

SAN Boot from HyperFlex iSCSI: UCS Servers Configuration Examples

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Configure](#)

[High-Level Overview of the SAN Boot from HyperFlex iSCSI](#)

[Configure HyperFlex](#)

[Configure UCS Standalone Server - CIMC](#)

[Configure UCS Manager](#)

[Configure IMM](#)

[Boot from HyperFlex iSCSI Target with MPIO](#)

[Windows OS Installation on iSCSI Boot LUN](#)

[Esxi OS Installation on iSCSI Boot LUN](#)

Introduction

This document describes the configuration of HyperFlex (HX), a standalone Unified Computing System (UCS) Server, UCS Server in Intersight Managed Mode (IMM), and UCS Managed Server in order to do a Storage Area Network (SAN) Boot from HyperFlex Internet Small Computer Systems Interface (iSCSI).

Contributed by Joost van der Made, Cisco TME, and Zayar Soe, Cisco Software QA Engineer.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- The UCS is initialized and configured; refer to the UCS configuration guide
- HyperFlex cluster is created; refer to HyperFlex configuration guide
- iSCSI network is configured with a VLAN; refer to HyperFlex configuration iSCSI configuration part (record those iSCSI target IP addresses, VLAN, Challenge Handshake Authentication Protocol (CHAP) information in order to use in this configuration guide)
- The Network Interface Controller or Card (NIC) cards must be Cisco Virtual Interface Card (VIC) 1300 or 1400 series

Components Used

The information in this document is based on these software and hardware versions:

- HyperFlex Data Platform (HXDP) 4.5(2a)
- UCS 220 M5 servers
- UCS firmware 4.1(3c)

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

A stateless server is easy in order to replace or expand, and it is only possible when the Boot Drive isn't local. In order to achieve this result, boot the server from a device outside the server, and SAN Boot makes this possible.

This document describes how to boot from iSCSI on HyperFlex using the Cisco UCS platform and how to troubleshoot. When this document talks about SAN Boot, the iSCSI Protocol is used in order to boot the server from a HyperFlex iSCSI Target Logical Unit Number (LUN). Fibre-channel connections are not part of this document.

In HXDP 4.5(2a) and later, the VIC1300 and VIC1400 are qualified to be iSCSI Initiators for HyperFlex iSCSI Targets. UCS Servers with this type of VICs can perform a SAN Boot from HyperFlex iSCSI.

This document explains the configuration of HyperFlex, a standalone UCS Server, UCS Server in IMM, and UCS Managed Server in order to do a SAN Boot from HyperFlex iSCSI. The last part covers the installation and configuration of Windows and ESXi Operating System (OS) with Multipath I/O (MPIO) Boot from SAN.

The target audience is UCS and HX administrators who have a basic understanding of UCS configuration, HX configuration, and OS installation.

Configure

High-Level Overview of the SAN Boot from HyperFlex iSCSI

HyperFlex iSCSI in a nutshell:

At the time of the configuration of the iSCSI network on the HyperFlex cluster, a HyperFlex iSCSI Cluster IP address is created. This address can be used in order to discover the Targets and LUN by the initiator. The HyperFlex Cluster determines which HyperFlex node connects. If there is a failure or one node is very busy, HyperFlex moves the Target to another node. A direct log-in from the initiator to a HyperFlex node is possible. In this case, the redundancy can be configured on the Initiator side.

The HyperFlex Cluster can consist of one or many HyperFlex Targets. Every Target has a unique iSCSI Qualified Name (IQN) and can have one or multiple LUNs, and these LUNs automatically get a LUN ID assigned.

The Initiator IQN is put in an **Initiator Group** linked to a HyperFlex Target where a LUN resides.

The Initiator Group can consist of one or more initiator IQNs. When an OS is already installed on a LUN, you can clone it and use it multiple times for a SAN Boot of different servers, which saves time.

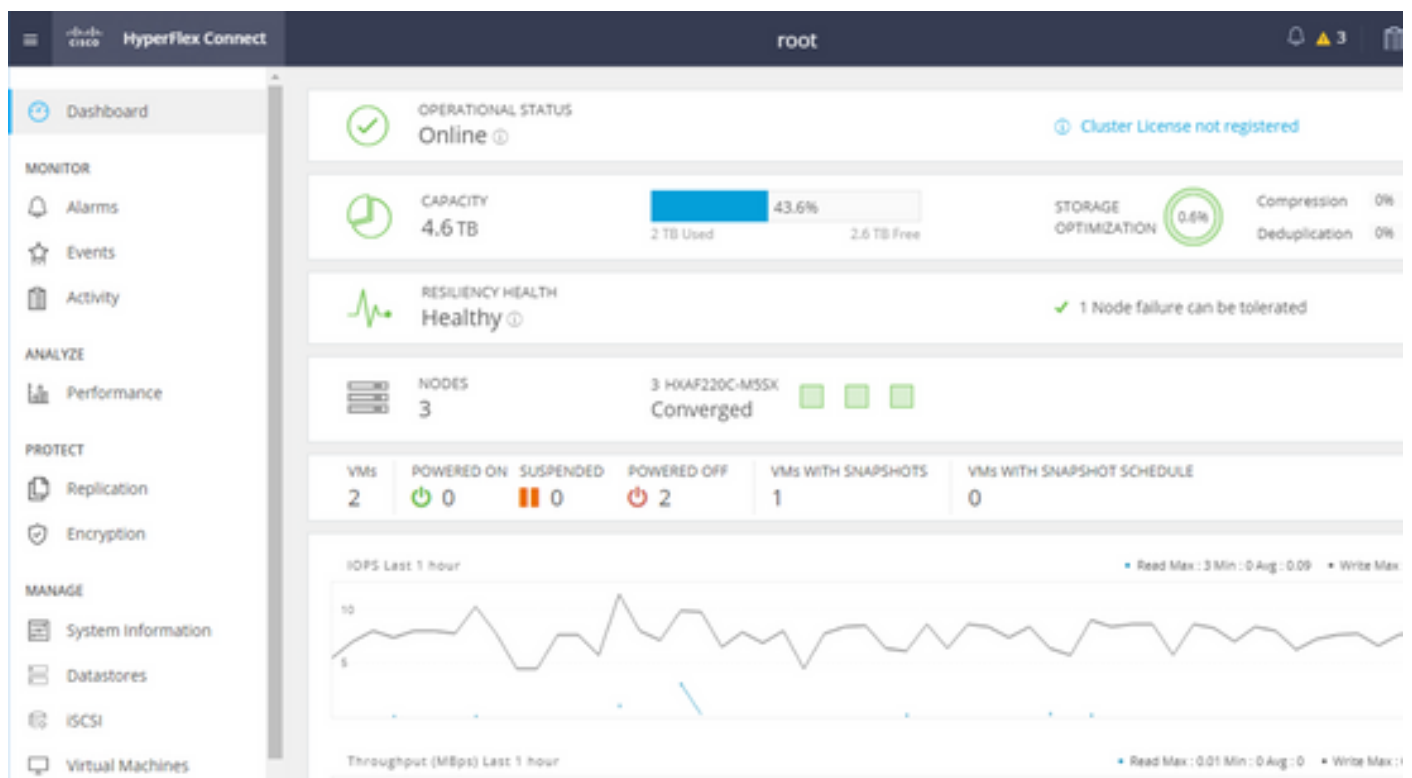
Note: A Windows OS cannot be cloned due to its behavior.

Configure HyperFlex

The configuration of HyperFlex for all three scenarios is the same. The IQN in the UCS Server configuration can be different compared to this section.

Prerequisites: Before you configure the steps mentioned in this document, a joint task must have already been done. These steps are not explained in this document. HyperFlex iSCSI Network is configured in HyperFlex. See [HyperFlex Admin Guide](#) about the steps.

Step 1. Open the HX-Connect, choose **iSCSI** as shown in this image:

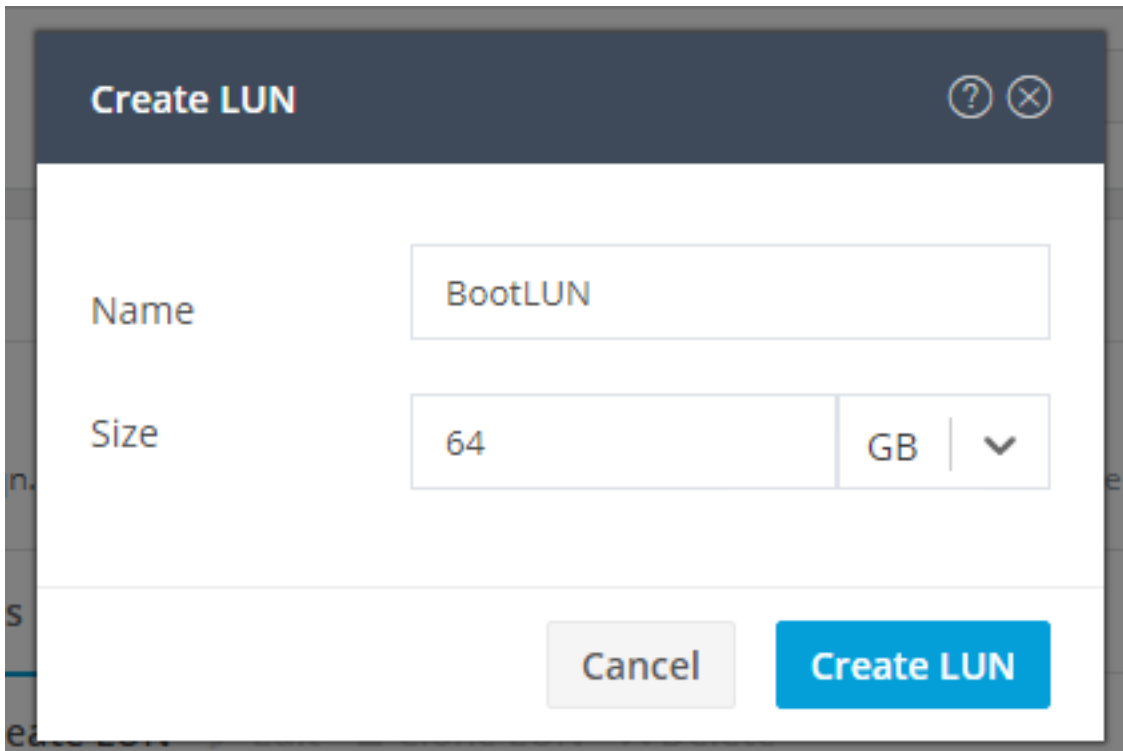


Step 2. Create a new HyperFlex iSCSI Target as shown in this image:

The screenshot shows the 'Create Target' configuration form. The title 'Create Target' is displayed in a dark header. Below the title, there is a text input field for 'Target Name' containing the value 'BFS'. At the bottom of the form, there is a checkbox labeled 'Enable CHAP authentication', which is currently unchecked.

