



Cisco UCS Manager Storage Management Guide, Release 4.1

First Published: 2020-02-20

Last Modified: 2021-07-29

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883



CONTENTS

PREFACE

Preface	xiii
Audience	xiii
Conventions	xiii
Related Cisco UCS Documentation	xv
Documentation Feedback	xv

CHAPTER 1

New and Changed Information	1
New and Changed Information	1

CHAPTER 2

Overview	3
Overview	3
Cisco UCS Manager User Documentation	3
Storage Options	4
Storage Design Considerations	6
Storage Configuration Sequence	6
Storage Protocols	7
The UCS Manager SAN Tab	7

CHAPTER 3

SAN Ports and Port Channels	9
Port Modes	9
Port Types	9
Server Ports	10
Configuring Server Ports	10
Reconfiguring a Port on a Fabric Interconnect	11
Enabling or Disabling a Port on a Fabric Interconnect	11
Unconfiguring a Port on a Fabric Interconnect	12

Appliance Ports	12
Configuring an Appliance Port	13
Modifying the Properties of an Appliance Port	14
FCoE and Fibre Channel Storage Ports	14
Configuring an Ethernet Port as an FCoE Storage Port	14
Configuring a Fibre Channel Storage Port	15
Restoring an Uplink Fibre Channel Port	15
FC Links Rebalancing	16
Converting FC Storage Port to FC Uplink Port	16
FCoE Uplink Ports	17
Configuring FCoE Uplink Ports	17
Unified Storage Ports	18
Configuring an Appliance Port as a Unified Storage Port	18
Unconfiguring a Unified Storage Port	19
Unified Uplink Ports	19
Configuring Unified Uplink Ports	20
Unconfiguring Unified Uplink Port	20
Policy-Based Port Error Handling	21
Configuring Error-Based Action	21
Fibre Channel Port Channels	22
Creating a Fibre Channel Port Channel	22
Enabling a Fibre Channel Port Channel	23
Disabling a Fibre Channel Port Channel	23
Adding Ports to and Removing Ports from a Fibre Channel Port Channel	23
Modifying the Properties of a Fibre Channel Port Channel	24
Deleting a Fibre Channel Port Channel	25
FCoE Port Channels	26
Creating an FCoE Port Channel	26
Deleting an FCoE Port Channel	26
Unified Uplink Port Channel	26

CHAPTER 4	Fibre Channel Zoning	29
	Information About Fibre Channel Zoning	29
	Information About Zones	29

Information About Zone Sets	30
Support for Fibre Channel Zoning in Cisco UCS Manager	30
Cisco UCS Manager-Based Fibre Channel Zoning	30
vHBA Initiator Groups	31
Fibre Channel Storage Connection Policy	31
Fibre Channel Active Zone Set Configuration	31
Switch-Based Fibre Channel Zoning	32
Guidelines and recommendations for Cisco UCS Manager-Based Fibre Channel Zoning	32
Configuring Fibre Channel Zoning	32
Creating a VSAN for Fibre Channel Zoning	33
Creating a New Fibre Channel Zone Profile	35
Deleting a Fibre Channel Zone Profile	37
Deleting a Fibre Channel User Zone	38
Fibre Channel Storage Connection Policies	38
Deleting a Fibre Channel Storage Connection Policy	38

CHAPTER 5**Named VSANs 39**

Named VSANs	39
Fibre Channel Uplink Trunking for Named VSANs	40
Guidelines and Recommendations for VSANs	40
Creating a Named VSAN	41
Creating a Storage VSAN	42
Deleting a Named VSAN	43
Changing the VLAN ID for the FCoE VLAN for a Storage VSAN	43
Enabling Fibre Channel Uplink Trunking	44
Disabling Fibre Channel Uplink Trunking	44

CHAPTER 6**SAN Pin Groups 47**

SAN Pin Groups	47
Creating a SAN Pin Group	47
Deleting a SAN Pin Group	48

CHAPTER 7**FC Identity Assignment 49**

Fibre Channel Identity	49
------------------------	----

CHAPTER 8**WWN Pools 51****WWN Pools 51**

Creating a WWNN Pool 52

Adding a WWN Block to a WWNN Pool 53

Deleting a WWN Block from a WWNN Pool 54

Adding a WWNN Initiator to a WWNN Pool 55

Deleting a WWPN Initiator from a WWPN Pool 55

Deleting a WWNN Pool 56

WWPN Pools 56

Creating a WWPN Pool 56

Adding a WWN Block to a WWPN Pool 58

Deleting a WWN Block from a WWPN Pool 58

Adding a WWPN Initiator to a WWPN Pool 59

Deleting a WWPN Initiator from a WWPN Pool 60

Deleting a WWPN Pool 60

WWxN Pools 61

Creating a WWxN Pool 61

Adding a WWN Block to a WWxN Pool 62

Deleting a WWN Block from a WWxN Pool 63

Deleting a WWxN Pool 63

CHAPTER 9**Storage-Related Policies 65****About vHBA Templates 65**

vHBA Template 65

Creating a vHBA Template 65

Binding a vHBA to a vHBA Template 67

Unbinding a vHBA from a vHBA Template 67

Deleting a vHBA Template 68

Fibre Channel Adapter Policies 68

Ethernet and Fibre Channel Adapter Policies 68

Creating a Fibre Channel Adapter Policy 71

Deleting a Fibre Channel Adapter Policy 76

About the Default vHBA Behavior Policy 77

Default vHBA Behavior Policy	77
Configuring a Default vHBA Behavior Policy	77
SPDM Security Policy	78
SPDM Security	78
Creating a SPDM Security Policy	79
Associating the Security Policy with a Server	80
Viewing the Fault Alert Settings	80
SAN Connectivity Policies	81
About the LAN and SAN Connectivity Policies	81
Privileges Required for LAN and SAN Connectivity Policies	81
Interactions between Service Profiles and Connectivity Policies	81
Creating a SAN Connectivity Policy	82
Creating a vHBA for a SAN Connectivity Policy	83
Deleting a vHBA from a SAN Connectivity Policy	83
Creating an Initiator Group for a SAN Connectivity Policy	83
Deleting an Initiator Group from a SAN Connectivity Policy	85
Deleting a SAN Connectivity Policy	85
Enabling Intel® Volume Management Device	85
Volume Management Device (VMD) Setup	85
Enabling VMD on UCS Manager	86
Enabling Volume Management Device (VMD) in Passthrough Mode	86
Volume Management Device (VMD) Passthrough Mode	86
Downloading VMD Drivers	87
Intel® Volume Management Device Drivers	87
Downloading the Linux VMD Drivers	88
Downloading the Windows VMD Drivers	89
Downloading the VMD Passthrough Drivers	90
Custom LED Status with VMD on NVMe	91

CHAPTER 10
SED Security Policies 95

Security Policies for Self-Encrypting Drives	95
Security Flags of the Controller and Disk	96
Managing Local Security Policies	96
Creating a Local Security Policy	96

- Modifying a Local Security Policy 97
- Inserting a Secured Disk into a Server with a Local Security Policy 97
- KMIP Client Certificate Policy 98
 - Creating a Global KMIP Client Certificate Policy 98
 - Creating a KMIP Client Certificate Policy for a Server 99
- Managing Remote Security Policies 100
 - Creating a Remote Security Policy 100
 - Modifying a Remote Security Policy 101
 - Modifying a Remote Security Key 101
 - Inserting a Secured Disk into a Server with a Remote Security Policy 102
- Enabling and Disabling Security on Disks 102
- Disabling Security on a Controller 102
- Unlocking a Locked Disk 103
- Erasing a Secure Foreign Configuration Disk 103
- Secure Data Deletion 104

CHAPTER 11

Storage Profiles 105

- Storage Profiles 105
- Cisco Boot Optimized M.2 RAID Controller 106
- Disk Groups and Disk Group Configuration Policies 106
 - Virtual Drives 107
 - Configuring a Disk Group Policy 108
- RAID Levels 112
- Automatic Disk Selection 113
- Supported LUN Modifications 114
- Unsupported LUN Modifications 114
- Disk Insertion Handling 115
 - Non-Redundant Virtual Drives 115
 - Redundant Virtual Drives with No Hot Spare Drives 116
 - Redundant Virtual Drives with Hot Spare Drives 116
 - Replacing Hot Spare Drives 116
 - Inserting Physical Drives into Unused Slots 116
- Virtual Drive Naming 117
- LUN Dereferencing 117

Controller Constraints and Limitations	117
Storage Profiles	118
Creating a Storage Profile	118
Creating a Specific Storage Profile	119
Deleting a Storage Profile	120
Local LUNs	120
Configuring Local LUNs	120
Displaying Details of All Local LUNs Inherited By a Service Profile	122
Deleting Local LUNs	123
LUN Set	123
LUN Set	123
Creating a LUN Set	124
Displaying the Details of a LUN Set	126
Deleting a LUN Set	128
PCH Controller Definitions	129
PCH SSD Controller Definition	129
Creating a Storage Profile PCH Controller Definition	130
Modifying a Service Profile PCH Controller Definition	135
Deleting a Storage Profile PCH Controller Definition	138
Migrating an M.2 Module	140
Replacing a Faulty M.2 Disk	142
Associating a Storage Profile with an Existing Service Profile	143
Configuring Storage Profiles	143
Importing Foreign Configurations for a RAID Controller on a Blade Server	143
Importing Foreign Configurations for a RAID Controller on a Rack Server	144
Configuring Local Disk Operations on a Blade Server	144
Configuring Local Disk Operations on a Rack Server	145
Configuring Local Disk Operations	145
Boot Policy for Local Storage	147
Configuring the Boot Policy for an Embedded Local LUN	148
Configuring the Boot Policy for an Embedded Local Disk	149
Local LUN Operations in a Service Profile	149
Preprovisioning a LUN Name	149
Claiming an Orphan LUN	150

Deploying and Undeploying a LUN 150
 Renaming a Service Profile Referenced LUN 151

CHAPTER 12

Mini Storage 153

Mini Storage 153
 Viewing Mini Storage Properties 153

CHAPTER 13

Configuring SD Card Support 155

FlexFlash Secure Digital Card Support 155
 FlexFlash FX3S Support 157
 Enabling FlexFlash SD Card Support 158
 Disabling FlexFlash SD Card Support 158
 Enabling Auto-Sync 159
 Formatting the SD Cards 159
 Resetting the FlexFlash Controller 160
 FlexUtil Secure Digital Card Support 160

CHAPTER 14

Direct Attached Storage 161

Direct Attached Storage 161
 Fibre Channel Switching Mode 162
 Configuring Fibre Channel Switching Mode 162
 Creating a Storage VSAN 163
 Creating a VSAN for Fibre Channel Zoning 164
 Configuring a Fibre Channel Storage Port 166
 Configuring Fibre Channel Zoning 167
 Creating a Fibre Channel Storage Connection Policy 168
 Creating a Service Profile with the Expert Wizard 169
 Associating a Service Profile with a Server or Server Pool 170
 Verifying Fibre Channel Zoning Configuration 171
 Troubleshooting Fibre Channel Zoning Configuration 172

CHAPTER 15

Storage Inventory 173

Local Disk Locator LED Status 173
 Toggling the Local Disk Locator LED On and Off 174

Custom LED Status with VMD on NVMe	174
NVMe-optimized M5 Servers	177
MSwitch Disaster Recovery	178
NVMe PCIe SSD Inventory	179
Viewing NVMe PCIe SSD Storage Inventory	179
Enabling Volume Management Device on UCS Storage	180
Enabling Intel® Volume Management Device	180
Volume Management Device (VMD) Setup	180
Enabling VMD on UCS Manager	181
Enabling Volume Management Device (VMD) in Passthrough Mode	181
Volume Management Device (VMD) Passthrough Mode	181
Configuring VMD Passthrough	182
Downloading VMD Drivers	182
Intel® Volume Management Device Drivers	182
Downloading the Linux VMD Drivers	184
Downloading the Windows VMD Drivers	184
Downloading the VMD Passthrough Drivers	185
Custom LED Status with VMD on NVMe	186
Enabling Intel® Volume Management Device	189
Volume Management Device (VMD) Setup	189
Enabling VMD on UCS Manager	189
Enabling Volume Management Device (VMD) in Passthrough Mode	190
Volume Management Device (VMD) Passthrough Mode	190
Configuring VMD Passthrough	190
Downloading VMD Drivers	191
Intel® Volume Management Device Drivers	191
Intel® Virtual RAID on CPU (VRoC) with VMD	192
Downloading the Linux VMD Drivers	192
Downloading the Windows VMD Drivers	193
Downloading the VMD Passthrough Drivers	194
Custom LED Status with VMD on NVMe	195
CHAPTER 16	
Cisco UCS S3260 System Storage Management	199
Storage Server Features and Components Overview	199

Cisco UCS S3260 Storage Management Operations	207
Disk Sharing for High Availability	208
Disk Zoning Policies	208
Creating a Disk Zoning Policy	209
Creating Disk Slots and Assigning Ownership	212
Associating Disk Zoning Policies to Chassis Profile	213
Disk Migration	214
Storage Enclosure Operations	215
Removing Chassis Level Storage Enclosures	215
Sas Expander Configuration Policy	216
Creating Sas Expander Configuration Policy	216
Deleting a Sas Expander Configuration Policy	218



Preface

- [Audience, on page xiii](#)
- [Conventions, on page xiii](#)
- [Related Cisco UCS Documentation, on page xv](#)
- [Documentation Feedback, on page xv](#)

Audience

This guide is intended primarily for data center administrators with responsibilities and expertise in one or more of the following:

- Server administration
- Storage administration
- Network administration
- Network security

Conventions

Text Type	Indication
GUI elements	GUI elements such as tab titles, area names, and field labels appear in this font . Main titles such as window, dialog box, and wizard titles appear in this font .
Document titles	Document titles appear in <i>this font</i> .
TUI elements	In a Text-based User Interface, text the system displays appears in <code>this font</code> .
System output	Terminal sessions and information that the system displays appear in <code>this font</code> .
CLI commands	CLI command keywords appear in this font . Variables in a CLI command appear in <i>this font</i> .
[]	Elements in square brackets are optional.

Text Type	Indication
{x y z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
<>	Nonprinting characters such as passwords are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.



Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the document.



Tip Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.



Timesaver Means *the described action saves time*. You can save time by performing the action described in the paragraph.



Caution Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.



Warning IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Cisco UCS Documentation

Documentation Roadmaps

For a complete list of all B-Series documentation, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/overview/guide/UCS_roadmap.html

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/overview/guide/ucs_rack_roadmap.html.

For information on supported firmware versions and supported UCS Manager versions for the rack servers that are integrated with the UCS Manager for management, refer to [Release Bundle Contents for Cisco UCS Software](#).

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to ucs-docfeedback@external.cisco.com. We appreciate your feedback.



CHAPTER 1

New and Changed Information

- [New and Changed Information](#), on page 1

New and Changed Information

This section provides information on new feature and changed behavior in Cisco UCS Manager, Release 4.1(3e).

Table 1: New Features and Changed Behavior in Cisco UCS Manager, Release 4.1(3e)

Feature	Description	Where Documented
Disable Local Storage RAID level	Added Disable Local Storage to support PCH SATA Controller (AHCI mode).	Creating a Storage Profile PCH Controller Definition , on page 130 and Modifying a Service Profile PCH Controller Definition , on page 135

This section provides information on new feature and changed behavior in Cisco UCS Manager, Release 4.1(2a).

Table 2: New Features and Changed Behavior in Cisco UCS Manager, Release 4.1(2a)

Feature	Description	Where Documented
Secure Data Deletion	Secure data deletion for compliance with Commission Regulation (EU) 2019/424 is now supported on some Cisco UCS servers.	Secure Data Deletion , on page 104

This section provides information on new feature and changed behavior in Cisco UCS Manager, Release 4.1(1a).

Table 3: New Features and Changed Behavior in Cisco UCS Manager, Release 4.1(1a)

Feature	Description	Where Documented
Virtual RAID on CPU (VRoC) with VMD	Cisco UCS Manager now supports Intel® Virtual RAID on CPU (VRoC), which allows you to create and manage RAID volumes within the BIOS of VMD-enabled Intel® NVMe SSD drives .	Volume Management Device (VMD) Setup, on page 85



CHAPTER 2

Overview

- [Overview, on page 3](#)
- [Cisco UCS Manager User Documentation, on page 3](#)
- [Storage Options, on page 4](#)
- [Storage Design Considerations, on page 6](#)
- [Storage Configuration Sequence, on page 6](#)
- [Storage Protocols, on page 7](#)
- [The UCS Manager SAN Tab, on page 7](#)

Overview

This guide describes how to configure the following storage management tasks:

- Ports and Port Channels
- Named VSANs
- SAN Pin Groups
- SAN Uplinks
- Pools
- FC Identity Assignment
- Storage-Related Policies
- Storage Profiles
- FlexFlash SD Card Support
- Direct Attached Storage
- Storage Inventory

Cisco UCS Manager User Documentation


Cisco UCS Manager offers you a new set of smaller, use-case based documentation described in the following table:

Guide	Description
Cisco UCS Manager Getting Started Guide	Discusses Cisco UCS architecture and Day 0 operations, including Cisco UCS Manager initial configuration and configuration best practices.
Cisco UCS Manager Administration Guide	Discusses password management, role-based access configuration, remote authentication, communication services, CIMC session management, organizations, backup and restore, scheduling options, BIOS tokens, and deferred deployments.
Cisco UCS Manager Infrastructure Management Guide	Discusses physical and virtual infrastructure components used and managed by Cisco UCS Manager.
Cisco UCS Manager Firmware Management Guide	Discusses downloading and managing firmware, upgrading through Auto Install, upgrading through service profiles, directly upgrading at endpoints using firmware auto sync, managing the capability catalog, deployment scenarios, and troubleshooting.
Cisco UCS Manager Server Management Guide	Discusses the new licenses, registering Cisco UCS domain with Cisco UCS Central, power capping, server boot, server profiles, and server-related policies.
Cisco UCS Manager Storage Management Guide	Discusses all aspects of storage management, such as SAN and VSAN in Cisco UCS Manager.
Cisco UCS Manager Network Management Guide	Discusses all aspects of network management, such as LAN and VLAN connectivity in Cisco UCS Manager.
Cisco UCS Manager System Monitoring Guide	Discusses all aspects of system and health monitoring, including system statistics in Cisco UCS Manager.
Cisco UCS S3260 Server Integration with Cisco UCS Manager	Discusses all aspects of management of UCS S-Series servers that are managed through Cisco UCS Manager.

Storage Options

The following are the UCS Manager storage options and the benefits of each.

Figure 1: Cisco UCS Manager Storage Options



- **Direct Attached Storage (DAS)**—This is the storage available inside a server and is directly connected to the system through the motherboard within a parallel SCSI implementation. DAS is commonly described as captive storage. Devices in a captive storage topology do not have direct access to the storage network and do not support efficient sharing of storage. To access data with DAS, a user must go through a front-end network. DAS devices provide little or no mobility to other servers and little scalability.

DAS devices limit file sharing and can be complex to implement and manage. For example, to support data backups, DAS devices require resources on the host and spare disk systems that other systems cannot use. The cost and performance of this storage depends upon the disks and RAID controller cards inside the servers. DAS is less expensive and is simple to configure; however, it lacks the scalability, performance, and advanced features provided by high-end storage.

- **Network Attached Storage (NAS)**—This storage is usually an appliance providing file system access. This storage could be as simple as an Network File System (NFS) or Common Internet File System (CIFS) share available to the servers. Typical NAS devices are cost-effective devices with not very high performance but have very high capacity with some redundancy for reliability. NAS is usually moderately expensive, simple to configure, and provides some advanced features; however, it also lacks scalability, performance, and advanced features provided by SAN.
- **Storage Area Network (SAN)**—A SAN is a specialized, high-speed network that attaches servers and storage devices. A SAN allows an any-to-any connection across the network by using interconnect elements, such as switches and directors. It eliminates the traditional dedicated connection between a server and storage, and the concept that the server effectively owns and manages the storage devices. It also eliminates any restriction to the amount of data that a server can access, currently limited by the number of storage devices that are attached to the individual server. Instead, a SAN introduces the flexibility of networking to enable one server or many heterogeneous servers to share a common storage utility. A network might include many storage devices, including disk, tape, and optical storage. Additionally, the storage utility might be located far from the servers that it uses. This type of storage

provides maximum reliability, expandability, and performance. The cost of SAN is also very high compared to other storage options.

SAN is the most resilient, highly scalable, and high performance storage; however, it is also the most expensive and complex to manage.

Storage Design Considerations

UCS storage physical connectivity has a slightly different design consideration as compared to LAN physical connectivity. The following are some design considerations for SAN connectivity:

- Northbound storage physical connectivity does not support virtual port channels (vPCs) like LAN connectivity.
- Port channels or trunking is possible to combine multiple storage uplink ports that provide physical link redundancy.
- Redundancy of storage resources is handled by the storage itself and varies from vendor to vendor.
- Connect storage through northbound Cisco storage devices, such as Nexus or MDS Fabric Switches.
- It is possible to connect storage directly to UCS Fabric Interconnects, which is recommended for small implementations because of the fabric interconnect physical ports consumption and increased processing requirements.
- Software configuration including VSANs and zoning is required for providing access to storage resources.

Storage Configuration Sequence

Follow the suggested sequence to configure a storage network:

1. Configure and enable server ports, uplink ports, and FC ports.
2. Create a management IP address pool (typically on the same subnet as the UCS Manager Admin IP address).
3. Create an UUID Pool, MAC Pool, WWNN Pool, WWPN Pool (or populate the corresponding "default" pools). Embed domain ID's. Use Fabric-specific Pools for MAC and WWPN (for example, Fabric-A, Fabric-B).
4. For SAN boot, create a unique "Boot Policy" for each storage array boot target.
5. Create VNIC templates (for example, eth0-A, eth1-B) that both draw from the above MAC Pool, and are associated with Fabric-A and Fabric-B respectively.
6. Create VHBA templates (for example, fc0-A, fc1-B) that both draw from the above WWPN Pool, and are associated with Fabric-A and Fabric-B respectively.
7. Create service profile templates that draw from all earlier established pools, policies and templates, as appropriate.
8. Instantiate the service profile from the template and associate the service profile to a given blade, or set the service profile template to associate with a particular server pool.

Storage Protocols

Fiber Channel, iSCSI, and Fiber Channel over Ethernet are protocols for SAN connectivity.

- **iSCSI**—An industry standard protocol for attaching various I/O peripherals such as printers, scanners, tape drives, and storage devices. The most common SCSI devices are disks and tape libraries.

SCSI is the core protocol to connect raw hard disk storage with the servers. To control remote storage with the SCSI protocol, different technologies are used as wrappers to encapsulate commands, such as FC and iSCSI.

Fiber Channel protocol provides the infrastructure to encapsulate the SCSI traffic and provided connectivity between computers and storage. FC operates at speeds of 2, 4, 8, and 16 Gbps.

- **Fiber Channel (FC)** consists of the following:
 - Hard disk arrays that provide raw storage capacity.
 - Storage processors to manage hard disks and provide storage LUNs and masking for the servers.
 - Fiber Channel Switches (also known as Fabric) that provide connectivity between storage processors and server HBAs.
 - Fiber Channel Host Bus Adapters: They are installed in the computer and provide connectivity to the SAN.

Fiber Channel identifies infrastructure components with World Wide Numbers (WWN) . WWNs are 64-bit addresses which uniquely identify the FC devices. Like MAC addresses, it has bits assigned to vendors to identify their devices. Each end device (like an HBA port) is given a World Wide Port Number (WWPN) and each connectivity device (like a Fabric switch) is given a World Wide Node Number (WWNN).

A Fiber Channel HBA used for connecting to a SAN is known as an initiator, and Fiber Channel SAN providing disks as LUNs is known as a target. The Fiber Channel protocol is different from Ethernet or TCP/IP protocols.

- **Fiber Channel over Ethernet (FCoE)** transport replaces the Fibre Channel cabling with 10 Gigabit Ethernet cables and provides lossless delivery over unified I/O. Ethernet is widely used in networking. With some advancement such as Data Center Ethernet (DCE) and Priority Flow Control (PFC) in Ethernet to make it more reliable for the datacenter, Fiber Channel is now also implemented on top of Ethernet. This implementation is known as FCoE.

The UCS Manager SAN Tab

From the SAN tab, you the UCS administrator can create, modify, and delete configuration elements related to SANs (FC, iSCSI) or direct attached FC/FCoE, NAS appliances, and communications.

The major nodes in this tab are the following:

- **SAN Cloud**—This node allows you to:
 - Configure SAN uplinks, including storage ports and port channels and SAN pin groups.
 - View the FC identity assignment

- Configure WWN Pools, including WWPN, WWxN, and WWxN, and iSCSI Qualified Name (IQN), pools.
- View the FSM details for a particular end point to determine if a task succeeded or failed and use the FSM to troubleshoot any failures.
- Monitor storage events and faults for health management.
- **Storage Cloud**—This node allows you to:
 - Configure storage FC links and storage FCoE interfaces (using SAN Storage Manager).
 - Configure VSAN settings.
 - Monitor SAN cloud events for health management.
- **Policies**—This node allows you to:
 - Configure threshold policies, classes, and properties and monitor events.
 - Configure threshold organization and sub-organization storage policies, including default VHBA, behavior, FC adaptor, LACP, SAN connectivity, SAN connector, and VHBA templates.
- **Pools**—This node allows you to configure pools defined in the system, including IQN, IQN suffix, WWNN, WWPN, and WWxN.
- **Traffic Monitoring Sessions**—This node allows you to configure port traffic monitoring sessions defined in the system.



CHAPTER 3

SAN Ports and Port Channels

- [Port Modes, on page 9](#)
- [Port Types, on page 9](#)
- [Server Ports, on page 10](#)
- [Reconfiguring a Port on a Fabric Interconnect, on page 11](#)
- [Enabling or Disabling a Port on a Fabric Interconnect, on page 11](#)
- [Unconfiguring a Port on a Fabric Interconnect, on page 12](#)
- [Appliance Ports, on page 12](#)
- [FCoE and Fibre Channel Storage Ports, on page 14](#)
- [FC Links Rebalancing, on page 16](#)
- [Converting FC Storage Port to FC Uplink Port, on page 16](#)
- [FCoE Uplink Ports, on page 17](#)
- [Unified Storage Ports, on page 18](#)
- [Unified Uplink Ports, on page 19](#)
- [Policy-Based Port Error Handling, on page 21](#)
- [Fibre Channel Port Channels, on page 22](#)
- [FCoE Port Channels, on page 26](#)
- [Unified Uplink Port Channel, on page 26](#)

Port Modes

The port mode determines whether a unified port on the fabric interconnect is configured to carry Ethernet or Fibre Channel traffic. You configure the port mode in Cisco UCS Manager. However, the fabric interconnect does not automatically discover the port mode.

Changing the port mode deletes the existing port configuration and replaces it with a new logical port. Any objects associated with that port configuration, such as VLANs and VSANS, are also removed. There is no restriction on the number of times you can change the port mode for a unified port.

Port Types

The port type defines the type of traffic carried over a unified port connection.

By default, unified ports changed to Ethernet port mode are set to the Ethernet uplink port type. Unified ports changed to Fibre Channel port mode are set to the Fibre Channel uplink port type. You cannot unconfigure Fibre Channel ports.

Changing the port type does not require a reboot.

Ethernet Port Mode

When you set the port mode to Ethernet, you can configure the following port types:

- Server ports
- Ethernet uplink ports
- Ethernet port channel members
- FCoE ports
- Appliance ports
- Appliance port channel members
- SPAN destination ports
- SPAN source ports



Note For SPAN source ports, configure one of the port types and then configure the port as SPAN source.

Fibre Channel Port Mode

When you set the port mode to Fibre Channel, you can configure the following port types:

- Fibre Channel uplink ports
- Fibre Channel port channel members
- Fibre Channel storage ports
- SPAN source ports



Note For SPAN source ports, configure one of the port types and then configure the port as SPAN source.

Server Ports

Configuring Server Ports

All of the port types listed are configurable on both the fixed module and expansion module, including the server ports.

This task describes only one method of configuring ports. You can also configure ports from a right-click menu, or in the LAN Uplinks Manager.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name* > **Fixed Module** > **Ethernet Ports**.
 - Step 3** Click a port under the **Ethernet Ports** node.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Actions** area, click **Reconfigure**.
 - Step 6** From the drop-down list, choose **Configure as Server Port**.
-

Reconfiguring a Port on a Fabric Interconnect

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
 - Step 3** Expand the node for the ports that you want to reconfigure.
 - Step 4** Click the port or ports that you want to reconfigure.
 - Step 5** In the **Work** pane, click the **General** tab.
 - Step 6** In the **Actions** area, click **Reconfigure**.
 - Step 7** From the drop-down list, choose which way you want the port reconfigured.
-

Example: Reconfiguring an Uplink Ethernet Port as a Server Port

1. Expand the **Ethernet Ports** node and select the port you want to reconfigure.
2. Follow steps 5 and 6 above.
3. From the drop-down list choose **Configure as Server Port**.

Enabling or Disabling a Port on a Fabric Interconnect

After you enable or disable a port on a fabric interconnect, wait for at least 1 minute before you re-acknowledge the chassis. If you re-acknowledge the chassis too soon, the pinning of server traffic from the chassis might not get updated with the changes to the port that you enabled or disabled.

You can enable or disable a port only when it is configured. If the port is unconfigured, the enable and disable options are not active.

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
 - Step 3** Expand the node for the ports that you want to enable or disable.
 - Step 4** Under the **Ethernet Ports** node, select a port.
 - Step 5** In the **Work** pane, click the **General** tab.
 - Step 6** In the **Actions** area, click **Enable Port** or **Disable Port**.
 - Step 7** If a confirmation dialog box displays, click **Yes**.
 - Step 8** Click **OK**.
-

Unconfiguring a Port on a Fabric Interconnect

Procedure

-
- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
 - Step 3** Expand the node for the ports that you want to unconfigure.
 - Step 4** Under the **Ethernet Ports** node, select a port.
 - Step 5** In the **Work** pane, click the **General** tab.
 - Step 6** In the **Actions** area, click **Unconfigure**.
 - Step 7** If a confirmation dialog box displays, click **Yes**.
 - Step 8** Click **OK**.
-

Appliance Ports

Appliance ports are only used to connect fabric interconnects to directly attached NFS storage.



-
- Note** When you create a new appliance VLAN, its IEEE VLAN ID is not added to the LAN Cloud. Therefore, appliance ports that are configured with the new VLAN remain down, by default, due to a pinning failure. To bring up these appliance ports, you have to configure a VLAN in the LAN Cloud with the same IEEE VLAN ID.
-

Cisco UCS Manager supports up to four appliance ports per fabric interconnect.

Configuring an Appliance Port

You can configure Appliance ports on either the fixed module or an expansion module.

This task describes only one method of configuring appliance ports. You can also configure appliance ports from the **General** tab for the port.



Note If you configure an appliance port when the uplink port is down, Cisco UCS Manager may display an error message stating that the appliance port has failed. This message is controlled by the **Action on Uplink Fail** option in the associated Network Control Policy.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the node for the ports that you want to configure.
- Step 4** Under the **Ethernet Ports** node, select a port.

If you want to reconfigure a server port, uplink Ethernet port, or FCoE storage port, expand the appropriate node.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Reconfigure**.
- Step 7** From the drop-down list, click **Configure as Appliance Port**.
- Step 8** If a confirmation dialog box displays, click **Yes**.
- Step 9** In the **Configure as Appliance Port** dialog box, complete the required fields.
- Step 10** In the **VLANs** area, do the following:
 - a) In the **Port Mode** field, click one of the following radio buttons to select the mode you want to use for the port channel:
 - **Trunk**—Cisco UCS Manager GUI displays the VLANs Table that lets you choose the VLANs you want to use.
 - **Access**—Cisco UCS Manager GUI displays the **Select VLAN** drop-down list that allows you to choose a VLAN to associate with this port or port channel.

With either mode, you can click the **Create VLAN** link to create a new VLAN.

Note If traffic for the appliance port needs to traverse the uplink ports, you must also define each VLAN used by this port in the LAN cloud. For example, you need the traffic to traverse the uplink ports if the storage is also used by other servers, or if you want to ensure that traffic fails over to the secondary fabric interconnect if the storage controller for the primary fabric interconnect fails.

 - b) If you clicked the **Trunk** radio button, complete the required fields in the VLANs table.
 - c) If you clicked the **Access** radio button, choose a VLAN from the **Select VLAN** drop-down list.

- Step 11** (Optional) If you want to add an endpoint, check the **Ethernet Target Endpoint** check box and specify the name and MAC address.
- Step 12** Click OK.
-

Modifying the Properties of an Appliance Port

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the node for the appliance port that you want to modify.
- Step 4** Expand **Ethernet Ports**.
- Step 5** Click the appliance port for which you want to modify the properties.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Actions** area, click **Show Interface**.
- You may need to expand the pane or use the scroll bars in the **Properties** dialog box to see all the fields.
- Step 8** In the **Properties** dialog box, modify the values as needed.
- Step 9** Click OK.
-

FCoE and Fibre Channel Storage Ports

Configuring an Ethernet Port as an FCoE Storage Port

You can configure FCoE storage ports on either the fixed module or an expansion module.

This task describes only one method of configuring FCoE storage ports. You can also configure FCoE storage ports from the **General** tab for the port.

Before you begin

The Fibre Channel switching mode must be set to Switching for these ports to be valid. The storage ports cannot function in end-host mode.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Depending upon the location of the ports you want to configure, expand one of the following:

- **Fixed Module**
- **Expansion Module**

- Step 4** Click one or more of the ports under the **Ethernet Ports** node.
If you want to reconfigure an uplink Ethernet port, server port, or appliance port, expand the appropriate node.
- Step 5** Right-click the selected port or ports and choose **Configure as FCoE Storage Port**.
On Cisco UCS 6454 Fabric Interconnects, ports 49-54 cannot be configured as FCoE storage ports.
On Cisco UCS 64108 Fabric Interconnects, port 97-108 cannot be configured as FCoE storage ports.
- Step 6** If a confirmation dialog box displays, click **Yes**.
- Step 7** Click **OK**.
-

Configuring a Fibre Channel Storage Port

This task describes only one method of configuring FC storage ports. You can also configure FC storage ports from the **General** tab for the port.

Before you begin

The Fibre Channel switching mode must be set to Switching for these ports to be valid. The storage ports cannot function in end-host mode.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the **Expansion Module** node.
- Step 4** Click one or more of the ports under the **FC Ports** node.
- Step 5** Right-click the selected port or ports and choose **Configure as FC Storage Port**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
- Step 7** Click **OK**.
-

Restoring an Uplink Fibre Channel Port

This task describes only one method of restoring an FC storage port to function as an uplink FC port. You can also reconfigure FC storage ports from the **General** tab for the port.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.

- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
 - Step 3** Expand the **Expansion Module** node.
 - Step 4** Click one or more of the ports under the **FC Ports** node.
 - Step 5** Right-click the selected port or ports and choose **Configure as Uplink Port**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
 - Step 7** Click **OK**.
-

FC Links Rebalancing

The FC uplinks balance automatically when FC Port Channels are utilized. To create FC Port Channels, refer to [Creating a Fibre Channel Port Channel, on page 22](#).

For the FC uplinks that are not members of the Port Channels (Individual ISLs), load balancing is done according to the FC uplinks balancing algorithm. For a vHBA of a host or service profile to choose an available FC uplink, when FC uplink trunking is disabled, the uplink and vHBA must belong to the same VSAN

For each vHBA, the algorithm searches for an FC uplink in the following order:

1. Least used FC uplink based on the number of vHBAs currently bound to the uplink.
2. If FC uplinks are equally balanced, then round robin is used.

This process continues for all the other vHBAs. The algorithm also considers other parameters such as pre-fip/fip adapters and number of flogis. You may not see the least-used component when there are less than six flogis.

After a port configuration or any other uplink state changes, if the traffic passing through the FC uplinks is no longer balanced, you can re-balance the traffic by resetting the vHBA(s) on each adapter and allow the load balancing algorithm to evaluate for the current state of the FC uplinks.

Converting FC Storage Port to FC Uplink Port

You can configure an FC Uplink port on either a fixed module or an expansion module.

This task describes only one method of configuring FC Uplink ports. You can also configure FC uplink ports from a right-click menu for the port.



Important For Cisco UCS 6400 Series, the fill pattern is greyed out and is automatically set to IDLE.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the node for the ports that you want to configure.

- Step 4** Under the **FC Ports** node, select any **Storage** port.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** From the **Actions** area, select **Configure as Uplink Port**.
- Step 7** If a confirmation dialog box displays, click **Yes**.
- Step 8** The Cisco UCS Manager GUI displays a success message.
- In the **Actions** area, **Configure as Uplink Port** becomes grayed out and **Configure as FC Storage Port** becomes active.
-

FCoE Uplink Ports

FCoE uplink ports are physical Ethernet interfaces between the fabric interconnects and the upstream Ethernet switch, used for carrying FCoE traffic. With this support the same physical Ethernet port can carry both Ethernet traffic and Fibre Channel traffic.

