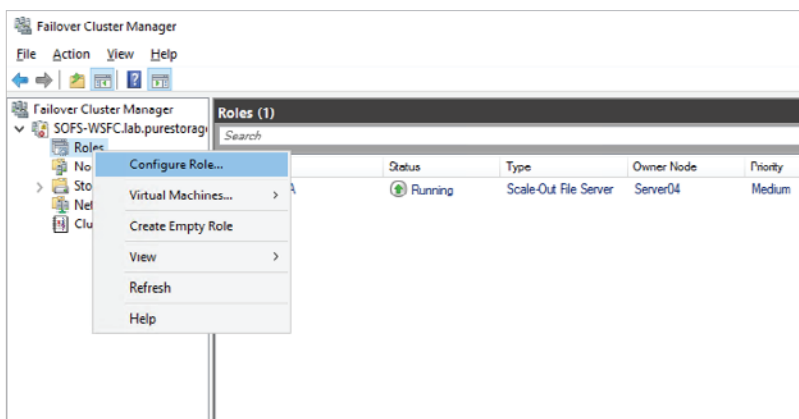


Figure 65. Configure a Role in Windows Server Failover Cluster

8. From the **Select Role** step, choose **File Server**. The file server role supports both the general purpose and scale-out application types of CAFS. Click **Next**.

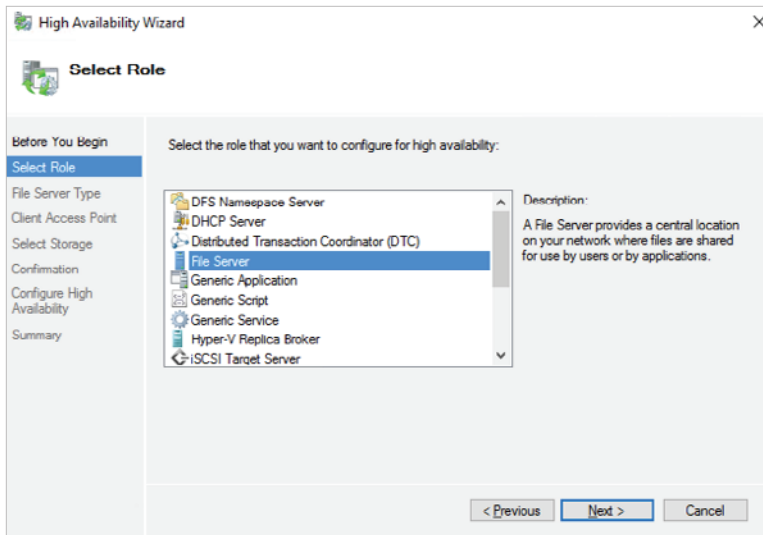


Figure 66. High Availability Wizard

9. From the **File Server Type** page, choose the option **File Server for General Use**. This option can be used for both Windows SMB-based file shares and NFS-based file shares. It also supports data deduplication, DFS replication, and data encryption. Click **Next**.

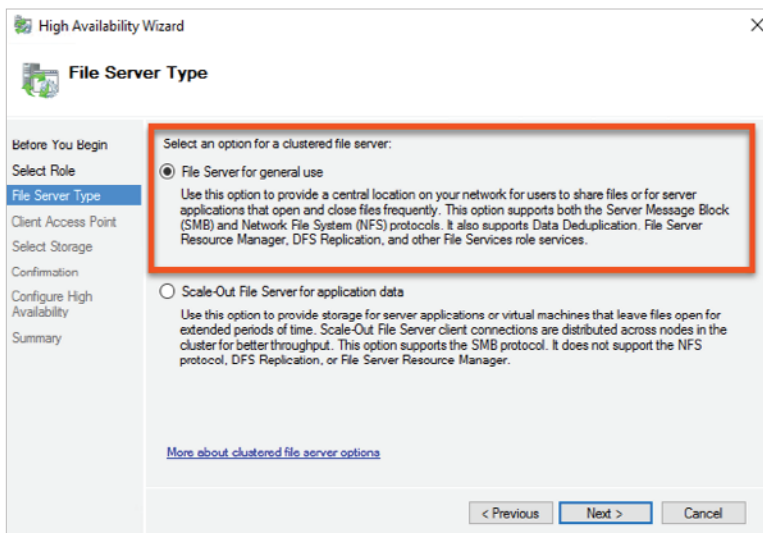


Figure 67. File Server Type

10. In the **Client Access Point** page, provide a server name and an IP address. Clients use the server name when they access the CAFS. This name will be registered in your DNS, and clients will use it like a server name. In addition, the general purpose CAFS also needs an IP address. Click **Next**.