In this configuration example, we don't use authentication. Give the **Target Name** a name without _ (Underscore) or other special characters. In these examples, CHAP authentication is not configured. For security reasons, it is possible in order to configure CHAP authentication. In the examples to install a Windows OS and ESXi on the BootFromSAN LUN, CHAP authentication is configured.


Step 3. Create a LUN within this Target as shown in this image:



The screenshot shows a 'Create LUN' dialog box. The title bar contains a question mark icon and a close icon. The main content area has two input fields: 'Name' with the text 'BootLUN' and 'Size' with the text '64' and a unit dropdown menu set to 'GB'. At the bottom of the dialog, there are two buttons: 'Cancel' and 'Create LUN'.

The name is just for your reference. Choose the appropriate size of the LUN. HyperFlex does not support LUN masking, and LUN IDs are generated automatically.

Step 4. Create an **Initiator Group** (IG) on HyperFlex with the IQN of the Initiator as shown in this image:



The screenshot shows a 'Create Initiator Group' dialog box. The title bar contains a question mark icon and a close icon. The main content area has three input fields: 'Name' with the text 'CIMCDemo', 'Initiators' (empty), and 'Initiator IQN' (empty). To the right of the 'Initiator IQN' field is an 'Add Initiators' button. Below the input fields is a list box containing the IQN 'iqn.2021-06.com.cisco.bootdemo:server1' with a trash icon and up/down arrows.

Decide a name for the IG. If you do not know the IQN of the Initiator at this moment, just add any valid IQN to this IG. Later you can delete it and add the correct initiator IQN name. Document the IG, in order to quickly find the Initiator name when you must change it.

In an IG, one or multiple Initiators IQN can be added.

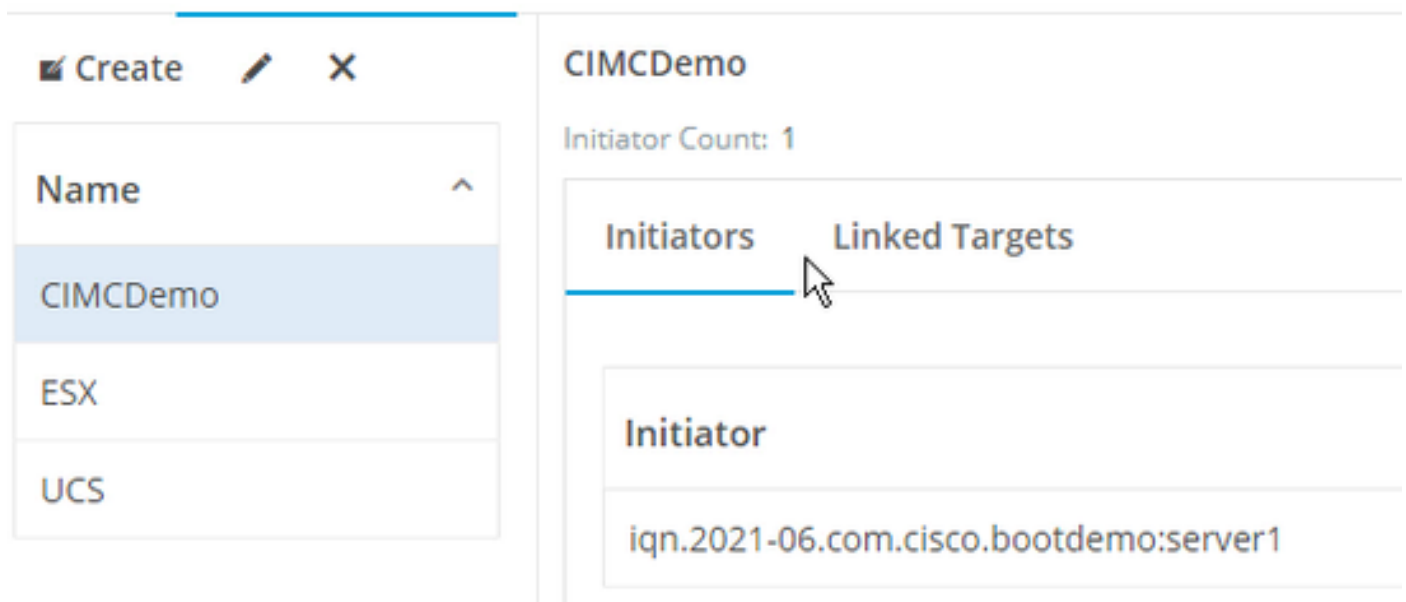
If the initiator is outside of the HyperFlex iSCSI subnet, run **hxcli iscsi allowlist add -p <ip address of the initiator>** command via the controller or HX WebCLI.

In order to verify if this IP address was added to the allowlist, run **hxcli iscsi allowlist show** command.

Step 4.1. Click **Create Initiator Group** as shown in this image:



Step 5. Link the IG with the HyperFlex Target. A HyperFlex Target with LUNs is created, and the IG is created. The last step for the HyperFlex configuration is to Link the Target with the IG. Choose the IG and select **Linked Targets** as shown in this image:



Select **Link** and choose the correct HyperFlex Target.

Step 5.1. Document HyperFlex Target IQN and LUN ID. Later, the HyperFlex Target IQN gets configured at the initiator. Choose the newly created Target and document the IQN. In this example, it is **iqn.1987-02.com.cisco.iscsi:CIMCDemoBoot** as shown in this image:

CIMCDemoBoot

IQN `iqn.1987-02.com.cisco.iscsi:CIMCDemoBoot` CHAP Authentication

LUNs		Linked Initiator Groups					
<input checked="" type="checkbox"/>	Create LUN	<input type="checkbox"/>	Edit	<input type="checkbox"/>	Clone LUN	<input type="checkbox"/>	Delete
	Name	^	LUN ID				
<input type="checkbox"/>	HDD1		LUN1				

The LUN ID at this target must also be documented and used later at the Initiator configuration. In this example, the LUN ID is **LUN1**.

If multiple targets are configured at the cluster, LUNs can have the same LUN ID on different HyperFlex Targets IQNs.

Configure UCS Standalone Server - CIMC


The server only has a Modular LAN-on-Motherboard (MLOM) with a network connection in this example. If there are multiple network adapters, please select the correct one. The procedure is the same as described here:

Prerequisites: Before you configure the steps mentioned in this document, a joint task must have already been done. These steps are not explained in this document.

- HyperFlex iSCSI Target, LUN, and IG are configured and linked
- CIMC is configured with an IP address that is reachable from a browser

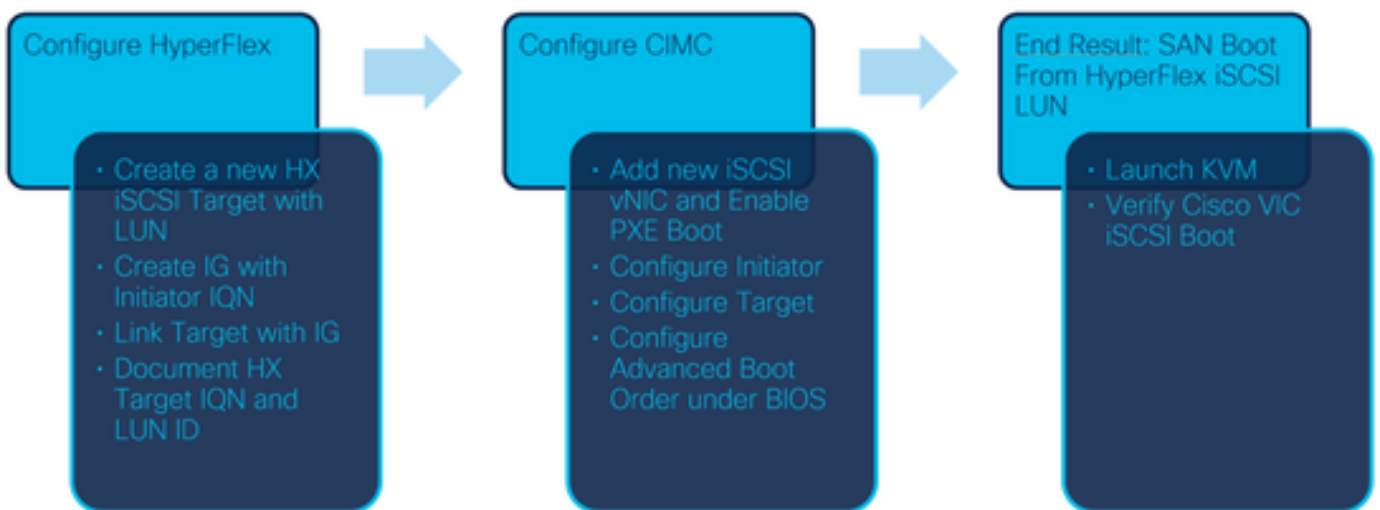
Network Diagram:

The physical network topology of the setup is as shown in this image:

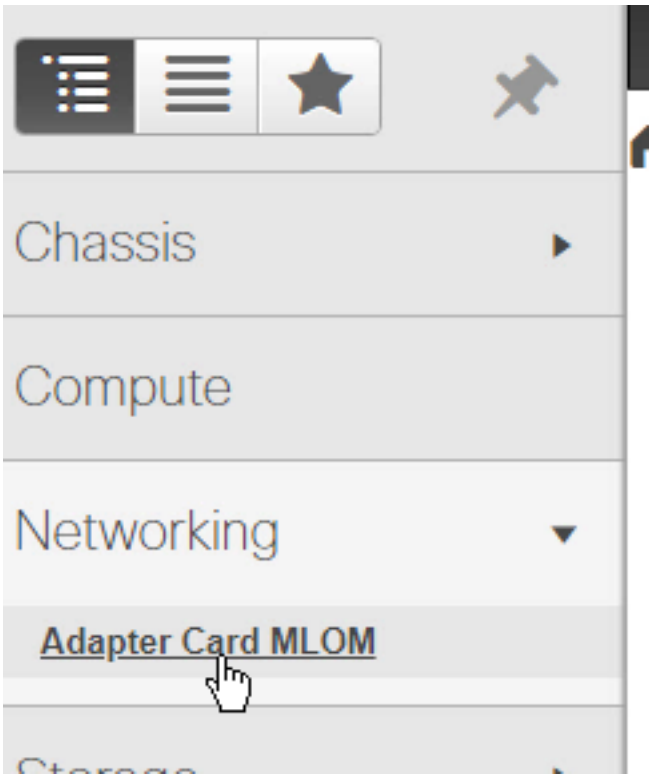


The UCS Standalone Server is connected via the MLOM to two Nexus switches. The two Nexus switches have a VPC connection to the Fabric Interconnect. Each HyperFlex node connects the network adapter to Fabric Interconnect A and B. For the SAN Boot, a Layer 2 iSCSI VLAN network is configured.