FCoE uplink ports connect to upstream Ethernet switches using the FCoE protocol for Fibre Channel traffic. This allows both the Fibre Channel traffic and Ethernet traffic to flow on the same physical Ethernet link.



Note FCoE uplinks and unified uplinks enable the multi-hop FCoE feature, by extending the unified fabric up to the distribution layer switch.

You can configure the same Ethernet port as any of the following:

- **FCoE uplink port**—As an FCoE uplink port for only Fibre Channel traffic.
- **Uplink port**—As an Ethernet port for only Ethernet traffic.
- **Unified uplink port**—As a unified uplink port to carry both Ethernet and Fibre Channel traffic.

Configuring FCoE Uplink Ports

You can configure an FCoE Uplink port on either a fixed module or an expansion module.

This task describes only one method of configuring FCoE Uplink ports. You can also configure FCoE uplink ports from a right-click menu or from the General tab for the port.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the node for the ports that you want to configure.
- Step 4** Under the **Ethernet Ports** node, select any **Unconfigured** port.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Reconfigure**.

- Step 7** From the drop down options, select **Configure as FCoE Uplink Port**.
- Step 8** If a confirmation dialog box displays, click **Yes**.
- Step 9** The Cisco UCS Manager GUI displays a success message.
In the **Properties** area, the **Role** changes to **Fcoe Uplink**.
-

Unified Storage Ports

Unified storage involves configuring the same physical port as both an Ethernet storage interface and an FCoE storage interface. You can configure any appliance port or FCoE storage port as a unified storage port. To configure a unified storage port, you must have the fabric interconnect in Fibre Channel switching mode.

In a unified storage port, you can enable or disable individual FCoE storage or appliance interfaces.

- In an unified storage port, if you do not specify a non-default VLAN for the appliance port, the FCoE-storage-native-vlan will be assigned as the native VLAN on the unified storage port. If the appliance port has a non-default native VLAN specified as native VLAN, this will be assigned as the native VLAN for the unified storage port.
- When you enable or disable the appliance interface, the corresponding physical port is enabled or disabled. So when you disable the appliance interface in unified storage, even if the FCoE storage is enabled, it goes down with the physical port.
- When you enable or disable the FCoE storage interface, the corresponding VFC is enabled or disabled. So when the FCoE storage interface is disabled in a unified storage port, the appliance interface will continue to function normally.

Configuring an Appliance Port as a Unified Storage Port

You can configure a unified storage port either from an appliance port or from an FCoE storage port. You can also configure the unified storage port from an unconfigured port. If you start from an unconfigured port, you will assign either an appliance configuration or an FCoE storage configuration to the port, and then will add another configuration to enable it as a unified storage port.



Important Make sure the fabric interconnect is in Fibre Channel switching mode.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Depending on the location of the ports you want to configure, expand one of the following:
- **Fixed Module**
 - **Expansion Module**

- Step 4** Under the **Ethernet Ports** node, select any the port that is already configured as an appliance port. In the **Work** pane, under the **General** tab, in the **Properties** area, the **Role** will show as **Appliance Storage**.
- Step 5** In the **Actions** area, click **Reconfigure**.
- Step 6** From the pop-up menu, select **Configure as FCoE Storage** port.
- Step 7** If a confirmation dialog box displays, click **Yes**.
- Step 8** The Cisco UCS Manager GUI displays a success message. In the **Properties** area, the **Role** changes to **Unified Storage**.
-

Unconfiguring a Unified Storage Port

You can unconfigure and remove both configurations from the unified connect port. Or you can unconfigure either of them and retain the other on the port.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the node for the ports that you want to unconfigure.
- Step 4** Under the **Ethernet Ports** node, select the port that you want to unconfigure.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Unconfigure**. You will see the following options:
- **Unconfigure FCoE Storage Port**
 - **Unconfigure Appliance Port**
 - **Unconfigure both**
- Step 7** Select one of the unconfigure options.
- Step 8** If a confirmation dialog box displays, click **Yes**.
- Step 9** The Cisco UCS Manager GUI displays a success message. In the **Properties** area, the **Role** changes to based on your unconfigure selection.
-

Unified Uplink Ports

When you configure an Ethernet uplink and an FCoE uplink on the same physical Ethernet port, it is called a unified uplink port. You can individually enable or disable either the FCoE or Ethernet interfaces independently.

- Enabling or disabling the FCoE uplink results in the corresponding VFC being enabled or disabled.
- Enabling or disabling an Ethernet uplink results in the corresponding physical port being enabled or disabled.

If you disable an Ethernet uplink, it disables the underlying physical port in a unified uplink. Therefore, even when the FCoE uplink is enabled, the FCoE uplink also goes down. But if you disable an FCoE uplink, only the VFC goes down. If the Ethernet uplink is enabled, it can still function properly in the unified uplink port.

Configuring Unified Uplink Ports

You can configure the unified uplink port from either of the following:

- From an existing FCoE uplink port or Ethernet uplink port
- From an unconfigured uplink port

You can configure the unified uplink port on either a fixed module or on an expansion module.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the node for the ports that you want to configure.
- Step 4** Under the **Ethernet Ports** node, select a port.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Properties** area, make sure the **Role** shows as **Fcoe Uplink**.
- Step 7** In the **Actions** area, click **Reconfigure**.
- Step 8** From the drop-down options, select **Configure as Uplink Port**.
- Step 9** If a confirmation dialog box displays, click **Yes**.
- Step 10** The Cisco UCS Manager GUI displays a success message.
- In the **Properties** area, the **Role** changes to **Unified Uplink**.
-

Unconfiguring Unified Uplink Port

You can unconfigure and remove both configurations from the unified uplink port. Or you can unconfigure either the FCoE configuration or Ethernet port configuration and retain the other on the port.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the node for the ports that you want to unconfigure.
- Step 4** Under the **Ethernet Ports** node, select the port you want to unconfigure.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Unconfigure**. Select one of the following options:
- **Unconfigure FCoE Uplink Port**

- **Unconfigure Uplink Port**
- **Unconfigure both**

- Step 7** If a confirmation dialog box displays, click **Yes**.
- Step 8** The Cisco UCS Manager GUI displays a success message. In the **Properties** area, the **Role** changes based on your unconfigure selection.
- Step 9** Click **Save Changes**.
-

Policy-Based Port Error Handling

If Cisco UCS Manager detects any errors on active network interface (NI) ports, and if the error-disable feature has been implemented, Cisco UCS Manager automatically disables the respective fabric interconnect port that is connected to the NI port that had errors. When a fabric interconnect port is error disabled, it is effectively shut down and no traffic is sent or received on that port.

The error-disable function serves two purposes:

- It lets you know which fabric interconnect port is error-disabled and that the connected NI Port has errors.
- It eliminates the possibility that this port can cause the failure of other ports other ports connected to the same Chassis/FEX. Such a failure can occur when the NI port has errors, which can ultimately cause serious network issues. The error-disable function helps prevent these situations.

Configuring Error-Based Action

Procedure

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Expand **Admin > All > Stats Management > fabric > Internal LAN > thr-policy-default > etherNiErrStats**.
- Step 3** Select a delta property.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** To implement an error-disable state on a fabric interconnect port, check the **Disable FI port when fault is raised** check box.
- Step 6** To enable auto recovery, in the **Enable Auto Recovery** field, select **Enable**.
- Step 7** To specify the time after which the port can automatically be re-enabled, in the **Time (in minutes)** field, type the desired value.
- Step 8** Click **Save Changes**.
-

Fibre Channel Port Channels

A Fibre Channel port channel allows you to group several physical Fibre Channel ports (link aggregation) to create one logical Fibre Channel link to provide fault-tolerance and high-speed connectivity. In Cisco UCS Manager, you create a port channel first and then add Fibre Channel ports to the port channel.



Note Fibre Channel port channels are not compatible with non-Cisco technology.

You can create up to four Fibre Channel port channels in each Cisco UCS domain with Cisco UCS 6200, 6300, , and 6400 Series Fabric Interconnects. Each Fibre Channel port channel can include a maximum of 16 uplink Fibre Channel ports.

You can create up to two Fibre Channel port channels in each Cisco UCS domain with Cisco UCS 6324 fabric interconnects. Each Fibre Channel port channel can include a maximum of four uplink Fibre Channel ports.

Ensure that the Fibre Channel port channel on the upstream NPIV switch is configured with its channel mode as **active**. If both the member port(s) and peer port(s) do not have the same channel mode configured, the port channel will not come up. When the channel mode is configured as **active**, the member ports initiate port channel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the port channel protocol, or responds with a nonnegotiable status, it defaults to the On mode behavior. The **active** port channel mode allows automatic recovery without explicitly enabling and disabling the port channel member ports at either end.

This example shows how to configure channel mode as active:

```
switch(config)# int po114
switch(config-if)# channel mode active
```

Creating a Fibre Channel Port Channel

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > SAN Cloud**.
- Step 3** Expand the node for the fabric where you want to create the port channel.
- Step 4** Right-click the **FC Port Channels** node and choose **Create Port Channel**.
- Step 5** In the **Set Port Channel Name** panel, specify the ID and name, then click **Next**.
- Step 6** In the **Add Ports** panel, specify the port channel admin speed, and add ports to the port channel.

Port Channel Admin Speed 1 Gbps and 2 Gbps are not available for Cisco UCS 6400 Series Fabric Interconnect. **Port Channel Admin Speed** 16 Gbps and 32 Gbps are available only for Cisco UCS 6400 Series Fabric Interconnect.

- Step 7** Click **Finish**.
-

Enabling a Fibre Channel Port Channel

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > SAN Cloud > Fabric > FC Port Channels**.
 - Step 3** Click the port channel you want to enable.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Actions** area, click **Enable Port Channel**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Disabling a Fibre Channel Port Channel

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > SAN Cloud > Fabric > FC Port Channels**.
 - Step 3** Click the port channel you want to disable.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Actions** area, click **Disable Port Channel**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Adding Ports to and Removing Ports from a Fibre Channel Port Channel

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > SAN Cloud > Fabric > FC Port Channels**.
- Step 3** Click the port channel to which you want to add or remove ports.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Add Ports**.
- Step 6** In the **Add Ports** dialog box, do one of the following:
 - To add ports, choose one or more ports in the **Ports** table, and then click the >> button to add the ports to the **Ports in the port channel** table.
 - To remove ports, choose one or more ports in the **Ports in the port channel** table, and then click the << button to remove the ports from the port channel and add them to the **Ports** table.

Step 7 Click **OK**.

Modifying the Properties of a Fibre Channel Port Channel



Note If you are connecting two Fibre Channel port channels, the admin speed for both port channels must match for the link to operate. If the admin speed for one or both of the Fibre Channel port channels is set to auto, Cisco UCS adjusts the admin speed automatically.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > SAN Cloud > Fabric > FC Port Channels**.
- Step 3** Click the port channel that you want to modify.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Properties** area, change the values in one or more of the following fields:

Name	Description
Name field	The user-defined name given to the port channel. This name can be between 1 and 16 alphanumeric characters.
VSAN drop-down list	The VSAN associated with the port channel.

Name	Description
Port Channel Admin Speed drop-down list	<p>The admin speed of the port channel. This can be:</p> <ul style="list-style-type: none"> • 1 Gbps <p>Note Not available for Cisco UCS 6400 Series Fabric Interconnect.</p> • 2 Gbps <p>Note Not available for Cisco UCS 6400 Series Fabric Interconnect.</p> • 4 Gbps • 8 Gbps • 16 Gbps <p>Note Available only for Cisco UCS 6400 Series Fabric Interconnect.</p> • 32 Gpbs <p>Note Available only for Cisco UCS 6400 Series Fabric Interconnect.</p> • auto <p>Note Not available for Cisco UCS 6400 Series Fabric Interconnect.</p>

Step 6 Click **Save Changes**.

Deleting a Fibre Channel Port Channel

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** Expand **SAN > SAN Cloud > Fabric > FC Port Channels**.
- Step 3** Right-click the port channel you want to delete and choose **Delete**.
- Step 4** If a confirmation dialog box displays, click **Yes**.

FCoE Port Channels

An FCoE port channel allows you to group several physical FCoE ports to create one logical FCoE port channel. At a physical level, the FCoE port channel carries FCoE traffic over an Ethernet port channel. So an FCoE port channel with a set of members is essentially an Ethernet port channel with the same members. This Ethernet port channel is used as a physical transport for FCoE traffic.

For each FCoE port channel, Cisco UCS Manager creates a VFC internally and binds it to an Ethernet port channel. FCoE traffic received from the hosts is sent over the VFC the same way as the FCoE traffic is sent over Fibre Channel uplinks.

Creating an FCoE Port Channel

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > SAN Cloud**.
 - Step 3** Expand the node for the fabric where you want to create the port channel.
 - Step 4** Right-click the **FCoE Port Channels** node and choose **Create FCoE Port Channel**.
 - Step 5** In the **Set Port Channel Name** panel of the **Create FCoE Port Channel** wizard, specify the ID and name, then click **Next**.
 - Step 6** In the **Add Ports** panel of the **Create FCoE Port Channel** wizard, specify the ports that you want to add.
 - Step 7** Click **Finish**.
-

Deleting an FCoE Port Channel

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** On the **SAN** tab, expand **SAN > SAN Cloud > Fabric > FCoE Port Channels**.
 - Step 3** Right-click the port channel you want to delete and choose **Delete**.
 - Step 4** If a confirmation dialog box displays, click **Yes**.
-

Unified Uplink Port Channel

When you create an Ethernet port channel and an FCoE port channel with the same ID, it is called a unified uplink port channel. When the unified port channel is created, a physical Ethernet port channel and a VFC are created on the fabric interconnect with the specified members. The physical Ethernet port channel is used to carry both Ethernet and FCoE traffic. The VFC binds FCoE traffic to the Ethernet port channel.

The following rules will apply to the member port sets of the unified uplink port channel:

- The Ethernet port channel and FCoE port channel on the same ID, must have the same set of member ports.
- When you add a member port channel to the Ethernet port channel, Cisco UCS Manager adds the same port channel to FCoE port channel as well. Similarly, adding a member to the FCoE port channel adds the member port to the Ethernet port channel.
- When you delete a member port from one of the port channels, Cisco UCS Manager automatically deletes the member port from the other port channel.

If you disable an Ethernet uplink port channel, it disables the underlying physical port channel in a unified uplink port channel. Therefore, even when the FCoE uplink is enabled, the FCoE uplink port channel also goes down. If you disable an FCoE uplink port channel, only the VFC goes down. If the Ethernet uplink port channel is enabled, it can still function properly in the unified uplink port channel.



CHAPTER 4

Fibre Channel Zoning

- [Information About Fibre Channel Zoning, on page 29](#)
- [Support for Fibre Channel Zoning in Cisco UCS Manager, on page 30](#)
- [Guidelines and recommendations for Cisco UCS Manager-Based Fibre Channel Zoning, on page 32](#)
- [Configuring Fibre Channel Zoning, on page 32](#)
- [Creating a VSAN for Fibre Channel Zoning, on page 33](#)
- [Creating a New Fibre Channel Zone Profile, on page 35](#)
- [Deleting a Fibre Channel Zone Profile, on page 37](#)
- [Deleting a Fibre Channel User Zone, on page 38](#)
- [Fibre Channel Storage Connection Policies, on page 38](#)

Information About Fibre Channel Zoning

Fibre Channel zoning allows you to partition the Fibre Channel fabric into one or more zones. Each zone defines the set of Fibre Channel initiators and Fibre Channel targets that can communicate with each other in a VSAN. Zoning also enables you to set up access control between hosts and storage devices or user groups.

The access and data traffic control provided by zoning does the following:

- Enhances SAN network security
- Helps prevent data loss or corruption
- Reduces performance issues

Information About Zones

A zone consists of multiple zone members and has the following characteristics:

- Members in a zone can access each other; members in different zones cannot access each other.
- Zones can vary in size.
- Devices can belong to more than one zone.
- A physical fabric can have a maximum of 8,000 zones.

Information About Zone Sets

Each zone set consists of one or more zones. You can use zone sets to enforce access control within the Fibre Channel fabric. In addition, zone sets provide you with the following advantages:

- Only one zone set can be active at any time.
- All zones in a zone set can be activated or deactivated as a single entity across all switches in the fabric.
- Changes to a zone set are not applied until the zone set has been activated. If you make changes to the active zone set, you must reactivate that zone set to apply the changes.
- A zone can be a member of more than one zone set.
- A switch in a zone can have a maximum of 500 zone sets.

Support for Fibre Channel Zoning in Cisco UCS Manager

Cisco UCS Manager supports switch-based Fibre Channel zoning and Cisco UCS Manager-based Fibre Channel zoning. You cannot configure a combination of zoning types in the same Cisco UCS domain. You can configure a Cisco UCS domain with one of the following types of zoning:

- Cisco UCS Manager-based Fibre Channel zoning—This configuration combines direct attach storage with local zoning. Fibre Channel or FCoE storage is directly connected to the fabric interconnects and zoning is performed in Cisco UCS Manager, using Cisco UCS local zoning. Any existing Fibre Channel or FCoE uplink connections need to be disabled. Cisco UCS does not currently support active Fibre Channel or FCoE uplink connections coexisting with the utilization of the UCS Local Zoning feature.
- Switch-based Fibre Channel zoning—This configuration combines direct attach storage with uplink zoning. The Fibre Channel or FCoE storage is directly connected to the fabric interconnects and zoning is performed externally to the Cisco UCS domain through an MDS or Nexus 5000 switch. This configuration does not support local zoning in the Cisco UCS domain.



Note Zoning is configured on a per-VSAN basis. You cannot enable zoning at the fabric level.

Cisco UCS Manager-Based Fibre Channel Zoning

With Cisco UCS Manager-based zoning, Cisco UCS Manager controls the Fibre Channel zoning configuration for the Cisco UCS domain, including creating and activating zones for all VSANs that you set up with this type of zoning. This type of zoning is also known as local zoning or direct attach storage with local zoning.



Note You cannot implement Cisco UCS Manager-based zoning if the VSAN is also configured to communicate with a VSAN on an upstream switch and includes Fibre Channel or FCoE uplink ports.

Supported Fibre Channel Zoning Modes

Cisco UCS Manager-based zoning supports the following types of zoning:

- Single initiator single target—Cisco UCS Manager automatically creates one zone for each vHBA and storage port pair. Each zone has two members. We recommend that you configure this type of zoning unless you expect the number of zones to exceed the maximum supported.
- Single initiator multiple targets—Cisco UCS Manager automatically creates one zone for each vHBA. We recommend that you configure this type of zoning if you expect the number of zones to reach or exceed the maximum supported.

vHBA Initiator Groups

vHBA initiator groups determine the Fibre Channel zoning configuration for all vHBAs in a service profile. Cisco UCS Manager does not include any default vHBA initiator groups. You must create vHBA initiator groups in any service profile that is to be assigned to servers included in a zone.

The configuration in a vHBA initiator group determines the following:

- The vHBAs included in the initiator group, which are sometimes referred to as vHBA initiators.
- A Fibre Channel storage connection policy, which includes the associated VSAN and the Fibre Channel target ports on the storage array.
- The type of Fibre Channel zoning to be configured for the vHBAs included in the group.

Fibre Channel Storage Connection Policy

The Fibre Channel storage connection policy contains a collection of target storage ports on storage arrays that you use to configure Cisco UCS Manager-based Fibre Channel zoning. You can create this policy under an organization or an initiator group.

The storage arrays in these zones must be directly connected to the fabric interconnects. The target storage ports on these arrays that you include in the Fibre Channel storage connection policy can be either Fibre Channel storage ports or FCoE storage ports. You use the WWN of a port to add it to the policy and to identify the port for the Fibre Channel zone.



Note Cisco UCS Manager does not create default Fibre Channel storage.

Fibre Channel Active Zone Set Configuration

In each VSAN that has been enabled for Fibre Channel zoning, Cisco UCS Manager automatically configures one zone set and multiple zones. The zone membership specifies the set of initiators and targets that are allowed to communicate with each other. Cisco UCS Manager automatically activates that zone set.

Cisco UCS Manager processes the user-configured vHBA initiator groups and their associated Fibre Channel storage connection policy to determine the desired connectivity between Fibre Channel initiators and targets. Cisco UCS Manager uses the following information to build pair-wise zone membership between initiators and targets:

- The port WWNs of the vHBA initiators derived from the vHBA initiator groups.

- The port WWNs of the storage array derived from the storage connection policy.

Switch-Based Fibre Channel Zoning

With switch-based zoning, a Cisco UCS domain inherits the zoning configuration from the upstream switch. You cannot configure or view information about your zoning configuration in Cisco UCS Manager. You have to disable zoning on a VSAN in Cisco UCS Manager to use switch-based zoning for that VSAN.

Guidelines and recommendations for Cisco UCS Manager-Based Fibre Channel Zoning

When you plan your configuration for Fibre Channel zoning, consider the following guidelines and recommendations:

Fibre Channel Switching Mode Must Be Switch Mode for Cisco UCS Manager Configurations

If you want Cisco UCS Manager to handle Fibre Channel zoning, the fabric interconnects must be in Fibre Channel Switch mode. You cannot configure Fibre Channel zoning in End-Host mode.

Symmetrical Configuration Is Recommended for High Availability

If a Cisco UCS domain is configured for high availability with two fabric interconnects, we recommend that both fabric interconnects are configured with the same set of VSANs.

Configuring Fibre Channel Zoning



Note This procedure provides a high level overview of the steps required to configure a Cisco UCS domain for Fibre Channel zoning that is controlled by Cisco UCS Manager. You must ensure that you complete all of the following steps.

Procedure

	Command or Action	Purpose
Step 1	If you have not already done so, disconnect the fabric interconnects in the Cisco UCS domain from any external Fibre Channel switches, such as an MDS.	
Step 2	If the Cisco UCS domain still includes zones that were managed by the external Fibre Channel switch, run the clear-unmanaged-fc-zone-all command on every affected VSAN to remove those zones.	This functionality is not currently available in the Cisco UCS Manager GUI. You must perform this step in the Cisco UCS Manager CLI.

	Command or Action	Purpose
Step 3	Configure the Fibre Channel switching mode for both fabric interconnects in Fibre Channel Switch mode.	You cannot configure Fibre Channel zoning in End-Host mode. See Configuring Fibre Channel Switching Mode , on page 162.
Step 4	Configure the Fibre Channel and FCoE storage ports that you require to carry traffic for the Fibre Channel zones.	See Configuring an Ethernet Port as an FCoE Storage Port , on page 14 and Configuring a Fibre Channel Storage Port , on page 15.
Step 5	Create one or more VSANs and enable Fibre Channel zoning on all VSANs that you require to carry traffic for the Fibre Channel zones.	For a cluster configuration, we recommend that you create the VSANs that you intend to include in a Fibre Channel zone in the SAN Uplinks Manager and use the common/global configuration to ensure they are accessible to both fabric interconnects. See Creating a VSAN for Fibre Channel Zoning , on page 33.
Step 6	Create one or more Fibre Channel storage connection policies.	You can perform this step when you configure Fibre Channel zoning in the service profiles, if you prefer. See Creating a VSAN for Fibre Channel Zoning , on page 33.
Step 7	Configure zoning in service profiles or service profile templates for servers that need to communicate through Fibre Channel zones.	Complete the following steps to complete this configuration: <ul style="list-style-type: none"> • Enable zoning in the VSAN or VSANs assigned to the vHBAs. See Creating a VSAN for Fibre Channel Zoning, on page 33 • Configure one or more vHBA initiator groups. See Creating a Service Profile with the Expert Wizard, on page 169.

Creating a VSAN for Fibre Channel Zoning



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, click the **SAN** node.

Step 3 In the **Work** pane, click the **SAN Uplinks Manager** link on the **SAN Uplinks** tab.

The SAN Uplinks Manager opens in a separate window.

Step 4 In the SAN Uplinks Manager, click the **VSAN** tab.

You can create the VSAN on any of the subtabs. However, if you use the **All** subtab, you can view all of the configured VSANs in the table.

Step 5 On the icon bar to the right of the table, click +.

If the + icon is disabled, click an entry in the table to enable it.

Step 6 In the **Create VSAN** dialog box, complete the following fields:

Name	Description
Name field	<p>The name assigned to the network.</p> <p>This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</p>
FC Zoning field	<p>Click the radio button to determine whether Cisco UCS Manager configures Fibre Channel zoning for the Cisco UCS domain. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The upstream switch handles Fibre Channel zoning, or Fibre Channel zoning is not implemented for the Cisco UCS domain. Cisco UCS Manager does not configure Fibre Channel zoning. • Enabled—Cisco UCS Manager configures and controls Fibre Channel zoning for the Cisco UCS domain. <p>Note If you enable Fibre Channel zoning through Cisco UCS Manager, do not configure the upstream switch with any VSANs that are being used for Fibre Channel zoning.</p>
Type radio button	<p>Click the radio button to determine how the VSAN should be configured. This can be one of the following:</p> <ul style="list-style-type: none"> • Common/Global—The VSAN maps to the same VSAN ID in all available fabrics. • Fabric A—The VSAN maps to the a VSAN ID that exists only in fabric A. • Fabric B—The VSAN maps to the a VSAN ID that exists only in fabric B. • Both Fabrics Configured Differently—The VSAN maps to a different VSAN ID in each available fabric. If you choose this option, Cisco UCS Manager GUI displays a VSAN ID field and a FCoE VLAN field for each fabric.

Name	Description
VSAN ID field	<p>The unique identifier assigned to the network.</p> <p>The ID can be between 1 and 4078, or between 4080 and 4093. 4079 is a reserved VSAN ID. In addition, if you plan to use FC end-host mode, the range between 3840 to 4079 is also a reserved VSAN ID range.</p>
FCoE VLAN field	<p>The unique identifier assigned to the VLAN used for Fibre Channel connections.</p> <p>VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:</p> <ul style="list-style-type: none"> • After an upgrade to Cisco UCS, Release 2.0—The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use. • After a fresh install of Cisco UCS, Release 2.0—The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049. <p>For FIP-capable, converged network adapters, such as the Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, the named VSAN must be configured with a named VLAN that is not the native VLAN for the FCoE VLAN ID. This configuration ensures that FCoE traffic can pass through these adapters.</p>

Step 7 Click **OK**.

Creating a New Fibre Channel Zone Profile

Perform the following procedure to create a new Fibre Channel Zone Profile.

Before you begin

Ensure that the VSAN is created for the Fiber Channel Zoning.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, click **Storage Cloud**.
- Step 3** Right-click **FC Zone Profiles** and choose **Create FC Zone Profile**.
- Step 4** In the **Create FC Zone Profile** dialog box, complete the following fields:

Field	Description
Name field	A name for the profile. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	This field is optional. Enter a brief description of the zone profile.
FC Zoning radio button	Select the desired state of the zone profile.
FC User Zones	Click the + icon on the right of the table to create FC User Zone. The Create FC User Zone window is displayed. Continue to next step for details to create FC user zone. Note If + icon is disabled, click an entry in the table to enable it.

Step 5 Complete the following fields in the **Create FC User Zone** dialog box:

Field	Description
Name field	A name for the FC Zone. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Path radio button	Click the radio button to determine how the VSAN should be configured. Following options are available: <ul style="list-style-type: none"> • Path A—The VSAN path to the VSAN ID that exists only in fabric A. • Path B—The VSAN path to the VSAN ID that exists only in fabric B.
Note	Perform one of the following to select a VSAN: <ul style="list-style-type: none"> • Select VSAN • Create VSAN • Create Storage VSAN

Field	Description
Select VSAN drop-down list	The unique identifier assigned to the VSANs that exists already in the network. The ID can be between 1 and 4078, or between 4080 and 4093. 4079 is a reserved VSAN ID. You cannot configure fibre channel zoning in end-host mode.
Create VSAN link	Click the + icon to create a new VSAN in a new window. This allows you to perform the same task as in Creating a VSAN for Fibre Channel Zoning, on page 33 . After creating the VSAN, you can return here and continue creating channel zone profile.
Create Storage VSAN link	Click the + icon to create a new storage VSAN in a new window. This allows you to perform the same task as in Creating a VSAN for Fibre Channel Zoning, on page 33 . After creating the VSAN, you can return here and continue creating channel zone profile.
Member WWPNs	Click the + icon on the right of the table to create World Wide Port Name (WWPN). The Create FC Zone Member window is displayed. Enter the WWPN for this zone.

Step 6 Click **OK** for **Create FC Zone Member** window.

Step 7 Click **OK** for **Create FC User Zone** window.

Step 8 Click **OK** for **Create FC Zone Profile** window.

The new Fibre Channel Zone Profile created is listed under **FC Zone Profiles**.

Deleting a Fibre Channel Zone Profile

Perform the following procedure to delete a Fibre Channel Zone Profile.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 On the **SAN** tab, click **Storage Cloud**.

Step 3 Click **FC Zone Profiles**.

Step 4 In the **Work** pane, right-click the name of the zone profile you wish to delete.

Step 5 Choose **Delete** and click **Yes** to confirm.

The zone profile is deleted from the system.

Deleting a Fibre Channel User Zone

Perform the following procedure to delete a Fibre Channel User Zone.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, click **Storage Cloud**.
- Step 3** Click **FC Zone Profiles**.
- Step 4** In the **Work** pane, select the zone profile.
- Step 5** Navigate to **FC User Zones**.
- Step 6** Right-click the name of the user zone you wish to delete.
- Step 7** Choose **Delete** and click **Yes** to confirm.

The user zone is deleted from the system.

Fibre Channel Storage Connection Policies

Deleting a Fibre Channel Storage Connection Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > Policies > *Organization_Name***.
 - Step 3** Expand the **Storage Connection Policies** node.
 - Step 4** Right-click the policy you want to delete and select **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-



CHAPTER 5

Named VSANs

- [Named VSANs, on page 39](#)
- [Fibre Channel Uplink Trunking for Named VSANs, on page 40](#)
- [Guidelines and Recommendations for VSANs, on page 40](#)
- [Creating a Named VSAN, on page 41](#)
- [Creating a Storage VSAN, on page 42](#)
- [Deleting a Named VSAN, on page 43](#)
- [Changing the VLAN ID for the FCoE VLAN for a Storage VSAN, on page 43](#)
- [Enabling Fibre Channel Uplink Trunking, on page 44](#)
- [Disabling Fibre Channel Uplink Trunking, on page 44](#)

Named VSANs

A named VSAN creates a connection to a specific external SAN. The VSAN isolates traffic to that external SAN, including broadcast traffic. The traffic on one named VSAN knows that the traffic on another named VSAN exists, but cannot read or access that traffic.

Like a named VLAN, the name that you assign to a VSAN ID adds a layer of abstraction that allows you to globally update all servers associated with service profiles that use the named VSAN. You do not need to reconfigure the servers individually to maintain communication with the external SAN. You can create more than one named VSAN with the same VSAN ID.

Named VSANs in Cluster Configurations

In a cluster configuration, a named VSAN can be configured to be accessible only to the Fibre Channel uplink ports on one fabric interconnect or to the Fibre Channel uplink ports on both fabric interconnects.

Named VSANs and the FCoE VLAN ID

You must configure each named VSAN with an FCoE VLAN ID. This property determines which VLAN is used for transporting the VSAN and its Fibre Channel packets.

For FIP-capable, converged network adapters, such as the Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, the named VSAN must be configured with a named VLAN that is not the native VLAN for the FCoE VLAN ID. This configuration ensures that FCoE traffic can pass through these adapters.

In the following sample configuration, a service profile with a vNIC and vHBA mapped to fabric A is associated with a server that has FIP capable, converged network adapters:

- The vNIC is configured to use VLAN 10.
- VLAN 10 is also designated as the native VLAN for the vNIC.
- The vHBA is configured to use VSAN 2.
- Therefore, VSAN 2 cannot be configured with VLAN 10 as the FCoE VLAN ID. VSAN 2 can be mapped to any other VLAN configured on fabric A.

Fibre Channel Uplink Trunking for Named VSANs

You can configure Fibre Channel uplink trunking for the named VSANs on each fabric interconnect. If you enable trunking on a fabric interconnect, all named VSANs in a Cisco UCS domain are allowed on all Fibre Channel uplink ports on that fabric interconnect.

Guidelines and Recommendations for VSANs

The following guidelines and recommendations apply to all named VSANs, including storage VSANs.

VSAN 4079 is a Reserved VSAN ID

Do not configure a VSAN as 4079. This VSAN is reserved and cannot be used in either FC switch mode or FC end-host mode.

If you create a named VSAN with ID 4079, Cisco UCS Manager marks that VSAN with an error and raises a fault.

Reserved VSAN Range for Named VSANs in FC Switch Mode

If you plan to use FC switch mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3040 to 4078.

VSANs in that range are not operational if the fabric interconnects are configured to operate in FC switch mode. Cisco UCS Manager marks that VSAN with an error and raises a fault.

Reserved VSAN Range for Named VSANs in FC End-Host Mode

If you plan to use FC end-host mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3840 to 4079.

VSANs in that range are not operational if the following conditions exist in a Cisco UCS domain:

- The fabric interconnects are configured to operate in FC end-host mode.
- The Cisco UCS domain is configured with Fibre Channel trunking or SAN port channels.

If these configurations exist, Cisco UCS Manager does the following:

1. Renders all VSANs with an ID in the range from 3840 to 4079 non-operational.
2. Raises a fault against the non-operational VSANs.
3. Transfers all non-operational VSANs to the default VSAN.

4. Transfers all vHBAs associated with the non-operational VSANs to the default VSAN.

If you disable Fibre Channel trunking and delete any existing SAN port channels, Cisco UCS Manager returns all VSANs in the range from 3840 to 4078 to an operational state and restores any associated vHBAs back to those VSANs.

Range Restrictions for Named VSAN IDs in FC Switch Mode

If you plan to use FC switch mode in a Cisco UCS domain, do not configure VSANs in the range from 3040 to 4078.

When a fabric interconnect operating in FC switch mode is connected to MDS as the upstream switch, VSANs configured in Cisco UCS Manager in the range from 3040 to 4078 and assigned as port VSANs cannot be created in MDS. This configuration results in a possible port VSAN mismatch.

Guidelines for FCoE VLAN IDs



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:

- After an upgrade to Cisco UCS, Release 2.0—The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use.
- After a fresh install of Cisco UCS, Release 2.0—The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.

Creating a Named VSAN



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

-
- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > SAN Cloud**.

Step 3 In the **Work** pane, click the **VSANs** tab.

Step 4 On the icon bar to the right of the table, click +.

If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VSAN** dialog box, complete the required.

Step 6 Click **OK**.

Cisco UCS Manager GUI adds the VSAN to one of the following **VSANs** nodes:

- The **SAN Cloud** > **VSANs** node for a storage VSAN accessible to both fabric interconnects.
- The **SAN Cloud** > *Fabric_Name* > **VSANs** node for a VSAN accessible to only one fabric interconnect.

Creating a Storage VSAN



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 On the **SAN** tab, expand **SAN** > **Storage Cloud**.

Step 3 In the **Work** pane, click the **VSANs** tab.

Step 4 On the icon bar to the right of the table, click +.

If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VSAN** dialog box, complete the required fields.

Step 6 Click **OK**.

Cisco UCS Manager GUI adds the VSAN to one of the following **VSANs** nodes:

- The **Storage Cloud** > **VSANs** node for a storage VSAN accessible to both fabric interconnects.
- The **Storage Cloud** > *Fabric_Name* > **VSANs** node for a VSAN accessible to only one fabric interconnect.

Deleting a Named VSAN

If Cisco UCS Manager includes a named VSAN with the same VSAN ID as the one you delete, the VSAN is not removed from the fabric interconnect configuration until all named VSANs with that ID are deleted.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, click the **SAN** node.
- Step 3** In the **Work** pane, click the **VSANs** tab.
- Step 4** Click one of the following subtabs, depending upon what type of VSAN you want to delete:

Subtab	Description
All	Displays all VSANs in the Cisco UCS domain.
Dual Mode	Displays the VSANs that are accessible to both fabric interconnects.
Switch A	Displays the VSANs that are accessible to only fabric interconnect A.
Switch B	Displays the VSANs that are accessible to only fabric interconnect B.

- Step 5** In the table, click the VSAN you want to delete.
You can use the **Shift** key or **Ctrl** key to select multiple entries.
- Step 6** Right-click the highlighted VSAN or VSANs and choose **Delete**.
- Step 7** If a confirmation dialog box displays, click **Yes**.

Changing the VLAN ID for the FCoE VLAN for a Storage VSAN



Caution Changing the VLAN ID of the FCoE VLAN for a storage VSAN causes a brief traffic outage. FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN results in a critical fault and traffic disrupt for all NICs and uplink ports using that FCoE VLAN. Ethernet traffic drops on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Changing the FCoE VLAN for the default VSAN or any configured VSAN under a global policy may result in storage disconnect or complete shut down.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, expand **SAN > Storage Cloud > VSANs**.

- Step 3** Choose the VSAN for which you want to modify the FCoE VLAN ID.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **FCoE VLAN** field, enter the desired VLAN ID.
 - Step 6** Click **Save Changes**.
-

Enabling Fibre Channel Uplink Trunking



Note If the fabric interconnects are configured for Fibre Channel end-host mode, enabling Fibre Channel uplink trunking renders all VSANs with an ID in the range from 3840 to 4079 non-operational.