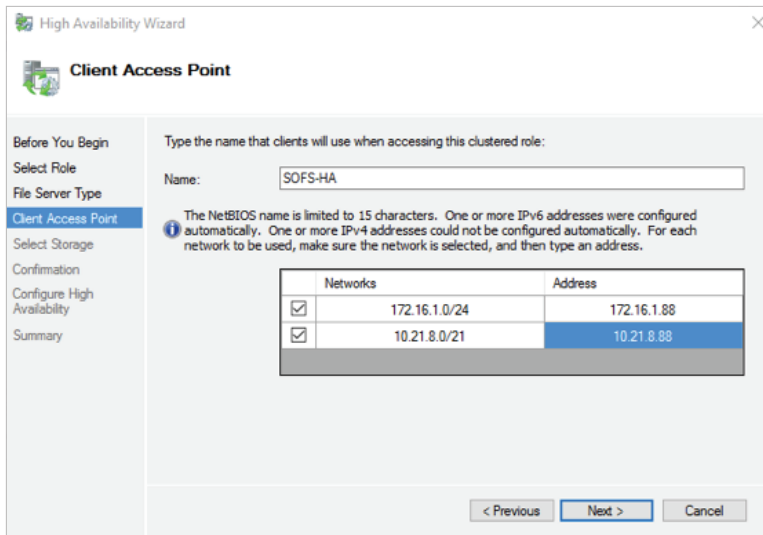


Figure 68. Client Access Point

11. From the **Select Storage** page, select the storage that will be available to the cluster. The selected storage must be listed under the cluster's storage node and designated as available storage. Click **Next**.

*Preassigned Clustered Shared Volumes (CSV) cannot be used for your general purpose Continuously Available File Shares.*