Workflow: The steps to be followed in order to configure SAN Boot from HyperFlex iSCSI LUN are as shown in this image:



Step 1. Configure the Network Adapter Card. Open CIMC in a browser and choose **Networking > Adapter Card MLOM** as shown in this image:



Step 2. Choose **vNICs** as shown in this image:



By default, there are already two vNICs configured. Leave them as they are as shown in this image:

Host Ethernet Interfaces

	Name	CDN	MAC Address
<input type="checkbox"/>	eth0	VIC-MLO...	E8:EB:34:9B:7F:60
<input type="checkbox"/>	eth1	VIC-MLO...	E8:EB:34:9B:7F:61

Step 3. Choose **Add vNIC** as shown in this image:

Add vNIC

▼ General

Name:	<input type="text" value="iscsi"/>	
CDN:	<input type="text"/>	
MTU:	<input type="text" value="1500"/>	(1500 - 9000)
Uplink Port:	<input type="text" value="0"/>	▼
MAC Address:	<input checked="" type="radio"/> Auto	
	<input type="radio"/>	<input type="text"/>
Class of Service:	<input type="text" value="0"/>	(0 - 6)
Trust Host CoS:	<input checked="" type="checkbox"/>	
PCI Order:	<input type="text" value="4"/>	(0 - 4)
Default VLAN:	<input type="radio"/> None	
	<input checked="" type="radio"/>	<input type="text" value="20"/>
VLAN Mode:	<input type="text" value="Access"/>	▼

This new vNIC transports the iSCSI traffic from the HyperFlex cluster to the UCS Server. In this example, the server has a Layer 2 iSCSI VLAN Connection. The VLAN is 20, and the VLAN Mode must be set to Access.

Step 3.1. Ensure **Enable PXE Boot** is checked as shown in this image:

Rate Limit: OFF
 (1 - 25000)

Channel Number: (1 - 1000)

PCI Link: 0 (0 - 1)

Enable NVGRE:

Enable VXLAN:

Geneve Offload:

Advanced Filter:

Port Profile:

Enable PXE Boot:

Enable VMQ:

Enable Multi Queue:

No. of Sub vNICs: 64 (1 - 64)

Enable iRFS:


Enable Uplink Failover:

Failback Timeout: 5 (0 - 600)

Step 3.2. Now you can add this vNIC. Use the **Add vNIC** option as shown in this image:



Step 4. On the left, choose the new created **iscsi** vNIC as shown in this image:



Step 4.1. Scroll down to the **iSCSI Boot Properties** and expand **Initiator** as shown in this image:

▼ Initiator

Name:	<input type="text" value="iqn.2021-06.com.cisco.bootdemo:server"/>	(0 - 222) chars
IP Address:	<input type="text" value="10.2.20.248"/>	
Subnet Mask:	<input type="text" value="255.255.255.0"/>	
Gateway:	<input type="text" value="10.2.20.254"/>	
Primary DNS:	<input type="text" value="10.2.20.16"/>	

The **Name** is the **IQN of the Initiator**. You can create your IQN as described in [RFC 3720](#). The **IP address** is which **IP address the UCS Server** gets for the iSCSI vNIC. This address must communicate with the HyperFlex iSCSI Cluster IP address. The HyperFlex Target doesn't have authentication so leave the rest blank as shown in this image:

Initiator Priority:	<input type="text" value="primary"/>	
Secondary DNS:	<input type="text"/>	
TCP Timeout:	<input type="text" value="15"/>	(0 - 255)
CHAP Name:	<input type="text"/>	(0 - 49) chars
CHAP Secret:	<input type="text"/>	(0 - 49) chars

Step 4.2. Configure the **Primary Target** as shown in this image:

▼ Primary Target

Name:	<input type="text" value="iqn.1987-02.com.cisco.iscsi:CIMCDemo"/>	(0 - 222) chars
IP Address:	<input type="text" value="10.3.112.13"/>	
TCP Port	3260	

The **Name** of the **Primary Target** is the **HyperFlex Target** which is linked to the IG with the IQN of this Initiator. The **IP address** is the **HyperFlex iSCSI Cluster IP address**.

Step 4.3. Ensure the **Boot LUN** is the correct one as shown in this image:

Boot LUN:	<input type="text" value="1"/>	(0 - 65535)
CHAP Name:	<input type="text"/>	(0 - 49) chars
CHAP Secret:	<input type="text"/>	(0 - 49) chars

Verification can be done in order to see the **LUN ID** of the LUN at the HyperFlex iSCSI Target. Click the **Save Changes and OK** button as shown in these images:

Save Changes



Do you want to save changes?

Settings will take effect upon the next server reset.

OK

Cancel

Prerequisites: Before you configure the steps mentioned in this document, a joint task must have already been done. A Service Profile was already created and assigned to a server. This step is not explained in this part of the document.

Step 1. Configure the CIMC Boot Order. Open the server CIMC and choose to **Compute as shown in this image:**

Chassis	▶
Compute	

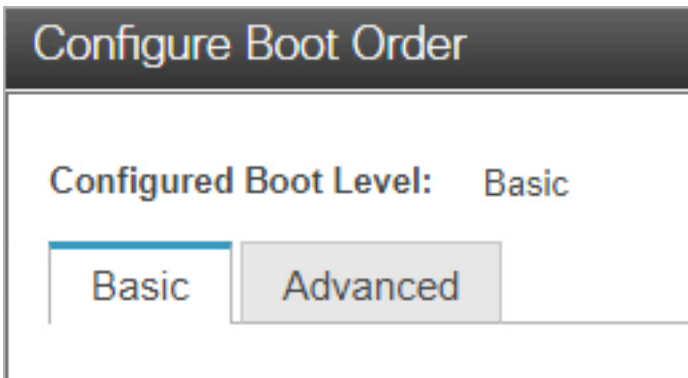
Step 1.1. Choose **BIOS>Configure Boot Order > Configure Boot order** as shown in these images:

Home / Compute / BIOS ★

BIOS	Remote Management	Troubleshooting	Power Policies
Enter BIOS Setup Clear BIOS CMOS Restore Manufacturing Custom Settings F			
Configure BIOS	Configure Boot Order	Configure BIOS Profile	

Configure Boot Order

Step 2. For iSCSI, we must use the **Advanced** Tab and choose **Add iSCSI Boot** as shown in these images:

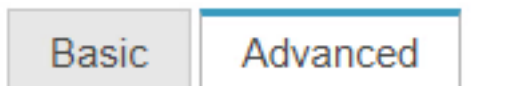


Configure Boot Order

Configured Boot Level: Basic

Basic Advanced

Configured Boot Level: Basic



Basic Advanced

Add Boot Device

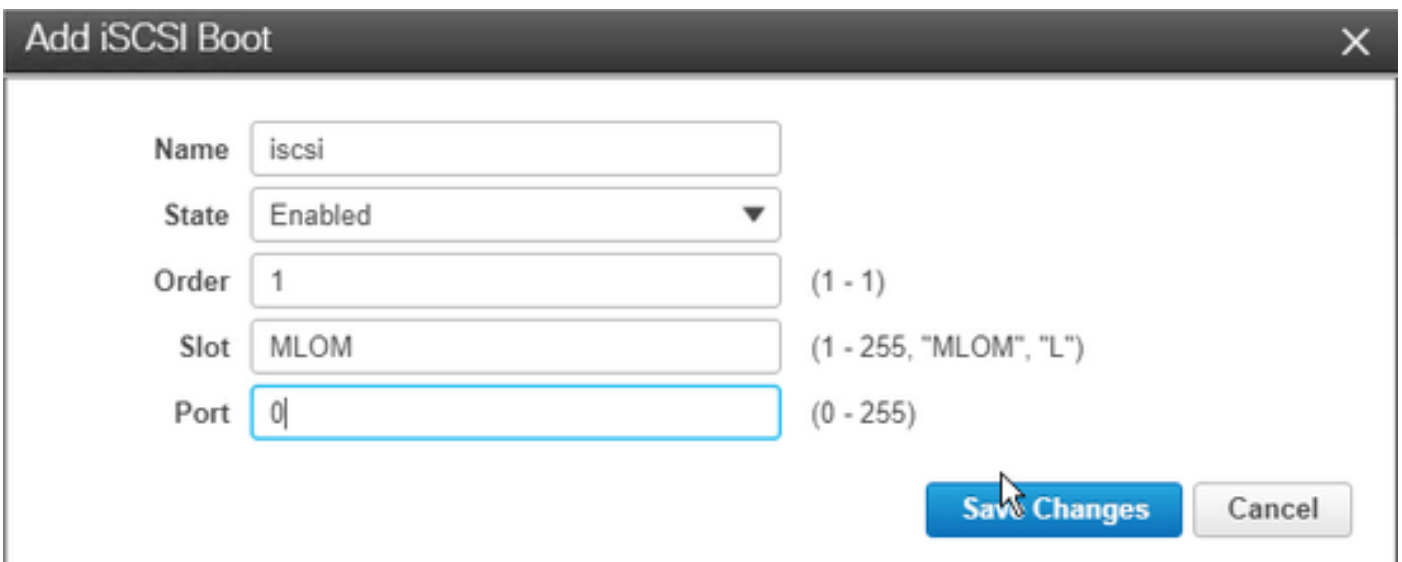
[Add Local HDD](#)