Note Before enabling VSAN trunking on a Fabric Interconnect, ensure that all host OS storage path redundancies are functioning. For more information on steps to monitor and ensure the Fibre Channel paths are recovered, see the [Verification that the Data Path is Ready](#) section. This should be followed to avoid an all paths down to the Fibre Channel Uplinks.

After confirmation, enable Fibre Channel Uplink Trunking on secondary Fabric Interconnect and wait until the secondary Fibre Channel VIF paths recover. Then move to enabling the primary Fabric Interconnect Fibre Channel Trunking after validating data paths.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > SAN Cloud**.
 - Step 3** Click the node for the fabric where you want to enable FC uplink trunking.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Actions** area, click **Enable FC Uplink Trunking**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Disabling Fibre Channel Uplink Trunking

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > SAN Cloud**.
- Step 3** Click the node for the fabric where you want to disable Fibre Channel uplink trunking.

- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Disable FC Uplink Trunking**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
-



CHAPTER 6

SAN Pin Groups

- [SAN Pin Groups, on page 47](#)

SAN Pin Groups

Cisco UCS uses SAN pin groups to pin Fibre Channel traffic from a vHBA on a server to an uplink Fibre Channel port on the fabric interconnect. You can use this pinning to manage the distribution of traffic from the servers.



Note In Fibre Channel switch mode, SAN pin groups are irrelevant. Any existing SAN pin groups will be ignored.

To configure pinning for a server, you must include the SAN pin group in a vHBA policy. The vHBA policy is then included in the service profile assigned to that server. All traffic from the vHBA will travel through the I/O module to the specified uplink Fibre Channel port.

You can assign the same pin group to multiple vHBA policies. As a result, you do not need to manually pin the traffic for each vHBA.



Important Changing the target interface for an existing SAN pin group disrupts traffic for all vHBAs which use that pin group. The fabric interconnect performs a log in and log out for the Fibre Channel protocols to re-pin the traffic.

Creating a SAN Pin Group

In a system with two fabric interconnects, you can associate the pin group with only one fabric interconnect or with both fabric interconnects.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > SAN Cloud**.

- Step 3** Right-click **SAN Pin Groups** and select **Create SAN Pin Group**.
- Step 4** Enter a unique name and description for the pin group.
- Step 5** To pin traffic for fabric interconnect A, do the following in the **Targets** area:
- Check the **Fabric A** check box.
 - Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the uplink Fibre Channel port you want to associate with the pin group.
- Step 6** To pin traffic for fabric interconnect B, do the following in the **Targets** area:
- Check the **Fabric B** check box.
 - Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the uplink Fibre Channel port you want to associate with the pin group.
- Step 7** Click **OK**.
-

What to do next

Include the pin group in a vHBA template.

Deleting a SAN Pin Group

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > SAN Cloud > SAN Pin Groups**.
- Step 3** Right-click the SAN pin group you want to delete and select **Delete**.
- Step 4** If a confirmation dialog box displays, click **Yes**.
-



CHAPTER 7

FC Identity Assignment

- [Fibre Channel Identity, on page 49](#)

Fibre Channel Identity

A Fibre Channel node and port must have a globally unique World Wide Number (WWN). In Cisco UCS WWNs are created as identity pools. A Fibre Channel node (a whole server, storage array) must have a World Wide Node Name (WWNN) and a Fibre Channel port must have a World Wide Port Name (WWPN). Both WWNNs and WWPNs are physical entities; hence, they have a 64-bit address.

The WWNN pool is created as one large pool for the Cisco UCS domain. You can use the default pool in the Cisco UCS Manager SAN tab. However, it is recommended to create a custom WWNN pool for that UCS domain.

A communicating device is a node. A host bus adapter in a server constitutes a Fibre Channel node. For servers and hosts, WWNN is unique for each host bus adapter (HBA). For a SAN switch, the WWNN is common for the chassis. For midrange storage, the WWNN is common for each controller unit. For enterprise storage, the WWNN is unique for the entire array.

Each server has a unique WWPN for each port of the HBA. For a SAN switch, the WWPN is available for each port in the chassis. For storage, each port has an individual.

The FC Identity Tab in Cisco UCS Manager displays the FC Identity of the devices in the Cisco UCS domain SAN cloud, including the:

- Selected device WWNN or WWPN identifier.
- Whether the identifier is assigned to a vHBA.
- vHBA to which the identifier is assigned.



CHAPTER 8

WWN Pools

- [WWN Pools, on page 51](#)
- [WWPN Pools, on page 56](#)
- [WWxN Pools, on page 61](#)

WWN Pools

A World Wide Name (WWN) pool is a collection of WWNs for use by the Fibre Channel vHBAs in a Cisco UCS domain. You create separate pools for the following:

- WW node names assigned to the vHBA
- WW port names assigned to the vHBA
- Both WW node names and WW port names



Important

A WWN pool can include only WWNNs or WWPNS in the ranges from 20:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPNS is converted to a MAC address. You cannot use WWPNS pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

If you use WWN pools in service profiles, you do not have to manually configure the WWNs that will be used by the server associated with the service profile. In a system that implements multi-tenancy, you can use a WWN pool to control the WWNs used by each organization.

You assign WWNs to pools in blocks.

WWNN Pools

A WWNN pool is a WWN pool that contains only WW node names. If you include a pool of WWNNs in a service profile, the associated server is assigned a WWNN from that pool.

WWPN Pools

A WWPN pool is a WWN pool that contains only WW port names. If you include a pool of WWPNs in a service profile, the port on each vHBA of the associated server is assigned a WWPN from that pool.

WWxN Pools

A WWxN pool is a WWN pool that contains both WW node names and WW port names. You can specify how many ports per node are created with WWxN pools. The pool size must be a multiple of *ports-per-node* + 1. For example, if you specify 7 ports per node, the pool size must be a multiple of 8. If you specify 63 ports per node, the pool size must be a multiple of 64.

You can use a WWxN pool whenever you select a WWNN or WWPN pool. The WWxN pool must be created before it can be assigned.

- For WWNN pools, the WWxN pool is displayed as an option in the **WWNN Assignment** drop-down list.
- For WWPN pools, choose **Derived** in the **WWPN Assignment** drop-down list.

Creating a WWNN Pool



Important

A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPN is converted to a MAC address. You cannot use WWPN pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools**.
- Step 3** Expand the node for the organization where you want to create the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **WWNN Pools** and select **Create WWNN Pool**.
- Step 5** In the **Define Name and Description** dialog box of the **Create WWNN Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the World Wide Node Name pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Description field	A description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Assignment Order field	This can be one of the following: <ul style="list-style-type: none"> • Default—Cisco UCS Manager selects a random identity from the pool. • Sequential—Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add WWN Blocks** page of the **Create WWNN Pool** wizard, click **Add**.

Step 8 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of <i>ports-per-node</i> + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 9 Click **OK**.

Step 10 Click **Finish**.

What to do next

Include the WWNN pool in a service profile and template.

Adding a WWN Block to a WWNN Pool



Important

A WWN pool can include only WWNNs or WWPNS in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPNS is converted to a MAC address. You cannot use WWPNS pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name** .
- Step 3** Expand the **WWNN Pools** node.
- Step 4** Right-click the WWNN pool to which you want to add a WWN block and select **Create WWN Block**.
- Step 5** In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of <i>ports-per-node</i> + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

- Step 6** Click **OK**.
-

Deleting a WWN Block from a WWNN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name > WWNN Pools > WWNN_Pool_Name** .
- Step 3** Right-click the WWN block that you want to delete and select **Delete**.
- Step 4** If a confirmation dialog box displays, click **Yes**.
-

Adding a WWNN Initiator to a WWNN Pool



Important A WWN pool can include only WWNNs or WWPNS in the ranges from 20:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPNS is converted to a MAC address. You cannot use WWPNS pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name**.
- Step 3** Expand the **WWNN Pools** node.
- Step 4** Right-click the WWNN pool to which you want to add a WWNN initiator and select **Create WWNN Initiator**.
- Step 5** In the **Create WWNN Initiator** dialog box, complete the following fields:

Name	Description
World Wide Name field	The WWN.
Name field	The name of the WWNN initiator. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the WWNN initiator. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

- Step 6** Click **OK**.

Deleting a WWPNS Initiator from a WWPNS Pool

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name**.
- Step 3** Expand the **WWPN Pools** node.

- Step 4** Choose the WWPN pool from which you want to delete a WWPN initiator.
 - Step 5** In the **Work** pane, click the **Initiators** tab.
 - Step 6** Right-click the initiator that you want to delete and choose **Delete**.
 - Step 7** If a confirmation dialog box displays, click **Yes**.
-

Deleting a WWNN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name**.
 - Step 3** Expand the **WWNN Pools** node.
 - Step 4** Right-click the WWNN pool you want to delete and select **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

WWPN Pools

Creating a WWPN Pool



Important

A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPN is converted to a MAC address. You cannot use WWPN pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN > Pools**.

Step 3 Expand the node for the organization where you want to create the pool.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click **WWPN Pools** and select **Create WWPN Pool**.

Step 5 In the **Define Name and Description** page of the **Create WWPN Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the World Wide Port Name pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Assignment Order field	This can be one of the following: <ul style="list-style-type: none"> • Default—Cisco UCS Manager selects a random identity from the pool. • Sequential—Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add WWN Blocks** page of the **Create WWPN Pool** wizard, click **Add**.

Step 8 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of <i>ports-per-node</i> + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 9 Click **OK**.

Step 10 Click **Finish**.

What to do next

Include the WWPN pool in a vHBA template.

Adding a WWN Block to a WWPN Pool



Important A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPN is converted to a MAC address. You cannot use WWPN pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name**.
- Step 3** Expand the **WWPN Pools** node.
- Step 4** Right-click the WWPN pool to which you want to add a WWN block and select **Create WWN Block**.
- Step 5** In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of <i>ports-per-node</i> + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

- Step 6** Click **OK**.

Deleting a WWN Block from a WWPN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.

- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name > WWPN Pools > WWPN_Pool_Name** .
- Step 3** Right-click the WWN block that you want to delete and select **Delete**.
- Step 4** If a confirmation dialog box displays, click **Yes**.

Adding a WWPN Initiator to a WWPN Pool



Important A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPN is converted to a MAC address. You cannot use WWPN pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name** .
- Step 3** Expand the **WWPN Pools** node.
- Step 4** Right-click the WWPN pool to which you want to add a WWPN initiator and select **Create WWPN Initiator**.
- Step 5** In the **Create WWPN Initiator** dialog box, complete the following fields:

Name	Description
World Wide Name field	The WWN.
Name field	The name of the WWPN initiator. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the WWPN initiator. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

- Step 6** If you want to add a SAN boot target, expand the **Boot Target** area and complete the following fields:

Name	Description
Boot Target WWPN field	The WWPN that corresponds to the location of the boot image.
Boot Target LUN field	The LUN that corresponds to the location of the boot image.

Step 7 Click **OK**.

Deleting a WWPN Initiator from a WWPN Pool

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name** .
- Step 3** Expand the **WWPN Pools** node.
- Step 4** Choose the WWPN pool from which you want to delete a WWPN initiator.
- Step 5** In the **Work** pane, click the **Initiators** tab.
- Step 6** Right-click the initiator that you want to delete and choose **Delete**.
- Step 7** If a confirmation dialog box displays, click **Yes**.
-

Deleting a WWPN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name** .
- Step 3** Expand the **WWPN Pools** node.
- Step 4** Right-click the WWPN pool you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.
-

WWxN Pools

Creating a WWxN Pool



Important A WWN pool can include only WWNNs or WWPNS in the ranges from 20:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPNS is converted to a MAC address. You cannot use WWPNS pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools**.
- Step 3** Expand the node for the organization where you want to create the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **WWxN Pools** and select **Create WWxN Pool**.
- Step 5** In the **Define Name and Description** page of the **Create WWxN Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the World Wide Port Name pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Max Ports per Node field	The maximum number of ports that can be assigned to each node name in this pool. You cannot change this value once the object has been saved.

Name	Description
Assignment Order field	This can be one of the following: <ul style="list-style-type: none"> • Default—Cisco UCS Manager selects a random identity from the pool. • Sequential—Cisco UCS Manager selects the lowest available identity from the pool.

Step 6

Click **Next**.

Step 7

In the **Add WWN Blocks** page of the **Create WWxN Pool** wizard, click **Add**.

Step 8

In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of <i>ports-per-node</i> + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 9

Click **OK**.

Step 10

Click **Finish**.

What to do next

Include the WWxN pool in a service profile and template.

Adding a WWN Block to a WWxN Pool

**Important**

A WWN pool can include only WWNNs or WWPNS in the ranges from 20:00:00:00:00:00:00 to 20:FF:00:FF:FF:FF:FF or from 50:00:00:00:00:00:00 to 5F:FF:00:FF:FF:FF:FF. All other WWN ranges are reserved. When fibre channel traffic is sent through the Cisco UCS infrastructure, the source WWPNS is converted to a MAC address. You cannot use WWPNS pool which can translate to a source multicast MAC address. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNS in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure**Step 1**

In the **Navigation** pane, click **SAN**.

Step 2

In the **SAN** tab, expand **SAN > Pools > Organization_Name**.

Step 3

Expand the **WWxN Pools** node.

Step 4 Right-click the WWxN pool to which you want to add a WWN block and select **Create WWN Block**.

Step 5 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of <i>ports-per-node</i> + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 6 Click **OK**.

Deleting a WWN Block from a WWxN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name > WWxN Pools > WWxN_Pool_Name**.
- Step 3** Right-click the WWN block that you want to delete and select **Delete**.
- Step 4** If a confirmation dialog box displays, click **Yes**.

Deleting a WWxN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** In the **SAN** tab, expand **SAN > Pools > Organization_Name** .
 - Step 3** Expand the **WWxN Pools** node.
 - Step 4** Right-click the WWxN pool you want to delete and select **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-



CHAPTER 9

Storage-Related Policies

- [About vHBA Templates, on page 65](#)
- [Fibre Channel Adapter Policies, on page 68](#)
- [About the Default vHBA Behavior Policy, on page 77](#)
- [SPDM Security Policy, on page 78](#)
- [SAN Connectivity Policies, on page 81](#)

About vHBA Templates

vHBA Template

This template is a policy that defines how a vHBA on a server connects to the SAN. It is also referred to as a vHBA SAN connectivity template.

You must include this policy in a service profile for it to take effect.

Creating a vHBA Template

Before you begin

This policy requires that one or more of the following resources already exist in the system:

- Named VSAN
- WWNN pool or WWPN pool
- SAN pin group
- Statistics threshold policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click the **vHBA Templates** node and choose **Create vHBA Template**.

Step 5 In the **Create vHBA Template** dialog box, complete the following fields:

Name	Description
Name field	The name of the virtual HBA template. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the template. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Fabric ID field	The name of the fabric interconnect that vHBAs created with this template are associated with.
Select VSAN drop-down list	The VSAN to associate with vHBAs created from this template.
Create VSAN link	Click this link if you want to create a VSAN.
Template Type field	This can be one of the following: <ul style="list-style-type: none"> • Initial Template—vHBAs created from this template are not updated if the template changes. • Updating Template—vHBAs created from this template are updated if the template changes.
Max Data Field Size field	The maximum size of the Fibre Channel frame payload bytes that the vHBA supports. Enter an integer between 256 and 2112. The default is 2048.
WWPN Pool drop-down list	The WWPN pool that a vHBA created from this template uses to derive its WWPN address.
QoS Policy drop-down list	The QoS policy that is associated with vHBAs created from this template.
Pin Group drop-down list	The SAN pin group that is associated with vHBAs created from this template.
Stats Threshold Policy drop-down list	The statistics collection policy that is associated with vHBAs created from this template.

Step 6 Click **OK**.

What to do next

Include the vHBA template in a service profile.

Binding a vHBA to a vHBA Template

You can bind a vHBA associated with a service profile to a vHBA template. When you bind the vHBA to a vHBA template, Cisco UCS Manager configures the vHBA with the values defined in the vHBA template. If the existing vHBA configuration does not match the vHBA template, Cisco UCS Manager reconfigures the vHBA. You can only change the configuration of a bound vHBA through the associated vHBA template. You cannot bind a vHBA to a vHBA template if the service profile that includes the vHBA is already bound to a service profile template.



Important If the vHBA is reconfigured when you bind it to a template, Cisco UCS Manager reboots the server associated with the service profile.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers** > **Service Profiles**.
 - Step 3** Expand the node for the organization that includes the service profile with the vHBA you want to bind.
If the system does not include multi-tenancy, expand the **root** node.
 - Step 4** Expand *Service_Profile_Name* > **vHBAs**.
 - Step 5** Click the vHBA you want to bind to a template.
 - Step 6** In the **Work** pane, click the **General** tab.
 - Step 7** In the **Actions** area, click **Bind to a Template**.
 - Step 8** In the **Bind to a vHBA Template** dialog box, do the following:
 - a) From the **vHBA Template** drop-down list, choose the template to which you want to bind the vHBA.
 - b) Click **OK**.
 - Step 9** In the warning dialog box, click **Yes** to acknowledge that Cisco UCS Manager may need to reboot the server if the binding causes the vHBA to be reconfigured.
-

Unbinding a vHBA from a vHBA Template

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.
- Step 3** Expand the node for the organization that includes the service profile with the vHBA you want to unbind.

If the system does not include multi-tenancy, expand the **root** node.

- Step 4** Expand *Service_Profile_Name* > vHBAs.
 - Step 5** Click the vHBA you want to unbind from a template.
 - Step 6** In the **Work** pane, click the **General** tab.
 - Step 7** In the **Actions** area, click **Unbind from a Template**.
 - Step 8** If a confirmation dialog box displays, click **Yes**.
-

Deleting a vHBA Template

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > Policies > Organization_Name**.
 - Step 3** Expand the **vHBA Templates** node.
 - Step 4** Right-click the vHBA template that you want to delete and choose **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

Fibre Channel Adapter Policies

Ethernet and Fibre Channel Adapter Policies

These policies govern the host-side behavior of the adapter, including how the adapter handles traffic. For example, you can use these policies to change default settings for the following:

- Queues
- Interrupt handling
- Performance enhancement
- RSS hash
- Failover in a cluster configuration with two fabric interconnects



Note For Fibre Channel adapter policies, the values displayed by Cisco UCS Manager may not match those displayed by applications such as QLogic SANsurfer. For example, the following values may result in an apparent mismatch between SANsurfer and Cisco UCS Manager:

- **Max LUNs Per Target**—SANsurfer has a maximum of 256 LUNs and does not display more than that number. Cisco UCS Manager supports a higher maximum number of LUNs. This parameter is applicable only for FC-Initiator.
- **Link Down Timeout**—In SANsurfer, you configure the timeout threshold for link down in seconds. In Cisco UCS Manager, you configure this value in milliseconds. Therefore, a value of 5500 ms in Cisco UCS Manager displays as 5s in SANsurfer.
- **Max Data Field Size**—SANsurfer has allowed values of 512, 1024, and 2048. Cisco UCS Manager allows you to set values of any size. Therefore, a value of 900 in Cisco UCS Manager displays as 512 in SANsurfer.
- **LUN Queue Depth**—The LUN queue depth setting is available for Windows system FC adapter policies. Queue depth is the number of commands that the HBA can send and receive in a single transmission per LUN. Windows Storport driver sets this to a default value of 20 for physical miniports and to 250 for virtual miniports. This setting adjusts the initial queue depth for all LUNs on the adapter. Valid range for this value is 1 - 254. The default LUN queue depth is 20. This feature only works with Cisco UCS Manager version 3.1(2) and higher. This parameter is applicable only for FC-Initiator.
- **IO TimeOut Retry**—When the target device does not respond to an IO request within the specified timeout, the FC adapter cancels the pending command then resends the same IO after the timer expires. The FC adapter valid range for this value is 1 - 59 seconds. The default IO retry timeout is 5 seconds. This feature only works with Cisco UCS Manager version 3.1(2) and higher.

Operating System Specific Adapter Policies

By default, Cisco UCS provides a set of Ethernet adapter policies and Fibre Channel adapter policies. These policies include the recommended settings for each supported server operating system. Operating systems are sensitive to the settings in these policies. Storage vendors typically require non-default adapter settings. You can find the details of these required settings on the support list provided by those vendors.



Important We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for an OS (instead of using the default adapter policy), you must use the following formulas to calculate values that work for that OS.

Depending on the UCS firmware, your driver interrupt calculations may be different. Newer UCS firmware uses a calculation that differs from previous versions. Later driver release versions on Linux operating systems now use a different formula to calculate the Interrupt Count. In this formula, the Interrupt Count is the maximum of either the Transmit Queue or the Receive Queue plus 2.

Interrupt Count in Linux Adapter Policies

Drivers on Linux operating systems use differing formulas to calculate the Interrupt Count, depending on the eNIC driver version. The UCS 3.2 release increased the number of Tx and Rx queues for the eNIC driver from 8 to 256 each.

Use one of the following strategies, according to your driver version.

For Linux drivers before the UCS 3.2 firmware release, use the following formula to calculate the Interrupt Count.

$$\begin{aligned} \text{Completion Queues} &= \text{Transmit Queues} + \text{Receive Queues} \\ \text{Interrupt Count} &= (\text{Completion Queues} + 2) \text{ rounded up to nearest power of } 2 \end{aligned}$$

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

$$\begin{aligned} \text{Completion Queues} &= 1 + 8 = 9 \\ \text{Interrupt Count} &= (9 + 2) \text{ rounded up to the nearest power of } 2 = 16 \end{aligned}$$

On drivers for UCS firmware release 3.2 and higher, the Linux eNIC drivers use the following formula to calculate the Interrupt Count.

$$\text{Interrupt Count} = (\# \text{Tx or Rx Queues}) + 2$$

For example:

$$\begin{aligned} \text{Interrupt Count } wq = 32, rq = 32, cq = 64 &- \text{ then Interrupt Count} = \text{Max}(32, 32) + 2 = 34 \\ \text{Interrupt Count } wq = 64, rq = 8, cq = 72 &- \text{ then Interrupt Count} = \text{Max}(64, 8) + 2 = 66 \\ \text{Interrupt Count } wq = 1, rq = 16, cq = 17 &- \text{ then Interrupt count} = \text{Max}(1, 16) + 2 = 18 \end{aligned}$$

Interrupt Count in Windows Adapter Policies

For Windows OS, the recommended adapter policy in UCS Manager for VIC 1400 series and above adapters is Win-HPN and if RDMA is used, the recommended policy is Win-HPN-SMB. For VIC 1400 series and above adapters, the recommended interrupt value setting is 512 and the Windows VIC driver takes care of allocating the required number of Interrupts.

For VIC 1300 and VIC 1200 series adapters, the recommended UCS Manager adapter policy is Windows and the Interrupt would be TX + RX + 2, rounded to closest power of 2. The maximum supported Windows queues is 8 for Rx Queues and 1 for Tx Queues.

Example for VIC 1200 and VIC 1300 series adapters:

$$\text{Tx} = 1, \text{Rx} = 4, \text{CQ} = 5, \text{Interrupt} = 8 (1 + 4 \text{ rounded to nearest power of } 2), \text{Enable RSS}$$

Example for VIC 1400 series and above adapters:

$$\text{Tx} = 1, \text{Rx} = 4, \text{CQ} = 5, \text{Interrupt} = 512, \text{Enable RSS}$$

NVMe over Fabrics using Fibre Channel

The NVM Express (NVMe) interface allows host software to communicate with a non-volatile memory subsystem. This interface is optimized for Enterprise non-volatile storage, which is typically attached as a register level interface to the PCI Express (PCIe) interface.

NVMe over Fabrics using Fibre Channel (FC-NVMe) defines a mapping protocol for applying the NVMe interface to Fibre Channel. This protocol defines how Fibre Channel services and specified Information Units (IUs) are used to perform the services defined by NVMe over a Fibre Channel fabric. NVMe initiators can access and transfer information to NVMe targets over Fibre Channel.

FC-NVMe combines the advantages of Fibre Channel and NVMe. You get the improved performance of NVMe along with the flexibility and the scalability of the shared storage architecture. Cisco UCS Manager Release 4.0(2) supports NVMe over Fabrics using Fibre Channel on UCS VIC 1400 Series adapters.

Starting with UCS Manager release 4.3(2b), NVMeoF using RDMA is supported on Cisco UCS VIC 14000 series adapters.

Cisco UCS Manager provides the recommended FC NVMe Initiator adapter policies in the list of pre-configured adapter policies. To create a new FC-NVMe adapter policy, follow the steps in the *Creating a Fibre Channel Adapter Policy* section.

NVMe over Fabrics Using RDMA

NVMe over Fabrics (NVMeoF) is a communication protocol that allows one computer to access NVMe namespaces available on another computer. NVMeoF is similar to NVMe, but differs in the network-related steps involved in using the NVMeoF storage devices. The commands for discovering, connecting, and disconnecting a NVMeoF storage device are integrated into the **nvme** utility provided in Linux..

The NVMeoF fabric that Cisco supports is RDMA over Converged Ethernet version 2 (RoCEv2). RoCEv2 is a fabric protocol that runs over UDP. It requires a no-drop policy.

The eNIC RDMA driver works in conjunction with the eNIC driver, which must be loaded first when configuring NVMeoF.

Cisco UCS Manager provides the default Linux-NVMe-RoCE adapter policy for creating NVMe RoCEv2 interfaces. Do not use the default Linux adapter policy. For complete information on configuring RoCEv2 over NVMeoF, refer to the *Cisco UCS Manager Configuration Guide for RDMA over Converged Ethernet (RoCE) v2*.

NVMeoF using RDMA is supported on M5 B-Series or C-Series Servers with Cisco UCS VIC 1400 Series adapters.

Starting with UCS Manager release 4.3(2b), NVMeoF using RDMA is supported on Cisco UCS VIC 14000 series adapters.

Creating a Fibre Channel Adapter Policy



Tip If the fields in an area do not display, click the **Expand** icon to the right of the heading.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **Adapter Policies** and choose **Create Fibre Channel Adapter Policy**.
- Step 5** Enter a name and description for the policy in the following fields:

Table 4:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 6 (Optional) In the **Resources** area, adjust the following values:

Name	Description
Transmit Queues field	The number of transmit queue resources to allocate. This value cannot be changed.
Ring Size field	The number of descriptors in each transmit queue. This parameter applies to Extended Link Services (ELS) and Common Transport (CT) fibre channel frames for generic services. It does not affect adapter performance. Enter an integer between 64 and 128. The default is 64.
Receive Queues field	The number of receive queue resources to allocate. This value cannot be changed.
Ring Size field	The number of descriptors in each receive queue. This parameter applies to Extended Link Services (ELS) and Common Transport (CT) fibre channel frames for generic services. It does not affect adapter performance. Enter an integer between 64 and 2048. The default is 64.
I/O Queues field	The number of IO queue resources the system should allocate. Enter an integer between 1 and 16. The default is 16.
Ring Size field	The number of descriptors in each I/O queue. Enter an integer between 64 and 512. The default is 512. Note The number of descriptors can affect the performance of the adapter, so we recommend that you do not change the default value.

Step 7 (Optional) In the **Options** area, adjust the following values:

Name	Description
FCP Error Recovery field	<p>Whether the system uses FCP Sequence Level Error Recovery (FC-TAPE) protocol for sequence level error recovery with tape devices. This enables or disables the Read Exchange Concise (REC) and Sequence Retransmission Request (SRR) functions on the VIC firmware. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—This is the default. • Enabled—You should select this option if your system is connected to one or more tape drive libraries. <p>Note This parameter only applies to a server with a Virtual Interface Card (VIC) adapter.</p>
Flogi Retries field	<p>The number of times that the system tries to log in to the fabric after the first failure.</p> <p>Enter any integer. To specify that the system continue to try indefinitely, enter infinite in this field. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter, or a converged network adapter.</p>
Flogi Timeout (ms) field	<p>The number of milliseconds that the system waits before it tries to log in again.</p> <p>Enter an integer between 1000 and 255000. The default is 4,000. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter or a converged network adapter.</p> <p>When a Flogi timeout value of 20 seconds or more is configured for a boot vHBA, it could lead to a SAN boot failure if the adapter does not receive an accept to the initial Flogi. For a boot-enabled vHBA, the recommended timeout values is 5 seconds or less.</p>
Plogi Retries field	<p>The number of times that the system tries to log into a port after the first failure.</p> <p>Enter an integer between 0 and 255. The default is 8. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter.</p>

Name	Description
Plogi Timeout (ms) field	<p>The number of milliseconds that the system waits before it tries to log in again.</p> <p>Enter an integer between 1000 and 255000. The default is 20,000. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>For an HBA that is going to be used to boot a Windows OS from SAN, the recommended value for this field is 4,000 ms.</p> <p>Note This parameter only applies to a server with a VIC adapter.</p> <p>When a Plogi timeout value of 20 seconds or more is configured for a boot vHBA, it could lead to a SAN boot failure if the adapter does not receive an accept to the initial Plogi. For a boot-enabled vHBA, the recommended timeout values is 5 seconds or less.</p>
Port Down Timeout (ms) field	<p>The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. This parameter is important for host multi-pathing drivers and it is one of the key indicators used for error processing.</p> <p>Enter an integer between 0 and 240000. The default is 30,000. For a server with a VIC adapter running ESX, the recommended value is 10,000.</p> <p>For a server with a port that is going to be used to boot a Windows OS from SAN, the recommended value for this field is 5000 milliseconds.</p> <p>We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter.</p>
IO Retry Timeout (seconds)	<p>The number of seconds that the FC adapter waits before aborting the pending command and resending the same IO. This happens when the network device does not responding to an IO request within the specified time.</p> <p>Enter an integer between 0 and 59 seconds. The default IO retry timeout is 5 seconds.</p>
Port Down IO Retry field	<p>The number of times an IO request to a port is returned because the port is busy before the system decides the port is unavailable.</p> <p>Enter an integer between 0 and 255. The default is 8. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter running Windows.</p>

Name	Description
Link Down Timeout (ms) field	<p>The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost.</p> <p>Enter an integer between 0 and 240000. The default is 30,000. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter running Windows.</p>
IO Throttle Count field	<p>The maximum number of data or control I/O operations that can be pending in the vHBA at one time. If this value is exceeded, the additional I/O operations wait in the queue until the number of pending I/O operations decreases and the additional operations can be processed.</p> <p>Note This parameter is not the same as the LUN queue depth, which is controlled by Cisco UCS Manager based on the operating system installed on the server.</p> <p>Enter an integer between 256 and 1024. The default is 256. We recommend you consult your storage array documentation for the optimal value for this parameter.</p>
Max LUNs Per Target field	<p>The maximum number of LUNs that the Fibre Channel driver will export or show. The maximum number of LUNs is usually controlled by the operating system running on the server.</p> <p>Enter an integer between 1 and 1024. The default value is 256. For servers running ESX or Linux, the recommended value is 1024.</p> <p>We recommend you consult your operating system documentation for the optimal value for this parameter.</p> <p>Note</p> <ul style="list-style-type: none"> • This parameter only applies to a server with a VIC adapter or a network adapter. • This parameter is applicable only for FC-Initiator.
LUN Queue Depth field	<p>The number of commands that the HBA can send and receive in a single transmission per LUN.</p> <p>Enter an integer between 1 and 254. The default LUN queue depth is 20.</p> <p>Note This parameter is applicable only for FC-Initiator.</p>

Name	Description
Interrupt Mode radio button	<p>The method used to send interrupts to the operating system from the driver. This can be one of the following:</p> <ul style="list-style-type: none"> • MSI-X—Message Signaled Interrupts (MSI) with the optional extension. We recommend that you select this option if the operating system on the server supports it. • MSI—MSI only. • INTx—PCI INTx interrupts. <p>Note This parameter only applies to a server with a VIC adapter or a network adapter running an operating system other than Windows. The Windows operating system ignores this parameter.</p>
vHBA Type radio button	<p>The vHBA type used in this policy. vHBAs supporting FC and FC-NVMe can now be created on the same adapter. The vHBA type used in this policy can be one of the following:</p> <ul style="list-style-type: none"> • FC Initiator—Legacy SCSI FC vHBA initiator • FC Target—vHBA that supports SCSI FC target functionality <p>Note This option is available as a Tech Preview.</p> <ul style="list-style-type: none"> • FC NVME Initiator—vHBA that is an FC NVME initiator, which discovers FC NVME targets and connects to them • FC NVME Target—vHBA that acts as an FC NVME target and provides connectivity to the NVME storage <p>Note This option is available as a Tech Preview.</p> <p>vHBA type is supported only on UCS VIC 1400 adapters .</p>

Step 8 Click **OK**.

Step 9 If a confirmation dialog box displays, click **Yes**.

Deleting a Fibre Channel Adapter Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN > Policies > Organization_Name**.

Step 3 Expand the **Fibre Channel Policies** node.

Step 4 Right-click the policy you want to delete and choose **Delete**.

- Step 5** If a confirmation dialog box displays, click **Yes**.
-

About the Default vHBA Behavior Policy

Default vHBA Behavior Policy

Default vHBA behavior policy allow you to configure how vHBAs are created for a service profile. You can choose to create vHBAs manually, or you can allow them to be created automatically.

You can configure the default vHBA behavior policy to define how vHBAs are created. This can be one of the following:

- **None**—Cisco UCS Manager does not create default vHBAs for a service profile. All vHBAs must be explicitly created.
- **HW Inherit**—If a service profile requires vHBAs and none have been explicitly defined, Cisco UCS Manager creates the required vHBAs based on the adapter installed in the server associated with the service profile.



Note If you do not specify a default behavior policy for vHBAs, **none** is used by default.

Configuring a Default vHBA Behavior Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies**.
- Step 3** Expand the **root** node.

You can configure only the default vHBA behavior policy in the root organization. You cannot configure the default vHBA behavior policy in a sub-organization.

- Step 4** Click **Default vHBA Behavior**.
- Step 5** On the **General Tab**, in the **Properties** area, click one of the following radio buttons in the **Action** field:
- **None**—Cisco UCS Manager does not create default vHBAs for a service profile. All vHBAs must be explicitly created.
 - **HW Inherit**—If a service profile requires vHBAs and none have been explicitly defined, Cisco UCS Manager creates the required vHBAs based on the adapter installed in the server associated with the service profile.

Step 6 Click **Save Changes**.

SPDM Security Policy

SPDM Security

Cisco UCS M6 Servers can contain mutable components that could provide vectors for attack against a device itself or use of a device to attack another device within the system. To defend against these attacks, the Security Protocol and Data Model (SPDM) Specification enables a secure transport implementation that challenges a device to prove its identity and the correctness of its mutable component configuration. This feature is supported on Cisco UCS C220 and C240 M6 Servers starting with in Cisco UCS Manager, Release 4.2(1d).



Note SPDM is currently not supported on the Cisco UCS C245 M6 Server.

SPDM defines messages, data objects, and sequences for performing message exchanges between devices over a variety of transport and physical media. It orchestrates message exchanges between Baseboard Management Controllers (BMC) and end-point devices over a Management Component Transport Protocol (MCTP). Message exchanges include authentication of hardware identities accessing the BMC. The SPDM enables access to low-level security capabilities and operations by specifying a managed level for device authentication, firmware measurement, and certificate management. Endpoint devices are challenged to provide authentication, and BMC authenticates the endpoints and only allows access for trusted entities.

The UCS Manager optionally allows uploads of external security certificates to BMC. A maximum of 40 SPDM certificates is allowed, including native internal certificates. Once the limit is reached, no more certificates can be uploaded. User uploaded certificates can be deleted but internal/default certificates cannot.

A SPDM security policy allows you to specify one of three Security level settings. Security can be set at one of the three levels listed below:

- Full Security:

This is the highest MCTP security setting. When you select this setting, a fault is generated when any endpoint authentication failure or firmware measurement failure is detected. A fault will also be generated if any of the endpoints do not support either endpoint authentication or firmware measurements.

- Partial Security (default):

When you select this setting, a fault is generated when any endpoint authentication failure or firmware measurement failure is detected. There will NOT be a fault generated when the endpoint doesn't support endpoint authentication or firmware measurements.

- No Security

When you select this setting, there will NOT be a fault generated for any failure (either endpoint measurement or firmware measurement failures).

You can also upload the content of one or more external/device certificates into BMC. Using a SPDM policy allows you to change or delete security certificates or settings as desired. Certificates can be deleted or replaced when no longer needed.

Certificates are listed in all user interfaces on a system.

Creating a SPDM Security Policy

This step creates a SPDM policy.



Note You can upload up to 40 SPDM certificates (including native certificates).