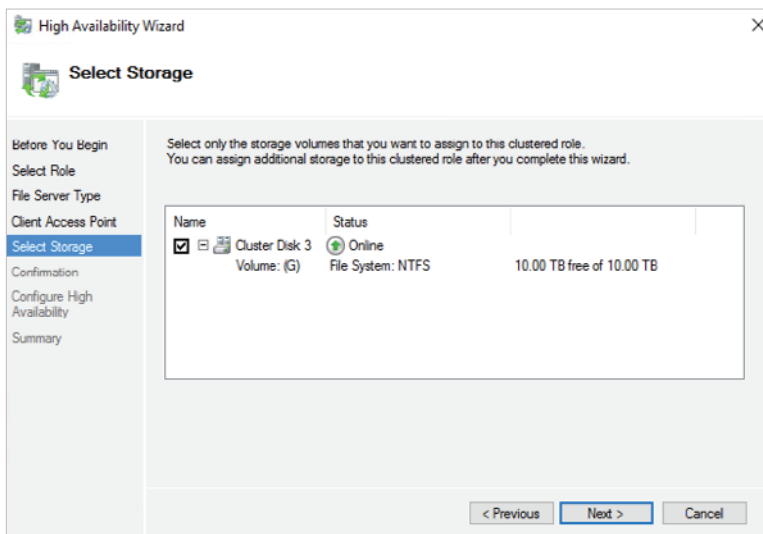


Figure 69. Select Storage

12. On the **Confirmation** page, review your settings. Click **Next** if you do not want to make further changes.

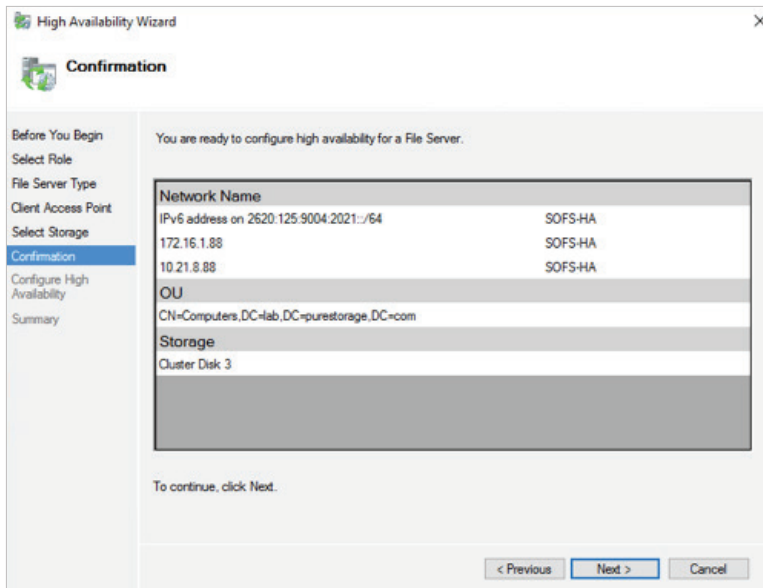


Figure 70. Confirmation for HA File Server creation

- The **Configure High Availability** dialog box opens and displays the status of the configuration process. When the process is complete, the **Summary** screen appears. Click **Finish**.

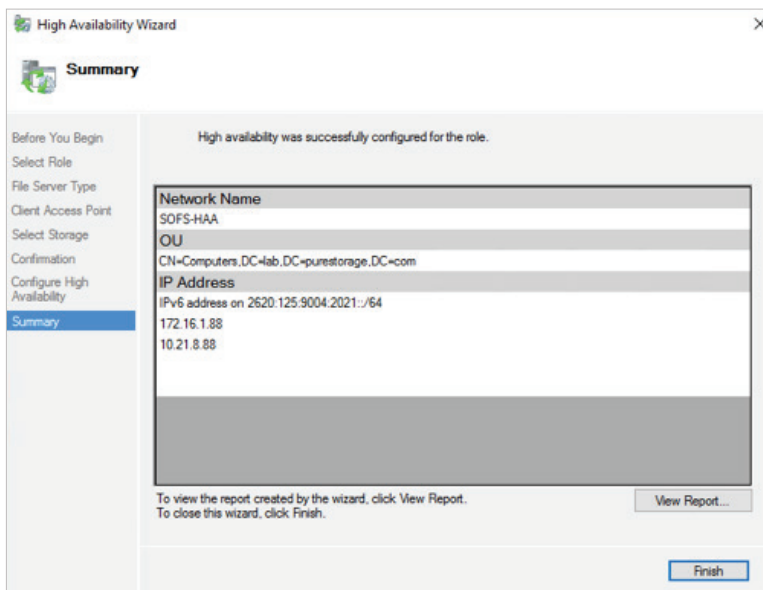


Figure 71. File Server completed

# CONCLUSION

Scalability, agility, reliability, performance, and management are just some of the core bullet points that IT managers and administrators require in the data center. Using the Scale-Out File Services Design Guide provides a reference for compute, networking, and storage platform that provides all the core bullet points mentioned, along with the ability to scale and grow as business demands require.

The Scale-Out File Services Design Guide can be implemented with any networking or compute hardware to build out a file share environment for small to large user basis, and also used for other application workloads.

# APPENDIX 1 – PURE STORAGE COMPONENTS

This section provides a detailed description of each infrastructure component.

## FLASHARRAY//M

FlashArray//M20, one of the mid-range models of FlashArray//M, is a solid state storage disk system that contains multiple flash memory drives instead of spinning hard disk drives. FlashArray//M can transfer data to and from solid state drives (SSDs) much faster than electromechanical disk drives.

FlashArray//M is powered by software that is purpose built for flash – delivering industry-leading 5:1 average data reduction, proven resiliency with non-disruptive operations and full performance, and disaster recovery & protection built-in. With the Evergreen™ Storage model, storage grows and evolves per business needs, providing value for a decade or more. Capacity, performance, and features can be upgraded non-disruptively without having to migrate anything.

With FlashArray//M, there is no need for any training or tuning – the installation and use of FlashArray//M is almost immediate. Additionally, the device can be managed from any iOS or Android device.