[Add PXE Boot](#)

[Add SAN Boot](#)

[Add iSCSI Boot](#)

Step 2.1. When you add iSCSI boot, the **Name** is for your reference. Ensure the **Order** is set to the lowest number, so it tries to boot from it first. The **Slot** in this example is **MLOM**. The port is set automatically to 0 as shown in this image:



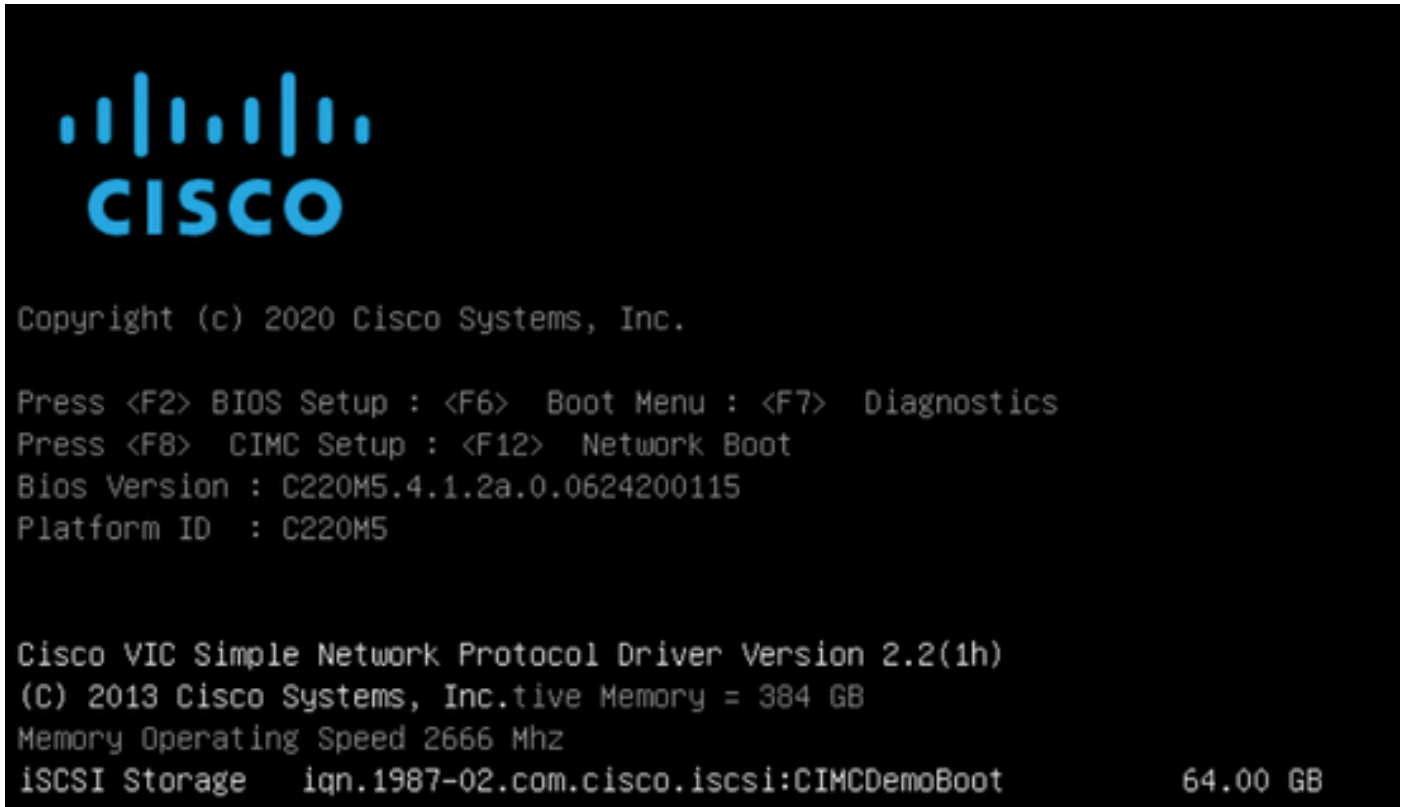
Add iSCSI Boot

Name	<input type="text" value="iscsi"/>	
State	<input type="text" value="Enabled"/>	▼
Order	<input type="text" value="1"/>	(1 - 1)
Slot	<input type="text" value="MLOM"/>	(1 - 255, "MLOM", "L")
Port	<input type="text" value="0"/>	(0 - 255)

Save Changes Cancel

Verify:

SAN Boot from HyperFlex iSCSI LUN. Reboot the server and verify that the BIOS sees the HyperFlex iSCSI LUN. When the Boot Order is set up correctly, it does a SAN boot from HyperFlex iSCSI LUN. On the BIOS screen, you see the **Cisco VIC Simple Network Protocol Driver**, and it shows the IQN of the HyperFlex Target LUN with the size of the LUN as shown in this image:




If the HyperFlex Target has multiple LUNs, they must be shown here.

When there is no OS installed on the LUN, you need to install it via vMedia or manually via the Keyboard, Video, Mouse (KVM).

Configure UCS Manager

Network Diagram:

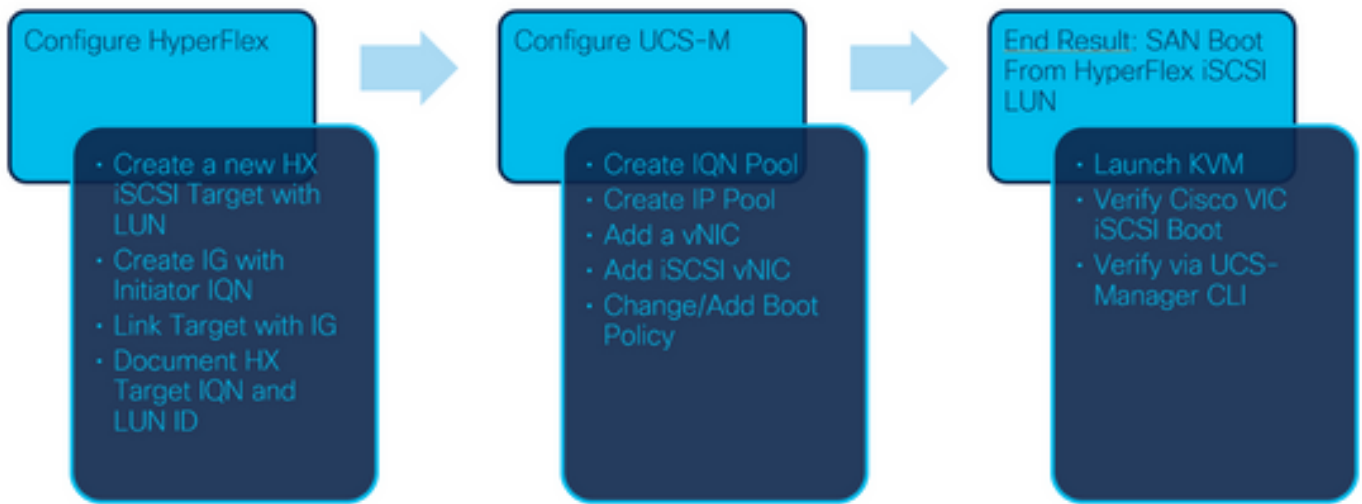
The physical network topology of the setup is as shown in this image:



The UCS Server is connected via Fabric Interconnects connected to the Nexus switches. The two Nexus switches have a VPC connection to the HyperFlex Fabric Interconnects. Each HyperFlex node connects the network adapter to Fabric Interconnect A and B. In this example, the iSCSI goes over different VLANs to show how you configure HyperFlex for this network situation. It is recommended that you eliminate Layer 3 routers and only use Layer 2 iSCSI VLANs to prevent this situation.

Workflow:

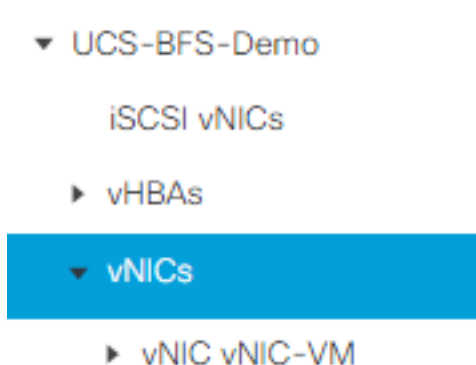
The steps to be followed in order to configure SAN Boot from HyperFlex iSCSI LUN are as shown in this image:



Step 1. Currently, there are no **iSCSI vNICs** configured on the Service Profile. There is only one entry under **vNICs** as shown in this image:



Step 1.1. Choose **vNICs** and click **Add** to add another vNIC for the iSCSI boot traffic as shown in these images:



vNICs

Name	MAC Address	Desired Order	Actual Order	Fabric ID
vNIC vNIC-...	00:25:B5:0D:50...	1	1	A

Advanced Filter Export Print

Delete Add

The **Name** is the name of the vNIC, and this name is needed later in the Boot Order Policy.

Step 1.2. Choose an already created MACPool. You can choose to have multiple vNICs for iSCSI over Fabric-A and Fabric-B or in order to select **Enable Failover**. In this example, the iSCSI vNIC is only connected over **Fabric A** as shown in this image:

Create vNIC

Name : vnic-iscsiboot

MAC Address

MAC Address Assignment: MACPool(17/24) ▼

[Create MAC Pool](#)

The MAC address will be automatically assigned from the selected pool.

The MAC address assignment change will be effective only after server reboot.

Use vNIC Template :

Fabric ID : Fabric A

Fabric B

Enable Failover

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash.

Step 1.3. Choose the VLAN which the iSCSI traffic is supposed to use. This example has the same iSCSI VLAN used by HyperFlex iSCSI Network as shown in this image:

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash.