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Go to **Policies**. Expand the root node.
- Step 3** Right-click **SPDM Certificate Policies** and select **Create SPDM Policy**.
- Step 4** Enter a name for this policy and select a **Fault Alert Setting** for the security level: **Disabled**, **Partial**, or **Full**.
- Full—If you select this option, then a fault is generated when there is any endpoint authentication failure for both supported and unsupported endpoints.
- Partial—If you select this option then a fault is generated when there is any endpoint authentication failure to only supported endpoints. No fault is generated when the endpoint does not support authentication.
- Disabled—If you select this option then no fault is generated for endpoint authentication failure for both supported and unsupported endpoints.
- The default is **Partial**.
- Note** To perform SPDM re-authentication and update the faults, Cisco IMC reboot or host reboot is required when the fault alert value is changed for an associated profile.
- Step 5** Click on **Add** in the **Create Policy** window. The **Add SPDM Certificate** window will open.
- Step 6** Name the certificate.
- UCS Manager supports only **Pem**certificates.
- Step 7** Paste the contents of the certificate into the Certificate field.
- Step 8** Click **OK** to add the certificate and return to the **Create SPDM Policy** window.
- You can add up to 40 certificates.
- Step 9** In the **Create SPDM Policy** menu, click **Okay**.
- After the SPDM policy is created, it will be listed immediately, along with its Alert setting, when you select **SPDM Certificate Policy** under the Server root Policies.
-

What to do next

Assign the Certificate to a Service Profile. The Service Profile must be associated with a server for it to take effect.

Associating the Security Policy with a Server

Before you begin

Create the SPDM security policy.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Go to **Service Profiles**. Expand the root node.
- Step 3** Select the Service Profile you want to associate with the Policy you created.
- a) On the **Policies** tab, scroll down and expand **SPDM Certificate Policy**. In the **SPDM Certificate Policy** dropdown, select the desired policy to associate with this Service Profile.
- Step 4** Click **OK**.
- The SPDM Policy will now be associated with the service profile.
-

What to do next

Check the fault alert level to make sure it is set to the desired setting.

Viewing the Fault Alert Settings

You can view the Fault Alert setting associated with a specific chassis.

Before you begin

Create a policy and associate it with a Service Profile.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Select a Rack-Mount Server.
- Step 3** On the **Inventory** tab, select **CIMC**.
- User uploaded certificates are listed and information for specific certificates can be selected and viewed.
-

SAN Connectivity Policies

About the LAN and SAN Connectivity Policies

Connectivity policies determine the connections and the network communication resources between the server and the LAN or SAN on the network. These policies use pools to assign MAC addresses, WWNs, and WWPNS to servers and to identify the vNICs and vHBAs that the servers use to communicate with the network.



Note We do not recommend that you use static IDs in connectivity policies, because these policies are included in service profiles and service profile templates and can be used to configure multiple servers.

Privileges Required for LAN and SAN Connectivity Policies

Connectivity policies enable users without network or storage privileges to create and modify service profiles and service profile templates with network and storage connections. However, users must have the appropriate network and storage privileges to create connectivity policies.

Privileges Required to Create Connectivity Policies

Connectivity policies require the same privileges as other network and storage configurations. For example, you must have at least one of the following privileges to create connectivity policies:

- admin—Can create LAN and SAN connectivity policies
- ls-server—Can create LAN and SAN connectivity policies
- ls-network—Can create LAN connectivity policies
- ls-storage—Can create SAN connectivity policies

Privileges Required to Add Connectivity Policies to Service Profiles

After the connectivity policies have been created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create connectivity policies.

Interactions between Service Profiles and Connectivity Policies

You can configure the LAN and SAN connectivity for a service profile through either of the following methods:

- LAN and SAN connectivity policies that are referenced in the service profile
- Local vNICs and vHBAs that are created in the service profile
- Local vNICs and a SAN connectivity policy
- Local vHBAs and a LAN connectivity policy

Cisco UCS maintains mutual exclusivity between connectivity policies and local vNIC and vHBA configuration in the service profile. You cannot have a combination of connectivity policies and locally created vNICs or vHBAs. When you include a LAN connectivity policy in a service profile, all existing vNIC configuration is erased, and when you include a SAN connectivity policy, all existing vHBA configuration in that service profile is erased.

Creating a SAN Connectivity Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **SAN Connectivity Policies** and choose **Create SAN Connectivity Policy**.
- Step 5** In the **Create SAN Connectivity Policy** dialog box, enter a name and optional description.
- Step 6** From the **WWNN Assignment** drop-down list in the **World Wide Node Name** area, choose one of the following:
- Choose **Select (pool default used by default)** to use the default WWN pool.
 - Choose one of the options listed under **Manual Using OUI** and then enter the WWN in the **World Wide Node Name** field.

You can specify a WWNN in the range from 20:00:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF:FF. You can click the **here** link to verify that the WWNN you specified is available.
 - Choose a WWN pool name from the list to have a WWN assigned from the specified pool. Each pool name is followed by two numbers in parentheses that show the number of WWNs still available in the pool and the total number of WWNs in the pool.
- Step 7** In the **vHBAs** table, click **Add**.
- Step 8** In the **Create vHBAs** dialog box, enter the name and optional description.
- Step 9** Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.
- Step 10** In the **Operational Parameters** area, choose the **Stats Threshold Policy**.
- Step 11** In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.
- Step 12** After you have created all the vHBAs you need for the policy, click **OK**.
-

What to do next

Include the policy in a service profile or service profile template.

Creating a vHBA for a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, expand **SAN > Policies > Organization_Name > San Connectivity Policies**.
- Step 3** Choose the policy for which you want to create a vHBA.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the table icon bar, click the + button.
- Step 6** In the **Create vHBAs** dialog box, enter the name and optional description.
- Step 7** Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.
- Step 8** In the **Operational Parameters** area, choose the **Stats Threshold Policy**.
- Step 9** In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.
- Step 10** Click **Save Changes**.
-

Deleting a vHBA from a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies > Organization_Name**.
- Step 3** Choose the policy from which you want to delete the vHBA.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **vHBAs** table, do the following:
- Click the vHBA that you want to delete.
 - On the icon bar, click **Delete**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
-

Creating an Initiator Group for a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.

- Step 2** Expand **SAN > Policies > Organization_Name**.
- Step 3** Choose the policy for which you want to create an initiator group.
- Step 4** In the **Work** pane, click the **vHBA Initiator Groups** tab.
- Step 5** In the table icon bar, click the + button.
- Step 6** In the **Create vHBA Initiator Group** dialog box, complete the following fields:

Name	Description
Name field	<p>The name of the vHBA initiator group.</p> <p>This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</p>
Description field	<p>A description of the group.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p>
Select vHBA Initiators table	<p>Check the check box in the Select column for each vHBA that you want to use.</p>
Storage Connection Policy drop-down list	<p>The storage connection policy associated with this vHBA initiator group. If you want to:</p> <ul style="list-style-type: none"> Use an existing storage connection policy, then choose that policy from the drop-down list. The Cisco UCS Manager GUI displays information about the policy and its FC target endpoints in the Global Storage Connection Policy area. <p>Create a new storage connection policy that will be globally available, then click the Create Storage Connection Policy link.</p> <ul style="list-style-type: none"> Create a local storage connection policy that is available only to this vHBA initiator group, then choose the Specific Storage Connection Policy option. The Cisco UCS Manager GUI displays the Specific Storage Connection Policy area that allows you to configure the local storage connection policy.
Create Storage Connection Policy link	<p>Click this link to create a new storage connection policy that will be available to all service profiles and service profile templates.</p>

- Step 7** Click **OK**.

Deleting an Initiator Group from a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > Policies > Organization_Name**.
 - Step 3** Choose the policy from which you want to delete the initiator group
 - Step 4** In the **Work** pane, click the **vHBA Initiator Groups** tab.
 - Step 5** In the table, do the following:
 - a) Click the initiator group that you want to delete.
 - b) On the icon bar, click **Delete**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Deleting a SAN Connectivity Policy

If you delete a SAN connectivity policy that is included in a service profile, it also deletes all vHBAs from that service profile and disrupts SAN data traffic for the server associated with the service profile.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > Policies > Organization_Name**.
 - Step 3** Expand the **SAN Connectivity Policies** node.
 - Step 4** Right-click the policy that you want to delete and choose **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

Enabling Intel® Volume Management Device

Volume Management Device (VMD) Setup

The Intel® Volume Management Device (VMD) is a tool that provides NVMe drivers to manage PCIe Solid State Drives attached to VMD-enabled domains. This includes Surprise hot-plug of PCIe drives and configuring blinking patterns to report status. PCIe Solid State Drive (SSD) storage lacks a standardized method to blink LEDs to represent the status of the device. With VMD, you can control LED indicators on both direct attached and switch attached PCIe storage using a simple command-line tool.

To use VMD, you must first enable VMD through a UCS Manager BIOS policy and set the UEFI boot options. Enabling VMD provides Surprise hot plug and optional LED status management for PCIe SSD storage that is attached to the root port. VMD Passthrough mode provides the ability to manage drives on guest VMs.

Enabling VMD also allows configuration of Intel® Virtual RAID on CPU (VRoC), a hybrid RAID architecture on Intel® Xeon® Scalable Processors. Documentation on the use and configuration of VRoC can be found at the Intel website.

IMPORTANT: VMD must be enabled in the UCS Manager BIOS settings before Operating System install. If enabled after OS installation, the server will fail to boot. This restriction applies to both standard VMD and VMD Passthrough. Likewise, once enabled, you cannot disable VMD without a loss of system function.

Enabling VMD on UCS Manager

To configure a BIOS and local boot Policy for VMD in UCS Manager, use the following procedure. The VMD platform default is disabled.



Note VMD must be enabled before OS installation.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 3** Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
 - Step 4** Go to **Policies > Root > BIOS Policies** and select the new policy.
 - Step 5** Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.
 - Step 6** Scroll down to **VMD Enable** and select **Enable**.
 - Step 7** Click **Save Changes** to enable VMD functions.
 - Step 8** In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode** and **Add NVMe** from the **Local Devices** menu. Click **Save Changes** to create the policy.
-

Enabling Volume Management Device (VMD) in Passthrough Mode

Volume Management Device (VMD) Passthrough Mode

The Intel® Volume Management Device (VMD) driver release package for Direct Device Assignment contains the Intel VMD UEFI Driver version for Direct Assign (PCIe PassThru) in VMware ESXi Hypervisor. The Intel VMD NVMe driver assists in the management of CPU-attached Intel PCIe NVMe SSDs.

The Intel VMD driver is required to enable the Direct Assign and discovery of the VMD physical addresses from a supported guest VM. Drivers are only provided for Passthrough mode for ESXi support of Red Hat Linux or Ubuntu. VMD Passthrough is enabled by configuring a UCS Manager BIOS policy before loading the Operating System. Once the Operating System has been loaded, you cannot enable or disable the VMD Passthrough option.



Note Passthrough mode is enabled by default, but you should always confirm that it is enabled before proceeding.

Configuring VMD Passthrough

Passthrough mode is only supported on ESXi drivers for Red Hat Linux or Ubuntu guest operating systems.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 3** Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
 - Step 4** Go to **Policies > Root > BIOS Policies** and select the new policy.
 - Step 5** Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.
 - Step 6** Scroll down to **VMD Enable** and select **Enable**.
 - Step 7** Click **Save Changes** to enable VMD functions.
 - Step 8** To finish enabling VMD Passthrough mode, select **Advanced** and **Intel Directed IO** from the submenus and scroll down to **Intel VT Directed IO**. Verify that the dropdown is set to **Enabled**. If not, set it.
 - Step 9** Click **Save Changes** to enable the VMD Passthrough policy.
 - Step 10** In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode**. Click **OK** to create the policy.
-

Downloading VMD Drivers

Intel® Volume Management Device Drivers

Intel® Volume Management Device (VMD) for NVMe enables drive management options using hardware logic inside the Intel Xeon processor. Specific drivers are available for the following operating systems:

- Linux
- Windows 2016, 2019
- VMWare



Note The latest VMWare drivers are available directly from the VMWare site. Following links in the VMWare driver download on the Cisco download site will take you directly to the VMWare login page.

For guest Operating Systems on ESXi, use VMD Passthrough mode. Supported Operating Systems for VMD Passthrough are:

- Red Hat Linux
- Ubuntu

To use the features of Intel VMD, you must:

- Enable VMD by creating a BIOS policy in the UCS Manager.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

- Install the appropriate VMD NVMe driver.
- Install the appropriate management tools for the driver package.
- Boot from UEFI.

Intel® Virtual RAID on CPU (VRoC) with VMD

Intel® Virtual RAID on CPU (VRoC) allows you to create and manage RAID volumes within the BIOS of VMD-enabled Intel NVMe SSD drives using hardware logic inside the Intel Xeon processor. More information on Intel VRoC can be found at: <https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html>.

The User Guides for Intel VRoC can be accessed at the direct link at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en

The Windows and Linux user documentation also contains information on how to configure Intel VRoC in the pre-boot environment. Creation of RAID volumes in VRoC is through the HII interface. The Windows documentation provides information on using the BIOS HII option to set up and configure RAID volumes in VRoC.

To use Intel VRoC, you must:

- Enable VMD in the BIOS settings
- Use UEFI boot mode
- Have sufficient drive resources to create the volume
- Use the BIOS HII option to set up and configure VRoC.

The Cisco implementation of Intel VRoC supports RAID 0 (striping), RAID 1 (mirroring), RAID 5 (striping with parity) and RAID 10 (combined mirroring and striping).

Downloading the Linux VMD Drivers

Complete these steps to download and install the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

-
- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.
- Step 2** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 3** Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.
- Step 4** Click on the latest release in the left panel.
- Note** The ISO image for VMD on blade servers is available from the 4.0(4f) release onward.
- Step 5** Click on **ISO image of UCS-related linux drivers only** and download the driver bundle.
- Step 6** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD > RHEL_{x.x}**.
- Step 7** Click on the version of Red Hat Linux that you wish to install.
- Step 8** Extract the contents of the folder. The folder contains both the driver package and associated documentation. Follow the installation procedure packaged with the drivers.
-

What to do next

The Intel® Virtual RAID on CPU (VRoC) Linux Software User Guide can be found with the user documentation at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en. It provides information on performing BIOS HII VRoC setup in the pre-boot environment, as well as how to install and use the programmable LED utility.

Downloading the Windows VMD Drivers

Complete these steps to download the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.
- Step 2** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 3** Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.
- Step 4** Click on the latest release in the left panel.
The ISO image for VMD is available from the 4.0(4f) release onward.
- Step 5** Click on **ISO image of UCS-related windows drivers only** and download the driver bundle.
- Step 6** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD > KIT_x_x_x_xxxx**.
- Step 7** Extract the contents of the folder.
- Step 8** Click on the entry for the kit and **KIT > Install**.
- Step 9** The folder contains both the driver package and associated documentation. Expand the zip file for **VROC_x_x_x_xxxxInstall**.
- Step 10** Follow the installation procedure packaged with the drivers.
-

What to do next

For setting up Intel® Virtual RAID on CPU (VRoC), refer to the online instructions at <https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html>.

Information on VRoC RAID features and management can be found in the *Windows Intel Virtual RAID on CPU Software User's Guide* at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/ssd-software/Windows_VROC_User_Guide.pdf.

Downloading the VMD Passthrough Drivers

Complete these steps to download and install the driver bundle for VMD Passthrough mode:



Note The VMD Passthrough driver bundle includes packages for both ESXi and Ubuntu.

Before you begin



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.

- Step 2** Search on **Servers - Unified Computing**.
- Step 3** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 4** Choose the UCS utilities from the Software Type selections: **Unified Computing System (UCS) Utilities**.
- Step 5** Click on the latest release in the left panel.

Note The ISO image for VMD is available from UCSM 4.0(4f) release onward.

- Step 6** Click on **ISO image of UCS-related vmware utilities only** and download the utilities bundle.

- Step 7** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD**.

The bundle provides both the driver installation package for the desired version of ESXi or VMD Direct Assign with Ubuntu, passthrough mode, and the Signed LED Offline bundle. Also included is a pdf that provides steps to configure an Ubuntu Virtual Machine in ESXi.

- Step 8** Click on either the version of ESXi that you wish to install or the zip file for Ubuntu.

For ESXi versions, Click on **ESXi_x > Direct Assign** and chose the desired zip file.

- Step 9** Extract the contents of the folder. Follow the installation procedure packaged with the driver software.

What to do next

Extract the contents of the LED management tools zip file. Install the management tools according to the instructions included with the driver package.

Before using the command line tools, the ESXi command line shell should be enabled from either the vSphere client or from the direct console of the ESXi host system.

Custom LED Status with VMD on NVMe

Once you have set up VMD, you can customize LED blinking patterns on PCIe NVMe drives. Information on LED customization can be found in the User Guides included in the driver packages.

LED Blinking

PCIe SSD drives lack a standard way to manage the LEDs that indicate drive status and health. Without this, there is a risk of removing the wrong drive, resulting in data loss. SSD drives have two indicators, the first being a green activity LED whose signals come directly from the SSD, and the second being a status LED whose signals come from the backplane. VMD manages only the status LEDs, not the activity LEDs.

LED Management only applies to NVMe and/or SATA drives. It does not support drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard .

LED Activity During Drive Hot-plug

VMD with NVMe supports Surprise hot-plugging. When a disk is hot-removed, then re-inserted into the same slot, the fault LED blinks for 10 seconds. This is expected behavior. The fail state is imposed on a slot's LEDs when the drive is removed, but the backplanes require the drive to be present in the slot for a LED to blink. Thus, the fail state exists once the drive is removed, but a LED blinks only when the new drive is inserted and discovered. The LED will return to normal once hot-plug event is handled.

Custom Blinking Patterns

VRoC with VMD allows you to perform basic LED management configuration of the status LEDs on compatible backplanes. Once the VMD NVMe driver is installed, you can install the VMD LED Management Tool, which lets you manage the LED through a command line interface. VMD allows you to customize LED blinking patterns on PCIe NVMe drives to better identify failing drives.

The tables below provide some brief guidelines for customized blinking on the various platforms. As individualized patterns are programmable, these tables provide only representative guidelines.

Table 5: LED Blinking Patterns: Windows

Status LED	Behavior	Options
"Activate LED"	Identifies a specific device in an enclosure by blinking the status LED of that drive in a designated pattern.	1-3600 seconds. Values outside this range default to 12 seconds. Default = 12 seconds
Drive Failure	Indicates a drive that is in a degraded or failed state by lighting the status LED of that device in a defined failure pattern.	The failure pattern is displayed until: <ul style="list-style-type: none"> • 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. • 2. From the time when a non-failed drive that is part of a RAID volume is removed, or the failed drive is identified and removed. It remains in failure state until a new drive is inserted into the same slot or the platform is rebooted. Default = Option 1
RAID volume Initialization or Verify and Repair Process	When a RAID volume is in Rebuild state, the status LEDs blink in the defined Rebuild pattern on either the specific drive being rebuilt or on the entire RAID volume that is being rebuilt.	Default = Enabled Can be: <ol style="list-style-type: none"> 1. Disabled (only on one drive) 2. Enabled (on all drives)
Managed unplug	During a managed hot unplug, the status LED of the managed drive blinks in the defined Locate pattern until the drive is physically ejected.	None. Enabled by default.

Status LED	Behavior	Options
RAID volume is migrating	During RAID volume migration, the status LEDs blink in the defined Rebuild pattern on all drives until the process is complete.	Default = Enabled Can be: 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs)
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 6: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller BLACKLIST	<code>ledmon</code> will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	Exclude controllers on the blacklist. Default = Support all controllers
RAID volume is initializing, verifying, or verifying and fixing BLINK_ON_INIT	Rebuild pattern on all drives in RAID volume (until initialization, verify, or verify and fix finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> scan interval INTERVAL	Defines the time interval between <code>ledmon sysfs scans</code> . The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is rebuilding REBUILD_BLINK_ON_ALL	Rebuild pattern on a single drive to which RAID volume rebuilds	1. False/Disabled (on one drive) 2. True/Enabled (on all drives) Default = False/Disabled
RAID volume is migrating BLINK_ON_MIGR	Rebuild pattern on all drives in RAID volume (until migration finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> debug level LOG_LEVEL	Corresponds with <code>-log-level</code> flag from <code>ledmon</code> .	Acceptable values are: quiet, error, warning, info, debug, all - 0 means 'quiet' and 5 means 'all' Default = 2
Set manage one RAID member or All RAID RAID_MEMBERS_ONLY	If the flag is set to <code>ledmon true</code> , will limit monitoring only to drives that are RAID members.	1. False / (all RAID member and PT) 2. True / (RAID member only) Default = False

Status LED	Behavior	Options
Limited scans only to specific controllers WHITELIST	<code>ledmon</code> limits changing the LED state to controllers listed on whitelist.	Limit changing LED state in whitelist controller. Default = No limit.

Table 7: LED Blinking Patterns: ESXi

Status LED	Behavior	Options
"Identify"	The ability to identify a specific device in an enclosure by blinking the status LED of that drive in the defined Locate pattern.	None. Default is Off.
"Off"	The ability to turn off the "Identify" LED once a specific device in an enclosure has been located.	None. Default is Off.



CHAPTER 10

SED Security Policies

- [Security Policies for Self-Encrypting Drives, on page 95](#)
- [Security Flags of the Controller and Disk, on page 96](#)
- [Managing Local Security Policies, on page 96](#)
- [KMIP Client Certificate Policy, on page 98](#)
- [Managing Remote Security Policies, on page 100](#)
- [Enabling and Disabling Security on Disks, on page 102](#)
- [Disabling Security on a Controller, on page 102](#)
- [Unlocking a Locked Disk, on page 103](#)
- [Erasing a Secure Foreign Configuration Disk, on page 103](#)
- [Secure Data Deletion, on page 104](#)

Security Policies for Self-Encrypting Drives

Self-Encrypting Drives (SEDs) have special hardware that encrypts incoming data and decrypts outgoing data in real-time. The data on the disk is always encrypted in the disk and stored in the encrypted form. The encrypted data is always decrypted on the way out of the disk. A media encryption key controls this encryption and decryption. This key is never stored in the processor or memory. Cisco UCS Manager supports SED security policies on Cisco UCS C-Series servers, B-Series M5 servers, and S-Series servers.

SEDs must be locked by providing a security key. The security key, which is also known as Key-Encryption Key or an authentication passphrase, is used to encrypt the media encryption key. If the disk is not locked, no key is required to fetch the data.

Cisco UCS Manager enables you to configure security keys locally or remotely. When you configure the key locally, you must remember the key. If you forget the key, it cannot be retrieved, and the data is lost. You can configure the key remotely by using a key management server (also known as KMIP server). This method addresses the issues related to safe-keeping and retrieval of the keys in the local management.

The encryption and decryption for SEDs is done through the hardware. Thus, it does not affect the overall system performance. SEDs reduce the disk retirement and redeployment costs through instantaneous cryptographic erasure. Cryptographic erasure is done by changing the media encryption key. When the media encryption key of a disk is changed, the data on the disk cannot be decrypted, and is immediately rendered unusable. With Cisco UCS Manager Release 3.1(3), SEDs offer disk theft protection for C-Series and S-Series servers. For HX servers, SEDs offer node theft protection. Cisco UCS Manager Release 4.0(2) extends the SED security policies to UCS B-Series M5 servers.

Security Flags of the Controller and Disk

Security flags indicate the current security status of the storage controller and disks.

The storage controller and disks have the following security flags:

- **Security Capable**—Indicates that the controller or disk is capable of supporting SED management.
- **Security Enable**—Indicates that the security-key is programmed on the controller or disk, and security is enabled on the device. This flag is set when you configure a security policy and associate it to a server, making the controller and disk secure. This flag is not set on an HX device.
- **Secured**—Indicates that the security-key is programmed on the controller or disk, and security is enabled on the HX device.

The following security flags are exclusive to storage disks:

- **Locked**—Indicates that the disk key does not match the key on the controller. This happens when you move disks across servers that are programmed with different keys. The data on a locked disk is inaccessible and the operating system cannot use the disk. To use this disk, you must either unlock the disk or secure erase the foreign configuration.
- **Foreign Secured**—Indicates that a secure disk is in foreign configuration. This happens when you unlock a locked disk with the right key, but the disk is in a foreign configuration state and the data on it is encrypted. To use this disk, you can either import the foreign configuration or clear the foreign config.

Managing Local Security Policies

Creating a Local Security Policy

Before you begin

You can create a local policy on a new or existing storage profile.

Procedure

-
- Step 1** In the **Navigation** pane, click **Storage > Storage Profiles**.
 - Step 2** Choose the storage profile where you want to create the policy.
 - Step 3** Click the **Security Policy** tab and then click **Create Security Policy** or right-click the storage profile and select **Create Security Policy**.
 - Step 4** Click the **Local Policy** option.
 - a) Enter **Key**.
The key must consist of 32 alphanumeric characters.
 - b) Click **OK**.
-

What to do next

The key thus created is associated to the storage profile for that server and is deployed under storage controller. To verify this, go to **Server ID > Inventory > Storage > Controller** and select a SAS storage controller. Go to the **General** tab and check whether the **Security** field shows as **drive security enable**.

Modifying a Local Security Policy

Procedure

- Step 1** In the **Navigation** pane, click **Storage > Storage Profiles**.
- Step 2** Choose the storage profile where you have created the policy.
- Step 3** Click the **Security Policy** tab.
- Step 4** (Optional) To modify the key for the local policy, in the **Local Policy** area:
- Enter a new security key for the database in the **Key** field.
 - Enter the current security key for the database in the **Deployed Key** field.
- Step 5** (Optional) To change the security policy from **Local Policy** to **Remote Policy**:
- Click the **Remote Policy** option.
 - Enter the primary server details in the **IP Address/Hostname** field.
 - (Optional) Enter the secondary server details in the **IP Address/Hostname** field.
 - (Optional) Enter the current security key for the database in the **Deployed Key** field.
 - (Optional) Enter the port number of the server in the **Port** field.
 - Enter the contents of the KMIP certificate in the **KMIP Server Public Certificate** field.
 - (Optional) Enter user credentials by clicking **Add Login Details**.
- Step 6** Click **Save Changes**.
-

Inserting a Secured Disk into a Server with a Local Security Policy

When you insert a secured disk into a server, one of the following will occur:

- The security-key on the drive matches that of the server and it automatically gets unlocked.
- The security-key on the disk is different from the security-key on the server. The disk will appear as a locked disk. You can do one of the following on a locked disk:
 - Erase the secure foreign configuration to delete all data on the disk.
 - Unlock the disk by providing the correct key of the disk. After unlocking the disk, the disk will be in the Foreign Secured state. You must immediately import or clear the foreign configuration for these disks.



Note If you unlock another set of disks before importing the foreign configuration for the current set of disks, the current set of disks become locked again and go in to the Locked state.

KMIP Client Certificate Policy

You can configure the key remotely by using a key management server, which is also known as KMIP server. You must create a KMIP client certificate policy before creating a remote policy. The hostname that is used for generating the certificate is the serial number of the KMIP server.

You can create a certificate policy from two separate scopes:

- Global scope—You can initially create a global certificate policy in this scope. Any modification of the certificate in this scope will not result in the regeneration of the certificate.
- Server scope—You can create or modify a certificate policy in this scope. This will result in a regeneration of the certificates. Such a certificate is specific to the server, and, for this server, overrides the global certificate.

After you create a KMIP client certificate policy, do one of the following:

- Copy the generated certificate to the KMIP Server.
- Use the generated Certificate Signing Request to get a CA-signed certificate. Copy this CA-signed certificate to the CIMC.

Creating a Global KMIP Client Certificate Policy

You can create a global KMIP client certificate policy.

The hostname that used to create the certificate when using this policy is the serial number of the server.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** In the **Work** pane, click the **Policies** tab.
- Step 3** Click the **Security** subtab.
- Step 4** Click **Create KMIP Client Cert Policy**.
- Step 5** In the **Create KMIP Client Cert Policy** dialog box that appears, enter the following information:

Name	Description
Country Code	The country code corresponding to the country in which the company resides. Enter two alphabetic characters in upper case.
State	The state or province in which the company requesting the certificate is headquartered. Enter up to 32 characters.

Name	Description
Locality	The city or town in which the company requesting the certificate is headquartered. Enter up to 32 characters.
Organization Name	The organization requesting the certificate. Enter up to 32 characters.
Organization Unit Name	The organizational unit. Enter up to 64 characters.
Email	The email address associated with the request.
Validity	The validity period of the certificate.

Step 6 Click **OK**.

Creating a KMIP Client Certificate Policy for a Server

You can create a KMIP client certificate policy for a server. This certificate is applicable only to the specific server, and overrides the global KMIP client certificate.

The hostname that used to create the certificate when using this policy is the serial number of the server.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** For C-Series and S-Series servers, expand **Equipment > Rack-Mounts > Servers > Server ID**.
- Step 3** For B-Series servers, expand **Equipment > Chassis > Chassis ID > Servers > Server ID**.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **Security** subtab.
- Step 6** Click **Create KMIP Client Cert Policy**.
- Step 7** In the **Create KMIP Client Cert Policy** dialog box that appears, enter the following information:

Name	Description
Country Code	The country code corresponding to the country in which the company resides. Enter two alphabetic characters in upper case.
State	The state or province in which the company requesting the certificate is headquartered. Enter up to 32 characters.

Name	Description
Locality	The city or town in which the company requesting the certificate is headquartered. Enter up to 32 characters.
Organization Name	The organization requesting the certificate. Enter up to 32 characters.
Organization Unit Name	The organizational unit. Enter up to 64 characters.
Email	The email address associated with the request.
Validity	The validity period of the certificate.

Step 8 Click **OK**.

Managing Remote Security Policies

Creating a Remote Security Policy

You can create a remote policy on a new or existing storage profile.

Before you begin

Ensure that you have created a KMIP client certificate policy.

Procedure

- Step 1** In the **Navigation** pane, click **Storage > Storage Profiles**.
- Step 2** Choose the storage profile where you want to create the policy.
- Step 3** Click the **Security Policy** tab and then click **Create Security Policy** or right-click the storage profile and select **Create Security Policy**.
- Step 4** Click the **Remote Policy** option.
- Enter the primary server details in the **IP Address/Hostname** field.
 - (Optional) Enter the secondary server details in the **IP Address/Hostname** field.
 - (Optional) Enter the port number of the server in the **Port** field.
 - Enter the contents of the KMIP certificate in the **KMIP Server Public Certificate** field.
 - (Optional) Enter user credentials by clicking **Add Login Details**.
 - Click **OK**.

A message that policy was created successfully is displayed.

What to do next

The key thus created is associated to the storage profile for that server and is deployed under storage controller. To verify this, go to **Server ID > Inventory > Storage > Controller** and select a SAS storage controller. Go to **General** tab and check whether the **Security** field shows as **drive security enable**.

Modifying a Remote Security Policy

Procedure

- Step 1** In the **Navigation** pane, click **Storage > Storage Profiles**.
- Step 2** Choose the storage profile where you have created the policy.
- Step 3** Click the **Security Policy** tab.
- Step 4** To modify the remote policy, in the **Remote Policy** area:
- Enter the primary server details in the **IP Address/Hostname** field.
 - (Optional) Enter the secondary server details in the **IP Address/Hostname** field.
 - (Optional) Enter the port number of the server in the **Port** field.
 - Enter the contents of the KMIP certificate in the **KMIP Server Public Certificate** field.
Save this certificate from the browser in base-64 format.
 - (Optional) Enter user credentials by clicking **Add Login Details**.
- Step 5** To change the security policy from **Remote Policy** to **Local Policy**:
- Click the **Local Policy** option.
 - Enter a new security key for the controller in the **Key** field.
- Step 6** Click **Save Changes**.
-

Modifying a Remote Security Key

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** For C-Series and S-Series servers, expand **Equipment > Rack-Mounts > Servers > Server ID**.
- Step 3** For B-Series servers, expand **Equipment > Chassis > Chassis ID > Servers > Server ID**.
- Step 4** In the **Work** area, click the **Inventory** tab.
- Step 5** Click the **Storage** subtab.
- Step 6** In the **Controllers** tab, select a SAS controller.

- Step 7** In the **General** tab, click **Modify Remote Key**.
-

Inserting a Secured Disk into a Server with a Remote Security Policy

When you insert a secured disk into a server with a remote security policy, the storage disk will appear as a locked disk. Do one of the following:

- Unlock the disk manually with the local key if the disk was previously locked using the local key.
- Unlock using the remote KMIP server.

When you move a secured disk from a server with a local security policy to a server with a remote security policy, the disk will come up as locked. Unlock the disk manually with the local key.

Enabling and Disabling Security on Disks

Before you begin

- To enable security on a disk, ensure that the disk is a JBOD.
- To secure erase a disk, the disk must be in an unconfigured good state.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** For C-Series and S-Series servers, expand **Equipment** > **Rack-Mounts** > **Servers** > **Server ID**.
- Step 3** For B-Series servers, expand **Equipment** > **Chassis** > **Chassis ID** > **Servers** > **Server ID**.
- Step 4** In the **Work** area, click the **Inventory** tab.
- Step 5** Click the **Storage** subtab.
- Step 6** In the **Disks** tab, select a disk.
- Step 7** In the **Details** area, click **Enable Encryption**.
- Step 8** To disable a secure disk, click **Secure Erase**.
-

Disabling Security on a Controller

Before you begin

You can disable security only on SAS controllers. To disable security on controller, you must first disable security on all the secure disks and delete all the secure virtual drives under the controller.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** For C-Series and S-Series servers, expand **Equipment > Rack-Mounts > Servers > Server ID**.
 - Step 3** For B-Series servers, expand **Equipment > Chassis > Chassis ID > Servers > Server ID**.
 - Step 4** In the **Work** area, click the **Inventory** tab.
 - Step 5** Click the **Storage** subtab.
 - Step 6** In the **Controllers** tab, select a SAS controller.
 - Step 7** In the **General** tab, click **Disable Security**.
-

Unlocking a Locked Disk

When the key of an SED does not match the key on the controller, it shows the disk as Locked, Foreign Secure. You must unlock the disks either by providing the security-key for that disk, or by using the remote KMIP server. After unlocking the disk, import or clear the foreign configuration.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment > Rack-Mounts > Servers > Server Number**.
 - Step 3** In the **Work** area, click the **Inventory** tab.
 - Step 4** Click the **Storage** subtab.
 - Step 5** In the **Controller** tab, select a SAS controller.
 - Step 6** To unlock a disk that is secured with a local security policy:
 - a) In the **General** tab, click **Unlock Disk**.
 - b) In the **Key** text box, provide the key that was used to lock the disk.
 - c) Click **OK**.
 - Step 7** To unlock a disk that is secured with a remote KMIP server, in the **General** tab, click **Unlock For Remote**.
-

After you unlock a locked disk, the security status of the disk will show as Foreign Secure.

What to do next

Import or clear foreign configuration.

Erasing a Secure Foreign Configuration Disk

You can erase a secure foreign configuration disk when you have a disk in locked state and you want to use the disk without accessing the existing data.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** For C-Series and S-Series servers, expand **Equipment > Rack-Mounts > Servers > Server ID**.
- Step 3** For B-Series servers, expand **Equipment > Chassis > Chassis ID > Servers > Server ID**.
- Step 4** In the **Work** area, click the **Inventory** tab.
- Step 5** Click the **Storage** subtab.
- Step 6** In the **Disks** tab, select a disk.
- Step 7** In the **General** tab, click **Secure Erase Foreign Configuration**.
-

Secure Data Deletion

The Commission Regulation (EU) 2019/424 requires that data be securely disposed of.

Secure data disposal is accomplished by using commonly available tools that erase the data from the various/drives, memory, and storage in the Cisco UCS servers and reset them to factory settings.

Secure data deletion for compliance with Commission Regulation (EU) 2019/424 is supported for the following Cisco UCS servers:

- Cisco UCS B200
- Cisco UCS B480
- Cisco UCS C125
- Cisco UCS C220
- Cisco UCS C240
- Cisco UCS C480
- Cisco UCS S3260

You must be familiar with what devices are present in your UCS server and run the appropriate tools for secure data deletion. In some cases, you may need to run multiple tools.

Full instructions on how to securely erase data are available at: <https://www.cisco.com/web/dofc/18794277.pdf>.



CHAPTER 11

Storage Profiles

- [Storage Profiles, on page 105](#)
- [Cisco Boot Optimized M.2 RAID Controller, on page 106](#)
- [Disk Groups and Disk Group Configuration Policies, on page 106](#)
- [RAID Levels, on page 112](#)
- [Automatic Disk Selection, on page 113](#)
- [Supported LUN Modifications, on page 114](#)
- [Unsupported LUN Modifications, on page 114](#)
- [Disk Insertion Handling, on page 115](#)
- [Virtual Drive Naming, on page 117](#)
- [LUN Dereferencing, on page 117](#)
- [Controller Constraints and Limitations, on page 117](#)
- [Storage Profiles, on page 118](#)
- [Configuring Storage Profiles, on page 143](#)

Storage Profiles

To allow flexibility in defining the number of storage disks, roles and usage of these disks, and other storage parameters, you can create and use storage profiles. A storage profile encapsulates the storage requirements for one or more service profiles. LUNs configured in a storage profile can be used as boot LUNs or data LUNs, and can be dedicated to a specific server. You can also specify a local LUN as a boot device. However, LUN resizing is not supported. The introduction of storage profiles allows you to do the following:

- Configure multiple virtual drives and select the physical drives that are used by a virtual drive. You can also configure the storage capacity of a virtual drive.
- Configure the number, type and role of disks in a disk group.
- Associate a storage profile with a service profile.