Figure 72. FlashArray//M20

The following table lists FlashArray//M20 specifications:

| FEATURE      | SPECIFICATION                                                                                                                                                                                             |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Capacity     | <ul style="list-style-type: none"><li>• 15 – 250+ TBs effective capacity</li><li>• 5 – 80 TBs raw capacity</li></ul>                                                                                      |
| Performance  | <ul style="list-style-type: none"><li>• Up to 200,000 32K IOPS @ &lt; 1ms average latency</li><li>• Up to 6 GB/s bandwidth</li></ul>                                                                      |
| Connectivity | <ul style="list-style-type: none"><li>• 2-port &amp; 4-port 16 Gb/s Fibre Channel (Onboard)</li><li>• 4 x 10 Gb/s Ethernet iSCSI (Configurable)</li><li>• 4 x 1 Gb/s Management and Replication</li></ul> |
| Physical     | <ul style="list-style-type: none"><li>• 3U</li><li>• 600 Watts (nominal draw)</li><li>• 95 lbs. (47.6 kg)</li><li>• 5.12" x 18.94" x 29.72" chassis</li></ul>                                             |

Table 1. FlashArray//M20 specifications

Check out complete FlashArray//M20 specifications with the following link.

<https://www.purestorage.com/products/flash-array-m/hardware-tech-spec-flash-array.html>

## PURITY OPERATING ENVIRONMENT

Purity implements advanced data reduction, storage management, and flash management features. All Purity features are built-in and included in the base cost of FlashArray//M.

- **Storage Software Built for Flash** – FlashCare technology virtualizes the entire pool of flash within the FlashArray, and allows Purity both to extend the life of, and ensure maximum performance from, consumer- grade MLC flash.
- **Granular and Adaptive** – Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity's data and flash management services to operate at the highest efficiency.
- **Best Data Reduction Available** – FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte aligned variable block size to enable effective reduction across a wide range of mixed workloads without tuning.
- **Highly Available and Resilient** – FlashProtect implements high availability, dual-parity RAID-3D, non- disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashArray during any failure or maintenance event.

- **Backup and Disaster Recovery Built In** – FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss locally and globally. All FlashProtect services are fully-integrated in the FlashArray and leverage the native data reduction capabilities.

## PURE1® – COMPONENT FEATURES

- **Pure1 Manage** – Combining local web-based management with cloud-based monitoring, Pure1 Manage allows you to manage your FlashArray wherever you are – with just a web browser or a smartphone mobile app.
- **Pure1 Connect** – A rich set of APIs, plugin-is, application connectors, and automation toolkits enable you to connect FlashArray//M to all your data center and cloud monitoring, management, and orchestration tools.
- **Pure1 Support** – FlashArray//M is constantly cloud-connected, enabling Pure Storage to deliver the most proactive support experience possible. Highly trained staff combined with big data analytics help resolve problems before they start.
- **Pure1 Collaborate** – Extend your development and support experience online, leveraging the Pure1 Collaborate community to get peer-based support, and to share tips, tricks, and scripts.

## APPENDIX 2 – CISCO COMPONENTS

The Cisco Unified Computing System™ (Cisco UCS) is a next-generation computing solution for small- to medium-sized businesses that unites compute, network, storage access, and virtualization into an organized structure designed to reduce total cost of ownership and introduce vastly improved infrastructure deployment mechanisms at scale. UCS incorporates a unified network fabric with scalable, modular, and powerful x86-architecture servers. With an innovative and proven design, Cisco UCS delivers an architecture that increases cost efficiency, agility, and flexibility beyond what traditional blade and rack-mount servers provide. Cisco makes organizations more effective by addressing the real problems that IT managers and executives face and solving them on a systemic level.



Figure 73. Cisco Unified Computing System

### **GREATER TIME-ON-TASK EFFICIENCY**

Automated configuration can change an IT organization's approach from reactive to proactive. The result is more time for innovation, less time spent on maintenance, and faster response times. These efficiencies allow IT staff more time to address strategic business initiatives. They also enable better quality of life for IT staff, which means higher morale and better staff retention – both critical elements for long-term efficiency.

Cisco UCS Manager is an embedded, model-based management system that allows IT administrators to set a vast range of server configuration policies, from firmware and BIOS settings to network and storage connectivity. Individual servers can be deployed in less time and with fewer steps than in traditional environments. Automation frees staff from tedious, repetitive, time-consuming chores – often the source of errors that cause downtime – making the entire data center more cost-effective.

### **EASIER SCALING**

Automation means rapid deployment, reduced opportunity cost, and better capital resource utilization. With Cisco UCS, rack-mount and blade servers can move from the loading dock and into production in a “plug-and-play” operation. Automatically configure blade servers using predefined policies simply by inserting the devices into an open blade chassis slot. Integrate rack-mount servers by connecting them to top-of-rack Cisco Nexus® fabric extenders. Since policies make configuration automated and repeatable, configuring 100 new servers is as straightforward as configuring one server, delivering agile, cost-effective scaling.

### **VIRTUAL BLADE CHASSIS**

With a separate network and separate management for each chassis, traditional blade systems are functionally an accidental architecture based on an approach that compresses all the components of a rack into each and every chassis. Such traditional blade systems are managed with multiple management tools that are combined to give the illusion of convergence for what is ultimately a more labor-intensive, error-prone,

and costly delivery methodology. Rack-mount servers are not integrated and must be managed separately or through additional tool sets, adding complexity, overhead, and the burden of more time.

## CISCO UCS 5108 CHASSIS

The Cisco UCS Mini chassis is a 6RU chassis that can accommodate up to 8 half-width blades. The chassis is a Cisco UCS 5108 Blade Server chassis that allows the two I/O bays at the rear of the chassis to accommodate Fabric Extenders such as the UCS 2208XP or UCS 6324 Fabric Interconnect modules. Cisco UCS Mini works on Intel Xeon processor E5-2600 with up to 36 CPU processors of 1.5 TB memory.

Having a unified fabric reduces the number of network interface cards (NICs), hosted bus adapters (HBAs), switches, and cables needed. Cisco UCS 5108 also offers support for one or two Cisco UCS 2100 Series or Cisco UCS 2200 I/O modules.



Figure 74. Cisco UCS 5108 Chassis

## CISCO UCS FABRIC INTERCONNECT 6248

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a core part of the Cisco Unified Computing System. Typically deployed in redundant pairs, the Cisco UCS 6248UP Fabric Interconnects provide uniform access to both networks and storage. It is highly scalable and flexible infrastructure.



Figure 75. Cisco UCS FI 6248

## CISCO NEXUS 5000

The Cisco® 5000 series switches are designed to meet the scaling demands of virtualized and cloud deployments. These switches support up to 2304 ports in a single management domain with Cisco FEX architecture. They support large buffers for congestion management, provide a hardware-based VXLAN (Layer 2, Layer 3,



and gateway), and are capable of network virtualization generic routing encapsulation (NVGRE). They deliver integrated Layer 3 services with large table sizes and buffers with 1-microsecond latency.



Figure 76. Cisco Nexus 5000 series switch

## CISCO UCS B200-M4 SERVERS

Cisco UCS B200 M4 server is a half-width blade server. Up to eight can reside in the 6-rack-unit (6RU) Cisco UCS 5108 Blade Server Chassis, offering one of the highest densities of servers per rack unit of blade chassis in the industry.

Cisco UCS B200 M4 is a density-optimized, half-width blade server that supports two CPU sockets for Intel E5-2600 v3 series CPUs and up to 24 DDR4 DIMMs. It supports one modular LOM (dedicated slot for Cisco's Virtual Interface Card) and one mezzanine adapter.



Figure 77: Cisco UCS B200 M4 Blade Server

In addition, Cisco UCS has the architectural advantage of not having to power and cool excess switches, NICs, and HBAs in each blade server chassis. Having a larger power budget per blade server provides uncompromised expandability and capabilities, as in the new Cisco UCS B200 M4 server with its leading memory-slot and drive capacity.

© 2017 Pure Storage, Inc. All rights reserved. Pure Storage, the "P" Logo, Pure1, and Evergreen Storage are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. Microsoft File Server and Microsoft Windows Storage Server are registered trademarks of Microsoft in the U.S. and other countries. Cisco UCS and Cisco Nexus are registered trademarks of Cisco in the U.S. and other countries. The Pure Storage product described in this documentation is distributed under a license agreement and may be used only in accordance with the terms of the agreement. The license agreement restricts its use, copying, distribution, decompilation, and reverse engineering. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

ps\_dg\_scale-out-file-services\_01



Pure Storage, Inc.

Twitter: [@purestorage](https://twitter.com/purestorage)

[www.purestorage.com](http://www.purestorage.com)

650 Castro Street, Suite #260  
Mountain View, CA 94041

T: 650-290-6088

F: 650-625-9667

Sales: [sales@purestorage.com](mailto:sales@purestorage.com)