VLANS VLAN Groups

Advanced Filter Export Print

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	gate_inside_10.2.6.0	<input type="radio"/>	6
<input type="checkbox"/>	Green-MGMT-10.2.12.0	<input type="radio"/>	12
<input checked="" type="checkbox"/>	hx-inband-iscsi-2014	<input checked="" type="radio"/>	2014
<input type="checkbox"/>	Infra_20	<input type="radio"/>	20
<input type="checkbox"/>	Orange-MGMT-10.2.21.0	<input type="radio"/>	21

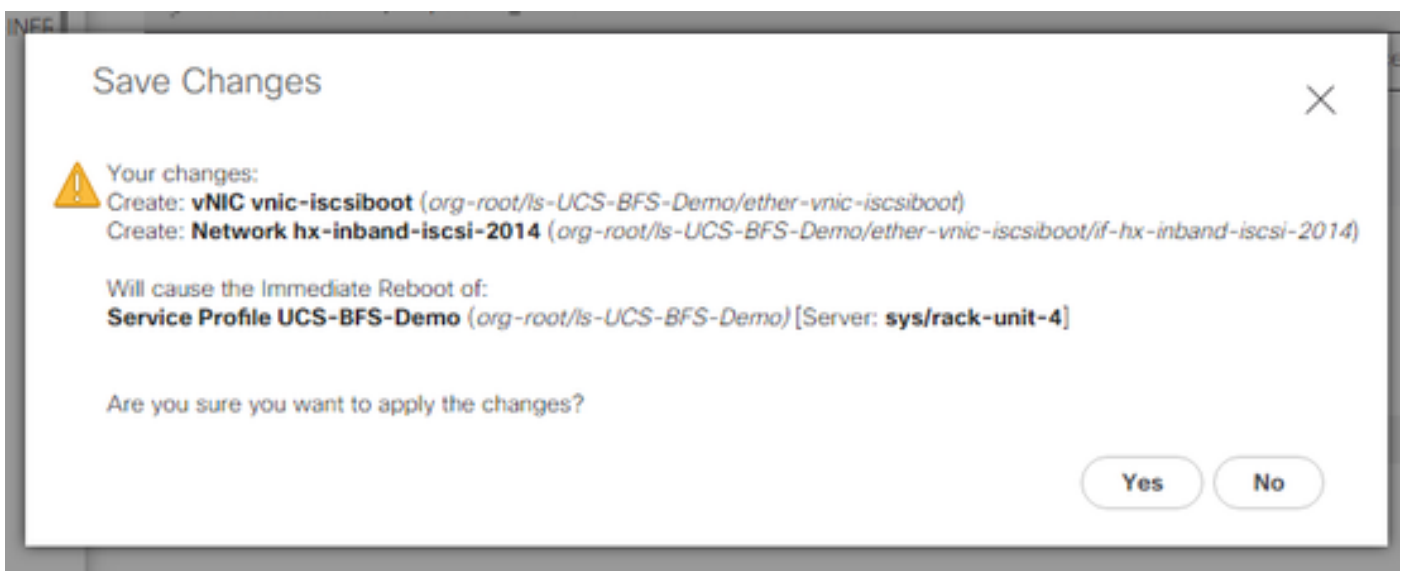
CDN Source : vNIC Name User Defined

MTU :

Note: Ensure this iSCSI VLAN is the Native VLAN. This is only a native VLAN from the server to the Fabric Interconnect and this VLAN doesn't have to be native outside of the Fabric Interconnects.

The best practice for iSCSI is to have Jumbo Frames, which have an MTU size of 9000. If you configure Jumbo Frames, ensure it is end-to-end Jumbo Frames. This includes the OS of the Initiator.

Step 1.4. Click **Save Changes** and **Yes** as shown in these images:



There are now two vNICs for the Service Profile.

Step 2. Add an iSCSI vNIC. Choose **iSCSI vNICs** and select **Add** as shown in these images:

▼ UCS-BFS-Demo

iSCSI vNICs

▶ vHBAs

▼ vNICs

▶ vNIC vnic-iscsiboot

▶ vNIC vNIC-VM

iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy
No data available		

⊕ Add ⊖ Delete ⓘ Modify

An iSCSI vNIC is now created.

[Spoiler](#)

Note: iSCSI vNIC is an iSCSI Boot Firmware Table (iBFT) placeholder for iSCSI boot configuration. It is not actual vNIC, and therefore a vNIC that underlies must be selected. Please do not assign a separate MAC address.

Note: iSCSI vNIC is an iSCSI Boot Firmware Table (iBFT) placeholder for iSCSI boot configuration. It is not actual vNIC, and therefore a vNIC that underlies must be selected. Please do not assign a separate MAC address.

Step 2.1. The **Name** is just an identifier. In VLAN, there is only a choice of one VLAN, which must have been the Native VLAN. Leave the **MAC address Assignment** in order to **Select(None used by default)** as shown in this image:

Create iSCSI vNIC

Name :

Overlay vNIC :

iSCSI Adapter Policy : [Create iSCSI Adapter Policy](#)

VLAN :

iSCSI MAC Address

MAC Address Assignment:

[Create MAC Pool](#)

Step 2.2. Change/Add Boot Policy. On the Service Profile, choose **Boot Order** as shown in this image:

[Servers](#) / [Service Profiles](#) / [root](#) / [Service Profile UCS-B...](#)

[<](#) [General](#) [Storage](#) [Network](#) [iSCSI vNICs](#) [vMedia Policy](#) [Boot Order](#)

Actions

[Modify Boot Policy](#)

Step 2.3. You can modify the Boot Policy only when other servers don't use this Boot Order Policy. In this example, a new Boot Policy is created. The Name is the name of this Boot Policy. If the BOOT LUN doesn't have any OS installed, ensure you choose, for instance, a **Remote CD-ROM**. This way, the OS can be installed via KVM Media. Click **Add iSCSI Boot** as shown in these images:

Modify Boot Policy

Boot Policy:

[Create Boot Policy](#)

Create Boot Policy

Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

WARNINGS:

The type (primary/secondary) does not indicate a boot order presence.

The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.

If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.

If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

+ Local Devices

+ CIMC Mounted vMedia

+ vNICs

+ vHBAs

- iSCSI vNICs

Add iSCSI Boot

+ EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu
No data available						

Move Up Move Down Delete

Set Uefi Boot Parameters

Step 2.4. The **iSCSI vNIC** is the name of the iSCSI vNIC that was created. Enter the same as shown in this image and click **OK**:

Add iSCSI Boot

iSCSI vNIC :

Step 3. The example in this step shows how to create one boot entry. A dual boot entry is possible with two vNICs. The iSCSI Target can still be the same. Because of Windows OS installation, there is a requirement that there must be a single boot entry or single path only at the time of installation. You have to return here and add after OS installation is done and MPIO is configured. This is covered in the section: MPIO.

Step 3.1. Select the Boot Policy that you just created and Expand **iSCSI** as shown in this image:

Modify Boot Policy

Boot Policy:

iSCSIBoot

Create Boot Policy

Name : **iSCSIBoot**
Description :
Reboot on Boot Order Change : **No**
Enforce vNIC/vHBA/iSCSI Name : **Yes**
Boot Mode : **Legacy**

WARNINGS:

The type (primary/secondary) does not indicate a boot order presence.

The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.

If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.

If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA/iS...	Type	LUN Name	WWN	Slot Number	Boot N
iSCSI	1						
iSCSI		bootiscsi_vnic	Primary				

Modify iSCSI vNIC

Set iSCSI Boot Parameters

Set USB Boot Parameters

If you do not see **Modify iSCSI vNIC**, the iSCSI vNIC was not the one you created.

Step 3.2. Choose **Set iSCSI Boot Parameters**. In this example, no authentication is used. The **Initiator Name Assignment** is via an IQN-Pool. This IQN-Pool can be created if it is not there. The **Initiator IP address Policy** is an IP Pool where the UCS Initiator gets its IP address. It can be created if there is no IP Pool created yet as shown in this image:

Set iSCSI Boot Parameters

Name : **bootiscsivnic**

Authentication Profile :

[Create iSCSI Authentication Profile](#)

Initiator Name

Initiator Name Assignment:

Initiator Name :

[Create IQN Suffix Pool](#)

[Reset Initiator Name](#)

The IQN will be assigned from the selected pool.
The available/total IQNs are displayed after the pool name.

Initiator Address

Initiator IP Address Policy:

IPv4 Address : **0.0.0.0**

Subnet Mask : **255.255.255.0**

Default Gateway : **0.0.0.0**

Primary DNS : **0.0.0.0**

Secondary DNS : **0.0.0.0**

[Create IP Pool](#)

The IP address will be automatically assigned from the selected pool.

Of course, it is possible to assign IP addresses manually.

Step 3.3. Scroll down and choose **iSCSI Static Target Interface** and click **Add** as shown in this image:

iSCSI Static Target Interface
 iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addre...
No data available				

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Step 3.4. The **iSCSI Target Name** is the HyperFlex iSCSI Target IQN documented at the time of the HyperFlex Target configuration. The **IPv4 address** is the HyperFlex iSCSI Cluster IP address. The **LUN ID** is the LUN ID that is documented at the time of the HyperFlex Target configuration as shown in this image:

Create iSCSI Static Target

iSCSI Target Name :

Priority :


Port :

Authentication Profile :
[Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

Step 3.5. Choose **OK** and **Yes** in order to modify the boot policy as shown in these images:



Modify Boot Policy



Your changes:

Modify: **Service Profile UCS-BFS-Demo** (*org-root/ls-UCS-BFS-Demo*)

Property: **Boot Policy Name**

Modify: **boot-vnic-bootiscsivnic** (*org-root/ls-UCS-BFS-Demo/iscsi-boot-params/boot-vnic-bootiscsivnic*)

Property: **MAC Pool**

Will cause the Immediate Reboot of:

Service Profile UCS-BFS-Demo (*org-root/ls-UCS-BFS-Demo*) [Server: **sys/rack-unit-4**]

Are you sure you want to apply the changes?

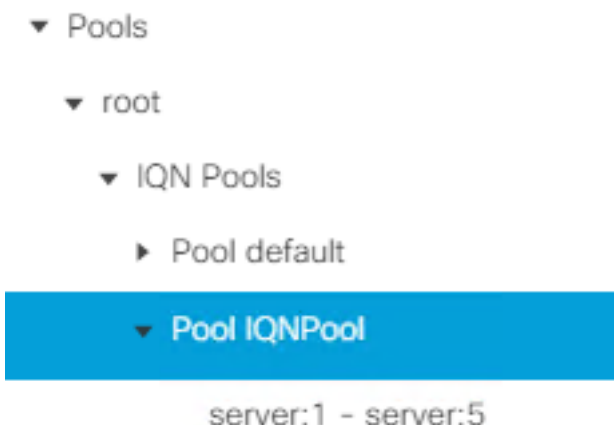
Press **Yes** to disregard the warning and submit changes, **No** to quit the wizard or **Cancel** to make changes to the current configuration.