You can create a storage profile both at an org level and at a service-profile level. A service profile can have a dedicated storage profile as well as a storage profile at an org level.

Cisco Boot Optimized M.2 RAID Controller

Beginning with 4.0(4a), Cisco UCS Manager supports Cisco boot optimized M.2 RAID controller (UCS-M2-HWRAID), which is based on Marvell® 88SE92xx PCIe to SATA 6Gb/s controller. It is supported on the following servers:

- Cisco UCS C220 M5 Server
- Cisco UCS C240 M5 Server
- Cisco UCS C480 M5 Server
- Cisco UCS B200 M5 Server
- Cisco UCS B480 M5 Server

The following M.2 drives are managed by the Cisco boot optimized M.2 RAID controller:

- 240GB M.2 6G SATA SSD
- 960GB M.2 6G SATA SSD

The Cisco boot optimized M.2 RAID controller supports only RAID1/JBOD (default - JBOD) mode and only UEFI boot mode.

Limitations of Cisco boot optimized M.2 RAID controller

- Existing LUN migration is not supported.
- **Local Disk Configuration** policy is not supported.
- The number of LUNs that can be created is limited to one because creating a single LUN uses the entire disk capacity.
- LUN is created using the **Local LUN** tab (see [Configuring Local LUNs, on page 120](#)) under storage profile and not using the controller definitions.
- You cannot mix different capacity M.2 drives.
- You cannot rename an orphan virtual drive on a blade or a rack server.

Disk Groups and Disk Group Configuration Policies

Servers in a chassis can use storage that is centralized in that chassis. You can select and configure the disks to be used for storage. A logical collection of these physical disks is called a disk group. Disk groups allow you to organize local disks. The storage controller controls the creation and configuration of disk groups.

A disk group configuration policy defines how a disk group is created and configured. The policy specifies the RAID level to be used for the disk group. It also specifies either a manual or an automatic selection of disks for the disk group, and roles for disks. You can use a disk group policy to manage multiple disk groups. However, a single disk group can be managed only by one disk group policy.

A hot spare is an unused extra disk that can be used by a disk group in the case of failure of a disk in the disk group. Hot spares can be used only in disk groups that support a fault-tolerant RAID level. In addition, a disk can be allocated as a global hot spare, which means that it can be used by any disk group.

Virtual Drives

A disk group can be partitioned into virtual drives. Each virtual drive appears as an individual physical device to the Operating System.

All virtual drives in a disk group must be managed by using a single disk group policy.

Configuration States

Indicates the configuration states of a virtual drive. Virtual drives can have the following configuration states:

- Applying—Creation of the virtual drive is in progress.
- Applied—Creation of the virtual drive is complete, or virtual disk policy changes are configured and applied successfully.
- Failed to apply—Creation, deletion, or renaming of a virtual drive has failed due to errors in the underlying storage subsystem.
- Orphaned—The service profile that contained this virtual drive is deleted or the service profile is no longer associated with a storage profile.



Note Orphaned LUNs cannot be used for booting OS. Although an image can be installed on these LUNs, booting from these drives will fail. To use any specific orphaned LUN, you must reassociate the storage profile, which will return it to the “Equipped” presence state. When there are orphaned LUNs with OS installed on it, and the boot policy associated with a service profile has Local LUN then OS booting will happen with any available orphaned LUNs. In case of multiple OS installed, there is no specific orphan LUN associated with any OS.

When there are orphaned LUNs with OS installed on it, and the boot policy associated with a service profile has Local LUN then OS booting will happen with any available orphaned LUNs. In case of multiple OS installed, there is no specific orphan LUN associated with any OS.

- Not in use—The service profile that contained this virtual drive is in the disassociated state.

Deployment States

Indicates the actions that you are performing on virtual drives. Virtual drives can have the following deployment states:

- No action—No pending work items for the virtual drive.
- Creating—Creation of the virtual drive is in progress.
- Deleting—Deletion of the virtual drive is in progress.
- Modifying—Modification of the virtual drive is in progress.

- Apply-Failed—Creation or modification of the virtual drive has failed.

Operability States

Indicates the operating condition of a virtual drive. Virtual drives can have the following operability states:

- Optimal—The virtual drive operating condition is good. All configured drives are online.
- Degraded—The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
- Cache-degraded—The virtual drive has been created with a write policy of **write back** mode, but the BBU has failed, or there is no BBU.



Note This state does not occur if you select the **always write back** mode.

- Partially degraded—The operating condition in a RAID 6 virtual drive is not optimal. One of the configured drives has failed or is offline. RAID 6 can tolerate up to two drive failures.
- Offline—The virtual drive is not available to the RAID controller. This is essentially a failed state.
- Unknown—The state of the virtual drive is not known.

Presence States

Indicates the presence of virtual drive components. Virtual drives have the following presence states:

- Equipped—The virtual drive is available.
- Mismatched—A virtual drive deployed state is different from its configured state.
- Missing—Virtual drive is missing.

Configuring a Disk Group Policy

You can configure the disks in a disk group policy automatically or manually.

Procedure

-
- Step 1** In the **Navigation** pane, click **Storage**.
 - Step 2** Expand **Storage > Storage Provisioning > Storage Policies**
 - Step 3** Expand the node for the organization where you want to create the disk group policy.
 - Step 4** Right-click **Disk Group Policies** in the organization and select **Create Disk Group Policy**.
 - Step 5** In the **Create Disk Group Policy** dialog box, specify the following:

Name	Description
Name field	The name of the policy

Name	Description
	This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved.
Description field	<p>A description of the policy. We recommend that you include information about where and when the policy should be used.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p>
RAID Level drop-down list	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • RAID 0 Striped • RAID 1 Mirrored <p>Note The Cisco boot optimized M.2 RAID controller (UCS-M2-HWRAID) supports only RAID1.</p> <ul style="list-style-type: none"> • RAID 5 Striped Parity • RAID 6 Striped Dual Parity • RAID 10 Mirrored and Striped <p>Note When you create a disk group with RAID 1 policy and configure four disks for it, a RAID 1E configuration is created internally by the storage controller.</p>

Step 6

Create LUNs using JBOD or UG drives under the following scenarios:

- a. When the drive state is UG and is in the disk group policy, and if Use JBOD is set to:
 - Yes—Both JBOD and UG drives can be used based on the drive slot ordering.
 - No—Only UG drives can be used.
- b. When drive state is JBOD and is in the disk group policy, and if Use JBOD is set to:
 - Yes—Both JBOD and UG drives can be used based on the drive slot ordering.
 - No—Only UG drives can be used.
- c. When the drive state is JBOD or UG and is in the disk group policy, and if Use JBOD is set to:
 - Yes—Both JBOD and UG drives can be used.
 - No—Only UG drives can be used.

Note The UCS Manager disk selection is based on the sequential slot number, irrespective of the drive state.

Step 7 To automatically configure the disks in a disk group policy, select **Disk Group Configuration (Automatic)** and specify the following:

Note If you have a setup with the Cisco Boot Optimized M.2 Raid Controller (UCS-M2-HWRAID), then go to [Step 8, on page 110](#).

Name	Description
Number of drives field	Specifies the number of drives for the disk group. The range for drives is from 0 to 24 drives. Unspecified is the default number of drives. When you select the number of drives as Unspecified , the number of drives will be selected according to the disk selection process.
Drive Type field	Drive type for the disk group. You can select: <ul style="list-style-type: none"> • HDD • SSD • Unspecified Unspecified is the default type of drive. When you select the drive type as Unspecified , the first available drive is selected. After this drive is selected, subsequent drives will be of a compatible type. For example, if the first was SSD, all subsequent drives would be SSD.
Number of Hot Spares field	Number of dedicated hot spares for the disk group. The range for dedicated hot spares is from 0 to 24 hot spares. Unspecified is the default number of dedicated hot spares. When you select the number of dedicated hot spares as Unspecified , the hot spares will be selected according to the disk selection process.
Min Drive Size field	Minimum drive size for the disk group. Only disks that match this criteria are available for selection. The range for minimum drive size is from 0 to 10240 GB. Unspecified is the default minimum drive size. When you select the minimum drive size as Unspecified , drives of all sizes will be available for selection.

Step 8 To manually configure the disks in a disk group policy, select **Disk Group Configuration (Manual)** and do the following:

- a) On the icon bar to the right of the table, click +
- b) In the **Create Local Disk Configuration Reference** dialog box, complete the following fields:

Name	Description
Slot field	Slot for which the local disk reference is configured. Note M.2 drives typically have Slot IDs = 253, 254. Additionally, verify the Slot IDs by navigating to Equipment > Server <i>servername</i> > Inventory > Storage > Disks

Name	Description
Role field	<p>Note If you have a setup with the Cisco Boot Optimized M.2 Raid Controller (UCS-M2-HWRAID), then select Normal (default). Selecting any other value results in configuration error.</p> <p>Role of the local disk in the disk group. You can select:</p> <ul style="list-style-type: none"> • Normal • Dedicated Hot Spare • Global Hot Spare
Span ID field	<p>Note If you have a setup with the Cisco Boot Optimized M.2 Raid Controller (UCS-M2-HWRAID), then this field does not apply. Leave the Span ID field as unspecified. Selecting any value results in configuration error.</p> <p>Specifies the ID of the span group to which the disk belongs. Disks belonging to a single span group can be treated as a single disk with a larger size. The values range from 0 to 8. For RAID-10, RAID-50, and RAID-60, minimum 2 spans are required and maximum 8 spans are supported. You can also set the Span ID as Unspecified, when spanning information is not required.</p>

Step 9

In the **Virtual Drive Configuration** area, specify the following:

- Note** If you have a setup with the Cisco Boot Optimized M.2 Raid Controller (UCS-M2-HWRAID), then:
- You can create only one virtual drive
 - For **Strip Size (KB)**, select **64KB** or **32KB**. Selecting any other value results in configuration error.
 - For **Access Policy**, **Read Policy**, **Write Cache Policy**, **IO Policy**, and **Drive Cache**, select **Platform Default**. Selecting any other value results in configuration error.

Name	Description
Strip Size (KB) field	Stripe size for a virtual drive. This can only be Platform Default .
Access Policy field	<p>Access policy for a virtual drive. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default • Read Write • Read Only • Blocked
Read Policy field	<p>Read policy for a virtual drive. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default • Read Ahead

Name	Description
	<ul style="list-style-type: none"> • Normal
Write Cache Policy field	Write-cache-policy for a virtual drive. This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Write Through • Write Back Good Bbu • Always Write Back
IO Policy field	I/O policy for a virtual drive. This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Direct • Cached
Drive Cache field	State of the drive cache. This can be one of the following: <ul style="list-style-type: none"> • Platform Default • No Change • Enable • Disable

All virtual drives in a disk group should be managed by using the same disk group policy.

Step 10

Click **OK**.

Note When you accept the virtual drive (VD) default values and associate the disk group policy to a service profile, you can modify the VD configuration after it is associated to a service profile. If you modify the VD default values from the WebBIOS to use the non-default values, a properties fault is not generated to verify the changed values.

RAID Levels

The RAID level of a disk group describes how the data is organized on the disk group for the purpose of ensuring availability, redundancy of data, and I/O performance.

The following are features provided by RAID:

- **Striping**—Segmenting data across multiple physical devices. This improves performance by increasing throughput due to simultaneous device access.
- **Mirroring**—Writing the same data to multiple devices to accomplish data redundancy.

- **Parity**—Storing of redundant data on an additional device for the purpose of error correction in the event of device failure. Parity does not provide full redundancy, but it allows for error recovery in some scenarios.
- **Spanning**—Allows multiple drives to function like a larger one. For example, four 20 GB drives can be combined to appear as a single 80 GB drive.

The supported RAID levels include the following:

- **RAID 0 Striped**—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- **RAID 1 Mirrored**—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **RAID 5 Striped Parity**—Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.

RAID 5 distributes parity data blocks among the disks that are part of a RAID-5 group and requires a minimum of three disks.

- **RAID 6 Striped Dual Parity**—Data is striped across all disks in the array and two sets of parity data are used to provide protection against failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.

Other than addition of a second parity block, RAID 6 is identical to RAID 5. A minimum of four disks are required for RAID 6.

- **RAID 10 Mirrored and Striped**—RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates through block-level striping. RAID 10 is mirroring without parity and block-level striping. A minimum of four disks are required for RAID 10.
- **RAID 50 Striped Parity and Striped**—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance.
- **RAID 60 Striped Dual Parity and Striped**—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.

Automatic Disk Selection

When you specify a disk group configuration, and do not specify the local disks in it, Cisco UCS Manager determines the disks to be used based on the criteria specified in the disk group configuration policy. Cisco UCS Manager can make this selection of disks in multiple ways.

When all qualifiers match for a set of disks, then disks are selected sequentially according to their slot number. Regular disks and dedicated hot spares are selected by using the lowest numbered slot.

The following is the disk selection process:

1. Iterate over all local LUNs that require the creation of a new virtual drive. Iteration is based on the following criteria, in order:
 - a. Disk type
 - b. Minimum disk size from highest to lowest

- c. Space required from highest to lowest
 - d. Disk group qualifier name, in alphabetical order
 - e. Local LUN name, in alphabetical order
2. Select regular disks depending on the minimum number of disks and minimum disk size. Disks are selected sequentially starting from the lowest numbered disk slot that satisfies the search criteria.



Note If you specify **Any** as the type of drive, the first available drive is selected. After this drive is selected, subsequent drives will be of a compatible type. For example, if the first drive was SATA, all subsequent drives would be SATA. Cisco UCS Manager Release 2.5 supports only SATA and SAS.

Cisco UCS Manager Release 2.5 does not support RAID migration.

3. Select dedicated hot spares by using the same method as normal disks. Disks are only selected if they are in an **Unconfigured Good** state.
4. If a provisioned LUN has the same disk group policy as a deployed virtual drive, then try to deploy the new virtual drive in the same disk group. Otherwise, try to find new disks for deployment.

Supported LUN Modifications

Some modifications that are made to the LUN configuration when LUNs are already deployed on an associated server are supported.

The following are the types of modifications that can be performed:

- Creation of a new virtual drive.
- Deletion of an existing virtual drive, which is in the orphaned state.
- Non-disruptive changes to an existing virtual drive. These changes can be made on an existing virtual drive without loss of data, and without performance degradation:
 - Policy changes. For example, changing the write cache policy.
 - Modification of boot parameters

The removal of a LUN will cause a warning to be displayed. Ensure that you take action to avoid loss of data.

Unsupported LUN Modifications

Some modifications to existing LUNs are not possible without destroying the original virtual drive and creating a new one. All data is lost in these types of modification, and these modifications are not supported.

Disruptive modifications to an existing virtual drive are not supported. The following are unsupported disruptive changes:

- Any supported RAID level change that can be handled through reconstruction. For example, RAID0 to RAID1.

- Increasing the size of a virtual drive through reconstruction.
- Addition and removal of disks through reconstruction.
- **Expand To Available** option is not supported for already deployed LUN.

Destructive modifications are also not supported. The following are unsupported destructive modifications:

- RAID-level changes that do not support reconstruction. For example, RAID5 to RAID1.
- Shrinking the size of a virtual drive.
- RAID-level changes that support reconstruction, but where there are other virtual drives present on the same drive group.
- Disk removal when there is not enough space left on the disk group to accommodate the virtual drive.
- Explicit change in the set of disks used by the virtual drive.

Disk Insertion Handling

When the following sequence of events takes place:

1. The LUN is created in one of the following ways:
 - a. You specify the slot specifically by using a local disk reference
 - b. The system selects the slot based on criteria specified by you
2. The LUN is successfully deployed, which means that a virtual drive is created, which uses the slot.
3. You remove a disk from the slot, possibly because the disk failed.
4. You insert a new working disk into the same slot.

The following scenarios are possible:

- [Non-Redundant Virtual Drives, on page 115](#)
- [Redundant Virtual Drives with No Hot Spare Drives, on page 116](#)
- [Redundant Virtual Drives with Hot Spare Drives, on page 116](#)
- [Replacing Hot Spare Drives, on page 116](#)
- [Inserting Physical Drives into Unused Slots, on page 116](#)

Non-Redundant Virtual Drives

For non-redundant virtual drives (RAID 0), when a physical drive is removed, the state of the virtual drive is **Inoperable**. When a new working drive is inserted, the new physical drive goes to an **Unconfigured Good** state.

For non-redundant virtual drives, there is no way to recover the virtual drive. You must delete the virtual drive and re-create it.

Redundant Virtual Drives with No Hot Spare Drives

For redundant virtual drives (RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60) with no hot spare drives assigned, virtual drive mismatch, virtual drive member missing, and local disk missing faults appear until you insert a working physical drive into the same slot from which the old physical drive was removed.

If the physical drive size is greater than or equal to that of the old drive, the storage controller automatically uses the new drive for the virtual drive. The new drive goes into the **Rebuilding** state. After rebuild is complete, the virtual drive goes back into the **Online** state.

Redundant Virtual Drives with Hot Spare Drives

For redundant virtual drives (RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60) with hot spare drives assigned, when a drive fails, or when you remove a drive, the dedicated hot spare drive, if available, goes into the **Rebuilding** state with the virtual drive in the **Degraded** state. After rebuilding is complete, that drive goes to the **Online** state.

Cisco UCSM raises a disk missing and virtual drive mismatch fault because although the virtual drive is operational, it does not match the physical configuration that Cisco UCSM expects.

if you insert a new disk in the slot with the disk missing, automatic copy back starts from the earlier hot spare disk to the newly inserted disk. After copy back, the hot spare disk is restored. In this state all faults are cleared.

If automatic copy back does not start, and the newly inserted disk remains in the **Unconfigured Good**, **JBOD**, or **Foreign Configuration** state, remove the new disk from the slot, reinsert the earlier hot spare disk into the slot, and import foreign configuration. This initiates the rebuilding process and the drive state becomes **Online**. Now, insert the new disk in the hot spare slot and mark it as hot spare to match it exactly with the information available in Cisco UCSM.

Replacing Hot Spare Drives

If a hot spare drive is replaced, the new hot spare drive will go to the **Unconfigured Good**, **Unconfigured Bad**, **JBOD**, or **Foreign Configuration** state.

Cisco UCSM will raise a virtual drive mismatch or virtual drive member mismatch fault because the hot spare drive is in a state different from the state configured in Cisco UCSM.

You must manually clear the fault. To do this, you must perform the following actions:

1. Clear the state on the newly inserted drive to **Unconfigured Good**.
2. Configure the newly inserted drive as a hot spare drive to match what is expected by Cisco UCSM.

Inserting Physical Drives into Unused Slots

If you insert new physical drives into unused slots, neither the storage controller nor Cisco UCSM will make use of the new drive even if the drive is in the **Unconfigured Good** state and there are virtual drives that are missing good physical drives.

The drive will simply go into the **Unconfigured Good** state. To make use of the new drive, you will need to modify or create LUNs to reference the newly inserted drive.

Virtual Drive Naming

When you use Cisco UCS Manager to create a virtual drive, Cisco UCS Manager assigns a unique ID that can be used to reliably identify the virtual drive for further operations. Cisco UCS Manager also provides the flexibility to provide a name to the virtual drive at the time of service profile association. Any virtual drive without a service profile or a server reference is marked as an orphan virtual drive.

In addition to a unique ID, a name is assigned to the drive. Names can be assigned in two ways:

- When configuring a virtual drive, you can explicitly assign a name that can be referenced in storage profiles.
- If you have not preprovisioned a name for the virtual drive, Cisco UCS Manager generates a unique name for the virtual drive.

You can rename an orphan virtual drive on a blade or a rack server that are not referenced by any service profile or server.



Note The renaming an orphan virtual drive is not supported for Cisco boot optimized M.2 Raid controller (UCS-M2-HWRAID).

LUN Dereferencing

A LUN is dereferenced when it is no longer used by any service profile. This can occur as part of the following scenarios:

- The LUN is no longer referenced from the storage profile
- The storage profile is no longer referenced from the service profile
- The server is disassociated from the service profile
- The server is decommissioned

When the LUN is no longer referenced, but the server is still associated, re-association occurs.

When the service profile that contained the LUN is disassociated, the LUN state is changed to **Not in use**.

When the service profile that contained the LUN is deleted, the LUN state is changed to **Orphaned**.

Controller Constraints and Limitations

- The following table provides the maximum supported virtual drives for servers:

Servers/Storage Controllers	Maximum Virtual Drives
UCSC-C240-M5, UCSC-C480-M5	32

Servers/Storage Controllers	Maximum Virtual Drives
UCS-S3260-M5, UCSC-C3X60-M4, UCSC-C3K-M4	64
UCSC-C240-M4, UCSC-C240-M3, UCSC-C24-M3	24
UCSB-MRAID12G	16
UCS-M2-HWRAID	2
For all other servers.	18

**Note**

- Storage controllers support the check max feature.
- When servers with multiple storage controllers are managed by the same storage profile, the maximum virtual drives are limited to the maximum value supported by the server.

- In Cisco UCS Manager Release 2.2(4), blade servers do not support drives with a block size of 4K, but rack-mount servers support such drives. If a drive with a block size of 4K is inserted into a blade server, discovery fails and the following error message appears: Unable to get Scsi Device Information from the system.
- In Cisco UCS Manager Release 3.1(2) and later releases, RAID Controller that does not support Out of band inventory (OOB) in C240 M4, M5 servers, display Operability as NA and Drive State as Unknown.

Storage Profiles

Creating a Storage Profile

You can create storage profile policies from the **Storage** tab in the **Navigation** pane. Additionally, you can also configure the default storage profile that is specific to a service profile from the **Servers** tab.

**Caution**

If you have a Cisco UCS blade or rack server with a default local disk configuration present in a Service Profile or Service Profile Template from an earlier release of UCS Manager and you upgrade to the 3.1 release and later releases, you can successfully create a Storage Profile with local LUNs in the same Service Profile or Service Profile Template if you change the Local Disk Configuration Default policy to **Any Configuration** instead of RAID level options in the local disk policy. The legacy LUN is thereafter part of the storage inventory.

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization where you want to create the storage profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the organization and select **Create Storage Profile**.
- Step 5** In the **Create Storage Profile** dialog box, specify the storage profile **Name**. You can provide an optional **Description** for this storage profile.
- Step 6** (Optional) In the **LUNs** area, create **Local LUNs** and add them to this storage profile.
See [Configuring Local LUNs, on page 120](#) for more information.
- Step 7** (Optional) In the **LUN Set** area, create **LUN Set** and add them to this storage profile.
See [Creating a LUN Set, on page 124](#) for more information.
- Step 8** In the **LUNs** area, create **Controller Definitions** and add them to this storage profile.
See [Creating a Storage Profile PCH Controller Definition, on page 130](#) for more information.
- Step 9** In the **LUNs** area, create **Security Policy** and add them to this storage profile.
See [Creating a Local Security Policy, on page 96](#) and [Creating a Remote Security Policy, on page 100](#) for more information.
- Step 10** Click **OK**.
-

Creating a Specific Storage Profile

Procedure

- Step 1** Expand **Servers > Service Profiles**.
- Step 2** Expand the node for the organization that contains the service profile for which you want to create a specific storage profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 3** Choose the service profile for which you want to create a specific storage profile.
- Step 4** In the **Work** pane, click the **Storage > LUN Configuration** tab.
- Step 5** In the **Actions** area, click **Modify Storage Profile**.
- Step 6** In the **Modify Storage Profile** dialog box, click the **Specific Storage Profile** tab.
- Step 7** Click **Create Specific Storage Profile**.
- Step 8** (Optional) In the **Specific Storage Profile** area, complete the **Description** field to set the description of the storage profile.

Each service profile can have only one specific storage profile. Hence, the name of this storage profile is provided by default.

Step 9 In the **Storage Items** area, **Create Local LUNs** and add them to this storage profile.

Step 10 Click **OK**.

Step 11 If a confirmation dialog box displays, click **Yes**.

Deleting a Storage Profile

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile that you want to delete.
- Step 4** Right-click the storage profile that you want to delete and select **Delete**.
- Step 5** Click **Yes** in the confirmation box that appears.

Local LUNs

Configuring Local LUNs

You can create local LUNs within a storage profile policy from the **Storage** tab in the **Navigation** pane. Additionally, you can also create local LUNs within the default storage profile that is specific to a service profile from the **Servers** tab.

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile within which you want to create a local LUN.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Create Local LUN**.
- Step 6** In the Create Local LUN dialog box, complete the following fields:

Name	Description
Create Local LUN option	(Appears when you create a local LUN) Selected by default when you create a local LUN.
Prepare Claim Local LUN option	(Appears when you create a local LUN) Select when you want to claim an orphan LUN.

Name	Description
Name field	<p>The name of the local LUN.</p> <p>This name can be between 1 and 10 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved.</p> <p>Note If the name given in Prepare Claim Local LUN is different from the name to be claimed, this LUN name and the Virtual drive name appearing in the LUN properties are different.</p>
Size (GB) field	<p>Size of this LUN in GB.</p> <p>Note You do not need to specify a LUN size while claiming an orphaned LUN.</p> <p>Note In a setup with the Cisco boot optimized M.2 Raid controller, this field is not grayed out. However, you do not have to populate this field. The system uses the full disk capacity to create the LUN, irrespective of the size specified.</p>
Fractional Size (MB) field	The fractional size of this LUN in MB.
Auto Deploy radio buttons	<p>Whether the local LUN should be automatically deployed or not. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto Deploy —Automatically deploys the local LUN. • No Auto Deploy —Does not automatically deploy the local LUN.
Expand To Available checkbox	<p>(Only available for rack and blade servers) Specifies that this LUN can be expanded to use the entire available disk group.</p> <p>For each drive group, only one LUN can use this option.</p> <p>Expand To Available option is not supported for already deployed LUN.</p>
Select Disk Group Configuration drop-down list	Chose the disk group configuration to be applied to this local LUN from the drop-down list.
Create Disk Group Policy link	Displays the Create Disk Group Policy dialog box to create a new disk group.

- Step 7** (Optional) Click **Create Disk Group Policy** to create a new disk group policy for this local LUN.
- Step 8** Click **OK**.

Displaying Details of All Local LUNs Inherited By a Service Profile

Storage profiles can be defined under org and as a dedicated storage profile under service profile. Thus, a service profile inherits local LUNs from both possible storage profiles. It can have a maximum of 2 such local LUNs. You can display the details of all local LUNs inherited by a service profile by using the following command:

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile that you want to display.
- Step 4** Choose the service profile whose inherited local LUNs you want to display.
- Step 5** In the **Work** pane, click the **Storage** tab.
- Step 6** Click the **Storage Profiles** subtab, and then click the **Local LUNs** tab.

Displays the following detailed information about all the local LUNs inherited by the specified service profile:

- **Name**—LUN name in the storage profile.
 - **RAID Level**—Summary of the RAID level of the disk group used.
 - **Size (MB)**—Size, in MB, of the LUN specified in the storage profile.
 - **Config State**—State of LUN configuration. The states can be one of the following:
 - **Applying**—Admin state is online, the LUN is associated with a server, and the virtual drive is being created.
 - **Applied**—Admin state is online, the LUN is associated with a server, and the virtual drive is created.
 - **Apply Failed**—Admin stage is online, the LUN is associated with a server, but the virtual drive creation failed.
 - **Not Applied**—The LUN is not associated with a server, or the LUN is associated with a service profile, but admin state is undeployed.
 - **Deploy Name**—The virtual drive name after deployment.
 - **LUN ID**—LUN ID.
 - **Drive State**—State of the virtual drive. The states are:
 - **Unknown**
 - **Optimal**
 - **Degraded**
 - **Inoperable**
 - **Partially Degraded**
-

Deleting Local LUNs

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile from which you want to delete a local LUN.
- Step 4** Expand **Local LUNs** for the storage profile that you want and select the LUN that you want to delete.
- Step 5** Right-click the LUN that you want to delete and select **Delete**.
A confirmation dialog box appears.
- Step 6** Click **Yes**.
-

LUN Set

LUN Set

Beginning with release 4.0(2a), Cisco UCS Manager provides the ability to configure a range of disk slots into individual RAID0 LUNs using LUN Set option.

The following guidelines should be considered while creating a LUN Set:

- Only SSD and HDD types of disks are allowed.
- Upto 60 disks are allowed in one range.
- You cannot add the same set of disks in range under two different LUN Set configurations.
- If a disk is set in the disk slot range of LUN Set, then you cannot configure the same disk set in Local LUN configuration under the same storage policy. Similarly, if a disk is set in Local LUN configuration, then you cannot use the same disk in the disk slot range of LUN Set.
- The server, in which the LUN Set is configured, should support OOB storage operations.
- You cannot configure a Local Disk Policy along with a Storage Policy in the same Service Profile.
- You cannot have the same name for a Local LUN and LUN Set.
- In S-series server PCH controllers, slots 201 and 202 do not support LUN Set.

Limitations of LUN Set

Cisco UCS Manager has the following limitations with LUN Set:

- You cannot claim orphaned Local LUNs into a LUN Set.
- Once created, you cannot modify a LUN Set. You should delete and create a new LUN Set with desired parameters.
- OS boot is not supported from LUN Set.

Creating a LUN Set

You can create LUN Set within a storage profile policy from the **Storage** tab in the **Navigation** pane. Additionally, you can also create LUN Set within the default storage profile that is specific to a service profile from the **Servers** tab.

Before you begin

Ensure that the disk set which you are going to use to create LUN Set is in **UnConfigured Good** or **JBOD** drive state



Note If the disk drive state is in **JBOD** state, then you may experience data loss if you use the same disk in the slot range.

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile within which you want to create a LUN Set.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Create LUN Set**.
- Step 6** In the Create LUN Set dialog box, complete the following fields:

Name	Description
Name field	The name of the LUN Set. This name can be between 1 and 10 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved.
RAID Level option	Currently Cisco UCS Manager supports only RAID 0 Striped option.
Disk Slot Range field	The slot range for the disk.

Name	Description
Strip Size (KB) drop-down list	<p>For striped virtual drives, the portion of the striped data segment that resides on each physical disk.</p> <ul style="list-style-type: none"> • Platform Default • 8KB • 16KB • 32KB • 64KB • 128KB • 256KB • 512KB • 1024KB
Access Policy option	<p>The type of access allowed. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default • Read Write • Read only • Blocked
Read Policy option	<p>The read-ahead cache mode. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default • Read Ahead • Normal
Write Cache Policy option	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default • Write Through • Write Back Good Bbu • Always Write Back

Name	Description
IO Policy option	This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Direct • Cached
Drive Cache option	This can be one of the following: <ul style="list-style-type: none"> • Platform Default • No Change • Enable • Disable
Security checkbox	Select this check box to secure a virtual drive.

Step 7 Click **OK**.

Displaying the Details of a LUN Set

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers > Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile that you want to display.

Step 4 Choose the service profile whose inherited local LUNs you want to display.

Step 5 In the **Work** pane, click the **Storage** tab.

Step 6 Click the **Storage Profiles** subtab, and then click the **LUN Set** tab.

Displays the following detailed information about all the LUN Set inherited by the specified service profile:

Table 8: LUN Set

Name	Description
Name column	The name of the LUN Set.
RAID Level option	Currently Cisco UCS Manager supports only RAID 0 Striped option.
Disk Slot Range field	The slot range for the disk.

Name	Description
Name field	<p>The name of the LUN Set.</p> <p>This name can be between 1 and 10 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved.</p>
RAID Level option	<p>Currently Cisco UCS Manager supports only RAID 0 Striped option.</p>
Disk Slot Range field	<p>The slot range for the disk.</p>
Strip Size (KB) drop-down list	<p>For striped virtual drives, the portion of the striped data segment that resides on each physical disk.</p> <ul style="list-style-type: none"> • Platform Default • 8KB • 16KB • 32KB • 64KB • 128KB • 256KB • 512KB • 1024KB
Access Policy option	<p>The type of access allowed. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default • Read Write • Read only • Blocked
Read Policy option	<p>The read-ahead cache mode. This can be one of the following:</p> <ul style="list-style-type: none"> • Platform Default • Read Ahead • Normal

Name	Description
Write Cache Policy option	This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Write Through • Write Back Good Bbu • Always Write Back
IO Policy option	This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Direct • Cached
Drive Cache option	This can be one of the following: <ul style="list-style-type: none"> • Platform Default • No Change • Enable • Disable
Security checkbox	Select this check box to secure a virtual drive.

Deleting a LUN Set

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile from which you want to delete a LUN Set.
- Step 4** Expand **LUN Set** for the storage profile that you want and select the LUN Set that you want to delete.
- Step 5** Right-click the LUN Set that you want to delete and select **Delete**.
A confirmation dialog box appears.
- Step 6** Click **Yes**.

PCH Controller Definitions

PCH SSD Controller Definition

Cisco UCS Manager Platform Controller Hub (PCH) Solid State Drive (SSD) Controller Definition provides a local storage configuration in storage profiles where you can configure all the disks in a single RAID or in a JBOD disk array.

The PCH Controller Definition configuration provides the following features:

- Ability to configure a single LUN RAID across two internal SSDs connected to the onboard PCH controller
- A way to configure the controller in two modes: AHCI (JBOD) and SWRAID (RAID).
- Ability to configure the PCH storage device in an Embedded Local LUN and Embedded Local Disk boot policy so precision control for boot order is achieved even with the presence of other bootable local storage devices in the server. Do not use the Local LUN or the Local JBOD options to boot from PCH disks
- Scrub policy support for the internal SSD drives. This is applicable only for the SWRAID mode. This does not apply for the AHCI and NORAIID of PCH Controller modes. *See the UCS Manager Server Management Guide.*
- Firmware upgrade support for the internal SSD drives. Disk firmware upgrade is supported only when the PCH Controller is in SWRAID mode. It is not supported for AHCI mode.

You can configure PCH controller SSDs in a storage profile policy. You can enable or disable protect configuration which saves the LUN configuration even after a service profile disassociation. You choose a controller mode. The PCH controller configuration supports only these two RAID options: RAID0 and RAID1. Use No RAID configuration option for AHCI mode where all the disks connected to the controller configured as JBOD disks. The configuration deployment happens as part of the storage profile association to a service profile process.

Cisco UCS Manager supports PCH managed internal SSDs in the following M4 servers:

- UCSC-C240-M4L
- UCSC-C240-M4SX

Cisco UCS Manager supports the following PCH managed SSDs on the M.2 card for all M5 servers:

- 240GB M.2 6G SATA SSD
- 960GB M.2 6G SATA SSD



Note You cannot have software RAID configuration in the controller definition and legacy boot mode configuration in boot policy together in M5 servers. Only UEFI boot mode is supported with software RAID configuration in the controller definition. This condition is applicable even if the drives are not used as boot drive.

For the PCH Controller Definition configuration in a Cisco UCS Manager boot policy two new devices exist to select: PCH LUN and PCH Disk. EmbeddedLocalLun represents the boot device in SWRAID mode and EmbeddedLocalDisk represent the boot devices in AHCI mode.

The system uses the same scrub policy is used to scrub supported SSDs. If the scrub is Yes, configured LUNs are destroyed as part of disassociation or re-discovery. If the scrub is No, configured LUNs are saved during disassociation and re-discovery.

Cisco UCS Manager supports firmware upgrade for the internal SSDs only when the PCH Controller is in SWRAID mode. It is not supported in AHCI mode.

FCH Controller Configuration

Fusion Controller Hub (FCH) SSD Controller Definition provides a local storage configuration in storage profiles for AMD based Cisco UCS C125 M5 Server. For AMD processor based servers, the PCH controller is referred to as FCH controller. The controller type remains as PCH in the Cisco UCS Manager GUI.

The FCH Controller works the same as PCH Controller except for the following differences:

- FCH is only in AHCI (JBOD) mode.



Note Cisco UCS Manager GUI shows the RAID support as **RAID0, RAID1**, however, Cisco UCS C125 M5 Server supports only AHCI mode.



Note You must re-acknowledge the server if you remove or insert disks managed by the PCH Controller.

- There are two FCH controllers:
 - First PCH controller manages SATA disks in the front panel (in the absence of a separate PCIe storage controller)
 - Second PCH controller manages the M.2 SSDs



Note For Cisco UCS C125 M5 Server, the PCH IDs are 3 and 4.



Note Further information and procedures related to PCH controller in this document are applicable to both Intel and AMD based servers.

Creating a Storage Profile PCH Controller Definition

The PCH Controller Definition provides a storage configuration in Storage Profiles where you can configure internal SSDs connected to a PCH controller. You create a name for the controller definition, specify whether you want the storage profile to retain the configuration even if the storage profile is disassociated from the service profile, and chose the RAID level to indicate the controller mode.

Procedure

- Step 1** In the **Navigation** pane, click **Storage > Storage Profiles**.
- Step 2** Choose the storage profile where you want to create the controller definition.
- Step 3** Click the **Controller Definitions** tab and then click **Add** at the bottom of the panel or right-click the storage profile and select **Create Controller Definition**.
- Step 4** In **Create Controller Definition** dialog box, configure the following information:

Name	Description
Name field	<p>The name of the storage controller.</p> <p>Note Once you save a PCH Controller Definition, you cannot modify the name from the General Tab Properties area.</p> <p>Enter up to 16 characters. You can use any alphanumeric characters. Special characters and spaces are not supported.</p>
Protect Configuration check box	<p>If checked, the storage profile retains the configuration even if the storage profile is disassociated from the service profile.</p> <p>Note If you disassociate the storage profile from a service profile with this option enabled, and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.</p>

Name	Description
RAID Level drop-down list	

Name	Description
	<p>This can be one of the following disk policy modes:</p> <ul style="list-style-type: none"> • Disable Local Storage—(Supported for PCH SSD Controller Definition) This disk policy mode is to disable the SATA AHCI Controller. This mode can be set only when disks are not present under the SATA AHCI controller. To re-enable this controller and to bring the controller back to its default value (AHCI), you can select No RAID or No Local Storage mode. • No Local Storage—(Supported for PCH SSD Controller Definition) For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk. • RAID 0 Striped—(Supported for PCH SSD Controller Definition) Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails. • RAID 1 Mirrored—(Supported for PCH SSD Controller Definition) Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives. • Any Configuration—(Supported for PCH SSD Controller Definition) For a server configuration that carries forward the local disk configuration without any changes. • No RAID—(Supported for PCH SSD Controller Definition) All the disks can be used individually without interdependency similar to JBOD disks. If you choose No RAID and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the No RAID mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the Inventory > Storage tab for the server. To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode. • RAID 5 Striped Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates. • RAID 6 Striped Dual Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.

Name	Description
	<ul style="list-style-type: none"> • RAID 10 Mirrored and Striped—(Not supported for PCH SSD Controller Definition) RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates. • RAID 50 Striped Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID 60 Striped Dual Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance. <p>Note Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.</p> <p>For RAID license information for a specific Cisco UCS server, see the Hardware Installation Guide for that server.</p>

Step 5

Click OK.

The new PCH Controller Definition appears in the navigation pane.

What to do next

For specific operating system software RAID driver installation procedures, see:

- *Installing LSI MegaSR Drivers for Windows and Linux* section in the [Cisco UCS C240 M4 Server Installation and Service Guide](#)
- *Installing LSI MegaSR Drivers For Windows and Linux* section in the [Cisco UCS C220 M5 Server Installation and Service Guide](#)
- *Installing LSI MegaSR Drivers For Windows and Linux* section in the [Cisco UCS C240 M5 Server Installation and Service Guide](#)
- *Installing LSI MegaSR Drivers For Windows and Linux* section in the [Cisco UCS C480 M5 Server Installation and Service Guide](#)



Note For Cisco UCS B200 M5 Server and Cisco UCS B480 M5 Server software RAID driver installation, follow the same procedure as any of the above M5 servers.

Modifying a Service Profile PCH Controller Definition

Before you begin

If you want to modify RAID level from **RAID 0 Striped** or **RAID 1 Mirrored** to **NO RAID**, then perform the following steps before starting the procedure:

1. Ensure that you have a scrub policy in the associated service profile. Refer *Creating a Service Profile with the Expert Wizard* in *Cisco UCS Manager Server Management Guide*.
2. Disassociate the server from the service profile. Refer *Disassociating a Service Profile from a Server or Server Pool* in *Cisco UCS Manager Server Management Guide*.

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** Expand **Storage Profiles** to the specific storage profile name that you want.
- Step 3** Expand **Controller Definitions** and click the specific controller definition that you want.
- Step 4** On the **General** tab, modify the following information:

Name	Description
Name field	<p>The name of the storage controller.</p> <p>Note Once you save a PCH Controller Definition, you cannot modify the name from the General Tab Properties area.</p> <p>Enter up to 16 characters. You can use any alphanumeric characters. Special characters and spaces are not supported.</p>
Protect Configuration check box	<p>If checked, the storage profile retains the configuration even if the storage profile is disassociated from the service profile.</p> <p>Note If you disassociate the storage profile from a service profile with this option enabled, and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.</p>

Name	Description
RAID Level drop-down list	

Name	Description
	<p>This can be one of the following disk policy modes:</p> <ul style="list-style-type: none"> • Disable Local Storage—(Supported for PCH SSD Controller Definition) This disk policy mode is to disable the SATA AHCI Controller. This mode can be set only when disks are not present under the SATA AHCI controller. To re-enable this controller and to bring the controller back to its default value (AHCI), you can select No RAID or No Local Storage mode. • No Local Storage—(Supported for PCH SSD Controller Definition) For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk. • RAID 0 Striped—(Supported for PCH SSD Controller Definition) Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails. • RAID 1 Mirrored—(Supported for PCH SSD Controller Definition) Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives. • Any Configuration—(Supported for PCH SSD Controller Definition) For a server configuration that carries forward the local disk configuration without any changes. • No RAID—(Supported for PCH SSD Controller Definition) All the disks can be used individually without interdependency similar to JBOD disks. If you choose No RAID and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the No RAID mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the Inventory > Storage tab for the server. To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode. • RAID 5 Striped Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates. • RAID 6 Striped Dual Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.

Name	Description
	<ul style="list-style-type: none"> • RAID 10 Mirrored and Striped—(Not supported for PCH SSD Controller Definition) RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates. • RAID 50 Striped Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID 60 Striped Dual Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance. <p>Note Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.</p> <p>For RAID license information for a specific Cisco UCS server, see the Hardware Installation Guide for that server.</p>

Step 5 Click OK.

The system displays whether it saved the modified PCH Controller Definition successfully.

What to do next

If you had disassociated the server from the service profile to modify RAID level from **RAID 0 Striped** or **RAID 1 Mirrored** to **NO RAID**, then perform the following steps:

1. For any M4 server, re-acknowledge the server. Refer *Reacknowledging a Rack-Mount Server* in *Cisco UCS Manager Server Management Guide*.
2. Associate the service profile to the server. Refer *Associating a Service Profile with a Server or Server Pool* in *Cisco UCS Manager Server Management Guide*.

Deleting a Storage Profile PCH Controller Definition

Procedure

Step 1 In the **Navigation** pane, click the **Storage** tab.

- Step 2** Expand **Storage Profiles**.
- Step 3** Expand **PCH Controller Definitions**.
- Step 4** In the **Navigation** pane, click the specific **Controller Definition** that you want to delete.
- Step 5** In the **General** tab **Actions** area, click **Delete**.
- Step 6** Confirm whether you want to delete the definition.
- The system displays whether it deleted the definition successfully. If not, see [PCH Controller Definition Configuration Troubleshooting](#), on page 139
- Step 7** If successfully deleted, click OK.
-

PCH Controller Definition Configuration Troubleshooting

PCH Controller Definition Creation

Unsuccessful PCH Controller Definition configuration exists under the following situations:

- You try to configure a Controller definition for an unsupported server model
- You try to use the legacy local disk configuration policy and also configures the PCH storage in storage profile
- You try to configure same controller using storage profile controller definition and also by using storage profile Local LUN configuration interface
- If the **Protect Configuration** checkbox is ON and you configured the RAID Type differently than the deployed configuration in SWRAID mode.
- If the **Protect Configuration** checkbox is ON and the RAID Type does not match the present controller mode.



Warning Any configuration change in the PCH storage configuration (like Controller mode change, RAID level change or controller qualifier change) for an already associated server triggers a PNUOS boot to happen causing a down time for the host OS.

Boot Policy

A configuration error occurs for any of the following cases:

- You select PCH Disk in boot policy but the primary or secondary target path slot number did not match with any of the inventoried internal SSD slot numbers.
- You select both PCH LUN and PCH Disk at the same time in the boot policy.

Firmware

For an incompatible software combination, there will not be any configuration error at the time of association. However, the storage configuration for the PCH SSD controller might fail or might not be deployed during association if you do not use the supported software combinations. Also, booting from the PCH SSD controller internal SSD might fail at the end of association for an incompatible software combination.

Migrating an M.2 Module

Migrating an M.2 module in SWRAID

Perform this procedure to migrate an M.2 module in SWRAID mode to a destination server:

Before you begin

Only UEFI boot mode is supported with software RAID configuration in the controller definition. This condition is applicable even if the drives are not used as boot drive. Ensure that the source and destination server boot mode is set to UEFI and controller definition is configured as same SWRAID (R0/R1).

Procedure

Step 1 Gracefully shut down the server.

Step 2 Physically remove the M.2 module.

The boot mode in the source server for SWRAID M.2 controller configuration in the source server has to be UEFI. Configure the boot policy of destination server with UEFI boot parameters under embedded disk.

Step 3 Insert the disk in the M.2 module in the destination server.

Step 4 Power on the server.

Step 5 Re-acknowledge the server.

Migrating an M.2 Module in AHCI Mode

Perform this procedure to migrate an M.2 module in NORAID mode to a destination server:

Before you begin

- If the source server is in legacy boot mode, ensure that the destination server is also in legacy boot mode and controller definition is configured as **NORAID**.
- If the source server is in UEFI boot mode, ensure that the destination server is also in UEFI boot mode and controller definition is configured as **NORAID**.

Procedure

Step 1 Gracefully shut down the server.

Step 2 Physically remove the M.2 module.

Step 3 Do one of the following:

- If the disk under M.2 controller had boot mode as UEFI on the source server, configure the boot policy of the destination server with UEFI boot parameters.
- If the disk under M.2 controller had boot mode as legacy on the source server, configure the boot policy of the destination server as legacy mode

Step 4 Insert the M.2 module in the destination server.

Step 5 Power on the server.

Step 6 Re-acknowledge the server.

Note If the disk is faulty, the server shows the disk status as **Not Detected**. Perform [Replacing a Faulty M.2 Disk, on page 142](#) to replace the faulty disk.

Migrating a SWRAID Disk

Perform this procedure to migrate a M.2 disk in SWRAID mode to a destination server:

Before you begin

Only UEFI boot mode is supported with software RAID configuration in the controller definition. This condition is applicable even if the drives are not used as boot drive. Ensure that the source and destination server boot mode is set to UEFI and controller definition is configured as same SWRAID (R0/R1).

Procedure

Step 1 Gracefully shut down the server.

Step 2 Physically remove the M.2 module and extract the disk.

If the disk is used as SWRAID in the source server the boot mode has to be UEFI and configure boot policy of destination server with UEFI boot parameters under embedded disk.

Step 3 Insert the disk in the M.2 module in the destination server.

Step 4 Power on the server.

Step 5 Re-acknowledge the server.

Note The **Drive State** of the disk should show as **Online**. If the disk is faulty, the sever fails to detect the disk or the **Drive State** shows as **BAD (or FAILED)** instead of **Online**. Perform [Replacing a Faulty M.2 Disk, on page 142](#) to replace the faulty disk.

Migrating a JBOD Disk in AHCI Mode

Perform this procedure to migrate a JBOD disk in NORAIID mode to a destination server:

Before you begin

- If the source server is in legacy boot mode, ensure that the destination server is also in legacy boot mode and controller definition is configured as **NORAIID**.
- If the source server is in UEFI boot mode, ensure that the destination server is also in UEFI boot mode and controller definition is configured as **NORAIID**.

Procedure

- Step 1** Gracefully shut down the server.
- Step 2** Physically remove the module and extract the M.2 disk.
- Step 3** Do one of the following:
- If the disk under M.2 controller had boot mode as UEFI on the source server, configure the boot policy of the destination server with UEFI boot parameters.
 - If the disk under M.2 controller had boot mode as legacy on the source server, configure the boot policy of the destination server as legacy mode
- Step 4** Insert the M.2 disk in the M.2 module on the destination server.
- Step 5** Power on the server.
- Step 6** Re-acknowledge the server.
-

Replacing a Faulty M.2 Disk

Perform this procedure to replace a faulty M.2 disk.

Before you begin

Ensure that the SWRAID controller definition is configured and the replacement disk formatted empty drive.

Procedure

- Step 1** Gracefully power down the server.
- Step 2** Physically remove the faulty M.2 drive. Use the **Serial Number** and **Disk Slot** to identify the faulty disk.
- Step 3** Insert the replacement M.2 drive.
- Step 4** Power on the server.
- Step 5** Wait for the disk to rebuild and then re-acknowledge the server.
- Note** SWRAID rebuild may take anywhere between 35 to 75 minutes depending on the disk size, disk speed, OS content, and other parameters.
- AHCI is a NORAIID configuration and hence rebuild is not applicable.
- Note** After replacing the faulty M.2 drive, the operability state and drive-state of the drive in other slot change to Degraded and Rebuilding. To bring back the drive to normal state, decommission and recommitment the blade.
-

Associating a Storage Profile with an Existing Service Profile

You can associate a storage profile with an existing service profile or a new service profile. See [Creating a Service Profile with the Expert Wizard, on page 169](#).

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers** > **Service Profiles**.
 - Step 3** Expand the node for the organization that contains the service profile that you want to associate with a storage profile.
 - Step 4** Choose the service profile that you want to associate with a storage profile.
 - Step 5** In the **Work** pane, click the **Storage** tab.
 - Step 6** Click the **LUN Configuration** subtab.
 - Step 7** In the **Actions** area, click **Modify Storage Profile**. The **Modify Storage Profile** dialog box appears.
 - Step 8** Click the **Storage Profile Policy** tab.
 - Step 9** To associate an existing storage profile with this service profile, select the storage profile that you want to associate from the **Storage Profile** drop-down list, and click **OK**. The details of the storage profile appear in the **Storage Items** area.
 - Step 10** To create a new storage profile and associate it with this service profile, click **Create Storage Profile**, complete the required fields, and click **OK**. [Creating a Storage Profile, on page 118](#) provides more information on creating a new storage profile.
 - Step 11** (Optional) To dissociate the service profile from a storage profile, select **No Storage Profile** from the **Storage Profile** drop-down list, and click **OK**.
-

Configuring Storage Profiles

Importing Foreign Configurations for a RAID Controller on a Blade Server

Before you begin

In a set up with Cisco boot optimized M.2 RAID controller, Cisco UCS Manager does not recognize which configuration to import if you connect two drives with different foreign configurations. You must first clear the configuration on one drive using the HII menu. For more information on clearing the configuration using the HII menu, see [Configuration Guides](#).

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Choose the server of the RAID controller for which you want to import foreign configurations.

- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **Controller** subtab.
- Step 6** In the **Actions** area, click **Import Foreign Configuration**.
-

Importing Foreign Configurations for a RAID Controller on a Rack Server

Before you begin

In a set up with Cisco boot optimized M.2 RAID controller, Cisco UCS Manager does not recognize which configuration to import if you connect two drives with different foreign configurations. You must first clear the configuration on one drive using the HII menu. For more information on clearing the configuration using the HII menu, see [Configuration Guides](#).

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Rack Mounts** > **Servers**.
- Note** For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure** *rack_enclosure_number* > **Servers**.
- Step 3** Choose the server of the RAID controller for which you want to import foreign configurations.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **Controller** subtab.
- Step 6** In the **Actions** area, click **Import Foreign Configuration**.
-

Configuring Local Disk Operations on a Blade Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Choose the server for which you want to configure local disk operations.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **Disks** subtab.
- Step 6** Right-click the disk that you want and select one of the following operations:
- **Clear Foreign Configuration State**—Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration.
 - **Set Unconfigured Good**—Specifies that the local disk can be configured.
 - **Set Prepare For Removal**—Specifies that the local disk is marked for removal from the chassis.

- **Set Undo Prepare For Removal**—Specifies that the local disk is no longer marked for removal from the chassis.
- **Mark as Dedicated Hot Spare**—Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives.
- **Remove Hot Spare**—Specifies that the local disk is no longer a hot spare.
- **Set JBOD to Unconfigured Good**—Specifies that the new local disk can be configured after being marked as Unconfigured Good.

Configuring Local Disk Operations on a Rack Server

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment > Rack Mounts > Servers**.

Note For Cisco UCS C125 M5 Servers, expand **Equipment > Rack Mounts > Enclosures > Rack Enclosure *rack_enclosure_number* > Servers**.

Step 3 Choose the server for which you want to configure local disk operations.

Step 4 In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.

Step 5 Click the **Disks** subtab.

Step 6 Right-click the disk that you want and select one of the following operations:

- **Clear Foreign Configuration State**—Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration.
- **Set Unconfigured Good**—Specifies that the local disk can be configured.
- **Set Prepare For Removal**—Specifies that the local disk is marked for removal.
- **Set Undo Prepare For Removal**—Specifies that the local disk is no longer marked for removal.
- **Mark as Dedicated Hot Spare**—Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives.
- **Remove Hot Spare**—Specifies that the local disk is no longer a hot spare.
- **Set JBOD to Unconfigured Good**—Specifies that the new local disk can be configured after being marked as **Unconfigured Good**.

Configuring Local Disk Operations

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment > Chassis > Chassis Number**

Step 3 In the **Work** pane, click the **Storage** tab.

- Step 4** Click the **Disks** subtab.
- Step 5** Right-click the disk that you want and select one of the following operations:
- **Clear Foreign Configuration State**—Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration.
 - **Set Unconfigured Good**—Specifies that the local disk can be configured.
 - **Set Prepare For Removal**—Specifies that the local disk is marked for removal from the chassis.
 - **Set Undo Prepare For Removal**—Specifies that the local disk is no longer marked for removal from the chassis.
 - **Mark as Dedicated Hot Spare**—Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives.
 - **Remove Hot Spare**—Specifies that the local disk is no longer a hot spare.
 - **Set JBOD to Unconfigured Good**—Specifies that the new local disk can be configured after being marked as Unconfigured Good.
-

Deleting an Orphan Virtual Drive

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > **Chassis Number**
- Step 3** In the **Work** pane, click the **Storage** tab.
- Step 4** Click the **LUNs** subtab.
- Step 5** Right-click the virtual drive that you want and select **Delete Orphaned LUN**.
A confirmation dialog box appears.
- Step 6** Click **Yes**.
-

Deleting an Orphan Virtual Drive on a Rack Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Rack Mounts** > **Servers**.
- Note** For Cisco UCS C125 M5 Servers, expand **Equipment** > **Rack Mounts** > **Enclosures** > **Rack Enclosure *rack_enclosure_number*** > **Servers**.
- Step 3** Choose the server for which you want to delete an orphan virtual drive.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **LUNs** subtab.
- Step 6** Right-click the virtual drive that you want and select **Delete Orphaned LUN**.

A confirmation dialog box appears.

- Step 7** Click **Yes**.
-

Renaming an Orphan Virtual Drive on a Blade Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Choose the server for which you want to rename an orphan virtual drive.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **LUNs** subtab.
- Step 6** Right-click the virtual drive that you want and select **Rename Referenced LUN**.
- Step 7** In the **Rename Referenced LUN** dialog box that appears, enter the new **LUN Name**.
- Step 8** Click **OK**.
-

Renaming an Orphan Virtual Drive on a Rack Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Rack Mounts > Servers**.
- Note** For Cisco UCS C125 M5 Servers, expand **Equipment > Rack Mounts > Enclosures > Rack Enclosure rack_enclosure_number > Servers**.
- Step 3** Choose the server for which you want to rename an orphan virtual drive.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **LUNs** subtab.
- Step 6** Right-click the virtual drive that you want and select **Rename Referenced LUN**.
- Step 7** In the **Rename Referenced LUN** dialog box that appears, enter the new **LUN Name**.
- Step 8** Click **OK**.
-

Boot Policy for Local Storage

You can specify the primary boot device for a storage controller as a local LUN or a JBOD disk. Each storage controller can have one primary boot device. However, in a storage profile, you can set only one device as the primary boot LUN.

Beginning with 4.0(4a), Cisco UCS Manager supports Cisco boot optimized M.2 Raid controller based off Marvell 88SE92xx PCIe to SATA 6Gb/s controller (UCS-M2-HWRAID). The controller supports only UEFI boot mode.

Local storage option in the boot policy supports the boot from the SATA drives in the Cisco boot optimized M.2 Raid controller.

Also, embedded local storage option in the boot policy supports the boot from the SATA drives in the Cisco boot optimized M.2 Raid controller. The primary and the secondary type boot specifically from the M.2 SATA drives.



Note Cisco UCS S3260 M3 server nodes do not support **Any** as an option for the **Local LUN Image Path** while adding a local LUN to the boot policy through the Cisco UCS Manager GUI. In the Cisco UCS Manager CLI, Cisco UCS S3260 M3 server nodes do not support the **local-any** command option.

Configuring the Boot Policy for an Embedded Local LUN



Note

- For a Cisco UCS S3260 M3 server node, before you migrate the boot policy from a local LUN or JBOD to an embedded local LUN or disk, you must delete the local storage configuration, wait for the association to complete and finally add the new local storage configuration. This will enable the boot operation from the PCH disks or LUNs.
- Specify one bootable LUN as either primary or secondary boot device. If you specify the bootable LUN as both primary and secondary boot devices, the boot policy will result in the service profile configuration error.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Policies**.
 - Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the root node.
 - Step 4** Select the boot policy that you want to configure.
 - Step 5** In the **Work** pane, click the **General** tab.
 - Step 6** Click the down arrows to expand the **Local Devices** area.
 - Step 7** Click **Add Embedded Local LUN** to configure the boot order of the local LUN.
 - Step 8** To configure the local LUN as the primary boot device, select **Primary**.
 - Step 9** In the **LUN Name** field, enter the name of the LUN to be configured as the primary boot device.
 - Step 10** Click **OK**.
-

Configuring the Boot Policy for an Embedded Local Disk



Note For UCSC-C125 server, if there is no separate PCIe storage controller, then do not configure boot policy for embedded local disk. Instead, use **Add Local Disk** option.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the root node.
- Step 4** Select the boot policy that you want to configure.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** Click the down arrows to expand the **Local Devices** area.
- Step 7** Click **Add Embedded Local Disk** to configure the local JBOD device as the primary boot device.
JBOD is supported only on the following servers:
- Cisco UCS B200 M3 blade server
 - Cisco UCS B260 M4 blade server
 - Cisco UCS B460 M4 blade server
 - Cisco UCS B200 M4 blade server
 - Cisco UCS C220 M4 rack-mount server
 - Cisco UCS C240 M4 rack-mount server
 - Cisco UCS C460 M4 rack-mount server
 - Cisco UCS S3260 M4 and M5 servers
- Step 8** In the **Disk Slot Number** field, enter the slot number of the JBOD disk to be configured as the primary boot device.
- Step 9** Click **OK**.
-

Local LUN Operations in a Service Profile

Preprovisioning a LUN Name

Preprovisioning a LUN name can be done only when the admin state of the LUN is **Undeployed**. If this LUN name exists and the LUN is orphaned, its is claimed by the service profile. If this LUN does not exist, a new LUN is created with the specified name.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles > Service_Profile_Name**.
 - Step 3** In the **Work** pane, click the **Storage** tab.
 - Step 4** Click the **LUN Configuration** tab.
 - Step 5** In the **Local LUNs** subtab, right-click the LUN for which you want to preprovision a LUN name and select **Pre-Provision LUN Name**.
 - Step 6** In the **Set Pre-Provision LUN Name** dialog box, enter the LUN name.
 - Step 7** Click **OK**.
-

Claiming an Orphan LUN

Claiming an orphan LUN can be done only when the admin state of the LUN is **Undeployed**. You can explicitly change the admin state of the LUN to **Undeployed** for claiming an orphan LUN.

If the LUN name is empty, set a LUN name before claiming it.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles > Service_Profile_Name**.
 - Step 3** In the **Work** pane, click the **Storage** tab.
 - Step 4** Click the **LUN Configuration** tab.
 - Step 5** In the **Local LUNs** subtab, right-click the LUN that you want to claim and select **Claim Orphan LUN**.
 - Step 6** In the **Claim Orphan LUN** dialog box that appears, select an orphaned LUN.
 - Step 7** Right-click the LUN and select **Set Admin State**.
 - Step 8** In the **Set Admin State** dialog box that appears, select **Undeployed** to undeploy a LUN and claim ownership.
 - Step 9** Click **OK**.
-

Deploying and Undeploying a LUN

You can deploy or undeploy a LUN. If the admin state of a local LUN is **Undeployed**, the reference of that LUN is removed and the LUN is not deployed.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles > Service_Profile_Name**.
- Step 3** In the **Work** pane, click the **Storage** tab.
- Step 4** Click the **LUN Configuration** tab.

- Step 5** In the **Local LUNs** subtab, right-click the LUN that you want to deploy or undeploy and select **Set Admin State**.
- Step 6** In the **Set Admin State** dialog box that appears, select **Online** to deploy a LUN or **Undeployed** to undeploy a LUN.
- Step 7** Click **OK**.
-

Renaming a Service Profile Referenced LUN

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles** > *Service_Profile_Name*.
- Step 3** In the **Work** pane, click the **Storage** tab.
- Step 4** Click the **LUN Configuration** tab.
- Step 5** In the **Local LUNs** subtab, right-click the LUN for which you want to rename the referenced LUN, and select **Rename Referenced LUN**.
- Step 6** In the **Rename Referenced LUN** dialog box that appears, enter the new name of the referenced LUN.
- Step 7** Click **OK**.
-



CHAPTER 12

Mini Storage

- [Mini Storage, on page 153](#)
- [Viewing Mini Storage Properties, on page 153](#)

Mini Storage

The mini storage slot is a new slot that is present on the Cisco UCS M5 blade and rack servers. This slot can be empty, populated with an SD storage module, or populated with an M.2 SATA module.



Note Cisco UCS Manager does not support micro-SD card.

The mini storage SD module consists of an in-built SD controller and two SD cardslots. These cards have RAID 1 capability.

The mini M.2 SATA module consists of two SATA slots. The PCH controller present on the server controls the SATA drives on this module.

Starting with 4.0(4a), Cisco UCS Manager supports Cisco boot optimized M.2 Raid controller based off Marvell 88SE92xx PCIe to SATA 6Gb/s controller (UCS-M2-HWRAID) for mini storage.

You can use Cisco UCS Manager to inventory and manage the mini storage modules.

Viewing Mini Storage Properties

Mini storage modules are supported only on M5 and higher servers.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Choose the server for which you want to view mini storage properties.
- Step 4** In the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Motherboard** subtab.

- Step 6** In the **Mini Storage** area, click the Expand icon to open that area.
Mini Storage properties are displayed.
-



CHAPTER 13

Configuring SD Card Support

- [FlexFlash Secure Digital Card Support, on page 155](#)
- [FlexUtil Secure Digital Card Support, on page 160](#)

FlexFlash Secure Digital Card Support

Overview

Cisco UCS B-Series, C-Series M3 and M4 and higher, and S-Series M4 servers support internal Secure Digital (SD) memory cards.

The SD cards are hosted by the Cisco Flexible Flash storage controller, a PCI-based controller which has two slots for SD cards. The cards contain a single partition called HV. When FlexFlash is enabled, Cisco UCS Manager displays the HV partition as a USB drive to both the BIOS and the host operating system.


You can populate one or both the SD card slots that are provided. If two SD cards are populated, you can use them in a mirrored mode.



Note Do not mix different capacity cards in the same server.

The SD cards can be used to store operating system boot images or other information. The following figure illustrates the SD card slots.

Figure 2: SD Card Slots



FlexFlash is disabled by default. You can enable FlexFlash in a local disk policy used in a service profile. When FlexFlash is enabled in a local disk policy, and the server is capable of supporting SD cards, the FlexFlash controller is enabled during service profile association. If a server is not capable of supporting SD cards or has an older CIMC version, a config failure message is displayed.

If you disable FlexFlash in a supported server, the Hypervisor or HV partition is immediately disconnected from the host. The FlexFlash controller will also be disabled as part of a related service profile disassociation.

The FlexFlash controller supports RAID-1 for dual SD cards. The FlexFlash scrub policy erases the HV partition in both cards, and brings the cards to a healthy RAID state.

You can configure new SD cards in a RAID pair and format them using one of the following methods:

- Format the SD cards. [Formatting the SD Cards, on page 159](#) provides detailed information.
- For an associated server, create a FlexFlash scrub policy and disassociate the service profile from the server. For an unassociated server, create a FlexFlash scrub policy and reacknowledge the server after modifying the default scrub policy.

The *Scrub Policy Settings* section in the *Cisco UCS Manager Server Management Guide* provides more details about the usage of the scrub policy.



Note Disable the scrub policy as soon as the pairing is complete.

To boot from the HV partition, the SD card must be present in the boot policy used in the service profile.

FlexFlash Firmware Management

The FlexFlash controller firmware is bundled as part of the CIMC image. When you upgrade the CIMC, if a newer firmware version is available for the FlexFlash controller, the controller can no longer be managed, and the FlexFlash inventory displays the **Controller State** as **Waiting For User Action** and the **Controller Health** as **Old Firmware Running**. To upgrade the FlexFlash controller firmware, you need to perform a board controller update. For more information, see the appropriate *Cisco UCS B-Series Firmware Management*

Guide, available at the following URL:

http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html.

Limitations for the Cisco Flexible Flash Storage Controller:

- The Cisco Flexible Flash storage controller only supports 16 GB, 32 GB, and 64 GB SD cards.



Note The 16 GB and 32 GB cards are supported only on the B200-M3 blade servers, and the 64 GB SD cards are supported only on the B200-M4 and M5 blade servers.

- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- Some Cisco UCS C-Series rack-mount servers have SD cards with four partitions: HV, HUU, SCU, and Drivers. Only the HV partition is visible in Cisco UCS Manager. You can migrate a four-partition SD card to a single HV partition card with a FlexFlash scrub policy.
- The FlexFlash controller does not support RAID-1 sync (mirror rebuild). If the SD cards are in a degraded RAID state, or if any metadata errors are reported by the controller, you must run the FlexFlash scrub policy to pair the cards for RAID. For more information about the FlexFlash scrub policy, see [Server-Related Policies](#). The following conditions might result in degraded RAID or metadata errors:
 - Inserting a new or used SD card in one slot, when the server already has an SD card populated in the second slot.
 - Inserting two SD cards from different servers.
- The server firmware version must be at 2.2(1a) or higher.

FlexFlash FX3S Support

Beginning with Release 2.2(3), Cisco UCS Manager allows additional FlexFlash support with the FX3S controller. The FX3S controller is present on the following servers:

- Cisco UCS B200 M4 and M5 blade server
- Cisco UCS C220 M4 and M5 rack server
- Cisco UCS C240 M4 and M5 rack server
- C480 M5 rack server
- C480 M5 ML blade server
- B480 M5 blade server
- Cisco UCS C125 M5 Server

FlexFlash operations with the FX3S control are similar to those with the Cisco Flexible Flash storage controller. FlexFlash is disabled by default, and is enabled using a local disk policy. You can also reset the controller, format the SD cards, and enable automatic synchronization of your paired SD cards.

The SD cards for the FX3S controller contain a single partition called Hypervisor.

Limitations for the Cisco FX3S Controller:

- The FX3S controller supports only 32 GB and 64 GB SD cards. 16 GB cards are not supported.
- The FX3S controller supports 128 GB cards on M5 blades and above.
- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- The server firmware version must be at 2.2(3a) or higher.

Enabling FlexFlash SD Card Support

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand **Local Disk Config Policies** and choose the local disk config policy for which you want to enable FlexFlash support.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **FlexFlash State** field, click the **Enable** radio button.
- Step 7** In the **FlexFlash Removable State** field, select the removable state. Click the **Yes** option if you need to define the Flex Flash SD card as removable. Check your hypervisor requirements for the required removable state setting.
- **No Change**(default) - Use this option if the hypervisor does not require a preset state for the SD card.
 - **No** - Use this option if the hypervisor requires a non-removable state for the SD card.
- Step 8** Click **Save Changes**.
-

Disabling FlexFlash SD Card Support

This procedure describes how to disable the FlexFlash capability in a local disk policy.

Procedure

- Step 1** In the **Navigation** pane, click on the **Servers** tab.
- Step 2** Select **Policies** from the **Filter** dropdown list.

- Step 3** Expand the **Local Disk Config Policies** tree.
 - Step 4** Highlight the policy for which you want to disable FlexFlash.
 - Step 5** Click the **Events** tab of the task pane, and select the **Disable** radio button next to **FlexFlash State**.
 - Step 6** If you need to replace the FlexFlash SD card, select the **Yes** radio button in the **FlexFlash Removable State** field. Make sure the SD cards are not in use before changing the FlexFlash Removable State.
 - Step 7** Click **Save Changes**.
-

Enabling Auto-Sync

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
 - Step 3** Click the server for which you want to enable auto-sync.
 - Step 4** In the **Work** pane, click the **Inventory** tab.
 - Step 5** Click the **Storage** subtab.
 - Step 6** In the **Actions** area, click **Enable Auto-sync**.
 - Step 7** In the **Enable Auto-sync** dialog box, choose the **Admin Slot Number** for the SD card that you want to use as the primary.
 - Step 8** Click **OK**.
-

Formatting the SD Cards

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
 - Step 3** Click the server for which you want to format the SD cards.
 - Step 4** In the **Work** pane, click the **Inventory** tab.
 - Step 5** Click the **Storage** subtab.
 - Step 6** In the **Actions** area, click **Format SD Cards**.
 - Step 7** Click **Yes** to format the SD cards.
-

Resetting the FlexFlash Controller

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
 - Step 3** Click the server for which you want to reset the FlexFlash controller.
 - Step 4** In the **Work** pane, click the **Inventory** tab.
 - Step 5** Click the **Storage** subtab.
 - Step 6** In the **Actions** area, click **Reset FlexFlash Controller**.
 - Step 7** Click **Yes** to reset the FlexFlash controller.
-

FlexUtil Secure Digital Card Support

The C-Series M5 Rack-Mount servers support a Micro-SD (FlexUtil) memory card for storage. UCS Manager however does not provide management support for Micro-SD card.



CHAPTER 14

Direct Attached Storage

- [Direct Attached Storage, on page 161](#)
- [Fibre Channel Switching Mode, on page 162](#)
- [Configuring Fibre Channel Switching Mode, on page 162](#)
- [Creating a Storage VSAN, on page 163](#)
- [Creating a VSAN for Fibre Channel Zoning, on page 164](#)
- [Configuring a Fibre Channel Storage Port, on page 166](#)
- [Configuring Fibre Channel Zoning, on page 167](#)

Direct Attached Storage

A typical Direct Attached Storage (DAS) system is made of a data storage device (for example enclosures holding a number of hard disk drives connected directly to a computer through a host bus adapter (HBA). Between those two points there is no network device (like a switch or router).

The main protocols used for DAS connections are ATA, SATA, eSATA, SCSI, SAS, USB, USB 3.0, IEEE 1394 and Fibre Channel.

Cisco UCS Manager allows you to have DAS with without the need for a SAN switch to push the zoning configuration.

The DAS configuration described assumes that the physical cables are already connected between the storage array ports and the Fabric Interconnects.

Cisco UCS 6400 Series Fabric Interconnects do not support 8 Gbps direct-attached FC connectivity (FC uplink ports or FC storage ports) without fill-pattern set to IDLE. When migrating to Cisco UCS 6400 Series Fabric Interconnects from Cisco UCS 6200 Series Fabric Interconnects, do one of the following:

- Use a SAN switch between the Cisco UCS 6400 Series Fabric Interconnect and the storage array with 8 GB FC connectivity.
- Upgrade the storage array to 16 GB or 32 GB FC connectivity.

Fibre Channel Switching Mode

The Fibre Channel switching mode determines how the fabric interconnect behaves as a switching device between the servers and storage devices. The fabric interconnect operates in either of the following Fibre Channel switching modes:

End-Host Mode

End-host mode allows the fabric interconnect to act as an end host to the connected fibre channel networks, representing all servers (hosts) connected to it through virtual host bus adapters (vHBAs). This behavior is achieved by pinning (either dynamically pinned or hard pinned) vHBAs to Fibre Channel uplink ports, which makes the Fibre Channel ports appear as server ports (N-ports) to the rest of the fabric. When in end-host mode, the fabric interconnect avoids loops by denying uplink ports from receiving traffic from one another.

End-host mode is synonymous with N Port Virtualization (NPV) mode. This mode is the default Fibre Channel Switching mode.



Note When you enable end-host mode, if a vHBA is hard pinned to an uplink Fibre Channel port and this uplink port goes down, the system cannot repin the vHBA, and the vHBA remains down.

Switch Mode

Switch mode is the traditional Fibre Channel switching mode. Switch mode allows the fabric interconnect to connect directly to a storage device. Enabling Fibre Channel switch mode is useful in Pod models where there is no SAN (for example, a single Cisco UCS domain that is connected directly to storage), or where a SAN exists (with an upstream MDS).

Switch mode is not the default Fibre Channel switching mode.



Note In Fibre Channel switch mode, SAN pin groups are irrelevant. Any existing SAN pin groups are ignored.

Configuring Fibre Channel Switching Mode



Important

When you change the Fibre Channel switching mode, Cisco UCS Manager logs you out and restarts the fabric interconnect. For a cluster configuration, Cisco UCS Manager restarts both fabric interconnects simultaneously in Cisco UCS Manager Release 3.1(1) and earlier releases. In Cisco UCS Manager Release 3.1(2), when the Fibre Channel switching mode is changed, the UCS fabric interconnects reload sequentially. In Cisco UCS Manager Release 3.1(3), and later releases, the subordinate fabric interconnect reboots first as a result of the change in switching mode. The primary fabric interconnect reboots only after you acknowledge it in **Pending Activities**. The primary fabric interconnect can take several minutes to complete the change in Fibre Channel switching mode and become system ready.



Note Reloading of both the fabric interconnects simultaneously will cause a system-wide downtime for approximately 10-15 minutes.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** In the **Work** pane, click the **General** tab.
- Step 4** In the **Actions** area of the **General** tab, click one of the following links:
- **Set Fibre Channel Switching Mode**
 - **Set Fibre Channel End-Host Mode**

The link for the current mode is dimmed.

- Step 5** In the dialog box, click **Yes**.
- Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager GUI.
-

Creating a Storage VSAN



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, expand **SAN** > **Storage Cloud**.
- Step 3** In the **Work** pane, click the **VSANs** tab.
- Step 4** On the icon bar to the right of the table, click +.
- If the + icon is disabled, click an entry in the table to enable it.
- Step 5** In the **Create VSAN** dialog box, complete the required fields.
- Step 6** Click **OK**.
- Cisco UCS Manager GUI adds the VSAN to one of the following **VSANs** nodes:

- The **Storage Cloud > VSANs** node for a storage VSAN accessible to both fabric interconnects.
- The **Storage Cloud > Fabric_Name > VSANs** node for a VSAN accessible to only one fabric interconnect.

Creating a VSAN for Fibre Channel Zoning



Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, click the **SAN** node.
- Step 3** In the **Work** pane, click the **SAN Uplinks Manager** link on the **SAN Uplinks** tab.
The SAN Uplinks Manager opens in a separate window.
- Step 4** In the SAN Uplinks Manager, click the **VSAN** tab.
You can create the VSAN on any of the subtabs. However, if you use the **All** subtab, you can view all of the configured VSANs in the table.
- Step 5** On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.
- Step 6** In the **Create VSAN** dialog box, complete the following fields:

Name	Description
Name field	The name assigned to the network. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
FC Zoning field	<p>Click the radio button to determine whether Cisco UCS Manager configures Fibre Channel zoning for the Cisco UCS domain. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The upstream switch handles Fibre Channel zoning, or Fibre Channel zoning is not implemented for the Cisco UCS domain. Cisco UCS Manager does not configure Fibre Channel zoning. • Enabled—Cisco UCS Manager configures and controls Fibre Channel zoning for the Cisco UCS domain. <p>Note If you enable Fibre Channel zoning through Cisco UCS Manager, do not configure the upstream switch with any VSANs that are being used for Fibre Channel zoning.</p>
Type radio button	<p>Click the radio button to determine how the VSAN should be configured. This can be one of the following:</p> <ul style="list-style-type: none"> • Common/Global—The VSAN maps to the same VSAN ID in all available fabrics. • Fabric A—The VSAN maps to the a VSAN ID that exists only in fabric A. • Fabric B—The VSAN maps to the a VSAN ID that exists only in fabric B. • Both Fabrics Configured Differently—The VSAN maps to a different VSAN ID in each available fabric. If you choose this option, Cisco UCS Manager GUI displays a VSAN ID field and a FCoE VLAN field for each fabric.
VSAN ID field	<p>The unique identifier assigned to the network.</p> <p>The ID can be between 1 and 4078, or between 4080 and 4093. 4079 is a reserved VSAN ID. In addition, if you plan to use FC end-host mode, the range between 3840 to 4079 is also a reserved VSAN ID range.</p>

Name	Description
FCoE VLAN field	<p>The unique identifier assigned to the VLAN used for Fibre Channel connections.</p> <p>VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:</p> <ul style="list-style-type: none"> • After an upgrade to Cisco UCS, Release 2.0—The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use. • After a fresh install of Cisco UCS, Release 2.0—The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049. <p>For FIP-capable, converged network adapters, such as the Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, the named VSAN must be configured with a named VLAN that is not the native VLAN for the FCoE VLAN ID. This configuration ensures that FCoE traffic can pass through these adapters.</p>

Step 7 Click **OK**.

Configuring a Fibre Channel Storage Port

This task describes only one method of configuring FC storage ports. You can also configure FC storage ports from the **General** tab for the port.

Before you begin

The Fibre Channel switching mode must be set to Switching for these ports to be valid. The storage ports cannot function in end-host mode.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
- Step 3** Expand the **Expansion Module** node.
- Step 4** Click one or more of the ports under the **FC Ports** node.
- Step 5** Right-click the selected port or ports and choose **Configure as FC Storage Port**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
- Step 7** Click **OK**.

Configuring Fibre Channel Zoning



Note This procedure provides a high level overview of the steps required to configure a Cisco UCS domain for Fibre Channel zoning that is controlled by Cisco UCS Manager. You must ensure that you complete all of the following steps.

Procedure

	Command or Action	Purpose
Step 1	If you have not already done so, disconnect the fabric interconnects in the Cisco UCS domain from any external Fibre Channel switches, such as an MDS.	
Step 2	If the Cisco UCS domain still includes zones that were managed by the external Fibre Channel switch, run the clear-unmanaged-fc-zone-all command on every affected VSAN to remove those zones.	This functionality is not currently available in the Cisco UCS Manager GUI. You must perform this step in the Cisco UCS Manager CLI.
Step 3	Configure the Fibre Channel switching mode for both fabric interconnects in Fibre Channel Switch mode.	You cannot configure Fibre Channel zoning in End-Host mode. See Configuring Fibre Channel Switching Mode , on page 162.
Step 4	Configure the Fibre Channel and FCoE storage ports that you require to carry traffic for the Fibre Channel zones.	See Configuring an Ethernet Port as an FCoE Storage Port , on page 14 and Configuring a Fibre Channel Storage Port , on page 15.
Step 5	Create one or more VSANs and enable Fibre Channel zoning on all VSANs that you require to carry traffic for the Fibre Channel zones.	For a cluster configuration, we recommend that you create the VSANs that you intend to include in a Fibre Channel zone in the SAN Uplinks Manager and use the common/global configuration to ensure they are accessible to both fabric interconnects. See Creating a VSAN for Fibre Channel Zoning , on page 33.
Step 6	Create one or more Fibre Channel storage connection policies.	You can perform this step when you configure Fibre Channel zoning in the service profiles, if you prefer. See Creating a VSAN for Fibre Channel Zoning , on page 33.
Step 7	Configure zoning in service profiles or service profile templates for servers that need to communicate through Fibre Channel zones.	Complete the following steps to complete this configuration: <ul style="list-style-type: none"> • Enable zoning in the VSAN or VSANs assigned to the vHBAs. See Creating a VSAN for Fibre Channel Zoning, on page 33

	Command or Action	Purpose
		<ul style="list-style-type: none"> Configure one or more vHBA initiator groups. See Creating a Service Profile with the Expert Wizard, on page 169.

Creating a Fibre Channel Storage Connection Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the **Storage Connection Policies** node and choose **Create Storage Connection Policy**.
- Step 5** In the **Create Storage Connection Policy** dialog box, complete the following fields:

Name	Description
Name field	<p>The name of the policy.</p> <p>This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</p>
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p>

- Step 6** In the **Zoning Type** field, click one of the following radio buttons:
- None**—Cisco UCS Manager does not configure Fibre Channel zoning.
 - Single Initiator Single Target**—Cisco UCS Manager automatically creates one zone for each vHBA and storage port pair. Each zone has two members. We recommend that you configure this type of zoning unless you expect the number of zones to exceed the maximum supported.
 - Single Initiator Multiple Targets**—Cisco UCS Manager automatically creates one zone for each vHBA. We recommend that you configure this type of zoning if you expect the number of zones to reach or exceed the maximum supported.
- Step 7** In the **FC Target Endpoints** table, click + on the icon bar to the right of the table.

If the + icon is disabled, click an entry in the table to enable it.

Step 8 In the **Create FC Target Endpoint** dialog box, complete the following fields and then click **OK**:

Name	Description
WWPN field	The WWPN (WWN) assigned to the physical target port on the Fibre Channel or FCoE storage array that the server uses to access the LUNs configured on the storage array.
Description field	A description of the target endpoint. We recommend that you include information about the port, LUNs, or storage array to which the target endpoint connects. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Path field	The fabric interconnect used for communications with the target endpoint.
Select VSAN drop-down list	The VSAN used for communications with the target endpoint.
Create VSAN link	Click this link if you want to create a VSAN.

Repeat this step until you have created all desired target endpoints for the policy.

Step 9 After you have created all desired target endpoints for the policy, click **OK**.

Creating a Service Profile with the Expert Wizard

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.
- Step 3** Expand the node for the organization where you want to create the service profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the organization and select **Create Service Profile (expert)**.
- Step 5** In the **Identify Service Profile** panel, specify the service profile **Name**, UUID assignment and click **Next**.
You can provide an optional description for this service profile. If the UUID is not available, you can also create a UUID Suffix Pool from this panel.
- Note** To create a service profile quickly, you can click **Finish** after specifying the name. Cisco UCS Manager creates a new service profile with the specified name and all system default values.
- Step 6** (Optional) In the **Networking** panel, specify the required information for the **Dynamic vNIC Connection Policy** and **LAN Connectivity** sections, then click **Next**.

You can create a dynamic vNIC connection policy and LAN connectivity policy from this panel.

Step 7 (Optional) In the **Storage** panel, specify the SAN configuration information such as, **Local Storage Policy**, **SAN Connectivity**, **WWNN** and **VSAN**, then click **Next**.

You can create a local disk configuration policy and SAN connectivity policy from this panel.

Step 8 (Optional) In the **Zoning** panel, specify the required zoning information, then click **Next**.

You can create the vHBA initiator groups from this panel.

Step 9 (Optional) In the **vNIC/vHBA Placement** panel, specify the placement method and PCI order, then click **Next**.

You can create a placement policy from this panel.

Step 10 (Optional) In the **Server Boot Order** panel, specify the **Boot Policy** from the drop-down list, then click **Next**.

You can create a boot policy from this panel.

Step 11 (Optional) In the **Maintenance Policy** panel, specify the maintenance policy, then click **Next**.

You can create a new maintenance policy and specify a maintenance schedule from this panel.

Step 12 (Optional) In the **Server Assignment** panel, specify the **Server Assignment** from the drop down list and the power state to apply on assignment, then click **Next**.

You can create a server pool or a host firmware package from this panel.

Step 13 (Optional) In the **Operational Policies** panel, specify the system operational information such as, **BIOS Configuration**, **External IPMI Management Configuration**, **Management IP Address**, **Monitoring Configuration (Thresholds)**, **Power Control Policy Configuration**, and **Scrub Policy**, then click **Finish**.

Note To set up an Outband IPv4 address or an Inband IPv4 or IPv6 address, click the respective tabs and complete the required fields.

If you do not find the policies you need for each of these configurations, you can create them from this panel.

Associating a Service Profile with a Server or Server Pool

Follow this procedure if you did not associate the service profile with a blade server or server pool when you created it, or to change the blade server or server pool with which a service profile is associated.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers > Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile that you want to associate with a new server or server pool.

If the system does not include multi tenancy, expand the **root** node.

Step 4 Right-click the service profile you want to associate with a server and select **Associate Service Profile**.

Step 5 In the **Associate Service Profile** dialog box, select one of the following options:

Option	Description
Server Pool	Select a server pool from the drop-down list. Cisco UCS Manager assigns a server from this pool to the service profile. Continue with Step 7.
Server	Navigate to the desired available server in the navigation tree and select the server which will be assigned to the service profile. Continue with Step 7.
Custom Server	Specifies the chassis and slot that contains the server that will be assigned to the service profile. If the server is not in the slot or is otherwise unavailable, the service profile will be associated with the server when it becomes available. Continue with Step 6.

Step 6 If you chose **Custom Server**, do the following:

- a) In the **Chassis Id** field, enter the number of the chassis where the selected server is located.
- b) In the **Server Id** field, enter the number of the slot where the selected server is located.

Step 7 If you want to restrict the migration of the service profile after it is associated with a server, check the **Restrict Migration** check box.

If you choose not to restrict migration, Cisco UCS Manager does not perform any compatibility checks on the new server before migrating the existing service profile. If the hardware of both servers are not similar, the association might fail.

Step 8 Click **OK**.

Verifying Fibre Channel Zoning Configuration

Verify that the zone configuration and zoneset activation works properly.

Procedure

- Step 1** In the **Navigation** pane, click the **Servers** tab.
- Step 2** On the **Servers** tab, expand **Servers > Service Profiles**.
- Step 3** Navigate to and click the service profile you created previously.
- Step 4** Click the **FC Zones** tab in the right pane.

Verify that:

- The initiator and target WWPN are in the same zone.
- The zone Admin State is Applied
- The Oper State is Active

Note UCS Manager automatically creates the zone name. The zone naming convention is ClusterName_FabricID_ZoneID_ServiceProfileName_InitiatorName.

Troubleshooting Fibre Channel Zoning Configuration

If you create a service profile, but cannot see the zones under the FC Zones tab, use this troubleshooting checklist:

- Is the zoning enabled on the intended VSAN?
- Is the service profile associated?
Zones are created only when the service profile is associated with the server.
- Is the correct storage connection policy selected under vHBA initiator groups?
- Is the correct VHBA added to the correct vHBA initiator group?
- Is the correct VSAN selected for the vHBAs?
- Are the correct VSAN and fabric selected under the storage connection policy?



CHAPTER 15

Storage Inventory

- [Local Disk Locator LED Status, on page 173](#)
- [Toggling the Local Disk Locator LED On and Off, on page 174](#)
- [Custom LED Status with VMD on NVMe, on page 174](#)
- [NVMe-optimized M5 Servers, on page 177](#)
- [NVMe PCIe SSD Inventory, on page 179](#)
- [Viewing NVMe PCIe SSD Storage Inventory, on page 179](#)
- [Enabling Volume Management Device on UCS Storage, on page 180](#)
- [Enabling Intel® Volume Management Device, on page 189](#)
- [Enabling Volume Management Device \(VMD\) in Passthrough Mode, on page 190](#)
- [Downloading VMD Drivers, on page 191](#)
- [Custom LED Status with VMD on NVMe, on page 195](#)

Local Disk Locator LED Status

The local disk locator LED is located on the slot where you insert the local disk. This LED identifies where a specific disk is inserted in a blade or rack server. The locator LED is useful for maintenance, when you need to remove a disk from among many disks in a server.

You can successfully turn on or off the local disk locator LED when:

- The server is powered on. UCS Manager generates an error if you attempt to turn the locator LED on or off when the server is powered off.
- The CIMC version is UCS Manager 3.1 or higher.
- The RAID controller supports the out-of-band (OOB) storage interface.

When Intel Volume Management Device (VMD) for NVMe is enabled, you can also configure blinking patterns for the LEDs on NVMe-managed devices to show drive status. VMD-enabled drives identified by a failure ID blink pattern can be hot-plugged without a system shutdown.

Toggling the Local Disk Locator LED On and Off

Before you begin

On and Off

- Ensure the server on which the disk is located is powered on. If the server is off, you are not able to turn on or off the local disk locator LED.

Procedure

- Step 1** In the **Navigation** pane, click the **Equipment** tab.
- Step 2** On the **Equipment** tab, expand **Equipment > Rack Mounts > Servers > Server Number**.
- For Rack-mounted servers, go to **Rack Mounts Servers Server Number**.
 - For Blade servers, go to **> Sensors > Storage Server Number**.
- Step 3** In the **Work** area, click the **Inventory > Storage > Disks** tabs.
The Storage Controller inventory appears.
- Step 4** Click a disk.
The disk details appear.
- Step 5** In the Actions area, click **Turn on Locator LED** or **Turn off Locator LED**.
The **Locator LED** state appears in the **Properties** area.
- Step 6** Click **Save Changes**.
-

Custom LED Status with VMD on NVMe

Once you have set up VMD, you can customize LED blinking patterns on PCIe NVMe drives. Information on LED customization can be found in the User Guides included in the driver packages.

LED Blinking

PCIe SSD drives lack a standard way to manage the LEDs that indicate drive status and health. Without this, there is a risk of removing the wrong drive, resulting in data loss. SSD drives have two indicators, the first being a green activity LED whose signals come directly from the SSD, and the second being a status LED whose signals come from the backplane. VMD manages only the status LEDs, not the activity LEDs.

LED Management only applies to NVMe and/or SATA drives. It does not support drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard .

LED Activity During Drive Hot-plug

VMD with NVMe supports Surprise hot-plugging. When a disk is hot-removed, then re-inserted into the same slot, the fault LED blinks for 10 seconds. This is expected behavior. The fail state is imposed on a slot's LEDs when the drive is removed, but the backplanes require the drive to be present in the slot for a LED to blink.

Thus, the fail state exists once the drive is removed, but a LED blinks only when the new drive is inserted and discovered. The LED will return to normal once hot-plug event is handled.

Custom Blinking Patterns

VRoC with VMD allows you to perform basic LED management configuration of the status LEDs on compatible backplanes. Once the VMD NVMe driver is installed, you can install the VMD LED Management Tool, which lets you manage the LED through a command line interface. VMD allows you to customize LED blinking patterns on PCIe NVMe drives to better identify failing drives.

The tables below provide some brief guidelines for customized blinking on the various platforms. As individualized patterns are programmable, these tables provide only representative guidelines.

Table 9: LED Blinking Patterns: Windows

Status LED	Behavior	Options
"Activate LED"	Identifies a specific device in an enclosure by blinking the status LED of that drive in a designated pattern.	1-3600 seconds. Values outside this range default to 12 seconds. Default = 12 seconds
Drive Failure	Indicates a drive that is in a degraded or failed state by lighting the status LED of that device in a defined failure pattern.	The failure pattern is displayed until: <ul style="list-style-type: none"> • 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. • 2. From the time when a non-failed drive that is part of a RAID volume is removed, or the failed drive is identified and removed. It remains in failure state until a new drive is inserted into the same slot or the platform is rebooted. Default = Option 1
RAID volume Initialization or Verify and Repair Process	When a RAID volume is in Rebuild state, the status LEDs blink in the defined Rebuild pattern on either the specific drive being rebuilt or on the entire RAID volume that is being rebuilt.	Default = Enabled Can be: <ol style="list-style-type: none"> 1. Disabled (only on one drive) 2. Enabled (on all drives)

Status LED	Behavior	Options
Managed unplug	During a managed hot unplug, the status LED of the managed drive blinks in the defined Locate pattern until the drive is physically ejected.	None. Enabled by default.
RAID volume is migrating	During RAID volume migration, the status LEDs blink in the defined Rebuild pattern on all drives until the process is complete.	Default = Enabled Can be: 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs)
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 10: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller BLACKLIST	<code>ledmon</code> will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	Exclude controllers on the blacklist. Default = Support all controllers
RAID volume is initializing, verifying, or verifying and fixing BLINK_ON_INIT	Rebuild pattern on all drives in RAID volume (until initialization, verify, or verify and fix finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> scan interval INTERVAL	Defines the time interval between <code>ledmon sysfs</code> scans. The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is rebuilding REBUILD_BLINK_ON_ALL	Rebuild pattern on a single drive to which RAID volume rebuilds	1. False/Disabled (on one drive) 2. True/Enabled (on all drives) Default = False/Disabled
RAID volume is migrating BLINK_ON_MIGR	Rebuild pattern on all drives in RAID volume (until migration finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> debug level LOG_LEVEL	Corresponds with <code>-log-level</code> flag from <code>ledmon</code> .	Acceptable values are: quiet, error, warning, info, debug, all - 0 means 'quiet' and 5 means 'all' Default = 2

Status LED	Behavior	Options
Set manage one RAID member or All RAID RAID_MEMBERS_ONLY	If the flag is set to <code>ledmon</code> true, will limit monitoring only to drives that are RAID members.	1. False / (all RAID member and PT) 2. True / (RAID member only) Default = False
Limited scans only to specific controllers WHITELIST	<code>ledmon</code> limits changing the LED state to controllers listed on whitelist.	Limit changing LED state in whitelist controller. Default = No limit.

Table 11: LED Blinking Patterns: ESXi

Status LED	Behavior	Options
"Identify"	The ability to identify a specific device in an enclosure by blinking the status LED of that drive in the defined Locate pattern.	None. Default is Off.
"Off"	The ability to turn off the "Identify" LED once a specific device in an enclosure has been located.	None. Default is Off.

NVMe-optimized M5 Servers

Beginning with 3.2(3a), Cisco UCS Manager supports the following NVMe-optimized M5 servers:

- **UCSC-C220-M5SN**—The PCIe MSwitch is placed in the dedicated MRAID slot for UCS C220 M5 servers. This setup supports up to 10 NVMe drives. The first two drives are direct-attached through the riser. The remaining eight drives are connected and managed by the MSwitch. This setup does not support any SAS/SATA drive combinations.
- **UCSC-C240-M5SN**—The PCIe MSwitch is placed in the riser-2 at slot-4 for UCS C240 M5 servers. The servers support up to 24 drives. Slots 1-8 are the NVMe drives connected and managed by the MSwitch. The servers also support up to two NVMe drives in the rear and are direct-attached through the riser. This setup supports SAS/SATA combination with the SAS/SATA drives from slots 9-24. These drives are managed by the SAS controller placed in the dedicated MRAID PCIe slot.
- **UCS-C480-M5**—UCS C480 M5 servers support up to three front NVMe drive cages, each supporting up to eight NVMe drives. Each cage has an interposer card, which contains the MSwitch. Each server can support up to 24 NVMe drives (3 NVMe drive cages x 8 NVMe drives). The servers also support a rear PCIe Aux drive cage, which can contain up to eight NVMe drives managed by an MSwitch placed in PCIe slot-10.

This setup does not support:

- a combination of NVMe drive cages and HDD drive cages
- a combination of the Cisco 12G 9460-8i RAID controller and NVMe drive cages, irrespective of the rear Auxiliary drive cage



Note The UCS C480 M5 PID remains same as in earlier release.



Note On B200 and B480 M5 blade servers, NVMe drives cannot be used directly with SAS controllers. Use an LSTOR-PT pass-through controller instead.

The following MSwitch cards are supported in NVMe optimized M5 servers:

- UCS-C480-M5 HDD Ext NVMe Card (UCSC-C480-8NVME)—Front NVMe drive cage with an attached interposer card containing the PCIe MSwitch. Each server supports up to three front NVMe drive cages and each cage supports up to 8 NVMe drives. Each server can support up to 24 NVMe drives (3 NVMe drive cages x 8 NVMe drives).
- UCS-C480-M5 PCIe NVMe Switch Card (UCSC-NVME-SC)—PCIe MSwitch card to support up to eight NVMe drives in the rear auxiliary drive cage inserted in PCIe slot 10.



Note Cisco UCS-C480-M5 servers support a maximum of 32 NVMe drives (24 NVMe drives in the front + 8 NVMe drives in the rear auxiliary drive cage)

- UCSC-C220-M5SN and UCSC-C240-M5SN do not have separate MSwitch PIDs. MSwitch cards for these servers are part of the corresponding NVMe optimized server.

MSwitch Disaster Recovery

You can recover a corrupted MSwitch and roll back to a previous working firmware.



Note If you have a setup with Cisco UCS C480 M5 Server, then MSwitch disaster recovery process can be performed only on one MSwitch at a time. If the disaster recovery process is already running for one MSwitch, then wait for it to complete. You can monitor the recovery status from FSM.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Rack-Mounts > Servers**.
- Step 3** Expand the server for the which contains the MSwitch.
- Step 4** In the **Work** pane, click **Inventory > Storage > Controller**.
- Step 5** Select the MSwitch which you want to recover.
- Step 6** Under the **General** tab, click **Disaster Recovery**.

Note Do not reset the server during the disaster recovery process.

Step 7 You can monitor the recovery status from FSM.

NVMe PCIe SSD Inventory

Cisco UCS Manager GUI discovers, identifies, and displays the inventory of Non-Volatile Memory Express (NVMe) Peripheral Component Interconnect Express (PCIe) SSD storage devices. You can view the health of the storage devices in the server. NVMe with PCIe SSD storage devices reduce latency, increased input/output operations per second (IOPS), and lower power consumption compared to SAS or SATA SSDs.



Note Virtual Controller for direct CPU-attached NVME drives will be shown in UCSM only when it detects the NVME drives and such NVME switch details are retained between drive insertions. These NVME switch entries will be removed only when the server is decommissioned.

Viewing NVMe PCIe SSD Storage Inventory

Procedure

- Step 1** In the **Navigation** pane, click the **Equipment** tab.
- Step 2** On the **Equipment** tab, expand **Equipment > Rack Mounts > Servers**.
- Step 3** Click the **Inventory** tab.
- Step 4** Do one of the following:
- Click the **Storage** tab.
You view the list of NVMe PCIe SSD storage devices named **Storage Controller NVME ID number**. You view the name, size, serial number, operating status, state and other details.
 - Click the NVMe PCIe SSD storage device.
You see the following inventory details:

Name	Description
ID	The NVMe PCIe SSD storage device configured on the server.
Model	The NVMe PCIe SSD storage device model.
Revision	The NVMe PCIe SSD storage device revision.
RAID Support	Whether the NVMe PCIe SSD storage device is RAID enabled.

Name	Description
OOB Interface Support	Whether the NVMe PCIe SSD storage device support out-of-band management .
PCIe Address	<p>The NVMe PCIe SSD storage device on the virtual interface card (VIC).</p> <p>Note PCIe Address does not come upon hot insertion of the NVMe card. To view this info, re-acknowledge the server.</p>
Number of Local Disks	The number of disks contained in the NVMe PCIe SSD storage device.
Rebuild Rate	Not applicable to NVMe PCIe SSD storage devices.
Vendor	The vendor that manufactured the NVMe PCIe SSD storage device.
PID	The NVMe PCIe SSD storage device product ID, also known as product name, model name, product number
Serial	The storage device serial number.

Enabling Volume Management Device on UCS Storage

Enabling Intel® Volume Management Device

Volume Management Device (VMD) Setup

The Intel® Volume Management Device (VMD) is a tool that provides NVMe drivers to manage PCIe Solid State Drives attached to VMD-enabled domains. This includes Surprise hot-plug of PCIe drives and configuring blinking patterns to report status. PCIe Solid State Drive (SSD) storage lacks a standardized method to blink LEDs to represent the status of the device. With VMD, you can control LED indicators on both direct attached and switch attached PCIe storage using a simple command-line tool.

To use VMD, you must first enable VMD through a UCS Manager BIOS policy and set the UEFI boot options. Enabling VMD provides Surprise hot plug and optional LED status management for PCIe SSD storage that is attached to the root port. VMD Passthrough mode provides the ability to manage drives on guest VMs.

Enabling VMD also allows configuration of Intel® Virtual RAID on CPU (VRoC), a hybrid RAID architecture on Intel® Xeon® Scalable Processors. Documentation on the use and configuration of VRoC can be found at the Intel website.

IMPORTANT: VMD must be enabled in the UCS Manager BIOS settings before Operating System install. If enabled after OS installation, the server will fail to boot. This restriction applies to both standard VMD and VMD Passthrough. Likewise, once enabled, you cannot disable VMD without a loss of system function.

Enabling VMD on UCS Manager

To configure a BIOS and local boot Policy for VMD in UCS Manager, use the following procedure. The VMD platform default is disabled.



Note VMD must be enabled before OS installation.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 3** Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
 - Step 4** Go to **Policies > Root > BIOS Policies** and select the new policy.
 - Step 5** Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.
 - Step 6** Scroll down to **VMD Enable** and select **Enable**.
 - Step 7** Click **Save Changes** to enable VMD functions.
 - Step 8** In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode** and **Add NVMe** from the **Local Devices** menu. Click **Save Changes** to create the policy.
-

Enabling Volume Management Device (VMD) in Passthrough Mode

Volume Management Device (VMD) Passthrough Mode

The Intel® Volume Management Device (VMD) driver release package for Direct Device Assignment contains the Intel VMD UEFI Driver version for Direct Assign (PCIe PassThru) in VMware ESXi Hypervisor. The Intel VMD NVMe driver assists in the management of CPU-attached Intel PCIe NVMe SSDs.

The Intel VMD driver is required to enable the Direct Assign and discovery of the VMD physical addresses from a supported guest VM. Drivers are only provided for Passthrough mode for ESXi support of Red Hat

Linux or Ubuntu. VMD Passthrough is enabled by configuring a UCS Manager BIOS policy before loading the Operating System. Once the Operating System has been loaded, you cannot enable or disable the VMD Passthrough option.



Note Passthrough mode is enabled by default, but you should always confirm that it is enabled before proceeding.

Configuring VMD Passthrough

Passthrough mode is only supported on ESXi drivers for Red Hat Linux or Ubuntu guest operating systems.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 3** Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
 - Step 4** Go to **Policies > Root > BIOS Policies** and select the new policy.
 - Step 5** Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.
 - Step 6** Scroll down to **VMD Enable** and select **Enable**.
 - Step 7** Click **Save Changes** to enable VMD functions.
 - Step 8** To finish enabling VMD Passthrough mode, select **Advanced** and **Intel Directed IO** from the submenus and scroll down to **Intel VT Directed IO**. Verify that the dropdown is set to **Enabled**. If not, set it.
 - Step 9** Click **Save Changes** to enable the VMD Passthrough policy.
 - Step 10** In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode**. Click **OK** to create the policy.
-

Downloading VMD Drivers

Intel® Volume Management Device Drivers

Intel® Volume Management Device (VMD) for NVMe enables drive management options using hardware logic inside the Intel Xeon processor. Specific drivers are available for the following operating systems:

- Linux
- Windows 2016, 2019
- VMWare



Note The latest VMWare drivers are available directly from the VMWare site. Following links in the VMWare driver download on the Cisco download site will take you directly to the VMWare login page.

For guest Operating Systems on ESXi, use VMD Passthrough mode. Supported Operating Systems for VMD Passthrough are:

- Red Hat Linux
- Ubuntu

To use the features of Intel VMD, you must:

- Enable VMD by creating a BIOS policy in the UCS Manager.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

- Install the appropriate VMD NVMe driver.
- Install the appropriate management tools for the driver package.
- Boot from UEFI.

Intel® Virtual RAID on CPU (VRoC) with VMD

Intel® Virtual RAID on CPU (VRoC) allows you to create and manage RAID volumes within the BIOS of VMD-enabled Intel NVMe SSD drives using hardware logic inside the Intel Xeon processor. More information on Intel VRoC can be found at: <https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html>.

The User Guides for Intel VRoC can be accessed at the direct link at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en

The Windows and Linux user documentation also contains information on how to configure Intel VRoC in the pre-boot environment. Creation of RAID volumes in VRoC is through the HII interface. The Windows documentation provides information on using the BIOS HII option to set up and configure RAID volumes in VRoC.

To use Intel VRoC, you must:

- Enable VMD in the BIOS settings
- Use UEFI boot mode
- Have sufficient drive resources to create the volume
- Use the BIOS HII option to set up and configure VRoC.

The Cisco implementation of Intel VRoC supports RAID 0 (striping), RAID 1 (mirroring), RAID 5 (striping with parity) and RAID 10 (combined mirroring and striping).

Downloading the Linux VMD Drivers

Complete these steps to download and install the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

-
- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.
 - Step 2** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
 - Step 3** Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.
 - Step 4** Click on the latest release in the left panel.
- Note** The ISO image for VMD on blade servers is available from the 4.0(4f) release onward.
- Step 5** Click on **ISO image of UCS-related linux drivers only** and download the driver bundle.
 - Step 6** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD > RHEL_{x.x}**.
 - Step 7** Click on the version of Red Hat Linux that you wish to install.
 - Step 8** Extract the contents of the folder. The folder contains both the driver package and associated documentation. Follow the installation procedure packaged with the drivers.
-

What to do next

The Intel® Virtual RAID on CPU (VRoC) Linux Software User Guide can be found with the user documentation at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en. It provides information on performing BIOS HII VRoC setup in the pre-boot environment, as well as how to install and use the programmable LED utility.

Downloading the Windows VMD Drivers

Complete these steps to download the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.
- Step 2** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 3** Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.
- Step 4** Click on the latest release in the left panel.
The ISO image for VMD is available from the 4.0(4f) release onward.
- Step 5** Click on **ISO image of UCS-related windows drivers only** and download the driver bundle.
- Step 6** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD > KIT_x_x_x_xxxx**.
- Step 7** Extract the contents of the folder.
- Step 8** Click on the entry for the kit and **KIT > Install**.
- Step 9** The folder contains both the driver package and associated documentation. Expand the zip file for **VROC_x_x_x_xxxxInstall**.
- Step 10** Follow the installation procedure packaged with the drivers.
-

What to do next

For setting up Intel® Virtual RAID on CPU (VRoC), refer to the online instructions at <https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html>.

Information on VRoC RAID features and management can be found in the *Windows Intel Virtual RAID on CPU Software User's Guide* at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/ssd-software/Windows_VROC_User_Guide.pdf.

Downloading the VMD Passthrough Drivers

Complete these steps to download and install the driver bundle for VMD Passthrough mode:



Note The VMD Passthrough driver bundle includes packages for both ESXi and Ubuntu.

Before you begin



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.

- Step 2** Search on **Servers - Unified Computing**.
- Step 3** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 4** Choose the UCS utilities from the Software Type selections: **Unified Computing System (UCS) Utilities**.
- Step 5** Click on the latest release in the left panel.

Note The ISO image for VMD is available from UCSM 4.0(4f) release onward.

- Step 6** Click on **ISO image of UCS-related vmware utilities only** and download the utilities bundle.
- Step 7** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD**.

The bundle provides both the driver installation package for the desired version of ESXi or VMD Direct Assign with Ubuntu, passthrough mode, and the Signed LED Offline bundle. Also included is a pdf that provides steps to configure an Ubuntu Virtual Machine in ESXi.

- Step 8** Click on either the version of ESXi that you wish to install or the zip file for Ubuntu.
- For ESXi versions, Click on **ESXi_x > Direct Assign** and chose the desired zip file.

- Step 9** Extract the contents of the folder. Follow the installation procedure packaged with the driver software.

What to do next

Extract the contents of the LED management tools zip file. Install the management tools according to the instructions included with the driver package.

Before using the command line tools, the ESXi command line shell should be enabled from either the vSphere client or from the direct console of the ESXi host system.

Custom LED Status with VMD on NVMe

Once you have set up VMD, you can customize LED blinking patterns on PCIe NVMe drives. Information on LED customization can be found in the User Guides included in the driver packages.

LED Blinking

PCIe SSD drives lack a standard way to manage the LEDs that indicate drive status and health. Without this, there is a risk of removing the wrong drive, resulting in data loss. SSD drives have two indicators, the first being a green activity LED whose signals come directly from the SSD, and the second being a status LED whose signals come from the backplane. VMD manages only the status LEDs, not the activity LEDs.

LED Management only applies to NVMe and/or SATA drives. It does not support drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard .

LED Activity During Drive Hot-plug

VMD with NVMe supports Surprise hot-plugging. When a disk is hot-removed, then re-inserted into the same slot, the fault LED blinks for 10 seconds. This is expected behavior. The fail state is imposed on a slot's LEDs when the drive is removed, but the backplanes require the drive to be present in the slot for a LED to blink. Thus, the fail state exists once the drive is removed, but a LED blinks only when the new drive is inserted and discovered. The LED will return to normal once hot-plug event is handled.

Custom Blinking Patterns

VRoC with VMD allows you to perform basic LED management configuration of the status LEDs on compatible backplanes. Once the VMD NVMe driver is installed, you can install the VMD LED Management Tool, which lets you manage the LED through a command line interface. VMD allows you to customize LED blinking patterns on PCIe NVMe drives to better identify failing drives.

The tables below provide some brief guidelines for customized blinking on the various platforms. As individualized patterns are programmable, these tables provide only representative guidelines.

Table 12: LED Blinking Patterns: Windows

Status LED	Behavior	Options
"Activate LED"	Identifies a specific device in an enclosure by blinking the status LED of that drive in a designated pattern.	1-3600 seconds. Values outside this range default to 12 seconds. Default = 12 seconds
Drive Failure	Indicates a drive that is in a degraded or failed state by lighting the status LED of that device in a defined failure pattern.	The failure pattern is displayed until: <ul style="list-style-type: none"> • 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. • 2. From the time when a non-failed drive that is part of a RAID volume is removed, or the failed drive is identified and removed. It remains in failure state until a new drive is inserted into the same slot or the platform is rebooted. Default = Option 1
RAID volume Initialization or Verify and Repair Process	When a RAID volume is in Rebuild state, the status LEDs blink in the defined Rebuild pattern on either the specific drive being rebuilt or on the entire RAID volume that is being rebuilt.	Default = Enabled Can be: <ol style="list-style-type: none"> 1. Disabled (only on one drive) 2. Enabled (on all drives)
Managed unplug	During a managed hot unplug, the status LED of the managed drive blinks in the defined Locate pattern until the drive is physically ejected.	None. Enabled by default.

Status LED	Behavior	Options
RAID volume is migrating	During RAID volume migration, the status LEDs blink in the defined Rebuild pattern on all drives until the process is complete.	Default = Enabled Can be: 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs)
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 13: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller BLACKLIST	<code>ledmon</code> will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	Exclude controllers on the blacklist. Default = Support all controllers
RAID volume is initializing, verifying, or verifying and fixing BLINK_ON_INIT	Rebuild pattern on all drives in RAID volume (until initialization, verify, or verify and fix finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> scan interval INTERVAL	Defines the time interval between <code>ledmon sysfs scans</code> . The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is rebuilding REBUILD_BLINK_ON_ALL	Rebuild pattern on a single drive to which RAID volume rebuilds	1. False/Disabled (on one drive) 2. True/Enabled (on all drives) Default = False/Disabled
RAID volume is migrating BLINK_ON_MIGR	Rebuild pattern on all drives in RAID volume (until migration finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> debug level LOG_LEVEL	Corresponds with <code>-log-level</code> flag from <code>ledmon</code> .	Acceptable values are: quiet, error, warning, info, debug, all - 0 means 'quiet' and 5 means 'all' Default = 2
Set manage one RAID member or All RAID RAID_MEMBERS_ONLY	If the flag is set to <code>ledmon true</code> , will limit monitoring only to drives that are RAID members.	1. False / (all RAID member and PT) 2. True / (RAID member only) Default = False

Status LED	Behavior	Options
Limited scans only to specific controllers WHITELIST	<code>ledmon</code> limits changing the LED state to controllers listed on whitelist.	Limit changing LED state in whitelist controller. Default = No limit.

Table 14: LED Blinking Patterns: ESXi

Status LED	Behavior	Options
"Identify"	The ability to identify a specific device in an enclosure by blinking the status LED of that drive in the defined Locate pattern.	None. Default is Off.
"Off"	The ability to turn off the "Identify" LED once a specific device in an enclosure has been located.	None. Default is Off.

Enabling Intel® Volume Management Device

Volume Management Device (VMD) Setup

The Intel® Volume Management Device (VMD) is a tool that provides NVMe drivers to manage PCIe Solid State Drives attached to VMD-enabled domains. This includes Surprise hot-plug of PCIe drives and configuring blinking patterns to report status. PCIe Solid State Drive (SSD) storage lacks a standardized method to blink LEDs to represent the status of the device. With VMD, you can control LED indicators on both direct attached and switch attached PCIe storage using a simple command-line tool.

To use VMD, you must first enable VMD through a UCS Manager BIOS policy and set the UEFI boot options. Enabling VMD provides Surprise hot plug and optional LED status management for PCIe SSD storage that is attached to the root port. VMD Passthrough mode provides the ability to manage drives on guest VMs.

Enabling VMD also allows configuration of Intel® Virtual RAID on CPU (VRoC), a hybrid RAID architecture on Intel® Xeon® Scalable Processors. Documentation on the use and configuration of VRoC can be found at the Intel website.

IMPORTANT: VMD must be enabled in the UCS Manager BIOS settings before Operating System install. If enabled after OS installation, the server will fail to boot. This restriction applies to both standard VMD and VMD Passthrough. Likewise, once enabled, you cannot disable VMD without a loss of system function.

Enabling VMD on UCS Manager

To configure a BIOS and local boot Policy for VMD in UCS Manager, use the following procedure. The VMD platform default is disabled.



Note VMD must be enabled before OS installation.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
 - Step 3** Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
 - Step 4** Go to **Policies > Root > BIOS Policies** and select the new policy.
 - Step 5** Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenu.
 - Step 6** Scroll down to **VMD Enable** and select **Enable**.
 - Step 7** Click **Save Changes** to enable VMD functions.
 - Step 8** In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode** and **Add NVMe** from the **Local Devices** menu. Click **Save Changes** to create the policy.
-

Enabling Volume Management Device (VMD) in Passthrough Mode

Volume Management Device (VMD) Passthrough Mode

The Intel® Volume Management Device (VMD) driver release package for Direct Device Assignment contains the Intel VMD UEFI Driver version for Direct Assign (PCIe PassThru) in VMware ESXi Hypervisor. The Intel VMD NVMe driver assists in the management of CPU-attached Intel PCIe NVMe SSDs.

The Intel VMD driver is required to enable the Direct Assign and discovery of the VMD physical addresses from a supported guest VM. Drivers are only provided for Passthrough mode for ESXi support of Red Hat Linux or Ubuntu. VMD Passthrough is enabled by configuring a UCS Manager BIOS policy before loading the Operating System. Once the Operating System has been loaded, you cannot enable or disable the VMD Passthrough option.



Note Passthrough mode is enabled by default, but you should always confirm that it is enabled before proceeding.

Configuring VMD Passthrough

Passthrough mode is only supported on ESXi drivers for Red Hat Linux or Ubuntu guest operating systems.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand the node for the organization where you want to create the policy.
If the system does not include multi tenancy, expand the **root** node.
- Step 3** Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
- Step 4** Go to **Policies > Root > BIOS Policies** and select the new policy.
- Step 5** Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.
- Step 6** Scroll down to **VMD Enable** and select **Enable**.
- Step 7** Click **Save Changes** to enable VMD functions.
- Step 8** To finish enabling VMD Passthrough mode, select **Advanced** and **Intel Directed IO** from the submenus and scroll down to **Intel VT Directed IO**. Verify that the dropdown is set to **Enabled**. If not, set it.
- Step 9** Click **Save Changes** to enable the VMD Passthrough policy.
- Step 10** In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode**. Click **OK** to create the policy.
-

Downloading VMD Drivers

Intel® Volume Management Device Drivers

Intel® Volume Management Device (VMD) for NVMe enables drive management options using hardware logic inside the Intel Xeon processor. Specific drivers are available for the following operating systems:

- Linux
- Windows 2016, 2019
- VMWare



Note The latest VMWare drivers are available directly from the VMWare site. Following links in the VMWare driver download on the Cisco download site will take you directly to the VMWare login page.

For guest Operating Systems on ESXi, use VMD Passthrough mode. Supported Operating Systems for VMD Passthrough are:

- Red Hat Linux
- Ubuntu

To use the features of Intel VMD, you must:

- Enable VMD by creating a BIOS policy in the UCS Manager.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

- Install the appropriate VMD NVMe driver.
- Install the appropriate management tools for the driver package.
- Boot from UEFI.

Intel® Virtual RAID on CPU (VRoC) with VMD

Intel® Virtual RAID on CPU (VRoC) allows you to create and manage RAID volumes within the BIOS of VMD-enabled Intel NVMe SSD drives using hardware logic inside the Intel Xeon processor. More information on Intel VRoC can be found at: <https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html>.

The User Guides for Intel VRoC can be accessed at the direct link at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en

The Windows and Linux user documentation also contains information on how to configure Intel VRoC in the pre-boot environment. Creation of RAID volumes in VRoC is through the HII interface. The Windows documentation provides information on using the BIOS HII option to set up and configure RAID volumes in VRoC.

To use Intel VRoC, you must:

- Enable VMD in the BIOS settings
- Use UEFI boot mode
- Have sufficient drive resources to create the volume
- Use the BIOS HII option to set up and configure VRoC.

The Cisco implementation of Intel VRoC supports RAID 0 (striping), RAID 1 (mirroring), RAID 5 (striping with parity) and RAID 10 (combined mirroring and striping).

Downloading the Linux VMD Drivers

Complete these steps to download and install the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.
- Step 2** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 3** Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.
- Step 4** Click on the latest release in the left panel.
- Note** The ISO image for VMD on blade servers is available from the 4.0(4f) release onward.
- Step 5** Click on **ISO image of UCS-related linux drivers only** and download the driver bundle.
- Step 6** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD > RHEL_{x.x}**.
- Step 7** Click on the version of Red Hat Linux that you wish to install.
- Step 8** Extract the contents of the folder. The folder contains both the driver package and associated documentation. Follow the installation procedure packaged with the drivers.
-

What to do next

The Intel® Virtual RAID on CPU (VRoC) Linux Software User Guide can be found with the user documentation at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en. It provides information on performing BIOS HII VRoC setup in the pre-boot environment, as well as how to install and use the programmable LED utility.

Downloading the Windows VMD Drivers

Complete these steps to download the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.



- Note** The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.
-

Procedure

- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.
- Step 2** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 3** Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.
- Step 4** Click on the latest release in the left panel.
- The ISO image for VMD is available from the 4.0(4f) release onward.

- Step 5** Click on **ISO image of UCS-related windows drivers only** and download the driver bundle.
- Step 6** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD > KIT_x_x_x_xxxx**.
- Step 7** Extract the contents of the folder.
- Step 8** Click on the entry for the kit and **KIT > Install**.
- Step 9** The folder contains both the driver package and associated documentation. Expand the zip file for **VROC_x_x_x_xxxxInstall**.
- Step 10** Follow the installation procedure packaged with the drivers.

What to do next

For setting up Intel® Virtual RAID on CPU (VRoC), refer to the online instructions at <https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html>.

Information on VRoC RAID features and management can be found in the *Windows Intel Virtual RAID on CPU Software User's Guide* at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/ssd-software/Windows_VROC_User_Guide.pdf.

Downloading the VMD Passthrough Drivers

Complete these steps to download and install the driver bundle for VMD Passthrough mode:



Note The VMD Passthrough driver bundle includes packages for both ESXi and Ubuntu.

Before you begin



Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

- Step 1** In a web browser, navigate to <https://software.cisco.com/download/home>.
- Step 2** Search on **Servers - Unified Computing**.
- Step 3** Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.
- Step 4** Choose the UCS utilities from the Software Type selections: **Unified Computing System (UCS) Utilities**.
- Step 5** Click on the latest release in the left panel.
- Note** The ISO image for VMD is available from UCSM 4.0(4f) release onward.
- Step 6** Click on **ISO image of UCS-related vmware utilities only** and download the utilities bundle.

- Step 7** When the driver bundle is downloaded, open it and select **Storage > Intel > VMD**.
- The bundle provides both the driver installation package for the desired version of ESXi or VMD Direct Assign with Ubuntu, passthrough mode, and the Signed LED Offline bundle. Also included is a pdf that provides steps to configure an Ubuntu Virtual Machine in ESXi.
- Step 8** Click on either the version of ESXi that you wish to install or the zip file for Ubuntu.
- For ESXi versions, Click on **ESXi_x > Direct Assign** and chose the desired zip file.
- Step 9** Extract the contents of the folder. Follow the installation procedure packaged with the driver software.
-

What to do next

Extract the contents of the LED management tools zip file. Install the management tools according to the instructions included with the driver package.

Before using the command line tools, the ESXi command line shell should be enabled from either the vSphere client or from the direct console of the ESXi host system.

Custom LED Status with VMD on NVMe

Once you have set up VMD, you can customize LED blinking patterns on PCIe NVMe drives. Information on LED customization can be found in the User Guides included in the driver packages.

LED Blinking

PCIe SSD drives lack a standard way to manage the LEDs that indicate drive status and health. Without this, there is a risk of removing the wrong drive, resulting in data loss. SSD drives have two indicators, the first being a green activity LED whose signals come directly from the SSD, and the second being a status LED whose signals come from the backplane. VMD manages only the status LEDs, not the activity LEDs.

LED Management only applies to NVMe and/or SATA drives. It does not support drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard .

LED Activity During Drive Hot-plug

VMD with NVMe supports Surprise hot-plugging. When a disk is hot-removed, then re-inserted into the same slot, the fault LED blinks for 10 seconds. This is expected behavior. The fail state is imposed on a slot's LEDs when the drive is removed, but the backplanes require the drive to be present in the slot for a LED to blink. Thus, the fail state exists once the drive is removed, but a LED blinks only when the new drive is inserted and discovered. The LED will return to normal once hot-plug event is handled.

Custom Blinking Patterns

VRoC with VMD allows you to perform basic LED management configuration of the status LEDs on compatible backplanes. Once the VMD NVMe driver is installed, you can install the VMD LED Management Tool, which lets you manage the LED through a command line interface. VMD allows you to customize LED blinking patterns on PCIe NVMe drives to better identify failing drives.

The tables below provide some brief guidelines for customized blinking on the various platforms. As individualized patterns are programmable, these tables provide only representative guidelines.

Table 15: LED Blinking Patterns: Windows

Status LED	Behavior	Options
"Activate LED"	Identifies a specific device in an enclosure by blinking the status LED of that drive in a designated pattern.	1-3600 seconds. Values outside this range default to 12 seconds. Default = 12 seconds
Drive Failure	Indicates a drive that is in a degraded or failed state by lighting the status LED of that device in a defined failure pattern.	The failure pattern is displayed until: <ul style="list-style-type: none"> • 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. • 2. From the time when a non-failed drive that is part of a RAID volume is removed, or the failed drive is identified and removed. It remains in failure state until a new drive is inserted into the same slot or the platform is rebooted. Default = Option 1
RAID volume Initialization or Verify and Repair Process	When a RAID volume is in Rebuild state, the status LEDs blink in the defined Rebuild pattern on either the specific drive being rebuilt or on the entire RAID volume that is being rebuilt.	Default = Enabled Can be: <ol style="list-style-type: none"> 1. Disabled (only on one drive) 2. Enabled (on all drives)
Managed unplug	During a managed hot unplug, the status LED of the managed drive blinks in the defined Locate pattern until the drive is physically ejected.	None. Enabled by default.
RAID volume is migrating	During RAID volume migration, the status LEDs blink in the defined Rebuild pattern on all drives until the process is complete.	Default = Enabled Can be: <ol style="list-style-type: none"> 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs)
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 16: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller BLACKLIST	<code>ledmon</code> will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	Exclude controllers on the blacklist. Default = Support all controllers
RAID volume is initializing, verifying, or verifying and fixing BLINK_ON_INIT	Rebuild pattern on all drives in RAID volume (until initialization, verify, or verify and fix finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> scan interval INTERVAL	Defines the time interval between <code>ledmon sysfs scans</code> . The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is rebuilding REBUILD_BLINK_ON_ALL	Rebuild pattern on a single drive to which RAID volume rebuilds	1. False/Disabled (on one drive) 2. True/Enabled (on all drives) Default = False/Disabled
RAID volume is migrating BLINK_ON_MIGR	Rebuild pattern on all drives in RAID volume (until migration finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> debug level LOG_LEVEL	Corresponds with <code>-log-level</code> flag from <code>ledmon</code> .	Acceptable values are: quiet, error, warning, info, debug, all - 0 means 'quiet' and 5 means 'all' Default = 2
Set manage one RAID member or All RAID RAID_MEMBERS_ONLY	If the flag is set to <code>ledmon true</code> , will limit monitoring only to drives that are RAID members.	1. False / (all RAID member and PT) 2. True / (RAID member only) Default = False
Limited scans only to specific controllers WHITELIST	<code>ledmon</code> limits changing the LED state to controllers listed on whitelist.	Limit changing LED state in whitelist controller. Default = No limit.

Table 17: LED Blinking Patterns: ESXi

Status LED	Behavior	Options
"Identify"	The ability to identify a specific device in an enclosure by blinking the status LED of that drive in the defined Locate pattern.	None. Default is Off.

Status LED	Behavior	Options
"Off"	The ability to turn off the "Identify" LED once a specific device in an enclosure has been located.	None. Default is Off.



CHAPTER 16

Cisco UCS S3260 System Storage Management

- [Storage Server Features and Components Overview](#), on page 199
- [Cisco UCS S3260 Storage Management Operations](#), on page 207
- [Disk Sharing for High Availability](#), on page 208
- [Storage Enclosure Operations](#), on page 215
- [Sas Expander Configuration Policy](#), on page 216

Storage Server Features and Components Overview

Storage Server Features

The following table summarizes the Cisco UCS S3260 system features:

Table 18: Cisco UCS S3260 System Features

Feature	Description
Chassis	Four rack unit (4RU) chassis
Processors	<ul style="list-style-type: none">• Cisco UCS S3260 M3 server nodes: Two Intel Xeon E5-2600 v2 Series processors inside each server node.• Cisco UCS S3260 M4 server nodes: Two Intel Xeon E5-2600 v4 Series processors inside each server node.• Cisco UCS S3260 M5 server nodes: Two Intel Skylake 2S-EP processors inside each server node.
Memory	Up to 16 DIMMs inside each server node.
Multi-bit error protection	This system supports multi-bit error protection.


Feature	Description
Storage	<p>The system has the following storage options:</p> <ul style="list-style-type: none"> • Up to 56 top-loading 3.5-inch drives • Up to four 3.5-inch, rear-loading drives in the optional drive expander module • Up to four 2.5-inch, rear-loading SAS solid state drives (SSDs) • One 2.5-inch, NVMe drive inside the server node <p>Note This is applicable for S3260 M4 servers only.</p> <ul style="list-style-type: none"> • Two 7 mm NVMe drive inside the server node <p>Note This is applicable for S3260 M5 servers only.</p> <ul style="list-style-type: none"> • Two 15 mm NVMe drive supported for IO Expander
Disk Management	<p>The system supports up to two storage controllers:</p> <ul style="list-style-type: none"> • One dedicated mezzanine-style socket for a Cisco storage controller card inside each server node
RAID Backup	<p>The supercap power module (SCPM) mounts to the RAID controller card.</p>
PCIe I/O	<p>The optional I/O expander provides two 8x Gen 3 PCIe expansion slots.</p> <p>Release 3.2(3) and later supports the following for S3260 M5 servers:</p> <ul style="list-style-type: none"> • Intel X550 dual-port 10GBase-T • Qlogic QLE2692 dual-port 16G Fiber Channel HBA • N2XX-AIPCI01 Intel X520 Dual Port 10Gb SFP+ Adapter

Feature	Description
Network and Management I/O	<p>The system can have one or two system I/O controllers (SIOCs). These provide rear-panel management and data connectivity.</p> <ul style="list-style-type: none"> • Two SFP+ 40 Gb ports each SIOC. • One 10/100/1000 Ethernet dedicated management port on each SIOC. <p>The server nodes each have one rear-panel KVM connector that can be used with a KVM cable, which provides two USB, one VGA DB-15, and one serial DB-9 connector.</p>
Power	Two or four power supplies, 1050 W each (hot-swappable and redundant as 2+2).
Cooling	<p>Four internal fan modules that pull front-to-rear cooling, hot-swappable. Each fan module contains two fans.</p> <p>In addition, there is one fan in each power supply.</p>

Front Panel Features

The following image shows the front panel features for the Cisco UCS S3260 system:

Figure 3: Front Panel Features




1	Operations panel	6	Temperature status LED
2	System Power button/LED	7	Power supply status LED
3	System unit identification button/LED	8	Network link activity LED

4	System status LED	9	Pull-out asset tag (not visible under front bezel)
5	Fan status LED	10	Internal-drive status LEDs

Rear Panel Features

The following image shows the rear panel features for the Cisco UCS S3260 system:

Figure 4: Front Panel Features



Disk Slots

1	<p>Server bay 1</p> <ul style="list-style-type: none"> • (Optional) I/O expander, as shown (with Cisco UCS S3260 M4 and M5 server node only) • (Optional) server node • (Optional) drive expansion module 	8	Not used at this time
---	--	---	-----------------------

2	<p>Server bay 2</p> <ul style="list-style-type: none"> • (Optional) server node (Cisco UCS S3260 M4 and M5 shown) • (Optional) drive expansion module 	9	Not used at this time
3	<p>System I/O controller (SIOC)</p> <ul style="list-style-type: none"> • SIOC 1 is required if you have a server node in server bay 1 • SIOC 2 is required if you have server node in server bay 2 	10	<p>Solid state drive bays (up to four 2.5-inch SAS SSDs)</p> <ul style="list-style-type: none"> • SSDs in bays 1 and 2 require a server node in server bay 1 • SSDs in bays 3 and 4 require a server node in server bay 2
4	<p>Power supplies (four, redundant as 2+2)</p>	11	<p>Cisco UCS S3260 M4 server node label (M4 SVRN)</p> <p>Note This label identifies a Cisco UCS S3260 M4 and M5 server node.</p> <p>The Cisco UCS S3260 M3 server node does not have a label.</p>
5	<p>40-Gb SFP+ ports (two on each SIOC)</p>	12	<p>KVM console connector (one each server node).</p> <p>Used with a KVM cable that provides two USB, one VGA, and one serial connector</p>

6	Chassis Management Controller (CMS) Debug Firmware Utility port (one each SIOC)	13	Server node unit identification button/LED
7	10/100/1000 dedicated management port, RJ-45 connector (one each SIOC)	14	Server node power button
		15	Server node reset button (resets chipset in the server node)

Storage Server Components

Server Nodes


The Cisco UCS S3260 system consists of one or two server nodes, each with two CPUs, DIMM memory of 128, 256, or 512 GB, and a RAID card up to 4 GB cache or a pass-through controller. The server nodes can be one of the following:

- Cisco UCS S3260 M3 Server Node
- Cisco UCS S3260 M4 Server Node—This node might include an optional I/O expander module that attaches to the top of the server node.
- Cisco UCS S3260 M5 Server Node—This node might include an optional I/O expander module that attaches to the top of the server node.

Disk Slots


The Cisco UCS S3260 chassis has 4 rows of 14 disk slots on the HDD motherboard and 4 additional disk slots on the HDD expansion tray. The following image shows the disk arrangement for the 56 top-accessible, hot swappable 3.5-inch 6 TB or 4 TB 7200 rpm NL-SAS HDD drives. A disk slot has two SAS ports and each is connected a SAS expander in the chassis.

Figure 5: Cisco UCS S3260 Top View



The following image shows the Cisco UCS S3260 chassis with the 4 additional disk slots on the HDD expansion tray.

Figure 6: Cisco UCS 3260 with the HDD expansion tray (Rear View)



If you have two server nodes with two SIOCs, you will have the following functionality:

1. The top server node works with the left SIOC (Server Slot1 with SIOC1).
2. The bottom server works with the right SIOC (Sever Slot 2 with SIOC2).

If you have one server node with two SIOCs, you can enable Server SIOC Connectivity functionality. Beginning with release 3.1(3), Cisco UCS S3260 system supports Server SIOC Connectivity functionality. Using this functionality, you can configure the data path through both the primary and auxiliary SIOCs when the chassis has single server and dual SIOCs set up.

SAS Expanders

The Cisco UCS S3260 system has two SAS expanders that run in redundant mode and connect the disks at the chassis level to storage controllers on the servers. The SAS expanders provide two paths between a storage controller, and hence enable high availability. They provide the following functionality:

- Manage the pool of hard drives.

- Disk zone configuration of the hard drives to storage controllers on the servers.

Beginning with release 3.2(3a), Cisco UCS Manager can enable single path access to disk by configuring single DiskPort per disk slot. This ensures that the server discovers only a single device and avoid a multi-path configuration.

The following table describes how the ports in each SAS expander are connected to the disks based on the type of deployment.

Port range	Connectivity
1-56	Top accessible disks
57-60	Disks in the HDD expansion tray.



Note The number of SAS uplinks between storage controller and SAS expander can vary based on the type of controller equipped in the server.

Storage Enclosures

A Cisco UCS S3260 system has the following types of storage enclosures:

Chassis Level Storage Enclosures

- **HDD motherboard enclosure**—The 56 dual port disk slots in the chassis comprise the HDD motherboard enclosure.
- **HDD expansion tray**—The 4 additional dual disk slots in the Cisco UCS S3260 system comprise the HDD expansion tray.



Note The HDD expansion tray is a field replaceable unit (FRU). The disks will remain unassigned upon insertion, and can be assigned to storage controllers. For detailed steps on how to perform disk zoning, see [Disk Zoning Policies, on page 208](#)

Server level Storage Enclosures

Server level storage enclosures are pre-assigned dedicated enclosures to the server. These can be one of the following:

- **Rear Boot SSD enclosure**—This enclosure contains two 2.5 inch disk slots on the rear panel of the Cisco UCS S3260 system. Each server has two dedicated disk slots. These disk slots support SATA SSDs.
- **Server board NVMe enclosure**—This enclosure contains one PCIe NVMe controller.



Note In the Cisco UCS S3260 system, even though disks can be physically present on the two types of enclosures described above, from the host OS all the disks are viewed as part of one SCSI enclosure. They are connected to SAS expanders that are configured to run as single SES enclosure.

Storage Controllers

Mezzanine Storage Controllers

The following table lists the storage controller type, firmware type, modes, sharing and OOB support for the various storage controllers.

Table 19:

Storage Controller Type	Firmware type	Modes	Sharing	OOB Support
UCSC-S3X60-R1GB	Mega RAID	HW RAID, JBOD	No	Yes
UCS-C3K-M4RAID	Mega RAID	HW RAID, JBOD	No	Yes
UCSC-S3X60-HBA	Initiator Target	Pass through	Yes	Yes
UCS-S3260-DHBA	Initiator Target	Pass through	Yes	Yes
UCS-S3260-DRAID	Mega RAID	HW RAID, JBOD	No	Yes

Other storage controllers

SW RAID Controller—The servers in the Cisco UCS S3260 system support two dedicated internal SSDs embedded into the PCIe riser that is connected to the SW RAID Controller. This controller is supported on the Cisco C3000 M3 servers.

NVMe Controller—This controller is used by servers in the Cisco UCS S3260 system for inventory and firmware updates of NVMe disks.

For more details about the storage controllers supported in the various server nodes, see the related service note:

- [Cisco UCS C3X60 M3 Server Node For Cisco UCS S3260 Storage Server Service Note](#)
- [Cisco UCS C3X60 M4 Server Node For Cisco UCS S3260 Storage Server Service Note](#)
- [Cisco UCS S3260 M5 Server Node For Cisco UCS S3260 Storage Server Service Note](#)

Cisco UCS S3260 Storage Management Operations

The following table summarizes the various storage management operations that you can perform with the Cisco UCS Manager integrated Cisco UCS S3260 system.

Operation	Description	See:
Disk Sharing for High Availability	The SAS expanders in the Cisco UCS S3260 system can manage the pool of drives at the chassis level. To share disks for high availability, perform the following: <ol style="list-style-type: none"> 1. Creating disk zoning policies. 2. Creating disk slots and assigning ownership. 3. Associating disks to chassis profile. 	"Disk Zoning Policies" section in this guide.
Storage Profiles, Disk Groups and Disk Group Configuration Policies	You can utilize Cisco UCS Manager's Storage Profile and Disk Group Policies for defining storage disks, disk allocation and management in the Cisco UCS S3260 system.	"Storage Profiles" section in the <i>Cisco UCS Manager Storage Management Guide, Release 3.1(2)</i> . <i>Cisco UCS Manager Storage Management Guide, Release 3.2</i> .
Storage Enclosure Operations	You can swap the HDD expansion tray with a server, or remove the tray if it was previously inserted.	"Removing Chassis Level Storage Enclosures" section in this guide.

Disk Sharing for High Availability

Disk Zoning Policies

You can assign disk drives to the server nodes using disk zoning. Disk zoning can be performed on the controllers in the same server or on the controllers on different servers. Disk ownership can be one of the following:

Unassigned

Unassigned disks are those not visible to the server nodes.

Dedicated

If this option is selected, you will need to set the values for the **Server** and **ControllerServer**, **Controller**, **Drive Path**, and **Slot Range** for the disk slot.



Note A disk is visible only to the assigned controller.

Beginning with release 3.2(3a), Cisco UCS Manager can enable single path access to disk by configuring single DiskPort per disk slot for Cisco UCS S3260 M5 and higher servers. Setting single path configuration ensures that the server discovers the disk drive only through a single drive path chosen in the configuration. Single path access is supported only for **Cisco UCS S3260 Dual Pass Through Controller** (UCS-S3260-DHBA)

Once single path access is enabled, you cannot downgrade to any release earlier than 3.2(3a). To downgrade, disable this feature and assign all the disk slots to both the disk ports by configuring disk path of the disk slots to **Path Both** in disk zoning policy.

Shared

Shared disks are those assigned to more than one controller. They are specifically used when the servers are running in a cluster configuration, and each server has its storage controllers in HBA mode.



Note Shared mode cannot be used under certain conditions when dual HBA controllers are used.

Chassis Global Hot Spare

If this option is selected, you will need to set the value for the **Slot Range** for the disk.



Important Disk migration and claiming orphan LUNs: To migrate a disk zoned to a server (Server 1) to another server (Server 2), you must mark the virtual drive (LUN) as transport ready or perform a hide virtual drive operation. You can then change the disk zoning policy assigned for that disk. For more information on virtual drive management, see the *Disk Groups and Disk Configuration Policies* section of the [Cisco UCS Manager Storage Management Guide](#).

Creating a Disk Zoning Policy

Procedure

- Step 1** In the Navigation pane, click **Chassis**.
- Step 2** Expand **Policies > root**.
- Step 3** Right-click **Disk Zoning Policies** and choose **Create Disk Zoning Policy**.
- Step 4** In the **Create Disk Zoning Policy** dialog box, complete the following:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p>
Preserve Config check box	<p>If this check box is selected, it preserves all configuration related information for the disks such as slot number, ownership, server assigned, controller assigned, and controller type.</p> <p>Note By default the Preserve Config check box remains unchecked.</p>

In the **Disk Zoning Information** area, complete the following:

Name	Field
Name column	The name for the disk slot.
Slot Number column	The slot number for the disk.

Name	Field
Ownership column	<p>The slot ownership value. This can be one of the following:</p> <ul style="list-style-type: none"> • Unassigned—This option is selected by default. You can set the slot number in the Slot Range field. • Dedicated—If this option is selected, you will need to set the values for the Server, Controller, Drive Path, and Slot Range for the disk slot. <p>Beginning with release 3.2(3a), Cisco UCS Manager can enable single path access to disk by configuring single DiskPort per disk slot. This ensures that the server discovers only a single device and avoid a multi-path configuration.</p> <p>Drive Path options are:</p> <ul style="list-style-type: none"> • Path Both (Default) - Drive path is zoned to both the SAS expanders. • Path 0 - Drive path is zoned to SAS expander 1. • Path 1 - Drive path is zoned to SAS expander 2. <ul style="list-style-type: none"> • Shared—If this option is selected, you will need to set the values for the Slot Range and controller information such as server assigned, controller assigned, and controller type for the disk slot. <p>Note Shared mode cannot be used under certain conditions when dual HBA controllers are used. To view the conditions for Shared mode for Dual HBA controller, see Table 20: Limitations for Shared Mode for Dual HBA Controller, on page 212.</p> <ul style="list-style-type: none"> • Chassis Global Hot Spare—If this option is selected, you will need to set the value for the Slot Range for the disk.
Assigned to Server column	The ID of the server that the disk is assigned.

Name	Field
Assigned to Controller column	The ID of the controller that the disk is assigned. Note In a Dual RAID setup, to migrate the disk from first controller to second, change the Assigned to Controller to the second controller.
Controller Type column	The type for the controller. If the disk is either dedicated or shared, the controller type is always SAS.

Table 20: Limitations for Shared Mode for Dual HBA Controller

Server	HDD Tray	Controller	Shared mode Support
Cisco UCS S3260	No	Dual HBA	Not Supported
Cisco UCS S3260	HDD Tray	Dual HBA	Not Supported
Pre-Provisioned	HDD Tray	Dual HBA	Not Supported

Creating Disk Slots and Assigning Ownership

After you create a disk zoning policy, you must create the disk slots, and assign ownership.

Procedure

-
- Step 1** In the Navigation pane, click **Chassis**.
 - Step 2** Expand **Policies > root > Disk Zoning Policies**, and select the disk zoning policy that you want to add disk slots.
 - Step 3** In the Work pane, under **Actions**, click **Add Slots to Policy**.
 - Step 4** In the **Add Slots to Policy** dialog box, complete the following:

Name	Description
Ownership check box	<p>The ownership for the disk slot. This can be one of the following:</p> <ul style="list-style-type: none"> • Unassigned—This option is selected by default. You can set the slot number in the Slot Range field. • Dedicated—If this option is selected, you will need to set the values for the Server, Controller, and Slot Range for the disk slot. • Shared—If this option is selected, you will need to set the values for the Slot Range and controller information such as server assigned, controller assigned, and controller type for the disk slot. <p>Note Shared mode cannot be used under certain conditions when dual HBA controllers are used. To view the conditions for Shared mode for Dual HBA controller, see Table 20: Limitations for Shared Mode for Dual HBA Controller, on page 212.</p> <ul style="list-style-type: none"> • Chassis Global Hot Spare—If this option is selected, you will need to set the value for the Slot Range for the disk.

Step 5 Click **OK**.

Associating Disk Zoning Policies to Chassis Profile

Procedure

- Step 1** In the **Navigation** pane, click **Chassis**.
- Step 2** Expand **Chassis > Chassis Profiles**.
- Step 3** Expand the node for the organization where you want to create the chassis profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the organization and select **Create Chassis Profile**.
- Step 5** In the **Identify Chassis Profile** page, specify the name for the chassis profile, and click **Next**.
- Step 6** (Optional) In the **Maintenance Policy** page, specify the name for the maintenance policy, and click **Next**.
- Step 7** In the **Chassis Assignment** page, select **Select existing Chassis** under **Chassis Assignment**, and then select the chassis that you want to associate with this chassis profile. Click **Next**.
- Step 8** In the **Disk Zoning** page, specify the disk zoning policy that you want to associate with this chassis profile.
- Step 9** Click **Finish**.

Disk Migration

Before you can migrate a disk zoned from one server to another, you must mark the virtual drive(LUN) as transport ready or perform a hide virtual drive operation. This will ensure that all references from the service profile have been removed prior to disk migration. For more information on virtual drives, please refer to the "virtual drives" section in the *Cisco UCS Manager Storage Management Guide, Release 3.1 Cisco UCS Manager Storage Management Guide, Release 3.2*



Note In a Dual RAID setup, to migrate the disk from first controller to second change the **Assigned to Controller** to the second controller in the disk zoning policy. Refer [Creating a Disk Zoning Policy, on page 209](#).

Procedure

- Step 1** In the Navigation pane, click **Equipment > Chassis > Servers**.
- Step 2** Choose the Sever where you want to perform disk migration.
- Step 3** In the Work pane, click the **Inventory** tab.
- Step 4** Click the **Storage** subtab.
- Step 5** Click the **LUNs** subtab.
- Step 6** Choose the storage controller where you want to prepare the virtual drives for migration to another server.
- Step 7** Choose the disk that you want to migrate.
- Step 8** In the **Actions** area, choose one of the following:

Name	Description
Rename	Click on this link to rename your disk.
Delete	Click on this link to delete your disk.
Set Transportation Ready	Click on this link for the safe migration of the virtual drive from one server to another. Note All virtual drives on a disk group must be marked as hidden before migrating or unassigning the disks from a server node.

Name	Description
ClearTransportation Ready	Click on this link to set the state of the virtual drive to no longer be transport ready.
Hide Virtual Drive	<p>Click on this option for the safe migration of the virtual drive from one server to another.</p> <p>Note All virtual drives on a disk group must be marked as hidden before migrating or unassigning the disks from a server node.</p>
Unhide Virtual Drive	Click on this link to unhide the virtual drive and enable IO operations.

Storage Enclosure Operations

Removing Chassis Level Storage Enclosures

You can remove the storage enclosure corresponding to HDD expansion tray in Cisco UCS Manager after it is physically removed. You cannot remove server level or any other chassis level storage enclosures.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Chassis > Servers > Storage Enclosures**.
- Step 3** Choose the storage enclosure that you want to remove.
- Step 4** In the **Actions** area, click **Remove Enclosure**.

Sas Expander Configuration Policy

Creating Sas Expander Configuration Policy

Procedure

	Command or Action	Purpose						
Step 1	In the Navigation pane, click Chassis .							
Step 2	Expand Chassis > Policies .							
Step 3	Expand the node for the organization where you want to create the policy.	If the system does not include multi tenancy, expand the root node.						
Step 4	Right-click Sas Expander Configuration Policies and choose Create Sas Expander Configuration Policy .							
Step 5	In the Create Sas Expander Configuration Policy dialog box, complete the following fields:	<table border="1"> <thead> <tr> <th>Name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Name field</td> <td>The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</td> </tr> <tr> <td>Description field</td> <td>A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</td> </tr> </tbody> </table>	Name	Description	Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.	Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Name	Description							
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.							
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).							

	Command or Action	Purpose	
		Name	Description
		6G-12G Mixed Mode field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Connection Management is disabled in this policy and the Sas Expander uses only 6G speeds even if 12G is available. • Enabled—Connection Management is enabled in this policy and it intelligently shifts between 6G and 12 G speeds based on availability. <p>After 6G-12G Mixed Mode is enabled, you cannot downgrade to any release earlier than 3.2(3a). To downgrade, disable this mode.</p> <ul style="list-style-type: none"> • No Change (Default)—Pre-existing configuration is retained. <p>Note Enabling or disabling 6G-12G Mixed Mode causes system reboot.</p> <p>6G-12G Mixed Mode field is available only for Cisco UCS S3260 M5 and higher servers.</p>
Step 6	Click OK .		

Deleting a Sas Expander Configuration Policy

Procedure

- Step 1** In the **Navigation** pane, click **Chassis**.
- Step 2** Expand **Chassis > Policies**.
- Step 3** Expand the node for the organization containing the pool.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand **Sas Expander Configuration Policies**.
- Step 5** Right-click the Sas Expander Configuration policy you want to delete and choose **Delete**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
-