Yes

No

Cancel

Step 4. Find Initiator IQN. The IQN of the UCS Initiator is not shown in the profile when this configuration is used. Navigate to **SAN** and choose the used **IQN Pools** as shown in this image:



Step 4.1. Note down the IQN of the Profile as shown in this image:

iqn.2021-07.local.hx.cluster:server:5 server:5 Yes org-root/ls-OrgBFS/iscsi-iscsivnic org-root/ls-iSCSIBoot/iscsi-iscsivnic

This Initiator name must be configured in HyperFlex **Initiator Group** linked to the HyperFlex

Target LUN where the server connects to the SAN Boot, as shown in this image:



Edit Initiator Group ? ×

Name:

Initiators:

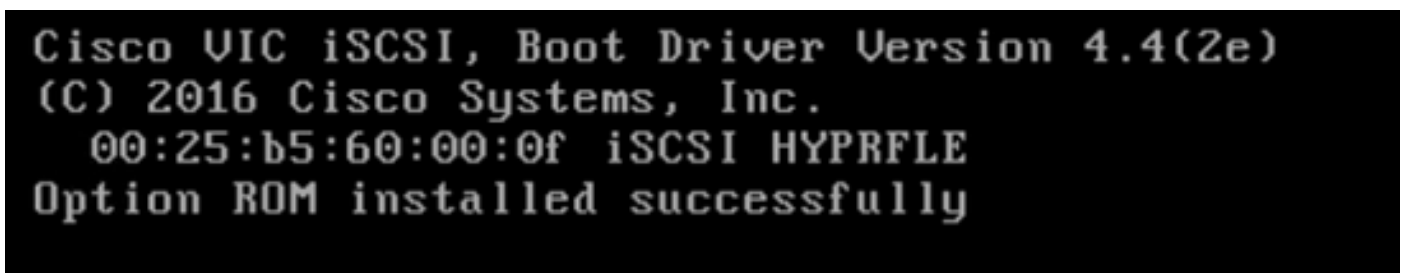
Initiator IQN: Add Initiators

🗑️ ⬆️ ⬆️


When you use a pool, the IQN name is not known in advance. If you create an IG with all the initiator IQNs, those Initiators can see all the same LUNs of the Target. This can be a situation that is not wanted.

Result:

SAN Boot from HyperFlex iSCSI LUN as shown in this image:



Install an OS on the Boot LUN if the LUN does not have any OS installed as shown in this image. The ESXi is installed on the LUN, and after installation, it boots from this LUN:



Troubleshoot iSCSI boot on UCS Manager CLI:

When there is a configuration error, you see **Initialize Error 1** as shown in this image:

```
Cisco VIC iSCSI, Boot Driver Version 4.4(2e)
(C) 2016 Cisco Systems, Inc.
Initialize error 1
```

Different causes can give this error. UCS Manager CLI can get more information about the Initialize Error. SSH to the UCS-Manager and log in. In our example, server 4 has the service profile, and there is only an MLOM present. This gives the value of **4/1**. Type the commands in the UCS Manager CLI as shown in this image:

```
FI-Charger-A# FI-Charger-A#
FI-Charger-A# connect adapter 4/1
adapter 0/4/1 # connect
No entry for terminal type "dumb";
using dumb terminal settings.
adapter 0/4/1 (top):1# attach-mcp
No entry for terminal type "dumb";
using dumb terminal settings.
adapter 0/4/1 (mcp):1# █
```

When you type **help**, you see a whole list of commands that are now possible. The commands for the iSCSI configuration are as shown in this image:

```
iscsi_get_config - Get iSCSI config
iscsi_ping - Test iSCSI connectivity
iscsi_show_eficfg - Show EFI Config
iscsi_show_ibft - Show iBFT posted
```

Problem 1: Ping Stats: In Progress

On the SSH session, type **iscsi_get_config** and check the output as shown in this image:

```

vnic iSCSI Configuration:
-----

vnic_id: 15
      host_id: 0
      link_state: Up

      Initiator Cfg:
      initiator_state: ISCSI_INITIATOR_READY
initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
      vlan: 0
      dhcp status: false
      IQN: iqn.2021-05.ucs.local:UCSServer:10
      IP Addr: 10.3.112.101
      Subnet Mask: 255.255.255.0
      Gateway: 10.3.112.254

      Target Cfg:
      Target Idx: 0
      State: ISCSI_TARGET_LOGIN
      Prev State: ISCSI_TARGET_DISABLED
      Target Error: ISCSI_TARGET_NO_ERROR
      IQN: iqn.1987-02.com.cisco.iscsi:UCS1SanBoot
      IP Addr: 10.3.12.13
      Port: 3260
      Boot Lun: 1
      Ping Stats: In Progress
adapter 0/1/1 (mcp):37#

```

The **Ping Stats** is **In Progress**. This means the initiator cannot ping the HyperFlex iSCSI Cluster IP address. Verify the network path from the initiator to the HyperFlex iSCSI Target. In our example, the Initiator iSCSI IP address is outside the iSCSI Subnet configured on the HyperFlex Cluster. The Initiator IP address must be added to the HyperFlex iSCSI Allowlist. SSH to the HyperFlex Cluster IP address and enter the command:

```

hxcli iscsi allowlist add -p <Initiator IP address>

```

In order to verify if the Initiator IP address is in the allowlist, use the command:

```

hxcli iscsi allowlist show

```

Problem 2: Target Error: "ISCSI_TARGET_LOGIN_ERROR"

On the SSH session, type **iscsi_get_config** and check the output as shown in this image:

```
adapter 0/1/1 (mcp):5# iscsi_get_config

vnic iSCSI Configuration:
-----

vnic_id: 15
      host_id: 0
      link_state: Up

      Initiator Cfg:
      initiator_state: ISCSI_INITIATOR_READY
initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
      vlan: 0
      dhcp status: false
      IQN: iqn.2021-05.ucs.local:UCSServer:10
      IP Addr: 10.3.112.101
      Subnet Mask: 255.255.255.0
      Gateway: 10.3.112.254

      Target Cfg:
      Target Idx: 0
      State: INVALID
      Prev State: ISCSI_TARGET_GET_SESSION_INFO
      Target Error: ISCSI_TARGET_LOGIN_ERROR
      IQN: HX
      IP Addr: 10.3.112.13
      Port: 3260
      Boot Lun: 0
      Ping Stats: Success (20.260ms)
adapter 0/1/1 (mcp):6# adapter 0/1/1 (mcp):6#
adapter 0/1/1 (mcp):6#
```

The **Target Error** is **ISCSI_TARGET_LOGIN_ERROR**. If authentication is used, verify the name and secrets. Ensure the Initiator IQN is in the HyperFlex Initiator Group and Linked to a Target.

Problem 3: Target Error: "ISCSI_TARGT_GET_HBT_ERROR"

On the SSH session, type **iscsi_get_config** and check the output as shown in this image:

```
vnic iSCSI Configuration:
-----

vnic_id: 15
      host_id: 0
      link_state: Up

      Initiator Cfg:
        initiator_state: ISCSI_INITIATOR_READY
        initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
          vlan: 0
          dhcp status: false
            IQN: iqn.2021-05.ucs.local:UCSServer:10
            IP Addr: 10.3.112.101
            Subnet Mask: 255.255.255.0
            Gateway: 10.3.112.254

          Target Cfg:
            Target Idx: 0
              State: INVALID
              Prev State: ISCSI_TARGET_GET_LUN_INFO
            Target Error: ISCSI_TARGET_GET_HBT_ERROR
              IQN: iqn.1987-02.com.cisco.iscsi:UCS1SanBoot
              IP Addr: 10.3.112.13
              Port: 3260
              Boot Lun: 0
            Ping Stats: Success (20.496ms)
adapter 0/1/1 (mcp):29#
```

The **Target Error** is **ISCSI_TARGET_GET_HBT_ERROR**. In the configuration of the **BOOT LUN**, a **wrong LUN ID** was used. In this case, the **BOOT LUN** was set to **0**, and it must have been assigned to **1**.

Working SAN Boot Configuration:

The SAN Boot From HyperFlex iSCSI must work when the iSCSI configuration is correct and you have the output as shown in this image:


```
adapter 0/1/1 (mcp):50# iscsi_get_config

vnic iSCSI Configuration:
-----

vnic_id: 15
      host_id: 0
      link_state: Up

      Initiator Cfg:
        initiator_state: ISCSI_INITIATOR_READY
initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
          vlan: 0
        dhcp status: false
          IQN: iqn.2021-05.ucs.local:UCSServer:10
          IP Addr: 10.3.112.101
        Subnet Mask: 255.255.255.0
          Gateway: 10.3.112.254

      Target Cfg:
        Target Idx: 0
          State: ISCSI_TARGET_READY
        Prev State: ISCSI_TARGET_DISABLED
        Target Error: ISCSI_TARGET_NO_ERROR
          IQN: iqn.1987-02.com.cisco.iscsi:UCS1SanBoot
          IP Addr: 10.3.112.13
          Port: 3260
        Boot Lun: 1
        Ping Stats: Success (20.598ms)

      Session Info:
        session_id: 0
        host_number: 0
        bus_number: 0
        target_id: 0
adapter 0/1/1 (mcp):51# adapter 0/1/1 (mcp):51#
adapter 0/1/1 (mcp):51# █
```


Configure IMM

Prerequisites:

- Fabric Interconnects are claimed in Intersight
- Intersight Server Profile is already created and attached to a server

Network Diagram:

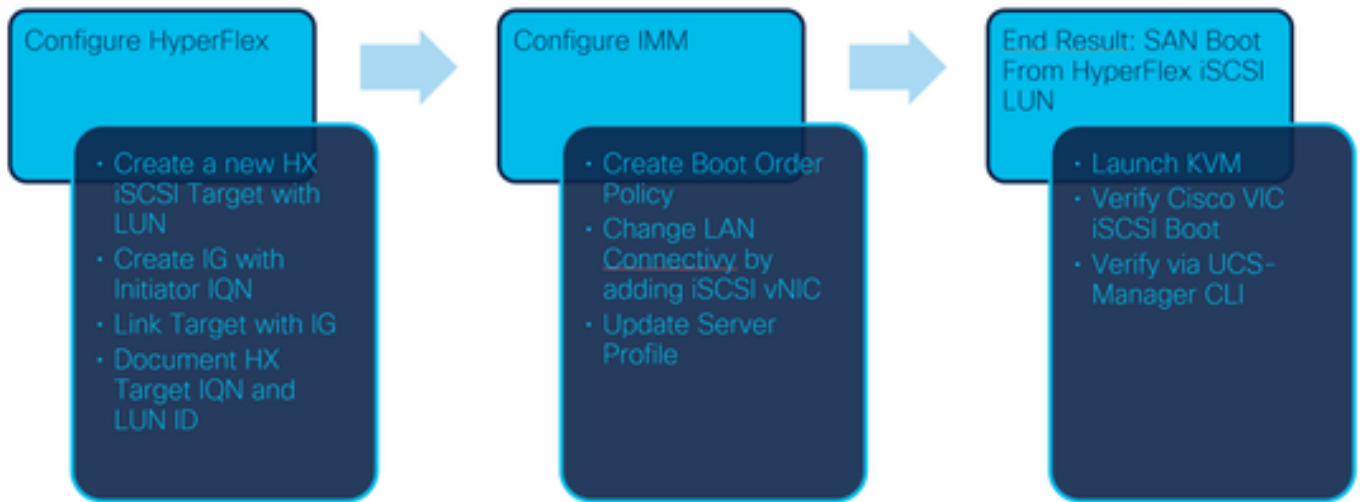
The physical network topology of the setup is as shown in this image:



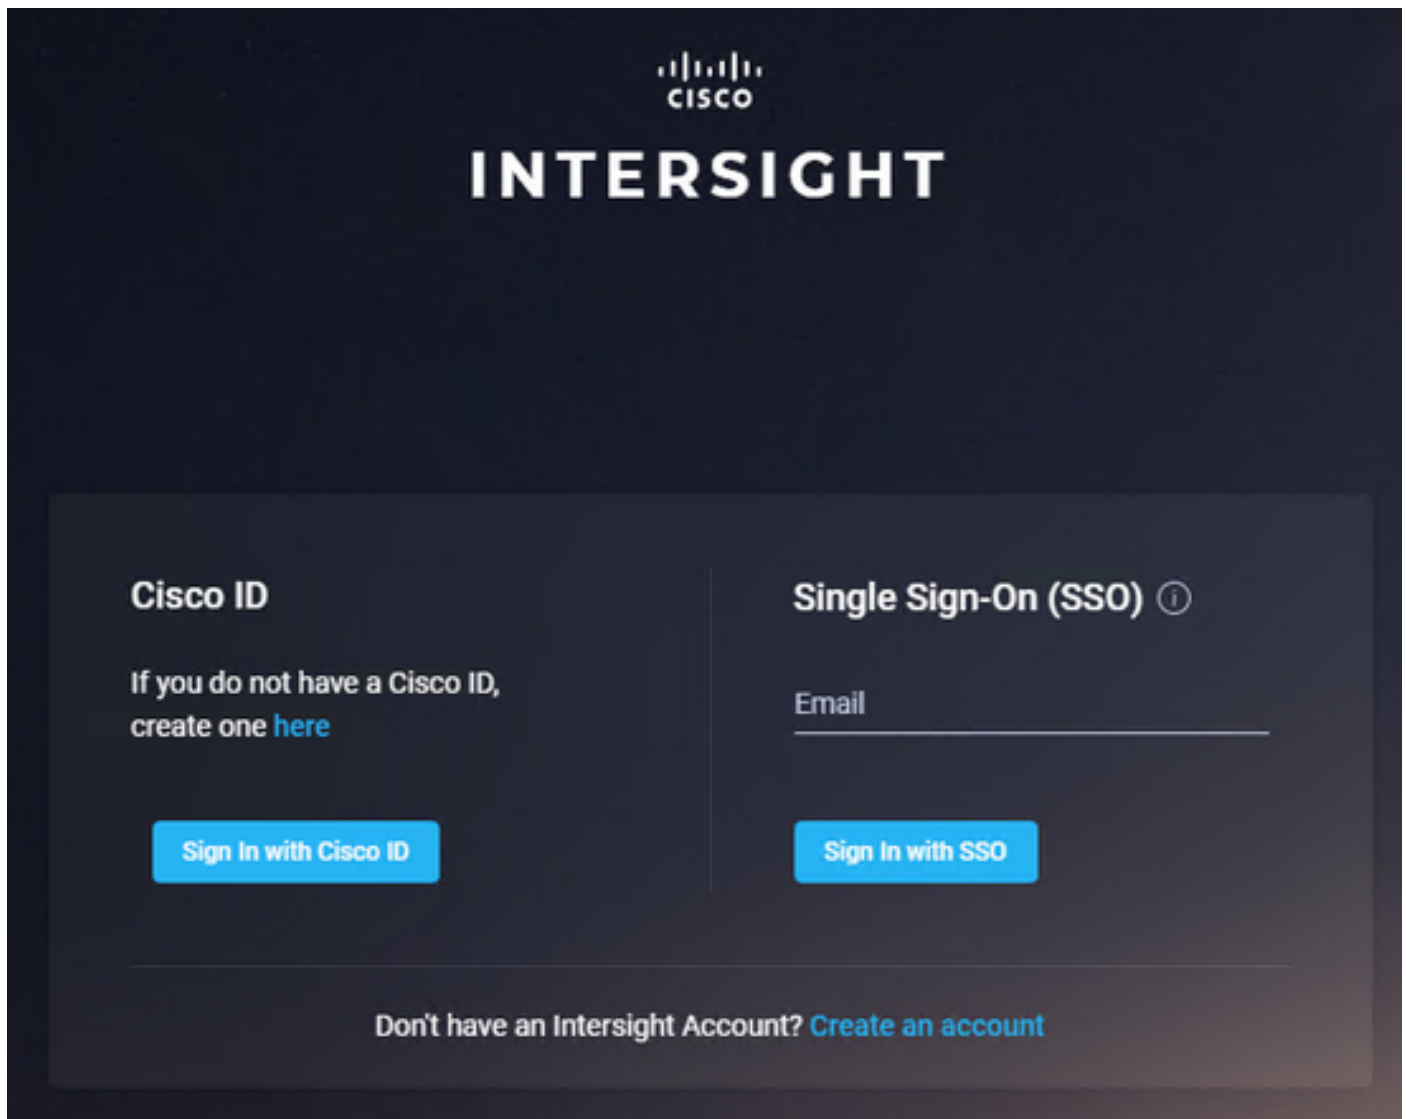
The UCS Server is in IMM and controlled via Intersight. The two Nexus switches have a VPC connection to the different pairs of Fabric Interconnect. Each HyperFlex node connects the network adapter to Fabric Interconnect A and B. A Layer 2 iSCSI VLAN network is configured without Layer 3 device delays for the SAN Boot.

Workflow:

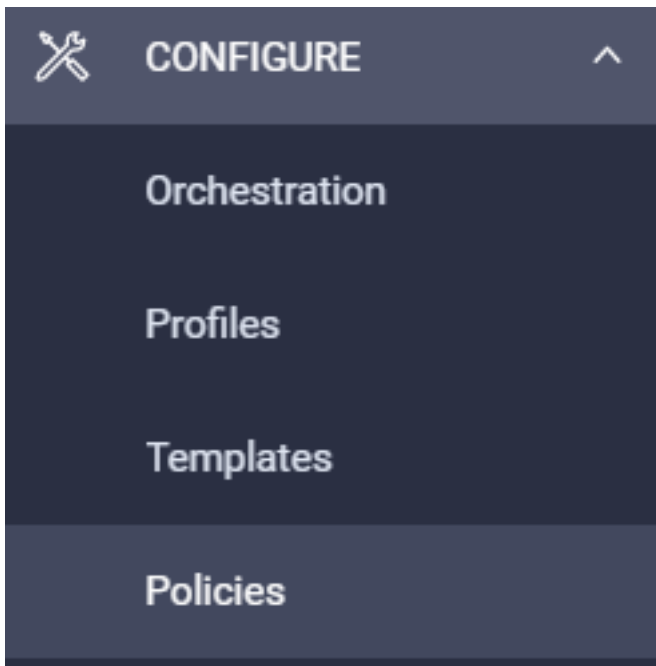
The steps to be followed in order to configure SAN Boot from HyperFlex iSCSI LUN are as shown in this image:



Step 1. In order to log in to Intersight, use <https://intersight.com> as shown in this image:



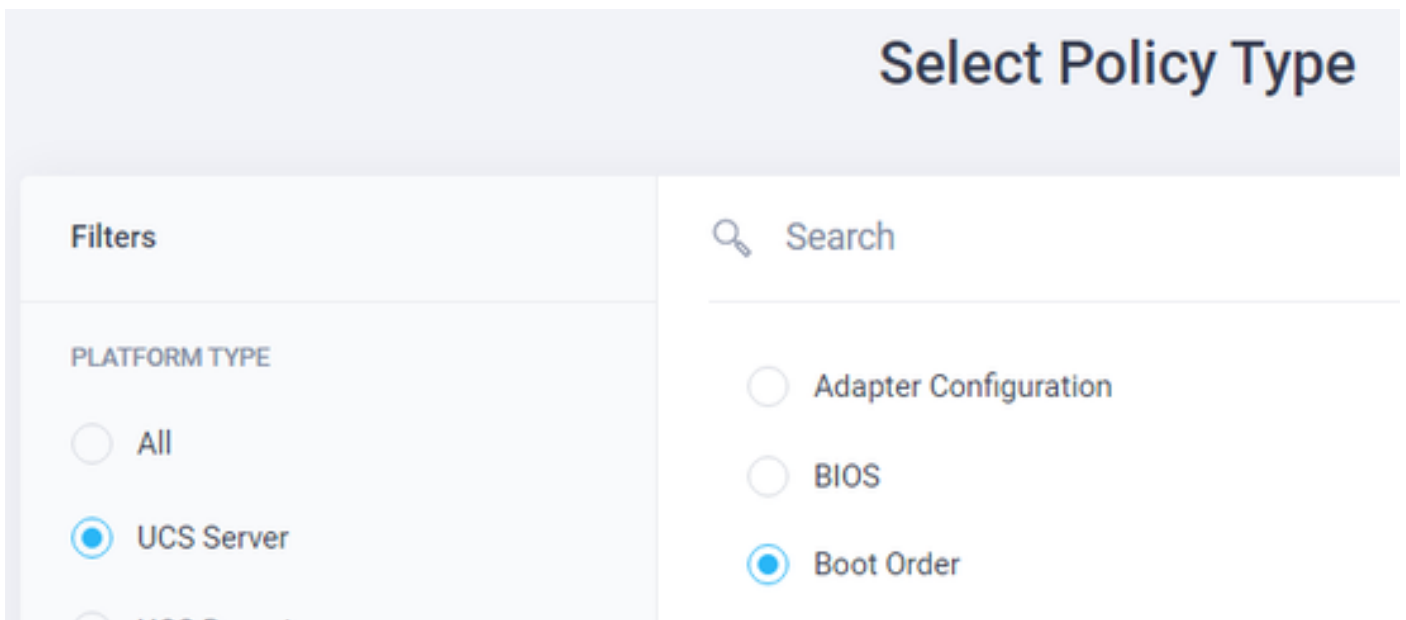
Step 2. Create a new Boot Order Policy. A new Boot Order policy is created for this server. Choose **Configure > Policies** as shown in this image:



Step 2.1. Click **Create Policy** at the right top corner as shown in this image:



Step 2.2. On the left, choose **UCS Server**. Choose **Boot Order** from the policies as shown in this image and click **Start**:



Step 2.3. In **Step 1.**, give it a unique **Name** as shown in this image and click **Next** :



Step 1

General

Add a name, description and tag for the policy.

Organization *

default



Name *

IMMBFSServer3

Set Tags

Description



<= 1024

Step 2.4. In **Step 2.**, choose **UCS Server (FI-Attached)**. For this example, leave the **Configured Boot Mode** at **Legacy**. Expand the **Add Boot Device** and select **iSCSI Boot** as shown in this image:



Step 2
Policy Details
Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

Configured Boot Mode

Legacy Unified Extensible Firmware Interface (UEFI)

Add Boot Device

- iSCSI Boot
- Local CDD
- Local Disk


Step 2.5. Give it a **Device Name** and an **Interface Name** as shown in this image:

iSCSI Boot (IMMBFSServer3) Enabled

Device Name *

Interface Name *

The name of the Interface Name must be documented, and it is used in order to create a new vNIC. Click **Create**, a pop up must show on the screen as shown in this image:

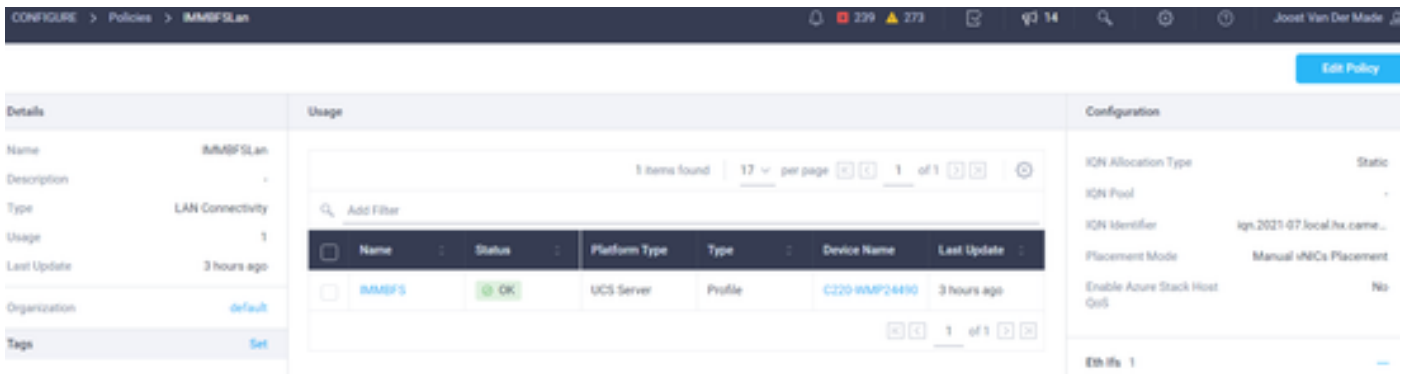


Step 3. Change LAN Connectivity. A new **LAN Connectivity** can be created. In this example, the current LAN Connectivity of the Server Profile is edited. Search for the user policy at the **Policies** overview as shown in this image:

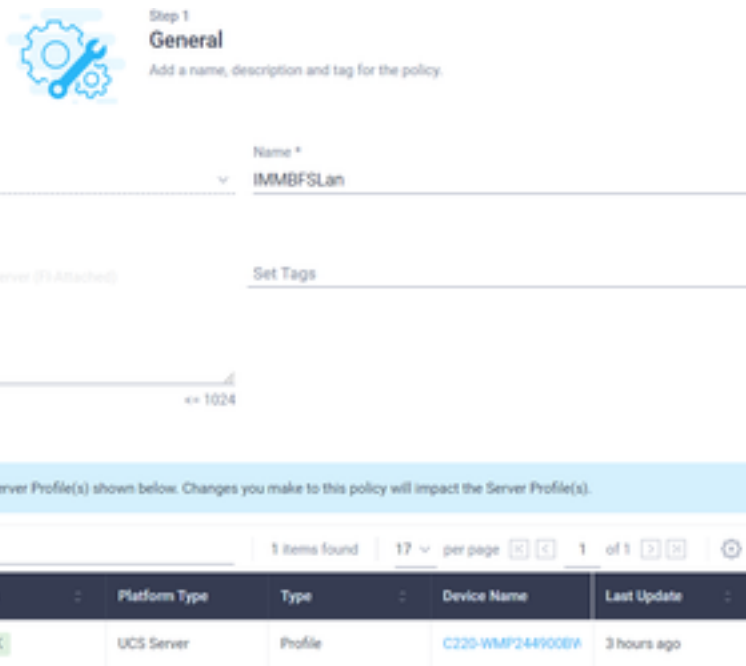
| Search IMMBFSLan Add Filter

<input type="checkbox"/>	Name	Platform Type	Type
<input type="checkbox"/>	IMMBFSLan	UCS Server	LAN Connectivity

Step 3.1. Choose **Edit Policy** as shown in this image:



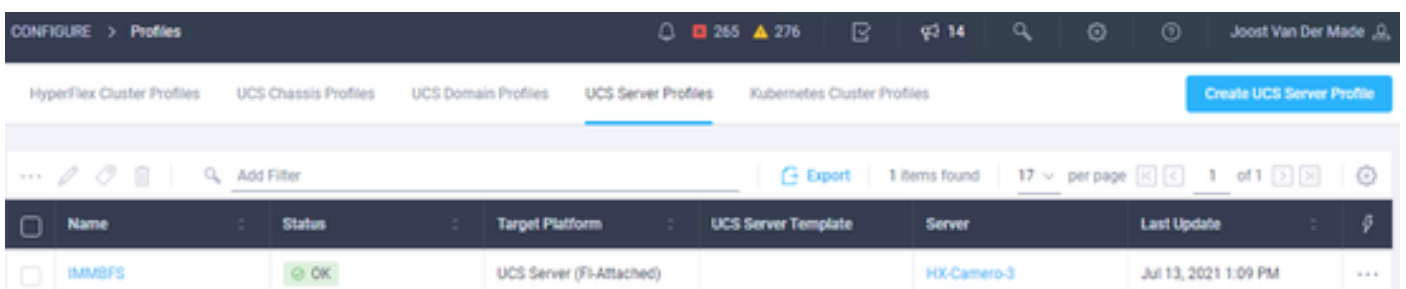
In this case, the Policy Name is **IMMBFSLan**. There is already a vNIC present in this configuration. Do not change anything in **Step 1**, as shown in this image and click **Next**:



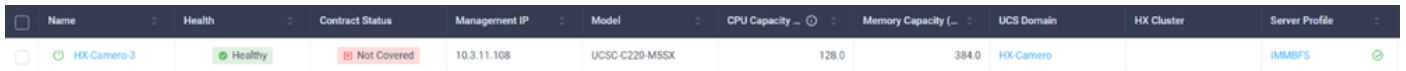
Step 3.2. In **Step 2.**, choose **Add vNIC** as shown in this image:



Step 4. Change the Server Profile. The LAN Connectivity policy was updated, and the **Boot Order** must be changed in this Server Profile. Choose **CONFIGURE > Profiles** from the left bar in order to locate the UCS server profile, as shown in this image:



The UCS Profile can be selected directly from the UCS Server as shown in this image:



The image shows a table of UCS servers with the following columns: Name, Health, Contract Status, Management IP, Model, CPU Capacity, Memory Capacity, UCS Domain, HX Cluster, and Server Profile. The data row shows a server named 'HX-Camero-3' with a 'Healthy' status, 'Not Covered' contract status, management IP '10.3.11.108', model 'UCSC-C220-M55X', CPU capacity '128.0', memory capacity '384.0', UCS domain 'HX-Camero', and server profile 'HMMBFS'.

Name	Health	Contract Status	Management IP	Model	CPU Capacity	Memory Capacity	UCS Domain	HX Cluster	Server Profile
HX-Camero-3	Healthy	Not Covered	10.3.11.108	UCSC-C220-M55X	128.0	384.0	HX-Camero		HMMBFS

The **Name** must be used in the **Boot Order** policy. The server only has one network adapter, the MLOM. This must be configured in **Slot ID**. Leave the **PCI Link** at 0. The **Switch ID** for this example is **A**, and the **PCI Order** is the number of the latest vNIC which is **1**. The **Ethernet Network Control Policy**, **Ethernet QoS**, and **Ethernet Adapter** can have default values. The best practice for iSCSI is to have an MTU of 9000, which can be configured in the **Ethernet QoS Policy**.

Step 4.1. Choose **Ethernet Network Group Policy > Select Policy** as shown in this image:

