



Cisco UCS Manager Storage Management Guide using the CLI, Release 3.2

First Published: 2017-08-18

Last Modified: 2019-08-27

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2017–2019 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

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

CHAPTER 1

New and Changed Information	1
New and Changed Information for This Release	1

CHAPTER 2

Overview	3
Overview	3
Storage Options	3
Storage Design Considerations	5
Storage Configuration Sequence	5
Storage Protocols	6
The UCS Manager SAN Tab	6

CHAPTER 3

SAN Ports and Port Channels	9
Port Modes	9
Port Types	9
Effect of Port Mode Changes on Data Traffic	10
FC Links Rebalancing	11
Configuring the Port Mode	12
Displaying Port Properties and Fibre Channel Statistics	14
Server Ports	15
Configuring a Server Port	15

Unconfiguring a Server Port	16
Unified Ports	16
Guidelines for Configuring Unified Ports	16
Cautions and Guidelines for Configuring Unified Uplink Ports and Unified Storage Ports	17
Beacon LEDs for Unified Ports	18
Configuring the Beacon LEDs for Unified Ports	19
Unified Ports on the Fabric Interconnect	20
Unified Storage Ports	20
Configuring a Unified Storage Port	20
Unified Uplink Ports	21
Configuring a Unified Uplink Port	21
Unified Uplink Port Channel	22
Configuring a Unified Uplink Port Channel	22
Cisco UCS Mini Scalability Ports	23
Configuring Scalability Ports	23
Appliance Ports	24
Configuring an Appliance Port	24
Assigning a Target MAC Address to an Appliance Port or Appliance Port Channel	26
Creating an Appliance Port	27
Mapping an Appliance Port to a Community VLAN	28
Unconfiguring an Appliance Port	29
FCoE Uplink Ports	30
Configuring a FCoE Uplink Port	30
Viewing FCoE Uplink Ports	31
Unconfiguring a FCoE Uplink Port	31
FCoE and Fibre Channel Storage Ports	32
Configuring a Fibre Channel Storage or FCoE Port	32
Unconfiguring a Fibre Channel Storage or FCoE Port	33
Restoring a Fibre Channel Storage Port Back to an Uplink Fibre Channel Port	33
Appliance Port Channels	34
Configuring an Appliance Port Channel	34
Unconfiguring an Appliance Port Channel	35
Enabling or Disabling an Appliance Port Channel	36
Adding a Member Port to an Appliance Port Channel	37

Deleting a Member Port from an Appliance Port Channel	37
Fibre Channel Port Channels	38
Configuring a Fibre Channel Port Channel	39
Unconfiguring a Fibre Channel Port Channel	40
Adding Channel Mode Active To The Upstream NPIV Fibre Channel Port Channel	40
Enabling or Disabling a Fibre Channel Port Channel	41
Adding a Member Port to a Fibre Channel Port Channel	42
Deleting a Member Port from a Fibre Channel Port Channel	43
FCoE Port Channels	43
Configuring a FCoE Port Channel	44
Adding a Member Port to a FCoE Uplink Port Channel	44
Adapter Port Channels	45
Viewing Adapter Port Channels	45
Event Detection and Action	46
Policy-Based Port Error Handling	47
Creating Threshold Definition	47
Configuring Error Disable on a Fabric Interconnect Port	48
Configuring Auto Recovery on a Fabric Interconnect Port	49
Viewing the Network Interface Port Error Counters	50
Fabric Port Channels	51
Load Balancing Over Ports	51
Cabling Considerations for Fabric Port Channels	52
Viewing Fabric Port Channels	52
Enabling or Disabling a Fabric Port Channel Member Port	53

CHAPTER 4
Fibre Channel Zoning 55

Information About Fibre Channel Zoning	55
Information About Zones	55
Information About Zone Sets	56
Support for Fibre Channel Zoning in Cisco UCS Manager	56
Cisco UCS Manager-Based Fibre Channel Zoning	56
vHBA Initiator Groups	57
Fibre Channel Storage Connection Policy	57
Fibre Channel Active Zone Set Configuration	57

Switch-Based Fibre Channel Zoning	58
Guidelines and recommendations for Cisco UCS Manager-Based Fibre Channel Zoning	58
Configuring Cisco UCS Manager Fibre Channel Zoning	58
Creating a VSAN for Fibre Channel Zoning	59
Creating a New Fibre Channel Zone Profile	60
Deleting a Fibre Channel Zone Profile	61
Deleting a Fibre Channel User Zone	62
Removing Unmanaged Zones from a VSAN Accessible to Both Fabric Interconnects	63
Removing Unmanaged Zones from a VSAN Accessible to One Fabric Interconnect	64
Configuring Fibre Channel Storage Connection Policies	65
Creating a Fibre Channel Storage Connection Policy	65
Deleting a Fibre Channel Storage Connection Policy	66

CHAPTER 5**Named VSANs 67**

Named VSANs	67
Fibre Channel Uplink Trunking for Named VSANs	68
Guidelines and Recommendations for VSANs	68
Creating a Named VSAN Accessible to Both Fabric Interconnects (Fibre Channel Uplink Mode)	70
Creating a Named VSAN Accessible to Both Fabric Interconnects (Fibre Channel Storage Mode)	71
Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Uplink Mode)	72
Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Storage Mode)	74
Deleting a Named VSAN	75
Changing the VLAN ID for the FCoE Native VLAN for a Named VSAN	76
Changing the VLAN ID for the FCoE Native VLAN for a Storage VSAN	76
Enabling or Disabling Fibre Channel Uplink Trunking	77

CHAPTER 6**SAN Pin Groups 79**

SAN Pin Groups	79
Configuring a SAN Pin Group	79
Configuring a FCoE Pin Group	80

CHAPTER 7**FC Identity Assignment 83**

Fibre Channel Identity	83
------------------------	----

CHAPTER 8	WWW Pools	85
	WWN Pools	85
	Creating a WWN Pool	86
	Deleting a WWN Pool	89

CHAPTER 9	Storage-Related Policies	91
	Configuring vHBA Templates	91
	vHBA Template	91
	Configuring a vHBA Template	91
	Deleting a vHBA Template	93
	Configuring Fibre Channel Adapter Policies	93
	Ethernet and Fibre Channel Adapter Policies	93
	Configuring a Fibre Channel Adapter Policy	95
	Deleting a Fibre Channel Adapter Policy	96
	Configuring the Default vHBA Behavior Policy	97
	Default vHBA Behavior Policy	97
	Configuring a Default vHBA Behavior Policy	97
	Configuring SAN Connectivity Policies	98
	About the LAN and SAN Connectivity Policies	98
	Privileges Required for LAN and SAN Connectivity Policies	99
	Interactions between Service Profiles and Connectivity Policies	99
	Creating a SAN Connectivity Policy	99
	Deleting a SAN Connectivity Policy	101
	Creating a vHBA for a SAN Connectivity Policy	101
	Deleting a vHBA from a SAN Connectivity Policy	104
	Creating an Initiator Group for a SAN Connectivity Policy	104
	Deleting an Initiator Group from a SAN Connectivity Policy	106

CHAPTER 10	Storage Profiles	107
	Storage Profiles	107
	Disk Groups and Disk Group Configuration Policies	107
	Virtual Drives	108
	RAID Levels	109

- Automatic Disk Selection 111
- Supported LUN Modifications 111
- Unsupported LUN Modifications 112
- Disk Insertion Handling 112
 - Non-Redundant Virtual Drives 113
 - Redundant Virtual Drives with No Hot Spare Drives 113
 - Redundant Virtual Drives with Hot Spare Drives 113
 - Replacing Hot Spare Drives 113
 - Inserting Physical Drives into Unused Slots 114
- Virtual Drive Naming 114
- LUN Dereferencing 114
- Controller Constraints and Limitations 115
- Configuring Storage Profiles 115
 - Configuring a Disk Group Policy 115
 - Setting the RAID Level 115
 - Automatically Configuring Disks in a Disk Group 116
 - Manually Configuring Disks in a Disk Group 118
 - Configuring Virtual Drive Properties 119
- Creating a Storage Profile 122
- Deleting a Storage Profile 123
- Creating Local LUNs 124
 - Reordering Local LUNs In a Storage Profile 126
 - Deleting Local LUNs In a Storage Profile 127
- Associating a Storage Profile with a Service Profile 128
- Displaying Details of All Local LUNs Inherited By a Service Profile 129
- Importing Foreign Configurations for a RAID Controller 132
- Configuring Local Disk Operations 132
- Configuring Virtual Drive Properties 134
 - Deleting an Orphaned Virtual Drive 136
 - Renaming an Orphaned Virtual Drive 139
- Configuring the Boot Policy for a Local LUN 140
- Local LUN Operations in a Service Profile 141
 - Preprovisioning a LUN Name or Claiming an Orphan LUN 141
 - Deploying and Undeploying a LUN 142

Renaming a Service Profile Referenced LUN 143

CHAPTER 11

Configuring SD Card Support 145

FlexFlash Secure Digital Card Support 145

Enabling or Disabling FlexFlash SD Card Support 147

FlexUtil Secure Digital Card Support 148

CHAPTER 12

Mini Storage 149

Mini Storage 149

Viewing Mini Storage Properties 149

Viewing the Storage Controller for the Mini Storage 150

CHAPTER 13

SED Security Policies 153

Security Policies for Self-Encrypting Drives 153

Security Flags of the Controller and Disk 154

Managing Local Security Policies 154

Creating a Local Security Policy 154

Modifying the Security Key of a Local Security Policy 155

Modifying the Security Policy from Local to Remote 156

Inserting a Secured Disk into a Server with a Local Security Policy 158

KMIP Client Certificate Policy 159

Creating a Global KMIP Client Certificate Policy 159

Creating a KMIP Client Certificate for a Server 160

Managing Remote Security Policies 162

Creating a Remote Security Policy 162

Modifying a Remote Security Key 164

Modifying the Security Policy from Remote to Local 165

Inserting a Secured Disk into a Server with a Remote Security Policy 166

Securing an Existing Virtual Drive 166

Enabling Security on a Disk 167

Erasing a Secure Disk 168

Disabling Security on a Controller 169

Unlocking a Locked Disk 169

Erasing a Secure Foreign Configuration Disk 170

Displaying the Security Flags of a Controller 171
 Displaying the Security Flags of a Local Disk 172
 Displaying the Security Flags of a Virtual Drive 173

CHAPTER 14

Cisco UCS S3260 System Storage Management 175
 Storage Server Features and Components Overview 175
 Cisco UCS S3260 Storage Management Operations 183
 Disk Sharing for High Availability 184
 Disk Zoning Policies 184
 Creating a Disk Zoning Policy 185
 Creating Disk Slots and Assigning Ownership 185
 Associating Disk Zoning Policies to Chassis Profile 187
 Disk Migration 188
 Storage Enclosure Operations 189
 Removing Chassis Level Storage Enclosures 189
 SAS Expander Configuration Policy 190
 Creating SAS Expander Configuration Policy 190
 Deleting a SAS Expander Configuration Policy 191



Preface

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

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](#).

Other Documentation Resources

Follow [Cisco UCS Docs on Twitter](#) to receive document update notifications.

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 for This Release, on page 1](#)

New and Changed Information for This Release

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to this guide or of all new features in this release.

Table 1: New Features and Changed Behavior in Cisco UCS Manager, Release 3.2(3)

Feature	Description	Where Documented
6G-12G Mixed Mode	6G-12G Mixed Mode is available for SAS Expander Configuration Policy.	Creating SAS Expander Configuration Policy, on page 190
Single Path support for UCS S3260 Dual Pass Through (UCS-S3260-DHBA) controller	In Cisco UCS S3260 M5 servers, Cisco UCS Manager enables single path access to disks by configuring a single DiskPort per disk slot.	Disk Zoning Policies, on page 184



CHAPTER 2

Overview

- [Overview, on page 3](#)
- [Storage Options, on page 3](#)
- [Storage Design Considerations, on page 5](#)
- [Storage Configuration Sequence, on page 5](#)
- [Storage Protocols, on page 6](#)
- [The UCS Manager SAN Tab, on page 6](#)

Overview

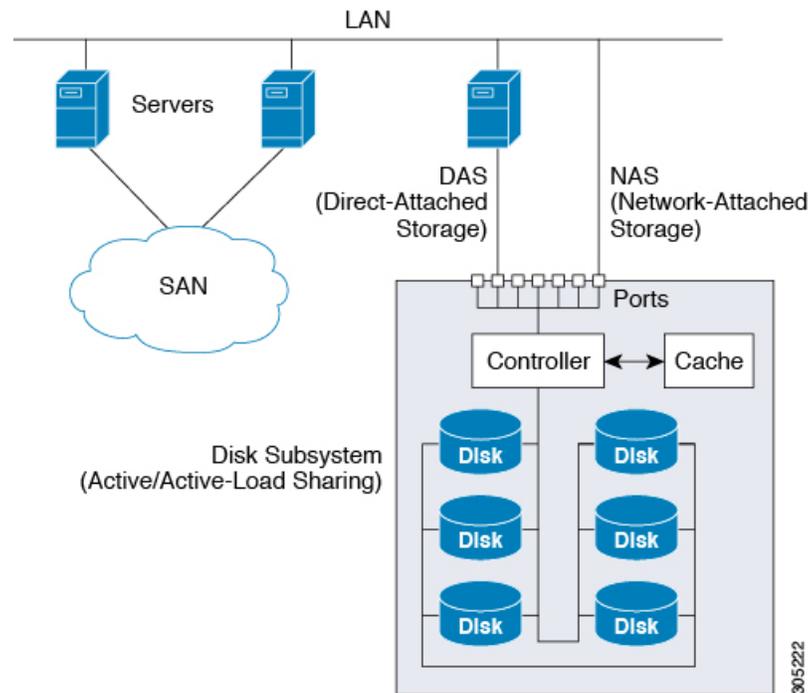
This guide describes how to configure the following storage management tasks:

- 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

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](#)
- [Effect of Port Mode Changes on Data Traffic, on page 10](#)
- [FC Links Rebalancing, on page 11](#)
- [Configuring the Port Mode, on page 12](#)
- [Displaying Port Properties and Fibre Channel Statistics, on page 14](#)
- [Server Ports, on page 15](#)
- [Unified Ports, on page 16](#)
- [Cisco UCS Mini Scalability Ports, on page 23](#)
- [Appliance Ports, on page 24](#)
- [FCoE Uplink Ports, on page 30](#)
- [FCoE and Fibre Channel Storage Ports, on page 32](#)
- [Appliance Port Channels, on page 34](#)
- [Fibre Channel Port Channels, on page 38](#)
- [FCoE Port Channels, on page 43](#)
- [Adapter Port Channels, on page 45](#)
- [Event Detection and Action, on page 46](#)
- [Fabric Port Channels, on page 51](#)

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
- FCoE Uplink ports
- SPAN source ports



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

Effect of Port Mode Changes on Data Traffic

Port mode changes can cause an interruption to the data traffic for the Cisco UCS domain. The length of the interruption and the traffic that is affected depend upon the configuration of the Cisco UCS domain and the module on which you made the port mode changes.



Tip To minimize the traffic disruption during system changes, form a Fibre Channel uplink port-channel across the fixed and expansion modules.

Impact of Port Mode Changes on an Expansion Module

After you make port mode changes on an expansion module, the module reboots. All traffic through ports on the expansion module is interrupted for approximately one minute while the module reboots.

Impact of Port Mode Changes on the Fixed Module in a Cluster Configuration

A cluster configuration has two fabric interconnects. After you make port changes to the fixed module, the fabric interconnect reboots. The impact on the data traffic depends upon whether or not you have configured the server vNICs to failover to the other fabric interconnect when one fails.

If you change the port modes on the expansion module of one fabric interconnect and then wait for that to reboot before changing the port modes on the second fabric interconnect, the following occurs:

- With server vNIC failover, traffic fails over to the other fabric interconnect and no interruption occurs.
- Without server vNIC failover, all data traffic through the fabric interconnect on which you changed the port modes is interrupted for approximately eight minutes while the fabric interconnect reboots.

If you change the port modes on the fixed modules of both fabric interconnects simultaneously, all data traffic through the fabric interconnects are interrupted for approximately eight minutes while the fabric interconnects reboot.

Impact of Port Mode Changes on the Fixed Module in a Standalone Configuration

A standalone configuration has only one fabric interconnect. After you make port changes to the fixed module, the fabric interconnect reboots. All data traffic through the fabric interconnect is interrupted for approximately eight minutes while the fabric interconnect reboots.

FC Links Rebalancing

The FC uplinks balance automatically when FC Port Channels are utilized. To create FC Port Channels, refer to [Configuring a Fibre Channel Port Channel, on page 39](#).

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.

Configuring the Port Mode



Caution Changing the port mode on either module can cause an interruption in data traffic because changes to the fixed module require a reboot of the fabric interconnect and changes on an expansion module require a reboot of that module .

If the Cisco UCS domain has a cluster configuration that is set up for high availability and servers with service profiles that are configured for failover, traffic fails over to the other fabric interconnect and data traffic is not interrupted when the port mode is changed on the fixed module.

In the Cisco UCS Manager CLI, there are no new commands to support Unified Ports. Instead, you change the port mode by scoping to the mode for the desired port type and then creating a new interface. When you create a new interface for an already configured slot ID and port ID, UCS Manager deletes the previously configured interface and creates a new one. If a port mode change is required because you configure a port that previously operated in Ethernet port mode to a port type in Fibre Channel port mode, UCS Manager notes the change.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope <i>port-type-mode</i>	<p>Enters the specified port type mode for one of the following port types:</p> <p>eth-server For configuring server ports.</p> <p>eth-storage For configuring Ethernet storage ports and Ethernet storage port channels.</p> <p>eth-traffic-mon For configuring Ethernet SPAN ports.</p> <p>eth-uplink For configuring Ethernet uplink ports.</p> <p>fc-storage For configuring Fibre Channel storage ports.</p> <p>fc-traffic-mon For configuring Fibre Channel SPAN ports.</p>

	Command or Action	Purpose
		fc-uplink For configuring Fibre Channel uplink ports and Fibre Channel uplink port channels.
Step 2	UCS-A <i>/port-type-mode</i> # scope fabric {a b}	Enters the specified port type mode for the specified fabric.
Step 3	UCS-A <i>/port-type-mode/fabric</i> # create interface slot-id port-id	Creates an interface for the specified port type. If you are changing the port type from Ethernet port mode to Fibre Channel port mode, or vice-versa, the following warning appears: Warning: This operation will change the port mode (from Ethernet to FC or vice-versa). When committed, this change will require the module to restart.
Step 4	Create new interfaces for other ports belonging to the Ethernet or Fibre Channel port block.	There are several restrictions that govern how Ethernet and Fibre Channel ports can be arranged on a fixed or expansion module. Among other restrictions, it is required that you change ports in groups of two. Violating any of the restrictions outlined in the Guidelines for Configuring Unified Ports section will result in an error.
Step 5	UCS-A <i>/port-type-mode/fabric/interface</i> # commit-buffer	Commits the transaction to the system configuration.

Based on the module for which you configured the port modes, data traffic for the Cisco UCS domain is interrupted as follows:

- Fixed module—The fabric interconnect reboots. All data traffic through that fabric interconnect is interrupted. In a cluster configuration that provides high availability and includes servers with vNICs that are configured for failover, traffic fails over to the other fabric interconnect and no interruption occurs. Changing the port mode for both sides at once results in both fabric interconnects rebooting simultaneously and a complete loss of traffic until both fabric interconnects are brought back up.

It takes about 8 minutes for the fixed module to reboot.

- Expansion module—The module reboots. All data traffic through ports in that module is interrupted.

It takes about 1 minute for the expansion module to reboot.

Example

The following example changes ports 3 and 4 on slot 1 from Ethernet uplink ports in Ethernet port mode to uplink Fibre Channel ports in Fibre Channel port mode:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create interface 1 3
```

```
Warning: This operation will change the port mode (from Ethernet to FC or vice-versa).
When committed, this change will require the fixed module to restart.
UCS-A /fc-uplink/fabric/interface* # up
UCS-A /fc-uplink/fabric* #create interface 1 4
Warning: This operation will change the port mode (from Ethernet to FC or vice-versa).
When committed, this change will require the fixed module to restart.
UCS-A /fc-uplink/fabric/interface* #commit-buffer
```

Displaying Port Properties and Fibre Channel Statistics

Procedure

	Command or Action	Purpose
Step 1	UCS-A /fabric-interconnect # connect nxos {a b}	Enters NX-OS mode for the fabric interconnect.
Step 2	UCS-A(nxos)# show interface fc <i>slot-id/port-id</i>	Displays the port properties and Fibre Channel statistics such as throughput rates and errors. Note The receive B2B credit remaining on a UCS 6400 Series fabric interconnect does not match the transmit B2B credit remaining on its peer switch. The receive B2B credit remaining parameter always remains at 64 because the credit is returned when the frame is released. On an MDS peer, the transmit B2B credit remaining parameter can go to 0. This results in the mismatch. You can run the show interface fc <i>slot-id/port-id</i> command on the peer switch to compare these parameters.

Example

The following example displays port properties and Fibre Channel statistics for UCS 6400 Series fabric interconnects:

```
UCS-A /fabric-interconnect # connect nxos a
UCS-A(nxos)# show interface fc 1/6
fc1/6 is trunking
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:06:00:de:fb:21:77:00
  Admin port mode is NP, trunk mode is on
  snmp link state traps are enabled
  Port mode is TNP
  Port vsan is 8
  Speed is 16 Gbps
  Transmit B2B Credit is 32
  Receive B2B Credit is 64
  Receive data field Size is 2112
```

```

Beacon is turned off
Belongs to san-port-channel 32
Trunk vsans (admin allowed and active) (1,5,7-8,40,120)
Trunk vsans (up) (1,5,7-8,40,120)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
5 minutes input rate 497578904 bits/sec,62197363 bytes/sec, 30981 frames/sec
5 minutes output rate 501679056 bits/sec,62709882 bytes/sec, 30319 frames/sec
430000799 frames input,863205473268 bytes
  0 discards,0 errors
  0 invalid CRC/FCS,0 unknown class
  0 too long,0 too short
423530360 frames output,876009587416 bytes
  0 discards,0 errors
  1 input OLS,1 LRR,9 NOS,0 loop inits
  1 output OLS,0 LRR, 8 NOS, 0 loop inits
  64 receive B2B credit remaining
  32 transmit B2B credit remaining
  0 low priority transmit B2B credit remaining
Last clearing of "show interface" counters :never

```

Server Ports

Configuring a Server Port

All of the port types listed are configurable on both the fixed and expansion module, including server ports, which are not configurable on the 6100 series fabric interconnect expansion module, but are configurable on the 6200 series fabric interconnect expansion module.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # create interface slot-num port-num	Creates an interface for the specified Ethernet server port.
Step 4	UCS-A /eth-server/fabric # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create an interface for Ethernet server port 4 on slot 1 of fabric B and commit the transaction:

```

UCS-A# scope eth-server
UCS-A /eth-server # scope fabric b
UCS-A /eth-server/fabric # create interface 1 4

```

```
UCS-A /eth-server/fabric* # commit-buffer
UCS-A /eth-server/fabric #
```

Unconfiguring a Server Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # delete interface slot-num port-num	Deletes the interface for the specified Ethernet server port.
Step 4	UCS-A /eth-server/fabric # commit-buffer	Commits the transaction to the system configuration.

Example

The following example unconfigures Ethernet server port 12 on slot 1 of fabric B and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope fabric b
UCS-A /eth-server/fabric # delete interface 1 12
UCS-A /eth-server/fabric* # commit-buffer
UCS-A /eth-server/fabric #
```

Unified Ports

Guidelines for Configuring Unified Ports

Consider the following guidelines and restrictions when configuring unified ports:

Hardware and Software Requirements

Unified ports are supported on the following:

- UCS 6300 Series Fabric Interconnects with Cisco UCS Manager Release 3.1 and later releases
- UCS 6200 Series Fabric Interconnects with Cisco UCS Manager Release 2.0 and later releases
- UCS 6324 Fabric Interconnects with Cisco UCS Manager Release 3.0 and later releases

Unified ports are not supported on 6100 series Fabric Interconnects, even if they are running Cisco UCS Manager, version 2.0.

Port Mode Placement

Because the Cisco UCS Manager GUI interface uses a slider to configure the port mode for unified ports on a fixed or expansion module, it automatically enforces the following restrictions which limits how port modes can be assigned to unified ports. When using the Cisco UCS Manager CLI interface, these restrictions are enforced when you commit the transaction to the system configuration. If the port mode configuration violates any of the following restrictions, the Cisco UCS Manager CLI displays an error:

- Ethernet ports must be grouped together in a block. For each module (fixed or expansion), the Ethernet port block must start with the first port and end with an even numbered port.
- Fibre Channel ports must be grouped together in a block. For each module (fixed or expansion), the first port in the Fibre Channel port block must follow the last Ethernet port and extend to include the rest of the ports in the module. For configurations that include only Fibre Channel ports, the Fibre Channel block must start with the first port on the fixed or expansion module.
- Alternating Ethernet and Fibre Channel ports is not supported on a single module.

Example of a valid configuration— Might include unified ports 1–16 on the fixed module configured in Ethernet port mode and ports 17–32 in Fibre Channel port mode. On the expansion module you could configure ports 1–4 in Ethernet port mode and then configure ports 5–16 in Fibre Channel mode. The rule about alternating Ethernet and Fibre Channel port types is not violated because this port arrangement complies with the rules on each individual module.

Example of an invalid configuration— Might include a block of Fibre Channel ports starting with port 16. Because each block of ports has to start with an odd-numbered port, you would have to start the block with port 17.



Note The total number of uplink Ethernet ports and uplink Ethernet port channel members that can be configured on each fabric interconnect is limited to 31. This limitation includes uplink Ethernet ports and uplink Ethernet port channel members configured on the expansion module.

Special Considerations for UCS Manager CLI Users

Because the Cisco UCS Manager CLI does not validate port mode changes until you commit the buffer to the system configuration, it is easy to violate the grouping restrictions if you attempt to commit the buffer before creating at least two new interfaces. To prevent errors, we recommend that you wait to commit your changes to the system configuration until you have created new interfaces for all of the unified ports changing from one port mode to another.

Committing the buffer before configuring multiple interfaces will result in an error, but you do not need to start over. You can continue to configure unified ports until the configuration satisfies the aforementioned requirements.

Cautions and Guidelines for Configuring Unified Uplink Ports and Unified Storage Ports

The following are cautions and guidelines to follow while working with unified uplink ports and unified storage ports:

- In an unified uplink port, if you enable one component as a SPAN source, the other component will automatically become a SPAN source.



Note If you create or delete a SPAN source under the Ethernet uplink port, Cisco UCS Manager automatically creates or deletes a SPAN source under the FCoE uplink port. The same happens when you create a SPAN source on the FCoE uplink port.

- You must configure a non default native VLAN on FCoE and unified uplink ports. This VLAN is not used for any traffic. Cisco UCS Manager will reuse an existing fcoe-storage-native-vlan for this purpose. This fcoe-storage-native-vlan will be used as a native VLAN on FCoE and unified uplinks.
- In an unified uplink port, if you do not specify a non default VLAN for the Ethernet uplink port the fcoe-storage-native-vlan will be assigned as the native VLAN on the unified uplink port. If the Ethernet port has a non default native VLAN specified as native VLAN, this will be assigned as the native VLAN for unified uplink port.
- When you create or delete a member port under an Ethernet port channel, Cisco UCS Manager automatically creates or deletes the member port under FCoE port channel. The same happens when you create or delete a member port in FCoE port channel.
- When you configure an Ethernet port as a standalone port, such as server port, Ethernet uplink, FCoE uplink or FCoE storage and make it as a member port for an Ethernet or FCoE port channel, Cisco UCS Manager automatically makes this port as a member of both Ethernet and FCoE port channels.
- When you remove the membership for a member port from being a member of server uplink, Ethernet uplink, FCoE uplink or FCoE storage, Cisco UCS Manager deletes the corresponding members ports from Ethernet port channel and FCoE port channel and creates a new standalone port.
- If you downgrade Cisco UCS Manager from release 2.1 to any of the prior releases, all unified uplink ports and port channels will be converted to Ethernet ports and Ethernet port channels when the downgrade is complete. Similarly, all the unified storage ports will be converted to appliance ports.
- For unified uplink ports and unified storage ports, when you create two interfaces, only one license is checked out. As long as either interface is enabled, the license remains checked out. The license will be released only if both the interfaces are disabled for a unified uplink port or a unified storage port.
- Cisco UCS 6100 series fabric interconnect switch can only support 1VF or 1VF-PO facing same downstream NPV switch.

Beacon LEDs for Unified Ports

Each port on the 6200 series fabric interconnect has a corresponding beacon LED. When the **Beacon LED** property is configured, the beacon LEDs illuminate, showing you which ports are configured in a given port mode.

You can configure the **Beacon LED** property to show you which ports are grouped in one port mode: either Ethernet or Fibre Channel. By default, the Beacon LED property is set to Off.



Note For unified ports on the expansion module, you can reset the **Beacon LED** property to the default value of **Off** during expansion module reboot.

Configuring the Beacon LEDs for Unified Ports

Complete the following task for each module for which you want to configure beacon LEDs.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fabric-interconnect {a b}	Enters fabric interconnect mode for the specified fabric.
Step 2	UCS-A /fabric # scope card slot-id	Enters card mode for the specified fixed or expansion module.
Step 3	UCS-A /fabric/card # scope beacon-led	Enters beacon LED mode.
Step 4	UCS-A /fabric/card/beacon-led # set admin-state {eth fc off}	Specifies which port mode is represented by illuminated beacon LED lights. eth All of the Unified Ports configured in Ethernet mode illuminate. fc All of the Unified Ports configured in Fibre Channel mode illuminate. off Beacon LED lights for all ports on the module are turned off.
Step 5	UCS-A /fabric/card/beacon-led # commit-buffer	Commits the transaction to the system configuration.

Example

The following example illuminates all of the beacon lights for Unified Ports in Ethernet port mode and commits the transaction:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric # scope card 1
UCS-A /fabric/card # scope beacon-led
UCS-A /fabric/card/beacon-led # set admin-state eth
UCS-A /fabric/card/beacon-led* # commit-buffer
UCS-A /fabric/card/beacon-led #
```

Unified Ports on the Fabric Interconnect

Unified ports are ports on the fabric interconnect that can be configured to carry either Ethernet or Fibre Channel traffic. These ports are not reserved. A Cisco UCS domain cannot use these ports until you configure them.



Note When you configure a port on a fabric interconnect, the administrative state is automatically set to enabled. If the port is connected to another device, this may cause traffic disruption. You can disable the port after configuring it.

Configurable beacon LEDs indicate which unified ports are configured for the selected port mode.

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, on either a fixed module or an expansion module. 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 a Unified Storage Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # create interface slot-num port-num	Creates an interface for the specified appliance port.
Step 4	UCS-A /eth-storage/fabric/interface* # commit buffer	Commits the transaction to the system configuration.

	Command or Action	Purpose
Step 5	UCS-A /eth-storage/fabric/interface* # scope fc-storage	Enters FC storage mode.
Step 6	UCS-A /fc-storage* # scope fabric {a b}	Enters Ethernet storage mode for the specific appliance port.
Step 7	UCS-A /fc-storage/fabric # create interface fcoe slot-num port-num	Adds FCoE storage port mode on the appliance port mode and creates a unified storage port..

Example

The following example creates an interface for an appliance port 2 on slot 3 of fabric A, adds fc storage to the same port to convert it as a unified port, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # create interface 3 2
UCS-A /eth-storage/fabric* # commit-buffer
UCS-A /eth-storage/fabric* # scope fc-storage
UCS-A /fc-storage*# scope fabric a
UCS-A /fc-storage/fabric* # create interface fcoe 3 2
UCS-A /fc-storage/fabric* # commit-buffer
UCS-A /fc-storage/fabric*
```

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 a Unified Uplink Port

To configure a unified uplink port, you will convert an existing FCoE uplink port as a unified port.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters Ethernet uplink fabric mode for the specified fabric.

	Command or Action	Purpose
Step 3	UCS-A /eth-uplink/fabric # create interface 15	Converts the FCoE uplink port as a unified port.
Step 4	UCS-A /eth-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a unified uplink port on an existing FCoE port:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric b
UCS-A /eth-uplink/fabric # create interface 1 5
UCS-A /eth-uplink/fabric/interface* # commit-buffer
UCS-A /eth-uplink/interface #
```

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.

Configuring a Unified Uplink Port Channel

To configure a unified uplink port channel, you will convert an existing FCoE uplink port channel as a unified port channel.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.

	Command or Action	Purpose
Step 2	UCS-A /eth-uplink # scope fabric {a b}	Enters Ethernet uplink fabric mode for the specified fabric.
Step 3	UCS-A /eth-uplink/fabric # create port-channel ID	Creates a port channel for the specified Ethernet uplink port.
Step 4	UCS-A /eth-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a unified uplink port channel on an existing FCoE port channel:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric b
UCS-A /eth-uplink/fabric # create port-channel 2
UCS-A /eth-uplink/fabric/port-channel* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Cisco UCS Mini Scalability Ports

The contains a scalability port as well as four unified ports. The scalability port is a 40GB QSFP+ breakout port that, with proper cabling, can support four 1G or 10G SFP+ ports. The scalability ports can be used as a licensed server port for supported Cisco UCS rack servers, an appliance port, or a FCoE port.

In the Cisco UCS Manager GUI, the scalability port is displayed as **Scalability Port 5** below the **Ethernet Ports** node. The individual breakout ports are displayed as **Port 1** through **Port 4**.

In the Cisco UCS Manager CLI, the scalability port is not displayed, but the individual breakout ports are displayed as **Br-Eth1/5/1** through **Br-Eth1/5/4**.

Configuring Scalability Ports

To configure ports, port channel members or SPAN members on the scalability port, scope into the scalability port first, then follow the steps for a standard unified port.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # scope aggr-interface slot-num port-num	Enters ethernet server fabric aggregate interface mode for the scalability port.

	Command or Action	Purpose
Step 4	UCS-A /eth-server/fabric/aggr-interface # show interface	Displays the interfaces on the scalability port.
Step 5	UCS-A /eth-server/fabric/aggr-interface # create interface slot-num port-num	Creates an interface for the specified Ethernet server port.
Step 6	UCS-A /eth-server/fabric/aggr-interface # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create an interface for Ethernet server port 3 on the fabric A scalability port and commit the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope fabric a
UCS-A /eth-server/fabric # scope aggr-interface 1 5
UCS-A /eth-server/fabric/aggr-interface # show interface
Interface:

Slot Id Aggr-Port ID Port Id Admin State Oper State State Reason
-----
      1          5      1 Enabled      Up
      1          5      2 Enabled      Up
      1          5      3 Enabled    Admin Down  Administratively Down
      1          5      4 Enabled    Admin Down  Administratively Down

UCS-A /eth-server/fabric/aggr-interface # create interface 1 3
UCS-A /eth-server/fabric/aggr-interface* # commit-buffer
UCS-A /eth-server/fabric/aggr-interface #
```

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.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # create interface slot-num port-num	Creates an interface for the specified appliance port.
Step 4	(Optional) UCS-A /eth-storage/fabric/interface # set portmode {access trunk}	<p>Specifies whether the port mode is access or trunk. By default, the mode is set to trunk.</p> <p>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.</p>
Step 5	(Optional) UCS-A /eth-storage/fabric/interface # set pingroupname pin-group name	Specifies the appliance pin target to the specified fabric and port, or fabric and port channel.
Step 6	(Optional) UCS-A /eth-storage/fabric/interface # set prio sys-class-name	<p>Specifies the QoS class for the appliance port. By default, the priority is set to best-effort.</p> <p>The sys-class-name argument can be one of the following class keywords:</p> <ul style="list-style-type: none"> • Fc—Use this priority for QoS policies that control vHBA traffic only. • Platinum—Use this priority for QoS policies that control vNIC traffic only. • Gold—Use this priority for QoS policies that control vNIC traffic only. • Silver—Use this priority for QoS policies that control vNIC traffic only. • Bronze—Use this priority for QoS policies that control vNIC traffic only. • Best Effort—Do not use this priority. It is reserved for the Basic Ethernet traffic lane. If you assign this priority to a QoS policy and configure another system class

	Command or Action	Purpose
		as CoS 0, Cisco UCS Manager does not default to this system class. It defaults to the priority with CoS 0 for that traffic.
Step 7	(Optional) UCS-A /eth-storage/fabric/interface # set adminspeed {10gbps 1 gbps}	Specifies the admin speed for the interface. By default, the admin speed is set to 10gbps.
Step 8	UCS-A /eth-storage/fabric/interface # commit buffer	Commits the transaction to the system configuration.

Example

The following example creates an interface for an appliance port 2 on slot 3 of fabric B, sets the port mode to access, pins the appliance port to a pin group called pingroup1, sets the QoS class to fc, sets the admin speed to 10 gbps, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric b
UCS-A /eth-storage/fabric # create interface 3 2
UCS-A /eth-storage/fabric* # set portmode access
UCS-A /eth-storage/fabric* # set pingroupname pingroup1
UCS-A /eth-storage/fabric* # set prio fc
UCS-A /eth-storage/fabric* # set adminspeed 10gbps
UCS-A /eth-storage/fabric* # commit-buffer
UCS-A /eth-storage/fabric #
```

What to do next

Assign a VLAN or target MAC address for the appliance port.

Assigning a Target MAC Address to an Appliance Port or Appliance Port Channel

The following procedure assigns a target MAC address to an appliance port. To assign a target MAC address to an appliance port channel, scope to the port channel instead of the interface.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.

	Command or Action	Purpose
Step 3	UCS-A /eth-storage/fabric # scope interface <i>slot-id port-id</i>	Enters Ethernet interface mode for the specified interface. Note To assign a target MAC address to an appliance port channel, use the scope port-channel command instead of scope interface .
Step 4	UCS-A /eth-storage/fabric/interface # create eth-target <i>eth-target name</i>	Specifies the name for the specified MAC address target.
Step 5	UCS-A /eth-storage/fabric/interface/eth-target # set mac-address <i>mac-address</i>	Specifies the MAC address in nn:nn:nn:nn:nn:nn format.

Example

The following example assigns a target MAC address for an appliance device on port 3, slot 2 of fabric B and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage* # scope fabric b
UCS-A /eth-storage/fabric* # scope interface 2 3
UCS-A /eth-storage/fabric/interface* # create eth-target macname
UCS-A /eth-storage/fabric/interface* # set mac-address 01:23:45:67:89:ab
UCS-A /eth-storage/fabric/interface* # commit-buffer
UCS-A /eth-storage/fabric #
```

The following example assigns a target MAC address for appliance devices on port channel 13 of fabric B and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage* # scope fabric b
UCS-A /eth-storage/fabric* # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel* # create eth-target macname
UCS-A /eth-storage/fabric/port-channel* # set mac-address 01:23:45:67:89:ab
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric #
```

Creating an Appliance Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A/eth-storage# create vlan <i>vlan-name</i> <i>vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet storage VLAN mode
Step 3	UCS-A/eth-storage/vlan# set sharing primary	Saves the changes.

	Command or Action	Purpose
Step 4	UCS-A/eth-storage/vlan# commit buffer	Commits the transaction to the system configuration.
Step 5	UCS-A/eth-storage# create vlan <i>vlan-name</i> <i>vlan-id</i>	Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet storage VLAN mode .
Step 6	UCS-A/eth-storage/vlan# set sharing community	Associates the primary VLAN to the secondary VLAN that you are creating.
Step 7	UCS-A/eth-storage/vlan# set pubnwnname <i>primary vlan-name</i>	Specifies the primary VLAN to be associated with this secondary VLAN.
Step 8	UCS-A/eth-storage/vlan# commit buffer	Commits the transaction to the system configuration.

Example

The following example creates an appliance port:

```
UCS-A# scope eth-storage
UCS-A/eth-storage# create vlan PRI600 600
UCS-A/eth-storage/vlan* # set sharing primary
UCS-A/eth-storage/vlan* # commit-buffer
UCS-A/eth-storage # create vlan COM602 602
UCS-A/eth-storage/vlan* # set sharing isolated
UCS-A/eth-storage/vlan* # set pubnwnname PRI600
UCS-A/eth-storage/vlan* # commit-buffer
```

Mapping an Appliance Port to a Community VLAN

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A/eth-storage# scope fabric { <i>a</i> <i>b</i> }	Enters Ethernet storage fabric interconnect mode for the specified fabric interconnect.
Step 3	UCS-A/eth-storage/fabric# create interface <i>slot-num port-num</i>	Creates an interface for the specified Ethernet server port.
Step 4	UCS-A/eth-storage/fabric/interface# exit	Exits from the interface. Note Ensure you commit the transaction after associating with the VLAN.
Step 5	UCS-A/eth-storage/fabric# exit	Exits from the fabric.
Step 6	UCS-A/eth-storage# scope vlan <i>vlan-name</i>	Enters the specified VLAN.

	Command or Action	Purpose
		Note Ensure community VLAN is created in the appliance cloud.
Step 7	UCS-A/eth-storage/vlan# create member-port <i>fabric slot-num port-num</i>	Creates the member port for the specified fabric, assigns the slot number, and port number and enters member port configuration.
Step 8	UCS-A/eth-storage/vlan/member-port# commit	Commits the transaction to the system configuration.

Example

The following example maps an appliance port to an community VLAN:

```
UCS-A# scope eth-storage
UCS-A/eth-storage# scope fabric a
UCS-A/eth-storage/fabric# create interface 1 22
UCS-A/eth-storage/fabric/interface*# exit
UCS-A/eth-storage/fabric*# exit
UCS-A/eth-storage*# scope vlan COM602
UCS-A/eth-storage/vlan*# create member-port a 1 22
UCS-A/eth-storage/vlan/member-port* commit
```

Unconfiguring an Appliance Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # delete eth-interface <i>slot-num port-num</i>	Deletes the interface for the specified appliance port.
Step 4	UCS-A /eth-storage/fabric # commit-buffer	Commits the transaction to the system configuration.

Example

The following example unconfigures appliance port 3 on slot 2 of fabric B and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric b
UCS-A /eth-storage/fabric # delete eth-interface 2 3
UCS-A /eth-storage/fabric* # commit-buffer
UCS-A /eth-storage/fabric #
```

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 a FCoE Uplink Port

All of the port types listed are configurable on both the fixed and expansion module, including server ports, which are not configurable on the 6100 series fabric interconnect expansion module, but are configurable on the 6200 series fabric interconnect expansion module.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters FC Uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters FC - Uplink mode for the specific fabric.
Step 3	UCS-A /fc-uplink/fabric # create fcoeinterface slot-numberport-number	Creates interface for the specified FCoE uplink port.
Step 4	UCS-A /fc-uplink/fabric/fabricinterface # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates an interface for FCoE uplink port 1 on slot 8 of fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create fcoeinterface 1 8
UCS-A /fc-uplink/fabric/fcoeinterface* # commit-buffer
UCS-A /fc-uplink/fabric/fcoeinterface #
```

Viewing FCoE Uplink Ports

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters FC Uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters FC - Uplink mode for the specific fabric.
Step 3	UCS-A /fc-uplink/fabric # show fcoeinterface	Lists the available interfaces.

Example

The following example displays the available FCoE uplink interfaces on fabric A:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # show fcoeinterface
FCoE Interface:

Slot Id      Port Id      Admin State Operational State Operational State Reason  Li
c State      Grace Prd
-----
-----
1           26 Enabled    Indeterminate
cense Ok           0

Fcoe Member Port:

Port-channel Slot  Port  Oper State      State Reason
-----
1           1    10 Sfp Not Present Unknown
1           1     3 Sfp Not Present Unknown
1           1     4 Sfp Not Present Unknown
1           1     6 Sfp Not Present Unknown
1           1     8 Sfp Not Present Unknown
2           1     7 Sfp Not Present Unknown
UCS-A /fc-uplink/fabric #
```

Unconfiguring a FCoE Uplink Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters FC Uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters FC - Uplink mode for the specific fabric.
Step 3	UCS-A /fc-uplink/fabric # delete fcoeinterface slot-numberport-number	Deletes the specified interface.
Step 4	UCS-A /fc-uplink/fabric/fabricinterface # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the FCoE uplink interface on port 1 on slot 8 of fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # delete fcoeinterface 1 8
UCS-A /fc-uplink/fabric/fcoeinterface* # commit-buffer
UCS-A /fc-uplink/fabric/fcoeinterface #
```

FCoE and Fibre Channel Storage Ports

Configuring a Fibre Channel Storage or FCoE Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope fabric {a b}	Enters Fibre Channel storage mode for the specified fabric.
Step 3	UCS-A /fc-storage/fabric # create interface {fc fcoe} slot-num port-num	Creates an interface for the specified Fibre Channel storage port.
Step 4	UCS-A /fc-storage/fabric # commit-buffer	Commits the transaction.

Example

The following example creates an interface for Fibre Channel storage port 10 on slot 2 of fabric A and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # scope fabric a
UCS-A /fc-storage/fabric* # create interface fc 2 10
UCS-A /fc-storage/fabric # commit-buffer
```

What to do next

Assign a VSAN.

Unconfiguring a Fibre Channel Storage or FCoE Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope fabric {a b}	Enters Fibre Channel storage mode for the specified fabric.
Step 3	UCS-A /fc-storage/fabric # delete interface {fc fcoe} slot-num port-num	Deletes the interface for the specified Fibre Channel or FCoE storage port.
Step 4	UCS-A /fc-storage/fabric # commit-buffer	Commits the transaction.

Example

The following example unconfigures Fibre Channel storage port 10 on slot 2 of fabric A and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # scope fabric a
UCS-A /fc-storage/fabric* # delete interface fc 2 10
UCS-A /fc-storage/fabric # commit-buffer
```

Restoring a Fibre Channel Storage Port Back to an Uplink Fibre Channel Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # create interface slot-num port-num	Creates an interface for the specified Fibre Channel uplink port.
Step 4	UCS-A /fc-uplink/fabric # commit-buffer	Commits the transaction.

Example

The following example creates an interface for Fibre Channel uplink port 10 on slot 2 of fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric* # create interface 2 10
UCS-A /fc-uplink/fabric # commit-buffer
```

Appliance Port Channels

An appliance port channel allows you to group several physical appliance ports to create one logical Ethernet storage link for the purpose of providing fault-tolerance and high-speed connectivity. In Cisco UCS Manager, you create a port channel first and then add appliance ports to the port channel. You can add up to eight appliance ports to a port channel.

Configuring an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage fabric mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # create port-channel port-num	Creates a port channel on the specified Ethernet storage port, and enters Ethernet storage fabric port channel mode.
Step 4	(Optional) UCS-A /eth-storage/fabric/port-channel # {enable disable}	Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	(Optional) UCS-A /eth-storage/fabric/port-channel # set name port-chan-name	Specifies the name for the port channel.
Step 6	(Optional) UCS-A /eth-storage/fabric/port-channel # set pingroupname pin-group name	Specifies the appliance pin target to the specified fabric and port, or fabric and port channel.
Step 7	(Optional) UCS-A /eth-storage/fabric/port-channel # set portmode {access trunk}	Specifies whether the port mode is access or trunk. By default, the mode is set to trunk.
Step 8	(Optional) UCS-A /eth-storage/fabric/port-channel # set prio sys-class-name	Specifies the QoS class for the appliance port. By default, the priority is set to best-effort. The sys-class-name argument can be one of the following class keywords: <ul style="list-style-type: none"> • Fc—Use this priority for QoS policies that control vHBA traffic only. • Platinum—Use this priority for QoS policies that control vNIC traffic only. • Gold—Use this priority for QoS policies that control vNIC traffic only.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Silver—Use this priority for QoS policies that control vNIC traffic only. • Bronze—Use this priority for QoS policies that control vNIC traffic only. • Best Effort—Do not use this priority. It is reserved for the Basic Ethernet traffic lane. If you assign this priority to a QoS policy and configure another system class as CoS 0, Cisco UCS Manager does not default to this system class. It defaults to the priority with CoS 0 for that traffic.
Step 9	(Optional) UCS-A /eth-storage/fabric/port-channel # set speed {1gbps 2gbps 4gbps 8gbps auto}	Specifies the speed for the port channel.
Step 10	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # create port-channel 13
UCS-A /eth-storage/fabric/port-channel* # enable
UCS-A /eth-storage/fabric/port-channel* # set name portchan13a
UCS-A /eth-storage/fabric/port-channel* # set pingroupname pingroup1
UCS-A /eth-storage/fabric/port-channel* # set portmode access
UCS-A /eth-storage/fabric/port-channel* # set prio fc
UCS-A /eth-storage/fabric/port-channel* # set speed 2gbps
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

Unconfiguring an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage fabric mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # delete port-channel <i>port-num</i>	Deletes the port channel from the specified Ethernet storage port.

	Command or Action	Purpose
Step 4	UCS-A /eth-storage/fabric # commit-buffer	Commits the transaction to the system configuration.

Example

The following example unconfigures the port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # delete port-channel 13
UCS-A /eth-storage/fabric* # commit-buffer
UCS-A /eth-storage/fabric #
```

Enabling or Disabling an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b}	Enters Ethernet storage mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # scope port-channel port-chan-name	Enters Ethernet storage port channel mode.
Step 4	UCS-A /eth-storage/fabric/port-channel # {enable disable}	Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example enables port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel* # enable
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

Adding a Member Port to an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage fabric mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # scope port-channel port-num	Enters Ethernet storage fabric port channel mode for the specified port channel.
Step 4	UCS-A /eth-storage/fabric/port-channel # create member-port slot-num port-num	Creates the specified member port from the port channel and enters Ethernet storage fabric port channel member port mode.
Step 5	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example adds the member port on slot 1, port 7 to the port channel on port 13 of fabric A and commits the transaction.

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel # create member-port 1 7
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

Deleting a Member Port from an Appliance Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-storage	Enters Ethernet storage mode.
Step 2	UCS-A /eth-storage # scope fabric {a b }	Enters Ethernet storage fabric mode for the specified fabric.
Step 3	UCS-A /eth-storage/fabric # scope port-channel port-num	Enters Ethernet storage fabric port channel mode for the specified port channel.
Step 4	UCS-A /eth-storage/fabric/port-channel # delete member-port slot-num port-num	Deletes the specified member port from the port channel.

	Command or Action	Purpose
Step 5	UCS-A /eth-storage/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes a member port from the port channel on port 13 of fabric A and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope fabric a
UCS-A /eth-storage/fabric # scope port-channel 13
UCS-A /eth-storage/fabric/port-channel # delete member-port 1 7
UCS-A /eth-storage/fabric/port-channel* # commit-buffer
UCS-A /eth-storage/fabric/port-channel #
```

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 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 NPV 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
```

Configuring 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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # create port-channel port-num	Creates a port channel on the specified Fibre Channel uplink port, and enters Fibre Channel uplink fabric port channel mode.
Step 4	(Optional) UCS-A /fc-uplink/fabric/port-channel # {enable disable}	Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	(Optional) UCS-A /fc-uplink/fabric/port-channel # set name port-chan-name	Specifies the name for the port channel.
Step 6	(Optional) UCS-A /fc-uplink/fabric/port-channel # set speed {1gbps 2gbps 4gbps 8gbps auto}	Specifies the speed for the port channel.
Step 7	UCS-A /fc-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates port channel 13 on fabric A, sets the name to portchan13a, enables the administrative state, sets the speed to 2 Gbps, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create port-channel 13
UCS-A /fc-uplink/fabric/port-channel* # enable
UCS-A /fc-uplink/fabric/port-channel* # set name portchan13a
UCS-A /fc-uplink/fabric/port-channel* # set speed 2gbps
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

Unconfiguring a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # delete port-channel port-num	Deletes the port channel on the specified Fibre Channel uplink port.
Step 4	UCS-A /fc-uplink/fabric # commit-buffer	Commits the transaction to the system configuration.

Example

The following example unconfigures port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # delete port-channel 13
UCS-A /fc-uplink/fabric* # commit-buffer
UCS-A /fc-uplink/fabric #
```

Adding Channel Mode Active To The Upstream NPIV Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # create port-channel port-num	Creates a port channel on the specified Fibre Channel uplink port, and enters Fibre Channel uplink fabric port channel mode.
Step 4	(Optional) UCS-A /fc-uplink/fabric/port-channel # {enable disable}	Enables or disables the administrative state of the port channel. The port channel is disabled by default.
Step 5	(Optional) UCS-A /fc-uplink/fabric/port-channel # set name port-chan-name	Specifies the name for the port channel.

	Command or Action	Purpose
Step 6	(Optional) UCS-A /fc-uplink/fabric/port-channel # scope <i>port-chan-name</i>	Specifies the name for the port channel.
Step 7	(Optional) UCS-A /fc-uplink/fabric/port-channel # channel mode { active }	Configures the channel-mode active on the upstream NPIV switch.
Step 8	UCS-A /fc-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example enables channel mode to active:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create port-channel 13
UCS-A /fc-uplink/fabric/port-channel* # enable
UCS-A /fc-uplink/fabric/port-channel* # set name portchan13a
UCS-A /fc-uplink/fabric/port-channel* # channel mode active
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel # exit
UCS-A /fc-uplink/fabric/ # show port-channel database

portchan13a
  Administrative channel mode is active
  Operational channel mode is active

UCS-A /fc-uplink/fabric/ #
```

Enabling or Disabling a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric { a b }	Enters Fibre Channel uplink mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # scope port-channel <i>port-chan-name</i>	Enters Fibre Channel uplink port channel mode.
Step 4	UCS-A /fc-uplink/fabric/port-channel # { enable disable }	Enables or disables the administrative state of the port channel. The port channel is disabled by default.

Example

The following example enables port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope port-channel 13
UCS-A /fc-uplink/fabric/port-channel* # enable
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

Adding a Member Port to a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # scope port-channel <i>port-num</i>	Enters Fibre Channel uplink fabric port channel mode for the specified port channel.
Step 4	UCS-A /fc-uplink/fabric/port-channel # create member-port <i>slot-num port-num</i>	Creates the specified member port from the port channel and enters Fibre Channel uplink fabric port channel member port mode.
Step 5	UCS-A /fc-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example adds the member port on slot 1, port 7 to port channel 13 on fabric A and commits the transaction.

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope port-channel 13
UCS-A /fc-uplink/fabric # create member-port 1 7
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

Deleting a Member Port from a Fibre Channel Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # scope port-channel port-num	Enters Fibre Channel uplink fabric port channel mode for the specified port channel.
Step 4	UCS-A /fc-uplink/fabric/port-channel # delete member-port slot-num port-num	Deletes the specified member port from the port channel.
Step 5	UCS-A /fc-uplink/fabric/port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes a member port from port channel 13 on fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope port-channel 13
UCS-A /fc-uplink/fabric # delete member-port 1 7
UCS-A /fc-uplink/fabric/port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/port-channel #
```

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.

Configuring a FCoE Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters FC Uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters FC - Uplink mode for the specific fabric.
Step 3	UCS-A /fc-uplink/fabric # create fcoe-port-channel number	Creates port channel for the specified FCoE uplink port.
Step 4	UCS-A /fc-uplink/fabric/fabricinterface # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates an interface for FCoE uplink port 1 on slot 4 of fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create fcoe-port-channel 4
UCS-A /fc-uplink/fabric/fcoe-port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/fcoe-port-channel #
```

Adding a Member Port to a FCoE Uplink Port Channel

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink fabric mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # scope fcoe-port-channel ID	Enters FCoE uplink port channel mode for the specified port channel.
Step 4	UCS-A /fc-uplink/fabric/fcoe-port-channel # create member-port slot-num port-num	Creates the specified member port from the port channel and enters FCoE uplink fabric port channel member port mode.

	Command or Action	Purpose
		<p>Note If the FCoE uplink port channel is a unified uplink port channel, you will get the following message:</p> <p>Warning: if this is a unified port channel then member will be added to the ethernet port channel of the same id as well.</p>
Step 5	UCS-A /fc-uplink/fabric/fcoe-port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example adds the member port on slot 1, port 7 to FCoE port channel 13 on fabric A and commits the transaction.

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope fcoe-port-channel 13
UCS-A /fc-uplink/fabric # create member-port 1 7
UCS-A /fc-uplink/fabric/fcoe-port-channel* # commit-buffer
UCS-A /fc-uplink/fabric/fcoe-port-channel #
```

Adapter Port Channels

An adapter port channel groups into one logical link all the physical links going from a Cisco UCS Virtual Interface Card (VIC) into an I/O.

Adapter port channels are created and managed internally by Cisco UCS Manager when it detects that the correct hardware is present. Adapter port channels cannot be configured manually. Adapter port channels are viewable using the Cisco UCS Manager GUI or the Cisco UCS Manager CLI.

Viewing Adapter Port Channels

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope iom {a b}	Enters chassis IOM mode for the specified IOM.
Step 3	UCS-A /chassis/iom # scope port group	Enters port group mode for the specified port group.
Step 4	UCS-A /chassis/iom/port group # show host-port-channel [detail expand]	Displays the adapter port channels on the specified chassis.

Example

This following example shows how to display information on host port channels within a port group mode:

```
UCS-A # scope chassis 1
UCS-A /chassis # scope iom a
UCS-A /chassis/iom # scope port group
UCS-A /chassis/iom/port group # show host-port-channel
```

Host Port channel:

Port Channel Id	Fabric ID	Oper State	State Reason
1289	B	Up	
1290	B	Up	
1306	B	Up	
1307	B	Up	
1309	B	Up	
1315	B	Up	

```
UCS-A /chassis/iom/port group #
```

Event Detection and Action

Cisco UCS Manager uses the statistics collection policy to monitor and trigger an alarm when there are faults in the network interface ports connected from the I/O module (IOM) to the fabric interconnect.

The error statistics for the network interface ports is called NiErrStats and consists of the following errors:

NiErrStats Error Name	Description
frameTx	Collects the TX_FRM_ERROR counter values.
tooLong	Collects the RX_TOOLONG counter values.
tooShort	Collects the sum of RX_UNDERSIZE and RX_FRAGMENT counter values.
Crc	Collects the sum of RX_CRERR_NOT_STOMPED and RX_CRCERR_STOMPED counter values.
inRange	Collects the RX_INRANGEERR counter values.



Note The network interface port statistics is collected only from active ports and the information is sent to Cisco UCS Manager.

Policy-Based Port Error Handling

If Cisco UCS Manager detects any errors on active NI ports, and if the error-disable feature is enabled, Cisco UCS Manager automatically disables the respective FI port that is connected to the NI port that had errors. When a FI 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 FI port is error-disabled and that the connected NI Port has errors.
- It eliminates the possibility that this port can cause other ports, which are connected to the same Chassis/FEX, to fail. 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.

Creating Threshold Definition

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope eth-server	Enters Ethernet storage mode.
Step 2	UCS-A/eth-server # scope stats-threshold-policy default	Enters statistics threshold policy mode.
Step 3	UCSA/eth-server/stats-threshold-policy # create class class-name	Creates the specified statistics threshold policy class and enters the organization statistics threshold policy class mode. To see a list of the available class name keywords, enter the create class ? command in organization threshold policy mode.
Step 4	UCS-A/eth-server/stats-threshold-policy/class # create property property-name	Creates the specified statistics threshold policy class property and enters the organization statistics threshold policy class property mode. To see a list of the available property name keywords, enter the create property ? command in organization threshold policy class mode.
Step 5	UCS-A/eth-server/stats-threshold-policy/class/property # set normal-value value	Specifies the normal value for the class property. The <i>value</i> format can vary depending on the class property being configured. To see the required format, enter the set normal-value ? command in organization statistics threshold policy class property mode.
Step 6	UCS-A/eth-server/stats-threshold-policy/class/property # create threshold-value {above-normal below-normal} {cleared condition critical info major minor warning}	Creates the specified threshold value for the class property and enters the organization statistics threshold policy class property threshold value mode.

	Command or Action	Purpose
Step 7	UCS-A/eth-server/stats-threshold-policy/class/property/threshold-value # set {deescalating escalating} value	Specifies the deescalating and escalating class property threshold value. The <i>value</i> format can vary depending on the class property threshold value being configured. To see the required format, enter the set deescalating ? or set escalating ? command in the organization statistics threshold policy class property threshold value mode.
Step 8	UCS-A/eth-server/stats-threshold-policy/class/property/threshold-value # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create a threshold definition:

```
UCS-A # scope eth-server
UCS-A /eth-server # scope stats-threshold-policy default
UCS-A /eth-server/stats-threshold-policy # create class ni-ether-error-stats
UCS-A /eth-server/stats-threshold-policy/class* # create property crc-delta
UCS-A /eth-server/stats-threshold-policy/class/property* # set normal-value 0
UCS-A /eth-server/stats-threshold-policy/class/property* # create threshold-value above-normal
major
UCS-A /eth-server/stats-threshold-policy/class/property/threshold-value* # set escalating
5
UCS-A /eth-server/stats-threshold-policy/class/property/threshold-value* # set deescalating
3
UCS-A /eth-server/stats-threshold-policy/class/property/threshold-value* # commit-buffer
```

Configuring Error Disable on a Fabric Interconnect Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope eth-server	Enters Ethernet storage mode.
Step 2	UCS-A/eth-server # scope stats-threshold-policy default	Enters statistics threshold policy mode.
Step 3	UCS-A/eth-server/stats-threshold-policy # scope class <i>class-name</i>	Enters the organization statistics threshold policy class mode for the specified statistics threshold policy class.
Step 4	UCS-A/eth-server/stats-threshold-policy/class # scope property <i>property-name</i>	Enters the organization statistics threshold policy class property mode for the specified statistics threshold policy class property.
Step 5	UCS-A/eth-server/stats-threshold-policy/class/property # set error-disable-fi-port {yes no}	Specifies the error disable state for the class property.

	Command or Action	Purpose
		Use the no option to disable error disable for the class property.
Step 6	UCS-A/eth-server/stats-threshold-policy/class/property* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to enable error disable on an FI port:

```
UCS-A # scope eth-server
UCS-A /eth-server # scope stats-threshold-policy default
UCS-A /eth-server/stats-threshold-policy # scope class ni-ether-error-stats
UCS-A /eth-server/stats-threshold-policy/class # scope property crc-delta
UCS-A /eth-server/stats-threshold-policy/class/property # set error-disable-fi-port yes
UCS-A /eth-server/stats-threshold-policy/class/property* # commit-buffer
```

Configuring Auto Recovery on a Fabric Interconnect Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope eth-server	Enters Ethernet storage mode.
Step 2	UCS-A/eth-server # scope stats-threshold-policy default	Enters statistics threshold policy mode.
Step 3	UCS-A/eth-server/stats-threshold-policy # scope class class-name	Enters the organization statistics threshold policy class mode for the specified statistics threshold policy class.
Step 4	UCS-A/eth-server/stats-threshold-policy/class # scope property property-name	Enters the organization statistics threshold policy class property mode for the specified statistics threshold policy class property.
Step 5	UCS-A/eth-server/stats-threshold-policy/class/property # set auto-recovery {enabled disabled}	Specifies the auto recovery state for the class property. Use the disabled option to disable auto recovery for the class property.
Step 6	UCS-A/eth-server/stats-threshold-policy/class/property* # set auto-recovery-time time	Specifies the time in minutes after which the port is automatically re-enabled. The auto recovery time can range from 0 minutes to 4294967295 minutes.
Step 7	UCS-A/eth-server/stats-threshold-policy/class/property* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to configure auto recovery on an FI port:

```
UCS-A # scope eth-server
UCS-A /eth-server # scope stats-threshold-policy default
UCS-A /eth-server/stats-threshold-policy # scope class ni-ether-error-stats
UCS-A /eth-server/stats-threshold-policy/class # scope property crc-delta
UCS-A /eth-server/stats-threshold-policy/class/property # set auto-recovery enabled
UCS-A /eth-server/stats-threshold-policy/class/property* # set auto-recovery-time 5
UCS-A /eth-server/stats-threshold-policy/class/property* # commit-buffer
```

Viewing the Network Interface Port Error Counters

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A/chassis # scope iom {a b}	Enters chassis IOM mode for the specified IOM.
Step 3	UCS-A/chassis/iom # scope port-group fabric	Enters the network interface port.
Step 4	UCS-A/chassis/iom/port-group # scope fabric-if <i>fabric-if number</i>	Enters the specified network interface port number.
Step 5	UCS-A/chassis/iom/port-group/fabric-if # show stats	Displays the error counters for the network interface port.

Example

The following example shows how to display the statistics for the network interface ports:

```
UCS-A # scope chassis 1
UCS-A/chassis # scope iom a
UCS-A/chassis/iom # scope port-group fabric
UCS-A/chassis/iom/port-group # scope fabric-if 1
UCS-A/chassis/iom/port-group/fabric-if # show stats
NI Ether Error Stats:
Time Collected: 2014-08-20T15:37:24:688
Monitored Object: sys/chassis-1/slot-1/fabric/port-1/ni-err-stats
Suspect: Yes
Crc (errors): 5000
Frame Tx (errors): 0
Too Long (errors): 0
Too Short (errors): 0
In Range (errors): 0
Thresholded: 0
```

Fabric Port Channels

Fabric port channels allow you to group several of the physical links from an IOM to a fabric interconnect into one logical link for redundancy and bandwidth sharing. As long as one link in the fabric port channel remains active, the fabric port channel continues to operate.

If the correct hardware is connected, fabric port channels are created by Cisco UCS Manager in the following ways:

- During chassis discovery according to the settings configured in the chassis discovery policy.
- After chassis discovery according to the settings configured in the chassis connectivity policy for a specific chassis.

For each IOM there is a single fabric port channel. Each uplink connecting an IOM to a fabric interconnect can be configured as a discrete link or included in the port channel, but an uplink cannot belong to more than one fabric port channel. For example, if a chassis with two IOMs is discovered and the chassis discovery policy is configured to create fabric port channels, Cisco UCS Manager creates two separate fabric port channels: one for the uplinks connecting IOM-1 and another for the uplinks connecting IOM-2. No other chassis can join these fabric port channels. Similarly, uplinks belonging to the fabric port channel for IOM-1 cannot join the fabric port channel for IOM-2.

Load Balancing Over Ports

Load balancing traffic among ports between IOMs and fabric interconnects uses the following criteria for hashing.

- For Ethernet traffic:
 - Layer 2 source and destination address
 - Layer 3 source and destination address
 - Layer 4 source and destination ports
- For FCoE traffic:
 - Layer 2 source and destination address
 - Source and destination IDs (SID and DID) and Originator Exchange ID (OXID)

In this example, a 2200 Series IOM module is verified by connecting iom *X* (where *X* is the chassis number).

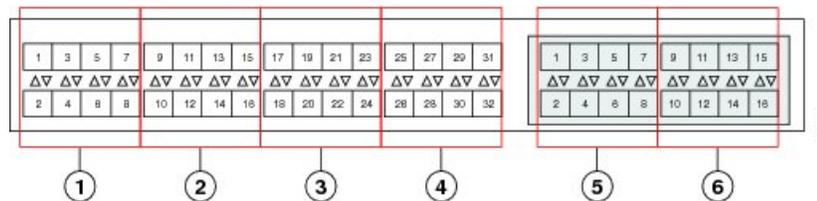
```
show platform software fwmctrl nifport
(....)
Hash Parameters:
 12_da: 1 12_sa: 1 12_vlan: 0
 13_da: 1 13_sa: 1
 14_da: 1 14_sa: 1
FCoE 12_da: 1 12_sa: 1 12_vlan: 0
FCoE 13_did: 1 13_sid: 1 13_oxid: 1
```

Cabling Considerations for Fabric Port Channels

When you configure the links between the Cisco UCS 2200 Series FEX and a Cisco UCS 6200 series fabric interconnect in fabric port channel mode, the available virtual interface namespace (VIF) on the adapter varies depending on where the FEX uplinks are connected to the fabric interconnect ports.

Inside the 6248 fabric interconnect there are six sets of eight contiguous ports, with each set of ports managed by a single chip. When all uplinks from an FEX are connected to a set of ports managed by a single chip, Cisco UCS Manager maximizes the number of VIFs used in service profiles deployed on the blades in the chassis. If uplink connections from an IOM are distributed across ports managed by separate chips, the VIF count is decreased.

Figure 2: Port Groups for Fabric Port Channels



Caution

Adding a second link to a fabric-port-channel port group is disruptive and will automatically increase the available amount of VIF namespace from 63 to 118. Adding further links is not disruptive and the VIF namespace stays at 118.



Caution

Linking a chassis to two fabric-port-channel port groups does not affect the VIF namespace unless it is manually acknowledged. The VIF namespace is then automatically set to the smaller size fabric port-channel port group usage (either 63 or 118 VIFs) of the two groups.

For high availability cluster-mode applications, we strongly recommend symmetric cabling configurations. If the cabling is asymmetric, the maximum number of VIFs available is the smaller of the two cabling configurations.

For more information on the maximum number of VIFs for your Cisco UCS environment, see the Configuration Limits document for your hardware and software configuration.

Viewing Fabric Port Channels

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.

	Command or Action	Purpose
Step 3	UCS-A /eth-server/fabric # show fabric-port-channel [detail expand]	Displays fabric port channels on the specified fabric interconnect.

Example

The following example displays information about configured fabric port channels on fabric interconnect A:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope fabric a
UCS-A /eth-server/fabric # show fabric-port-channel
Fabric Port Channel:
  Port Channel Id Chassis Id Admin State Oper State      State Reason
  -----
                1025 1      Enabled   Failed        No operational members
                1026 2      Enabled   Up
UCS-A /eth-server/fabric #
```

Enabling or Disabling a Fabric Port Channel Member Port

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-server	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # scope fabric {a b}	Enters Ethernet server fabric mode for the specified fabric.
Step 3	UCS-A /eth-server/fabric # scope fabric-port-channel port-chan-id	Enters Ethernet server fabric, fabric port channel mode for the specified fabric.
Step 4	UCS-A /eth-server/fabric/fabric-port-channel # scope member-port slot-id port-id	Enters Ethernet server fabric, fabric port channel mode for the specified member port.
Step 5	UCS-A /eth-server/fabric/fabric-port-channel # {enable disable}	Enables or disables the specified member port.
Step 6	UCS-A /eth-server/fabric/fabric-port-channel # commit-buffer	Commits the transaction to the system configuration.

Example

The following example disables fabric channel member port 1 31 on fabric port channel 1025 and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope fabric a
UCS-A /eth-server/fabric # scope fabric-port-channel 1025
UCS-A /eth-server/fabric/fabric-port-channel # scope member-port 1 31
```

```
UCS-A /eth-server/fabric/fabric-port-channel/member-port # disable
UCS-A /eth-server/fabric/fabric-port-channel/member-port* # commit-buffer
UCS-A /eth-server/fabric/fabric-port-channel/member-port #
```



CHAPTER 4

Fibre Channel Zoning

- [Information About Fibre Channel Zoning, on page 55](#)
- [Support for Fibre Channel Zoning in Cisco UCS Manager, on page 56](#)
- [Guidelines and recommendations for Cisco UCS Manager-Based Fibre Channel Zoning, on page 58](#)
- [Configuring Cisco UCS Manager Fibre Channel Zoning, on page 58](#)
- [Creating a VSAN for Fibre Channel Zoning, on page 59](#)
- [Creating a New Fibre Channel Zone Profile, on page 60](#)
- [Deleting a Fibre Channel Zone Profile, on page 61](#)
- [Deleting a Fibre Channel User Zone, on page 62](#)
- [Removing Unmanaged Zones from a VSAN Accessible to Both Fabric Interconnects, on page 63](#)
- [Removing Unmanaged Zones from a VSAN Accessible to One Fabric Interconnect, on page 64](#)
- [Configuring Fibre Channel Storage Connection Policies, on page 65](#)

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 Cisco UCS Manager 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.
Step 4	Configure the Fibre Channel and FCoE storage ports that you require to carry traffic for the Fibre Channel zones.	
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 Fibre Channel storage mode and accessible to both fabric interconnects.
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.
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. • Configure one or more vHBA initiator groups.

Creating a VSAN for Fibre Channel Zoning

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # create vsan { <i>VSAN_Name</i> } { <i>VSAN_ID</i> } { <i>FCoE_VLAN_ID</i> }	Enter the following: <ul style="list-style-type: none"> • <i>VSAN_Name</i>- 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. • <i>VSAN_ID</i>- The unique identifier assigned to the network. The ID can be between 1 and 4078, or between 4080 and 4093. 4079 is a reserved VSAN ID.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • FCoE_VLAN_ID- The unique identifier assigned to the VLAN used for Fibre Channel connections. The ID can be between 1 and 4029, or between 4048 and 4093. 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: <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 UCSCNA 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 3	UCS-A /fc-uplink # commit-buffer	

Example

The following examples create a VSAN named TestVsan and commit the changes in the system:

```
UCS-A # scope fc-uplink
UCS-A /fc-uplink # create vsan TestVsan 2 30
UCS-A /fc-uplink/vsan* # commit-buffer
UCS-A /fc-uplink/vsan #
```

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

	Command or Action	Purpose
Step 1	UCS-A # scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # create fc-zone-profile <i>Profile_Name</i>	Creates a Fibre Channel profile with the specified name.
Step 3	UCS-A /fc-storage/fc-zone-profile * # create fc-user-zone <i>Zone_Name</i>	Enters Fibre Channel zone profile mode and creates the specified Fibre Channel zone.
Step 4	UCS-A /fc-storage/fc-zone-profile/fc-user-zone* # set path {A B}	Sets the Fibre Channel zone path.
Step 5	UCS-A /fc-storage/fc-zone-profile/fc-user-zone* # set vsan <i>VSAN_Name</i>	Sets the Fibre Channel zone to the named VSAN.
Step 6	UCS-A /fc-storage/fc-zone-profile/fc-user-zone* # create member <i>wwpn</i>	Creates the WWPN for the Fibre Channel zone profile.
Step 7	UCS-A /fc-storage/fc-zone-profile/fc-user-zone* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create FC Zoning Policy named myProfile:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # create fc-zone-profile myProfile
UCS-A /fc-storage/fc-zone-profile* # create fc-user-zone myZone
UCS-A /fc-storage/fc-zone-profile/fc-user-zone* # set path A
UCS-A /fc-storage/fc-zone-profile/fc-user-zone* # set vsan test
UCS-A /fc-storage/fc-zone-profile/fc-user-zone* # create member 20:c2:11:25:b5:00:00:7f
UCS-A /fc-storage/fc-zone-profile/fc-user-zone/member* # commit-buffer
```

Deleting a Fibre Channel Zone Profile

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope fc-storage	Enters Fibre Channel storage mode.

	Command or Action	Purpose
Step 2	UCS-A /fc-storage # delete fc-zone-profile <i>Profile_Name</i>	Deletes a Fibre Channel profile with the specified name.
Step 3	UCS-A /fc-storage* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to delete an FC Zone Profile named myProfile:

```
UCS-A # scope fc-storage
UCS-A /fc-storage # delete fc-zone-profile myProfile
UCS-A /fc-storage* # commit-buffer
UCS-A /fc-storage #
```

Deleting a Fibre Channel User Zone

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope fc-zone-profile <i>Profile_Name</i>	Enters the specified Fibre Channel profile.
Step 3	UCS-A /fc-storage/fc-zone-profile # delete fc-user-zone <i>Userzone_Name</i>	Deletes the specified Fibre Channel User Zone.
Step 4	UCS-A /fc-storage/fc-zone-profile* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to delete an FC User Zone Profile named myZone:

```
UCS-A # scope fc-storage
UCS-A /fc-storage # scope fc-zone-profile myProfile
UCS-A /fc-storage/fc-zone-profile # delete fc-user-zone myZone
UCS-A /fc-storage/fc-zone-profile* # commit-buffer
UCS-A /fc-storage #
```

Removing Unmanaged Zones from a VSAN Accessible to Both Fabric Interconnects

After you disconnect the external Fibre Channel switch, the Fibre Channel zones that were managed by that switch might not be cleared from the Cisco UCS domain. This procedure removes those zones from each VSAN in the Cisco UCS domain so that you can configure Fibre Channel zoning in Cisco UCS.

Before you begin

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.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink mode for the specified fabric interconnect.
Step 3	UCS-A /fc-uplink/fabric # scope vsan vsan-name	Enters VSAN mode for the specified named VSAN.
Step 4	UCS-A /fc-uplink/fabric/vsan # clear-unmanaged-fc-zones-all	Clears all unmanaged Fibre Channel zones from the specified named VSAN. If desired, you can repeat Steps 2 through 4 to remove unmanaged zones from all VSANs that are accessible to the specified fabric interconnect before you commit the buffer.
Step 5	UCS-A /fc-uplink/fabric/vsan # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to remove unmanaged zones from a named VSAN accessible to fabric interconnect A and commit the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope vsan finance
UCS-A /fc-uplink/fabric/vsan # clear-unmanaged-fc-zones-all
UCS-A /fc-uplink/fabric/vsan* # commit-buffer
UCS-A /fc-uplink #
```

Removing Unmanaged Zones from a VSAN Accessible to One Fabric Interconnect

After you disconnect the external Fibre Channel switch, the Fibre Channel zones that were managed by that switch might not be cleared from the Cisco UCS domain. This procedure removes those zones from each VSAN in the Cisco UCS domain so that you can configure Fibre Channel zoning in Cisco UCS.

Before you begin

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.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope vsan vsan-name	Enters VSAN mode for the specified named VSAN.
Step 3	UCS-A /fc-uplink/vsan # clear-unmanaged-fc-zones-all	Clears all unmanaged Fibre Channel zones from the specified named VSAN. If desired, you can repeat steps 2 and 3 to remove unmanaged zones from all VSANs that are accessible to both fabric interconnects before you commit the buffer.
Step 4	UCS-A /fc-uplink/vsan # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to remove unmanaged zones from a named VSAN and commit the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope vsan finance
UCS-A /fc-uplink/vsan # clear-unmanaged-fc-zones-all
UCS-A /fc-uplink/vsan* # commit-buffer
UCS-A /fc-uplink #
```

Configuring Fibre Channel Storage Connection Policies

Creating a Fibre Channel Storage Connection Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # create storage-connection-policy <i>policy-name</i>	Creates a storage connection policy with the specified policy name, and enters organization storage connection policy mode.
Step 3	UCS-A /org # set zoning-type { none simt sist }	<ul style="list-style-type: none"> • 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 4	UCS-A /org/storage-connection-policy # create storage-target <i>wwpn</i>	Creates a storage target endpoint with the specified WWPN, and enters storage target mode.
Step 5	UCS-A /org/storage-connection-policy/storage-target # set target-path { a b }	Specifies which fabric interconnect is used for communications with the target endpoint.
Step 6	UCS-A /org/storage-connection-policy/storage-target # set target-vsan <i>vsan</i>	Specifies which VSAN is used for communications with the target endpoint.
Step 7	UCS-A /org/storage-connection-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example configures a Fibre Channel storage connection policy in the root organization named scPolicyZone1, using fabric interconnect A and the default VSAN, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create storage-connection-policy scPolicyZone1
UCS-A /org/storage-connection-policy* set zoning-type sist
UCS-A /org/storage-connection-policy* # create storage-target 20:10:20:30:40:50:60:70
UCS-A /org/storage-connection-policy/storage-target* # set target-path a
UCS-A /org/storage-connection-policy/storage-target* # set target-vsan default
UCS-A /org/storage-connection-policy* # commit-buffer
UCS-A /org/storage-connection-policy #
```

Deleting a Fibre Channel Storage Connection Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # delete storage-connection-policy <i>policy-name</i>	Deletes the specified storage connection policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the storage connection policy named scPolicyZone1 from the root organization and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete san-connectivity-policy scPolicyZone1
UCS-A /org* # commit-buffer
UCS-A /org #
```



CHAPTER 5

Named VSANs

- [Named VSANs, on page 67](#)
- [Fibre Channel Uplink Trunking for Named VSANs, on page 68](#)
- [Guidelines and Recommendations for VSANs, on page 68](#)
- [Creating a Named VSAN Accessible to Both Fabric Interconnects \(Fibre Channel Uplink Mode\), on page 70](#)
- [Creating a Named VSAN Accessible to Both Fabric Interconnects \(Fibre Channel Storage Mode\), on page 71](#)
- [Creating a Named VSAN Accessible to One Fabric Interconnect \(Fibre Channel Uplink Mode\), on page 72](#)
- [Creating a Named VSAN Accessible to One Fabric Interconnect \(Fibre Channel Storage Mode\), on page 74](#)
- [Deleting a Named VSAN, on page 75](#)
- [Changing the VLAN ID for the FCoE Native VLAN for a Named VSAN, on page 76](#)
- [Changing the VLAN ID for the FCoE Native VLAN for a Storage VSAN, on page 76](#)
- [Enabling or Disabling Fibre Channel Uplink Trunking, on page 77](#)

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 Accessible to Both Fabric Interconnects (Fibre Channel Uplink Mode)



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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # create vsan <i>vsan-name</i> <i>vsan-id fcoe-id</i>	Creates the specified named VSAN, specifies the VSAN name, VSAN ID and FCoE VLAN ID, and enters Fibre Channel uplink VSAN mode. <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.
Step 3	UCS-A /fc-uplink/vsan # set fc-zoning { disabled enabled }	Configures Fibre Channel zoning for the VSAN, as follows: <ul style="list-style-type: none"> • disabled—The upstream switch configures and controls the Fibre Channel zoning or Fibre Channel zoning is not implemented on this VSAN. • enabled—Cisco UCS Manager configures and controls Fibre Channel zoning.
Step 4	UCS-A /fc-uplink/vsan # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a named VSAN for both fabric interconnects, names the VSAN accounting, assigns the VSAN ID 2112, assigns the FCoE VLAN ID 4021, enables the VSAN for Cisco UCS Manager-based Fibre Channel zoning, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink* # create vsan accounting 2112 4021
UCS-A /fc-uplink/vsan # set fc-zoning enabled
UCS-A /fc-uplink/vsan* # commit-buffer
UCS-A /fc-uplink/vsan #
```

Creating a Named VSAN Accessible to Both Fabric Interconnects (Fibre Channel Storage Mode)



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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # create vsan <i>vsan-name</i> <i>vsan-id</i> <i>fc-oe-id</i>	Creates the specified named VSAN, specifies the VSAN name, VSAN ID, and FCoE VLAN ID, and enters Fibre Channel storage VSAN mode. <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.

	Command or Action	Purpose
Step 3	UCS-A /fc-storage/vsan # create member-port {fc fcoe} {a b} slot-id port-id	Creates a member port; specifies whether the port type, fabric, slot ID and port ID.
Step 4	UCS-A /fc-storage/vsan # set fc-zoning {disabled enabled}	Configures Fibre Channel zoning for the VSAN, as follows: <ul style="list-style-type: none"> • disabled—The upstream switch configures and controls the Fibre Channel zoning or Fibre Channel zoning is not implemented on this VSAN. • enabled—Cisco UCS Manager configures and controls Fibre Channel zoning.
Step 5	UCS-A /fc-storage/vsan # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a named VSAN, names the VSAN finance, assigns the VSAN ID 3955, assigns the FCoE VLAN ID 4021, creates a member port and assigns it to member port A, slot 1 port 40, enables the VSAN for Cisco UCS Manager-based Fibre Channel zoning, and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage/ # create VSAN finance 3955 4021
UCS-A /fc-storage/vsan # create member-port fcoe a 1 40
UCS-A /fc-storage/vsan # set fc-zoning enabled
UCS-A /fc-storage/vsan/member-port* # commit-buffer
UCS-A /fc-storage/vsan/member-port #
```

Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Uplink Mode)



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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.

	Command or Action	Purpose
Step 2	UCS-A /fc-uplink # scope fabric {a b}	Enters Fibre Channel uplink fabric interconnect mode for the specified fabric interconnect (A or B).
Step 3	UCS-A /fc-uplink/fabric # create vsan <i>vsan-name vsan-id fcoe-id</i>	Creates the specified named VSAN, specifies the VSAN name, VSAN ID, and FCoE VLAN ID, and enters Fibre Channel uplink VSAN mode. <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.
Step 4	UCS-A /fc-uplink/vsan # set fc-zoning {disabled enabled}	Configures Fibre Channel zoning for the VSAN, as follows: <ul style="list-style-type: none"> • disabled—The upstream switch configures and controls the Fibre Channel zoning or Fibre Channel zoning is not implemented on this VSAN. • enabled—Cisco UCS Manager configures and controls Fibre Channel zoning.
Step 5	UCS-A /fc-uplink/fabric/vsan # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a named VSAN for fabric interconnect A, names the VSAN finance, assigns the VSAN ID 3955, assigns the FCoE VLAN ID 2221, enables the VSAN for Cisco UCS Manager-based Fibre Channel zoning, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # create vsan finance 3955 2221
UCS-A /fc-uplink/vsan # set fc-zoning enabled
UCS-A /fc-uplink/fabric/vsan* # commit-buffer
UCS-A /fc-uplink/fabric/vsan #
```

Creating a Named VSAN Accessible to One Fabric Interconnect (Fibre Channel Storage Mode)



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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # scope fabric {a b}	Enters Fibre Channel storage mode for the specified fabric interconnect.
Step 3	UCS-A /fc-storage/fabric # create vsan vsan-name vsan-id fcoe-id	Creates the specified named VSAN, specifies the VSAN name, VSAN ID, and FCoE VLAN ID, and enters Fibre Channel storage VSAN mode. <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.
Step 4	UCS-A /fc-storage/fabric/vsan # create member-port {fc fcoe} {a b} slot-id port-id	Creates a member port on the specified VSAN.
Step 5	UCS-A /fc-storage/vsan # set fc-zoning {disabled enabled}	Configures Fibre Channel zoning for the VSAN, as follows: <ul style="list-style-type: none"> • disabled—The upstream switch configures and controls the Fibre Channel zoning or Fibre Channel zoning is not implemented on this VSAN.

	Command or Action	Purpose
		<ul style="list-style-type: none"> enabled—Cisco UCS Manager configures and controls Fibre Channel zoning.
Step 6	UCS-A /fc-storage/fabric/vsan # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a named VSAN on fabric A, names the VSAN finance, assigns the VSAN ID 3955, assigns the FCoE VLAN ID 2221, creates a member port and assigns the it to member port A, slot 1 port 40, and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage/ # scope fabric a
UCS-A /fc-storage/fabric # create VSAN finance 3955 2221
UCS-A /fc-storage/fabric/vsan # create member-port a 1 40
UCS-A /fc-storage/fabric/vsan # set fc-zoning enabled
UCS-A /fc-storage/fabric/vsan/member-port* # commit-buffer
UCS-A /fc-storage/fabric/vsan/member-port #
```

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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # delete vsan <i>vsan-name</i>	Deletes the specified named VSAN.
Step 3	UCS-A /fc-uplink # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to delete a named VSAN and commit the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # delete vsan finance
UCS-A /fc-uplink* # commit-buffer
UCS-A /fc-uplink #
```

Changing the VLAN ID for the FCoE Native VLAN for 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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope vsan vsan-name	Enters VSAN mode for the specified named VSAN.
Step 3	UCS-A /fc-uplink/vsan # set fcoe-vlan fcoe-vlan-id	Sets the unique identifier assigned to the VLAN used for Fibre Channel connections.
Step 4	UCS-A /fc-uplink/vsan # commit-buffer	Commits the transaction to the system configuration.

Example

The following example changes the VLAN ID for the FCoE Native VLAN on a named VSAN called finance to 4000 and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope vsan finance
UCS-A /fc-uplink/vsan # set fcoe-vlan 4000
UCS-A /fc-uplink/vsan* # commit-buffer
UCS-A /fc-uplink/vsan #
```

Changing the VLAN ID for the FCoE Native VLAN for 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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-storage	Enters Fibre Channel storage mode.
Step 2	UCS-A /fc-storage # set fcoe-storage-native-vlan <i>fcoe-id</i>	Sets the unique identifier assigned to the VLAN used for Fibre Channel connections.
Step 3	UCS-A /fc-storage # commit-buffer	Commits the transaction to the system configuration.

Example

The following example changes the VLAN ID for the FCoE Native VLAN on a storage VSAN called finance to 4000 and commits the transaction:

```
UCS-A# scope fc-storage
UCS-A /fc-storage # set fcoe-storage-native-vlan 4000
UCS-A /fc-storage* # commit-buffer
UCS-A /fc-storage #
```

Enabling or Disabling 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.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # scope fabric {a b }	Enters Fibre Channel uplink mode for the specified fabric.
Step 3	UCS-A /fc-uplink/fabric # set uplink-trunking {enabled disabled }	Enables or disables uplink trunking.
Step 4	UCS-A /fc-uplink/fabric # commit-buffer	Commits the transaction to the system configuration.

Example

The following example enables Fibre Channel uplink trunking for fabric A and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # set uplink-trunking enabled
UCS-A /fc-uplink/fabric* # commit-buffer
UCS-A /fc-uplink/fabric #
```



CHAPTER 6

SAN Pin Groups

- [SAN Pin Groups, on page 79](#)
- [Configuring a SAN Pin Group, on page 79](#)
- [Configuring a FCoE Pin Group, on page 80](#)

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.

Configuring 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

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters Fibre Channel uplink mode.
Step 2	UCS-A /fc-uplink # create pin-group <i>pin-group-name</i>	Creates a Fibre Channel (SAN) pin group with the specified name, and enters Fibre Channel uplink pin group mode.
Step 3	(Optional) UCS-A /fc-uplink/pin-group # set descr <i>description</i>	Provides a description for the pin group. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	(Optional) UCS-A /fc-uplink/pin-group # set target {a b dual} port <i>slot-num / port-num</i>	Sets the Fibre Channel pin target to the specified fabric and port.
Step 5	UCS-A /fc-uplink/pin-group # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a SAN pin group named fcpingroup12, provides a description for the pin group, sets the pin group target to slot 2, port 1, and commits the transaction:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # create pin-group fcpingroup12
UCS-A /fc-uplink/pin-group* # set descr "This is my pin group #12"
UCS-A /fc-uplink/pin-group* # set target a port 2/1
UCS-A /fc-uplink/pin-group* # commit-buffer
UCS-A /fc-uplink/pin-group #
```

What to do next

Include the pin group in a vHBA template.

Configuring a FCoE Pin Group

You can create a FCoE pin group, and specify the FCoE uplink port as the pin group target.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope fc-uplink	Enters FC uplink mode.

	Command or Action	Purpose
Step 2	UCS-A /fc-uplink # create pin-group fcoepingroup	Creates a FCoE pin group with the specified name, and enters FCoE uplink pin group mode.
Step 3	UCS-A /fc-uplink/pin-group # set target a fcoe-port 1/8	Sets FCoE port 1/8 as the target port for this pin group.
Step 4	UCS-A /fc-uplink/pin-group # commit-buffer	Commits the transaction to the system configuration.

Example

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # create pin-group fcoepingroup
UCS-A /fc-uplink/pin-group* #set target a fcoe-port 1/8
UCS-A /fc-uplink/pin-group* # commit-buffer
UCS-A /fc-uplink/pin-group #
```




CHAPTER 7

FC Identity Assignment

- [Fibre Channel Identity, on page 83](#)

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 85](#)
- [Creating a WWN Pool, on page 86](#)
- [Deleting a WWN Pool, on page 89](#)

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 20:00:00:00:00:00:00:00 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

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 WWN 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # create wwn-pool <i>wwn-pool-name</i> { node-and-port-wwn-assignment node-wwn-assignment port-wwn-assignment }	Creates a WWN pool with the specified name and purpose, and enters organization WWN pool mode. This can be one of the following: <ul style="list-style-type: none"> • node-and-port-wwn-assignment—Creates a WWxN pool that includes both world wide node names (WWNNs) and world wide port names (WWPNs). • node-wwn-assignment—Creates a WWNN pool that includes only WWNNs. • port-wwn-assignment—Creates a WWPN pool that includes only WWPNs.

	Command or Action	Purpose
		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.
Step 3	(Optional) UCS-A /org/wwn-pool # set descr <i>description</i>	Provides a description for the WWN pool. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/wwn-pool # set assignmentorder { default sequential }	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 5	UCS-A /org/wwn-pool # set max-ports-per-node { 15-ports-per-node 3-ports-per-node 31-ports-per-node 63-ports-per-node 7-ports-per-node }	For WWxN pools, specify the maximum number of ports that can be assigned to each node name in this pool. The default value is 3-ports-per-node . Note The pool size for WWxN pools must be a multiple of <i>ports-per-node</i> + 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.
Step 6	UCS-A /org/wwn-pool # create block <i>first-wwn last-wwn</i>	Creates a block (range) of WWNs, and enters organization WWN pool block mode. You must specify the first and last WWN in the block using the form <i>nn:nn:nn:nn:nn:nn:nn:nn</i> , with the WWNs separated by a space. Note A WWN pool can contain more than one WWN block. To create multiple WWN blocks, you must enter multiple create block commands from organization WWN pool mode.
Step 7	UCS-A /org/wwn-pool/block # exit	Exits organization WWN pool block mode.

	Command or Action	Purpose
Step 8	UCS-A /org/wwn-pool # create initiator <i>wwn</i> <i>wwn</i>	Creates a single initiator for a WWNN or WWPN pool, and enters organization WWN pool initiator mode. You must specify the initiator using the form <i>nn:nn:nn:nn:nn:nn:nn</i> . Note A WWNN or WWPN pool can contain more than one initiator. To create multiple initiators, you must enter multiple create initiator commands from organization WWN pool mode.
Step 9	UCS-A /org/wwn-pool/initiator # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create a WWNN pool named sanpool, provide a description for the pool, specify a block of WWNs and an initiator to be used for the pool, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # create wwn-pool sanpool node-wwn-assignment
UCS-A /org/wwn-pool* # set descr "This is my WWNN pool"
UCS-A /org/wwn-pool* # create block 20:00:00:25:B5:00:00:00 20:00:00:25:B5:00:00:01
UCS-A /org/wwn-pool/block* # exit
UCS-A /org/wwn-pool* # create initiator 23:00:00:05:AD:1E:02:00
UCS-A /org/wwn-pool/initiator* # commit-buffer
UCS-A /org/wwn-pool/initiator #
```

The following example shows how to create a WWxN pool named sanpool, provide a description for the pool, specify seven ports per node, specify a block of eight WWNs to be used for the pool, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # create wwn-pool sanpool node-and-port-wwn-assignment
UCS-A /org/wwn-pool* # set descr "This is my WWxN pool"
UCS-A /org/wwn-pool* # set max-ports-per-node 7-ports-per-node
UCS-A /org/wwn-pool* # create block 20:00:00:25:B5:00:00:00 20:00:00:25:B5:00:00:08
UCS-A /org/wwn-pool/block* # commit-buffer
UCS-A /org/wwn-pool/block #
```

What to do next

- Include the WWPN pool in a vHBA template.
- Include the WWNN pool in a service profile and template.
- Include the WWxN pool in a service profile and template.

Deleting a WWN 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # delete wwn-pool <i>pool-name</i>	Deletes the specified WWN pool.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to delete the WWN pool named pool4 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # delete wwn-pool pool4
UCS-A /org* # commit-buffer
UCS-A /org #
```




CHAPTER 9

Storage-Related Policies

- [Configuring vHBA Templates, on page 91](#)
- [Configuring Fibre Channel Adapter Policies, on page 93](#)
- [Configuring the Default vHBA Behavior Policy, on page 97](#)
- [Configuring SAN Connectivity Policies, on page 98](#)

Configuring 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.

Configuring a vHBA Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create vhma-templ <i>vhba-templ-name</i> [fabric {a b}] [fc-if <i>vsan-name</i>]	Creates a vHBA template and enters organization vHBA template mode.
Step 3	(Optional) UCS-A /org/vhma-templ # set descr <i>description</i>	Provides a description for the vHBA template.
Step 4	(Optional) UCS-A /org/vhma-templ # set fabric {a b}	Specifies the fabric to use for the vHBA. If you did not specify the fabric when creating the vHBA template in Step 2, then you have the option to specify it with this command.

	Command or Action	Purpose
Step 5	(Optional) UCS-A /org/vhba-templ # set fc-if <i>vsan-name</i>	Specifies the Fibre Channel interface (named VSAN) to use for the vHBA template. If you did not specify the Fibre Channel interface when creating the vHBA template in Step 2, you have the option to specify it with this command.
Step 6	UCS-A /org/vhba-templ # set max-field-size <i>size-num</i>	Specifies the maximum size of the Fibre Channel frame payload (in bytes) that the vHBA supports.
Step 7	UCS-A /org/vhba-templ # set pin-group <i>group-name</i>	Specifies the pin group to use for the vHBA template.
Step 8	UCS-A /org/vhba-templ # set qos-policy <i>mac-pool-name</i>	Specifies the QoS policy to use for the vHBA template.
Step 9	UCS-A /org/vhba-templ # set stats-policy <i>policy-name</i>	Specifies the server and server component statistics threshold policy to use for the vHBA template.
Step 10	UCS-A /org/vhba-templ # set type { initial-template updating-template }	Specifies the vHBA template update type. If you do not want vHBA instances created from this template to be automatically updated when the template is updated, use the initial-template keyword; otherwise, use the updating-template keyword to ensure that all vHBA instances are updated when the vHBA template is updated.
Step 11	UCS-A /org/vhba-templ # set wwpn-pool <i>pool-name</i>	Specifies the WWPN pool to use for the vHBA template.
Step 12	UCS-A /org/vhba-templ # commit-buffer	Commits the transaction to the system configuration.

Example

The following example configures a vHBA template and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create vhba template VhbaTempFoo
UCS-A /org/vhba-templ* # set descr "This is a vHBA template example."
UCS-A /org/vhba-templ* # set fabric a
UCS-A /org/vhba-templ* # set fc-if accounting
UCS-A /org/vhba-templ* # set max-field-size 2112
UCS-A /org/vhba-templ* # set pin-group FcPinGroup12
UCS-A /org/vhba-templ* # set qos-policy policy34foo
UCS-A /org/vhba-templ* # set stats-policy ServStatsPolicy
UCS-A /org/vhba-templ* # set type updating-template
UCS-A /org/vhba-templ* # set wwpn-pool SanPool7
UCS-A /org/vhba-templ* # commit-buffer
UCS-A /org/vhba-templ #
```

Deleting a vHBA Template

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete vhma-templ <i>vhba-templ-name</i>	Deletes the specified vHBA template.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the vHBA template named VhbaTempFoo and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete vhma template VhbaTempFoo
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring 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.
- **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.
- **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.

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.

Completion Queues = Transmit Queues + Receive Queues

Interrupt Count = (Completion Queues + 2) rounded up to nearest power of 2

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

Completion Queues = 1 + 8 = 9

Interrupt Count = (9 + 2) rounded up to the nearest power of 2 = 16

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

Interrupt Count = (#Tx or Rx Queues) + 2

For example:

Interrupt Count wq = 32, rq = 32, cq = 64 - then Interrupt Count = Max(32, 32) + 2 = 34

Interrupt Count wq = 64, rq = 8, cq = 72 - then Interrupt Count = Max(64, 8) + 2 = 66

Interrupt Count wq = 1, rq = 16, cq = 17 - then Interrupt count = Max(1, 16) + 2 = 18

Configuring a Fibre Channel Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create fc-policy <i>policy-name</i>	Creates the specified Fibre Channel adapter policy and enters organization Fibre Channel policy mode.
Step 3	(Optional) UCS-A /org/fc-policy # set descr <i>description</i>	Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	(Optional) UCS-A /org/fc-policy # set error-recovery { fc-error-recovery { disabled enabled } link-down-timeout <i>timeout-msec</i> port-down-io-retry-count <i>retry-count</i> port-down-timeout <i>timeout-msec</i> }	Configures the Fibre Channel error recovery.
Step 5	(Optional) UCS-A /org/fc-policy # set interrupt mode { intx msi msi-x }	Configures the driver interrupt mode.

	Command or Action	Purpose
Step 6	(Optional) UCS-A /org/fc-policy # set port { io-throttle-count <i>throttle-count</i> max-luns <i>max-num</i> }	Configures the Fibre Channel port.
Step 7	(Optional) UCS-A /org/fc-policy # set port-f-logi { retries <i>retry-count</i> timeout <i>timeout-msec</i> }	Configures the Fibre Channel port fabric login (FLOGI).
Step 8	(Optional) UCS-A /org/fc-policy # set port-p-logi { retries <i>retry-count</i> timeout <i>timeout-msec</i> }	Configures the Fibre Channel port-to-port login (PLOGI).
Step 9	(Optional) UCS-A /org/fc-policy # set recv-queue { count <i>count</i> ring-size <i>size-num</i> }	Configures the Fibre Channel receive queue.
Step 10	(Optional) UCS-A /org/fc-policy # set scsi-io { count <i>count</i> ring-size <i>size-num</i> }	Configures the Fibre Channel SCSI I/O.
Step 11	(Optional) UCS-A /org/fc-policy # set trans-queue ring-size <i>size-num</i> }	Configures the Fibre Channel transmit queue.
Step 12	UCS-A /org/fc-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example configures a Fibre Channel adapter policy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create fc-policy FcPolicy42
UCS-A /org/fc-policy* # set descr "This is a Fibre Channel adapter policy example."
UCS-A /org/fc-policy* # set error-recovery error-detect-timeout 2500
UCS-A /org/fc-policy* # set port max-luns 4
UCS-A /org/fc-policy* # set port-f-logi retries 250
UCS-A /org/fc-policy* # set port-p-logi timeout 5000
UCS-A /org/fc-policy* # set recv-queue count 1
UCS-A /org/fc-policy* # set scsi-io ring-size 256
UCS-A /org/fc-policy* # set trans-queue ring-size 256
UCS-A /org/fc-policy* # commit-buffer
UCS-A /org/fc-policy #
```

Deleting a Fibre Channel Adapter Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # delete fc-policy <i>policy-name</i>	Deletes the specified Fibre Channel adapter policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the Fibre Channel adapter policy named FcPolicy42 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete fc-policy FcPolicy42
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring 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

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode.
Step 2	UCS-A/org # scope vhma-beh-policy	Enters default vHBA behavior policy mode.

	Command or Action	Purpose
Step 3	UCS-A/org/vhba-beh-policy # set action { hw-inherit [<i>template_name name</i>] none }	Specifies the default vHBA behavior policy. This can be one of the following: <ul style="list-style-type: none"> • 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. If you specify hw-inherit, you can also specify a vHBA template to create the vHBAs. • none—Cisco UCS Manager does not create default vHBAs for a service profile. All vHBAs must be explicitly created.
Step 4	UCS-A/org/vhba-beh-policy # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to set the default vHBA behavior policy to **hw-inherit**.

```
UCS-A # scope org /
UCS-A/org # scope vhba-beh-policy
UCS-A/org/vhba-beh-policy # set action hw-inherit
UCS-A/org/vhba-beh-policy* # commit-buffer
UCS-A/org/vhba-beh-policy #
```

Configuring 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter <i>/</i> as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # create san-connectivity-policy <i>policy-name</i>	Creates the specified SAN connectivity policy, and enters organization network control policy mode. 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.
Step 3	(Optional) UCS-A /org/lan-connectivity-policy # set descr <i>policy-name</i>	Adds a description to the policy. We recommend that you include information about where and how the policy should be used. 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 4	UCS-A /org/service-profile # set identity { dynamic-uuid { <i>uuid</i> derived } dynamic-wwnn { <i>wwnn</i> derived } uuid-pool <i>pool-name</i> wwnn-pool <i>pool-name</i> }	Specifies how the server acquires a UUID or WWNN. You can do one of the following: <ul style="list-style-type: none"> • Create a unique UUID in the form <i>nnnnnnnnn-nnnn-nnnn-nnnnnnnnnnnnn</i> • Derive the UUID from the one burned into the hardware at manufacture • Use a UUID pool • Create a unique WWNN in the form <i>hh : hh : hh : hh : hh : hh : hh</i> • Derive the WWNN from one burned into the hardware at manufacture • Use a WWNN pool
Step 5	UCS-A /org/lan-connectivity-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create a SAN connectivity policy named SanConnect242 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # create san-connectivity-policy SanConnect242
UCS-A /org/san-connectivity-policy* # set descr "SAN connectivity policy"
UCS-A /org/san-connectivity-policy* # set identity wwnn-pool SanPool7
```

```
UCS-A /org/san-connectivity-policy* # commit-buffer
UCS-A /org/san-connectivity-policy #
```

What to do next

Add one or more vHBAs and/or initiator groups to this SAN connectivity policy.

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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # delete san-connectivity-policy <i>policy-name</i>	Deletes the specified SAN connectivity policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to delete a SAN connectivity policy named SanConnect52 from the root organization and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # delete san-connectivity-policy SanConnect52
UCS-A /org* # commit-buffer
UCS-A /org #
```

Creating a vHBA for a SAN Connectivity Policy

If you are continuing from [Creating a SAN Connectivity Policy, on page 99](#), begin this procedure at Step 3.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope san-connectivity-policy <i>policy-name</i>	Enters SAN connectivity policy mode for the specified SAN connectivity policy.

	Command or Action	Purpose
Step 3	UCS-A /org/san-connectivity-policy # create vhma <i>vhba-name</i> [fabric { a b }] [fc-if <i>fc-if-name</i>]	Creates a vHBA for the specified SAN connectivity policy and enters vHBA mode. 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.
Step 4	UCS-A /org/san-connectivity-policy/vhba # set adapter-policy <i>policy-name</i>	Specifies the adapter policy to use for the vHBA.
Step 5	UCS-A /org/san-connectivity-policy/vhba # set identity { dynamic-wwpn { <i>wwpn</i> derived } wwpn-pool <i>wwn-pool-name</i> }	Specifies the WWPN for the vHBA. You can set the storage identity using one of the following options: <ul style="list-style-type: none"> • Create a unique WWPN in the form <i>hh:hh:hh:hh:hh:hh:hh:hh</i>. You can specify a WWPN 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. • If you want the WWPN to be compatible with Cisco MDS Fibre Channel switches, use the WWPN template 20:00:00:25:B5:XX:XX:XX. • Derive the WWPN from one burned into the hardware at manufacture. • Assign a WWPN from a WWN pool.
Step 6	UCS-A /org/san-connectivity-policy/vhba # set max-field-size <i>size-num</i>	Specifies the maximum size of the Fibre Channel frame payload (in bytes) that the vHBA supports. Enter an integer between 256 and 2112. The default is 2048.
Step 7	UCS-A /org/san-connectivity-policy/vhba # set order { <i>order-num</i> unspecified }	Specifies the PCI scan order for the vHBA.
Step 8	UCS-A /org/san-connectivity-policy/vhba # set pers-bind { disabled enabled }	Disables or enables persistent binding to Fibre Channel targets.
Step 9	UCS-A /org/san-connectivity-policy/vhba # set pin-group <i>group-name</i>	Specifies the SAN pin group to use for the vHBA.

	Command or Action	Purpose
Step 10	UCS-A /org/san-connectivity-policy/vhba # set qos-policy <i>policy-name</i>	Specifies the QoS policy to use for the vHBA.
Step 11	UCS-A /org/san-connectivity-policy/vhba # set stats-policy <i>policy-name</i>	Specifies the statistics threshold policy to use for the vHBA.
Step 12	UCS-A /org/san-connectivity-policy/vhba # set template-name <i>policy-name</i>	Specifies the vHBA template to use for the vHBA. If you choose to use a vHBA template for the vHBA, you must still complete all of the configuration not included in the vHBA template, including Steps 4, 7, and 8.
Step 13	UCS-A /org/san-connectivity-policy/vhba # set vcon {1 2 3 4 any}	Assigns the vHBA to one or all virtual network interface connections.
Step 14	UCS-A /org/san-connectivity-policy/vhba # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to configure a vHBA for a SAN connectivity policy named SanConnect242 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope san-connectivity-policy SanConnect242
UCS-A /org/san-connectivity-policy* # create vhba vhba3 fabric a
UCS-A /org/san-connectivity-policy/vhba* # set adapter-policy AdaptPol2
UCS-A /org/san-connectivity-policy/vhba* # set identity wwpn-pool SanPool7
UCS-A /org/san-connectivity-policy/vhba* # set max-field-size 2112
UCS-A /org/san-connectivity-policy/vhba* # set order 0
UCS-A /org/san-connectivity-policy/vhba* # set pers-bind enabled
UCS-A /org/san-connectivity-policy/vhba* # set pin-group FcPinGroup12
UCS-A /org/san-connectivity-policy/vhba* # set qos-policy QosPol5
UCS-A /org/san-connectivity-policy/vhba* # set stats-policy StatsPol2
UCS-A /org/san-connectivity-policy/vhba* # set template-name SanConnPol3
UCS-A /org/san-connectivity-policy/vhba* # set vcon any
UCS-A /org/san-connectivity-policy/vhba* # commit-buffer
UCS-A /org/san-connectivity-policy/vhba #
```

What to do next

If desired, add another vHBA or an initiator group to the SAN connectivity policy. If not, include the policy in a service profile or service profile template.

Deleting a vHBA from a SAN Connectivity Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope san-connectivity-policy <i>policy-name</i>	Enters SAN connectivity policy mode for the specified SAN connectivity policy.
Step 3	UCS-A /org/san-connectivity-policy # delete vHBA <i>vhba-name</i>	Deletes the specified vHBA from the SAN connectivity policy.
Step 4	UCS-A /org/san-connectivity-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to delete a vHBA named vHBA3 from a SAN connectivity policy named SanConnect242 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope san-connectivity-policy SanConnect242
UCS-A /org/san-connectivity-policy # delete vHBA vHBA3
UCS-A /org/san-connectivity-policy* # commit-buffer
UCS-A /org/san-connectivity-policy #
```

Creating an Initiator Group for a SAN Connectivity Policy

If you are continuing from [Creating a SAN Connectivity Policy, on page 99](#), begin this procedure at Step 3.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope san-connectivity-policy <i>policy-name</i>	Enters SAN connectivity policy mode for the specified SAN connectivity policy.
Step 3	UCS-A /org/san-connectivity-policy # create initiator-group <i>group-name</i> fc	Creates the specified initiator group for Fibre Channel zoning and enters initiator group mode. 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

	Command or Action	Purpose
		cannot change this name after the object is saved.
Step 4	UCS-A /org/san-connectivity-policy/initiator-group # create initiator <i>vhba-name</i>	Creates the specified vHBA initiator in the initiator group. If desired, repeat this step to add a second vHBA initiator to the group.
Step 5	UCS-A /org/san-connectivity-policy/initiator-group # set storage-connection-policy <i>policy-name</i>	Associates the specified storage connection policy with the SAN connectivity policy. Note This step assumes that you want to associate an existing storage connection policy to associate with the SAN connectivity policy. If you do, continue with Step 10. If you want to create a local storage definition for this policy instead, continue with Step 6.
Step 6	UCS-A /org/san-connectivity-policy/initiator-group/storage-connection-def # create storage-target <i>wwpn</i>	Creates a storage target endpoint with the specified WWPN, and enters storage target mode.
Step 7	UCS-A /org/san-connectivity-policy/initiator-group/storage-connection-def # set target-path {a b}	Specifies which fabric interconnect is used for communications with the target endpoint.
Step 8	UCS-A /org/san-connectivity-policy/initiator-group/storage-connection-def # set target-vsan <i>vsan</i>	Specifies which VSAN is used for communications with the target endpoint.
Step 9	UCS-A /org/san-connectivity-policy/initiator-group # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to configure an initiator group named `initGroupZone1` with two initiators for a SAN connectivity policy named `SanConnect242`, configure a local storage connection policy definition named `scPolicyZone1`, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope san-connectivity-policy SanConnect242
UCS-A /org/san-connectivity-policy # create initiator-group initGroupZone1 fc
UCS-A /org/san-connectivity-policy/initiator-group* # set zoning-type sist
UCS-A /org/san-connectivity-policy/initiator-group* # create initiator vhba1
UCS-A /org/san-connectivity-policy/initiator-group* # create initiator vhba2
UCS-A /org/san-connectivity-policy/initiator-group* # create storage-connection-def
scPolicyZone1
UCS-A /org/san-connectivity-policy/initiator-group/storage-connection-def* # create
storage-target
```

```

20:10:20:30:40:50:60:70
UCS-A /org/san-connectivity-policy/initiator-group/storage-connection-def/storage-target*
# set
target-path a
UCS-A /org/san-connectivity-policy/initiator-group/storage-connection-def/storage-target*
# set
target-vsan default
UCS-A /org/san-connectivity-policy/initiator-group* # commit-buffer
UCS-A /org/san-connectivity-policy/initiator-group #

```

What to do next

If desired, add another initiator group or a vHBA to the SAN connectivity policy. If not, include the policy in a service profile or service profile template.

Deleting an Initiator Group from a SAN Connectivity Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope san-connectivity-policy <i>policy-name</i>	Enters SAN connectivity policy mode for the specified SAN connectivity policy.
Step 3	UCS-A /org/san-connectivity-policy # delete initiator-group <i>group-name</i>	Deletes the specified initiator group from the SAN connectivity policy.
Step 4	UCS-A /org/san-connectivity-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to delete an initiator group named `initGroup3` from a SAN connectivity policy named `SanConnect242` and commit the transaction:

```

UCS-A# scope org /
UCS-A /org # scope san-connectivity-policy SanConnect242
UCS-A /org/san-connectivity-policy # delete initiator-group initGroup3
UCS-A /org/san-connectivity-policy* # commit-buffer
UCS-A /org/san-connectivity-policy #

```



CHAPTER 10

Storage Profiles

- [Storage Profiles, on page 107](#)
- [Disk Groups and Disk Group Configuration Policies, on page 107](#)
- [RAID Levels, on page 109](#)
- [Automatic Disk Selection, on page 111](#)
- [Supported LUN Modifications, on page 111](#)
- [Unsupported LUN Modifications, on page 112](#)
- [Disk Insertion Handling, on page 112](#)
- [Virtual Drive Naming, on page 114](#)
- [LUN Dereferencing, on page 114](#)
- [Controller Constraints and Limitations, on page 115](#)
- [Configuring Storage Profiles, on page 115](#)

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.

Disk Groups and Disk Group Configuration Policies

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.

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.
- 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.

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:

- 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 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**— (Not supported for PCH SSD Controller Definition) 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**— (Not supported for PCH SSD Controller Definition) 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**— (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.

**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. For RAID license information for a specific Cisco UCS server, see the Hardware Installation Guide for that server.

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:
 1. Disk type
 2. Minimum disk size from highest to lowest
 3. Space required from highest to lowest
 4. Disk group qualifier name, in alphabetical order
 5. 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.

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:
 1. You specify the slot specifically by using a local disk reference
 2. 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 113](#)

- [Redundant Virtual Drives with No Hot Spare Drives, on page 113](#)
- [Redundant Virtual Drives with Hot Spare Drives, on page 113](#)
- [Replacing Hot Spare Drives, on page 113](#)
- [Inserting Physical Drives into Unused Slots, on page 114](#)

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 UCSM to create a virtual drive, UCSM assigns a unique ID that can be used to reliably identify the virtual drive for further operations. UCSM 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, UCSM generates a unique name for the virtual drive.

You can rename virtual drives that are not referenced by any service profile or server.

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

- For Cisco UCS C240, C220, C24, and C22 servers, each storage controller supports 24 virtual drives. For all other servers, each storage controller supports 16 virtual drives.
- 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 and M5 servers, display Operability as NA and Drive State as Unknown.

Configuring Storage Profiles

Configuring a Disk Group Policy

You can choose to configure a disk group policy through automatic or manual disk selection. Configuring a disk group involves the following:

1. [Setting the RAID Level, on page 115](#)
2. [Automatically Configuring Disks in a Disk Group, on page 116](#) or [Manually Configuring Disks in a Disk Group, on page 118](#)
3. [Configuring Virtual Drive Properties, on page 119](#)

Setting the RAID Level

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org# create disk-group-config-policy <i>disk-group-name</i>	Creates a disk group configuration policy with the specified name and enters disk group configuration policy mode.
Step 3	UCS-A /org/disk-group-config-policy* # set raid-level <i>raid-level</i>	Specifies the RAID level for the disk group configuration policy. The RAID levels that you can specify are: <ul style="list-style-type: none"> • raid-0-striped • raid-1-mirrored

	Command or Action	Purpose
		<ul style="list-style-type: none"> • raid-10-mirrored-and-striped • raid-5-striped-parity • raid-6-striped-dual-parity • raid-50-striped-parity-and-striped • raid-60-striped-dual-parity-and-striped
Step 4	UCS-A /org/disk-group-config-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to set the RAID level for a disk group configuration policy.

```
UCS-A# scope org
UCS-A /org # create disk-group-config-policy raid5policy
UCS-A /org/disk-group-config-policy* # set raid-level raid-5-striped-parity
UCS-A /org/disk-group-config-policy* # commit-buffer
```

What to do next

Automatically or manually configure disks as part of the disk group configuration policy.

Automatically Configuring Disks in a Disk Group

You can allow UCSM to automatically select and configure disks in a disk group.

When you create a disk group with RAID 1 policy and configure four disks for it, a RAID1E configuration is created internally by the storage controller.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org# enter disk-group-config-policy <i>disk-group-name</i>	Enters disk group configuration policy mode for the specified disk group name.
Step 3	UCS-A /org/disk-group-config-policy* # enter disk-group-qual	Enters disk group qualification mode. In this mode, UCSM automatically configures disks as part of the specified disk group.
Step 4	UCS-A /org/disk-group-config-policy/disk-group-qual* # set drive-type <i>drive-type</i>	Specifies the drive type for the disk group. You can select: <ul style="list-style-type: none"> • HDD

	Command or Action	Purpose
		<ul style="list-style-type: none"> • SSD • Unspecified <p>Note If you specify Unspecified 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 was SSD, all subsequent drives would be SSD.</p>
Step 5	UCS-A /org/disk-group-config-policy/disk-group-qual* # set min-drive-size <i>drive-size</i>	<p>Specifies the minimum drive size for the disk group. Only disks that match this criteria will be available for selection.</p> <p>The range for minimum drive size is from 0 to 10240 GB. You can also set the minimum drive size as Unspecified. If you set the minimum drive size as Unspecified, drives of all sizes will be available for selection.</p>
Step 6	UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-ded-hot-spares <i>hot-spare-num</i>	<p>Specifies the number of dedicated hot spares for the disk group.</p> <p>The range for dedicated hot spares is from 0 to 24 hot spares. You can also set the number of dedicated hot spares as Unspecified. If you set the number of dedicated hot spares as Unspecified, the hot spares will be selected according to the disk selection process.</p>
Step 7	UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-drives <i>drive-num</i>	<p>Specifies the number of drives for the disk group.</p> <p>The range for drives is from 0 to 24 drives for Cisco UCS C240, C220, C24, and C22 servers. For all other servers, the limit is 16 drives per server. You can also set the number of drives as Unspecified. If you set the number of drives as Unspecified, the number of drives will be selected according to the disk selection process.</p>
Step 8	UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-glob-hot-spares <i>hot-spare-num</i>	<p>Specifies the number of global hot spares for the disk group.</p> <p>The range for global hot spares is from 0 to 24 hot spares. You can also set the number of global hot spares as Unspecified. If you set the number of global hot spares as Unspecified, the global hot spares will be</p>

	Command or Action	Purpose
		selected according to the disk selection process.
Step 9	UCS-A /org/disk-group-config-policy/disk-group-qual* # set use-remaining-disks {no yes}	Specifies whether the remaining disks in the disk group policy should be used or not. The default value for this command is no .
Step 10	UCS-A /org/disk-group-config-policy/disk-group-qual* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to automatically configure disks for a disk group configuration policy.

```
UCS-A# scope org
UCS-A /org # enter disk-group-config-policy raid5policy
UCS-A /org/disk-group-config-policy* # enter disk-group-qual
UCS-A /org/disk-group-config-policy/disk-group-qual* # set drive-type hdd
UCS-A /org/disk-group-config-policy/disk-group-qual* # set min-drive-size 1000
UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-ded-hot-spares 2
UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-drives 7
UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-glob-hot-spares 2
UCS-A /org/disk-group-config-policy/disk-group-qual* # set use-remaining-disks no
UCS-A /org/disk-group-config-policy/disk-group-qual* # commit-buffer

UCS-A# scope org
UCS-A /org # enter disk-group-config-policy raid5policy
UCS-A /org/disk-group-config-policy* # enter disk-group-qual
UCS-A /org/disk-group-config-policy/disk-group-qual* # set drive-type ssd
UCS-A /org/disk-group-config-policy/disk-group-qual* # set min-drive-size 1000
UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-ded-hot-spares 2
UCS-A /org/disk-group-config-policy/disk-group-qual* # set num-drives 7
UCS-A /org/disk-group-config-policy/disk-group-qual* # commit-buffer
```

What to do next

Configure Virtual Drives.

Manually Configuring Disks in a Disk Group

You can manually configure disks for a disk group.

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.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org# enter disk-group-config-policy <i>disk-group-name</i>	Enters disk group configuration policy mode for the specified disk group name.
Step 3	UCS-A /org/disk-group-config-policy* # create local-disk-config-ref <i>slot-num</i>	Creates a local disk configuration reference for the specified slot and enters local disk configuration reference mode.
Step 4	UCS-A /org/disk-group-config-policy/local-disk-config-ref *# set role <i>role</i>	Specifies the role of the local disk in the disk group. You can select: <ul style="list-style-type: none"> • ded-hot-spare: Dedicated hot spare • glob-hot-spare: Global hot spare • normal
Step 5	UCS-A /org/disk-group-config-policy/local-disk-config-ref *# set span-id <i>span-id</i>	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. You can also set the Span ID as Unspecified when spanning information is not required. <p>Note In Cisco UCS Release 2.5, you can have a maximum of 4 span groups.</p>
Step 6	UCS-A /org/disk-group-config-policy/local-disk-config-ref *# commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to manually configure disks for a disk group configuration policy.

```
UCS-A# scope org
UCS-A /org # enter disk-group-config-policy raid5policy
UCS-A /org/disk-group-config-policy* # create local-disk-config-ref 1
UCS-A /org/disk-group-config-policy/local-disk-config-ref *# set role ded-hot-spare
UCS-A /org/disk-group-config-policy/local-disk-config-ref* # set span-id 1
UCS-A /org/disk-group-config-policy/local-disk-config-ref *# commit-buffer
```

What to do next

Configure Virtual Drive Properties.

Configuring Virtual Drive Properties

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

If you try to associate to a server that does not support these properties, a configuration error will be generated.

Only the following storage controllers support these properties:

- LSI 6G MegaRAID SAS 9266-8i
- LSI 6G MegaRAID SAS 9271-8i
- LSI 6G MegaRAID 9265-8i
- LSI MegaRAID SAS 2208 ROMB
- LSI MegaRAID SAS 9361-8i

For the LSI MegaRAID SAS 2208 ROMB controller, these properties are supported only in the B420-M3 blade server. For the other controllers, these properties are supported in multiple rack servers.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org# scope disk-group-config-policy <i>disk-group-name</i>	Enters disk group configuration policy mode for the specified disk group name.
Step 3	UCS-A /org/disk-group-config-policy* # create virtual-drive-def	Creates a virtual drive definition and enters the virtual drive definition mode.
Step 4	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set access-policy <i>policy-type</i>	Specifies the access policy. This can be one of the following: <ul style="list-style-type: none"> • blocked • platform-default • read-only: • read-write
Step 5	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set drive-cache <i>state</i>	Specifies the state of the drive cache. This can be one of the following: <ul style="list-style-type: none"> • enable • disable • no-change • platform-default <p>Important In Cisco UCS Release 2.5, the drive cache state cannot be changed. It will remain as platform-default, irrespective of the drive cache state that you select.</p>

	Command or Action	Purpose
Step 6	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set io-policy <i>policy-type</i>	Specifies the I/O policy. This can be one of the following: <ul style="list-style-type: none"> • cached • direct • platform-default
Step 7	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set read-policy <i>policy-type</i>	Specifies the read policy. This can be one of the following: <ul style="list-style-type: none"> • normal • platform-default • read-ahead
Step 8	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set strip-size <i>strip-size</i>	Specifies the strip size. This can be one of the following: <ul style="list-style-type: none"> • 64 KB • 128 KB • 256 KB • 512 KB • 1024 KB • platform-default
Step 9	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set write-cache-policy <i>policy-type</i>	Specifies the write-cache-policy. This can be one of the following: <ul style="list-style-type: none"> • always-write-back • platform-default • write-back-good-bbu • write-through
Step 10	UCS-A /org/disk-group-config-policy/virtual-drive-def* # commit-buffer	Commits the transaction to the system configuration.
Step 11	UCS-A /org/disk-group-config-policy/virtual-drive-def* # show	Displays the configured virtual drive properties.

Example

This example shows how to configure virtual disk properties:

```
UCS-A# scope org
UCS-A /org # scope disk-group-config-policy raid0policy
UCS-A /org/disk-group-config-policy # create virtual-drive-def
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set access-policy read-write
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set drive-cache enable
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set io-policy cached
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set read-policy normal
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set strip-size 1024
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set write-cache-policy write-through
UCS-A /org/disk-group-config-policy/virtual-drive-def* # commit-buffer
UCS-A /org/disk-group-config-policy/virtual-drive-def # show

Virtual Drive Def:
  Strip Size (KB): 1024KB
  Access Policy: Read Write
  Read Policy: Normal
  Configured Write Cache Policy: Write Through
  IO Policy: Cached
  Drive Cache: Enable
UCS-A /org/disk-group-config-policy/virtual-drive-def #
```

What to do next

Create a Storage Profile

Creating a Storage Profile

You can create a storage profile at the org level and at the service-profile level.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # create storage-profile <i>storage-profile-name</i>	Creates a storage profile with the specified name at the org level and enters storage-profile configuration mode.
Step 3	UCS-A /org/storage-profile* # commit-buffer	Commits the transaction to the system configuration.
Step 4	(Optional) UCS-A /org* # enter service-profile <i>service-profile-name</i>	Enters the specified service profile.
Step 5	(Optional) UCS-A /org/service-profile* # create storage-profile-def	Creates a storage profile at the service-profile level.

	Command or Action	Purpose
Step 6	UCS-A /org/service-profile/storage-profile-def* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to create a storage profile at the org level.

```
UCS-A# scope org
UCS-A /org # create storage-profile stp2
UCS-A /org/storage-profile* # commit-buffer
```

This example shows how to create a storage profile at the service-profile level.

```
UCS-A# scope org
UCS-A /org* # enter service-profile sp1
UCS-A /org/service-profile* # create storage-profile-def
UCS-A /org/service-profile/storage-profile-def* # commit-buffer
```

What to do next

Create Local LUNs

Deleting a Storage Profile

You can delete a storage profile that was created at the org level or at the service-profile level.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # delete storage-profile <i>storage-profile-name</i>	Deletes the storage profile with the specified name at the org level.
Step 3	(Optional) UCS-A /org # scope service-profile <i>service-profile-name</i>	Enters the specified service profile.
Step 4	(Optional) UCS-A /org/service-profile # delete storage-profile-def	Deletes the dedicated storage profile at the service-profile level.

Example

This example shows how to delete a storage profile at the org level.

```
UCS-A # scope org
UCS-A /org # delete storage-profile stor1
```

This example shows how to delete a storage profile at the service-profile level.

```
UCS-A # scope org
UCS-A /org # scope service-profile sp1
UCS-A /org/service-profile # delete storage-profile-def
```

Creating Local LUNs

You can create local LUNs within a storage profile at the org level and within a dedicated storage profile at the service-profile level.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # enter storage-profile <i>storage-profile-name</i>	Enters storage-profile mode for the specified storage profile.
Step 3	UCS-A /org/storage-profile* # create local-lun <i>lun-name</i>	Creates a local LUN with the specified name.
Step 4	UCS-A /org/storage-profile/local-lun* # set auto-deploy { auto-deploy no-auto-deploy }	Specifies whether the LUN should be auto-deployed or not.
Step 5	UCS-A /org/storage-profile/local-lun* # set disk-policy-name <i>disk-policy-name</i>	Specifies the name of the disk policy name for this LUN.
Step 6	UCS-A /org/storage-profile/local-lun* # set order <i>order-num</i>	Specifies the order of this LUN. The order can range from 1 to 64. You can also specify that the order should be lowest-available for the system to automatically assign the lowest available order to the LUN. Multiple LUNs referenced by a storage profile must have unique names and unique orders.
Step 7	UCS-A /org/storage-profile/local-lun* # set expand-to-avail { no yes }	Specifies whether the LUN should be expanded to the entire available disk group. For each service profile, only one LUN can be configured to use this option.
Step 8	UCS-A /org/storage-profile/local-lun* # set size <i>size</i>	Specifies the size of this LUN in GB. The size can range from 1 GB to 10240 GB. Note You do not need to specify a LUN size while claiming an orphaned LUN.

	Command or Action	Purpose
Step 9	UCS-A /org/storage-profile/local-lun* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to configure a local LUN within a storage profile at the org level.

```
UCS-A# scope org
UCS-A /org # enter storage-profile stp2
UCS-A /org/storage-profile* # create local-lun lun2
UCS-A /org/storage-profile/local-lun* # set disk-policy-name dpn2
UCS-A /org/storage-profile/local-lun* # set order 2
UCS-A /org/storage-profile/local-lun* # set size 1000
UCS-A /org/storage-profile/local-lun* # commit-buffer
```

```
UCS-A# scope org
UCS-A /org # enter storage-profile stp2
UCS-A /org/storage-profile* # create local-lun lun2
UCS-A /org/storage-profile/local-lun* # set auto-deploy no-auto-deploy
UCS-A /org/storage-profile/local-lun* # set disk-policy-name dpn2
UCS-A /org/storage-profile/local-lun* # set expand-to-avail yes
UCS-A /org/storage-profile/local-lun* # set size 1000
UCS-A /org/storage-profile/local-lun* # commit-buffer
```

This example shows how to configure a local LUN within a dedicated storage profile at the service-profile level.

```
UCS-A# scope org
UCS-A /org* # enter service-profile stp1
UCS-A /org/service-profile* # enter storage-profile-def
UCS-A /org/service-profile/storage-profile-def # create local-lun lun1
UCS-A /org/service-profile/storage-profile-def/local-lun* # set disk-policy-name dpn1
UCS-A /org/service-profile/storage-profile-def/local-lun* # set order 1
UCS-A /org/service-profile/storage-profile-def/local-lun* # set size 1000
UCS-A /org/service-profile/storage-profile-def/local-lun* # commit-buffer

UCS-A# scope org
UCS-A /org # enter service-profile spl
UCS-A /org/service-profile* # enter storage-profile-def
UCS-A /org/service-profile/storage-profile-def # create local-lun lun1
UCS-A /org/service-profile/storage-profile-def/local-lun* # set auto-deploy no-auto-deploy
UCS-A /org/service-profile/storage-profile-def/local-lun* # set disk-policy-name dpn1
UCS-A /org/service-profile/storage-profile-def/local-lun* # set expand-to-avail yes
UCS-A /org/service-profile/storage-profile-def/local-lun* # set size 1000
UCS-A /org/service-profile/storage-profile-def/local-lun* # commit-buffer
```

What to do next

Associate a Storage Profile with a Service Profile

Reordering Local LUNs In a Storage Profile

You can use the **set order** command to change the local LUN visibility order to the server. This operation will reboot the server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # enter storage-profile <i>storage-profile-name</i>	Enters storage-profile mode for the specified storage profile.
Step 3	UCS-A /org/storage-profile # enter local-lun <i>lun-name</i>	Enters local-lun mode for the specified local LUN.
Step 4	UCS-A /org/storage-profile/local-lun* # set disk-policy-name <i>disk-policy-name</i>	Specifies the name of the disk policy name for this LUN.
Step 5	UCS-A /org/storage-profile/local-lun # set order <i>order-num</i>	Specifies the order of this LUN. The order can range from 1 to 64. You can also specify that the order should be lowest-available for the system to automatically assign the lowest available order to the LUN.
Step 6	UCS-A /org/storage-profile/local-lun* # set size <i>size</i>	Specifies the size of this LUN in GB. The size can range from 1 GB to 10240 GB.
Step 7	UCS-A /org/storage-profile/local-lun # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to reorder a local LUN within a storage profile at the org level.

```
UCS-A# scope org
UCS-A /org # enter storage-profile stp1
UCS-A /org/storage-profile* # enter local-lun lun1
UCS-A /org/storage-profile/local-lun* # set disk-policy-name dpn1
UCS-A /org/storage-profile/local-lun* # set order 1
UCS-A /org/storage-profile/local-lun* # set size 10
UCS-A /org/storage-profile/local-lun* # exit
UCS-A /org/storage-profile* # enter local-lun lun2
UCS-A /org/storage-profile/local-lun* # set disk-policy-name dpn2
UCS-A /org/storage-profile/local-lun* # set order 2
UCS-A /org/storage-profile/local-lun* # set size 10
UCS-A /org/storage-profile/local-lun* # exit
UCS-A /org/storage-profile* # commit-buffer

UCS-A /org/storage-profile # show configuration

    enter storage-profile stp1
        enter local-lun lun1
```

```

        set auto-deploy auto-deploy
        set disk-policy-name dpn1
        set order 1
        set size 10
    exit
    enter local-lun lun2
        set auto-deploy auto-deploy
        set disk-policy-name dpn2
        set order 2
        set size 10
    exit
    set descr ""
exit

UCS-A /org/storage-profile # enter local-lun lun1
UCS-A /org/storage-profile/local-lun # set order 2
UCS-A /org/storage-profile/local-lun* # exit
UCS-A /org/storage-profile* # enter local-lun lun2
UCS-A /org/storage-profile/local-lun* # set order 1
UCS-A /org/storage-profile/local-lun* # exit
UCS-A /org/storage-profile* # commit-buffer
UCS-A /org/storage-profile # show configuration

    enter storage-profile stp1
        enter local-lun lun1
            set auto-deploy auto-deploy
            set disk-policy-name dpn1
            set order 2
            set size 10
        exit
        enter local-lun lun2
            set auto-deploy auto-deploy
            set disk-policy-name dpn2
            set order 1
            set size 10
        exit
        set descr ""
    exit

```

Deleting Local LUNs In a Storage Profile

When a LUN is deleted, the corresponding virtual drive is marked as orphan after the virtual drive reference is removed from the server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # enter storage-profile <i>storage-profile-name</i>	Enters storage-profile mode for the specified storage profile.
Step 3	(Optional) UCS-A /org/storage-profile* # show local-lun	Displays the local LUNs in the specified storage profile.

	Command or Action	Purpose
Step 4	UCS-A /org/storage-profile* # delete local-lun <i>lun-name</i>	Deletes the specified LUN.
Step 5	UCS-A /org/storage-profile* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to delete a LUN in a storage profile.

```
UCS-A # scope org
UCS-A /org # enter storage-profile stp2
UCS-A /org/storage-profile # show local-lun
```

Local SCSI LUN:

LUN Name	Size (GB)	Order	Disk Policy Name	Auto Deploy
-----	-----	-----	-----	-----
luna	1	2	raid0	Auto Deploy
lunb	1	1	raid0	Auto Deploy

```
UCS-A /org/storage-profile # delete local-lun luna
UCS-A /org/storage-profile* # commit-buffer
UCS-A /org/storage-profile* # show local-lun
```

Local SCSI LUN:

LUN Name	Size (GB)	Order	Disk Policy Name	Auto Deploy
-----	-----	-----	-----	-----
lunb	1	1	raid0	Auto Deploy

Associating a Storage Profile with a Service Profile

A storage profile created under org can be referred by multiple service profiles, and a name reference in service profile is needed to associate the storage profile with a service profile.



Important

Storage profiles can be defined under org and under service profile (dedicated). Hence, a service profile inherits local LUNs from both possible storage profiles. A service profile can have a maximum of two such local LUNs.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>service-profile-name</i>	Enters the specified service profile mode.
Step 3	UCS-A /org/service-profile # set storage-profile-name <i>storage-profile-name</i>	Associates the specified storage profile with the service profile. Note To dissociate the service profile from a storage profile, use the set storage-profile-name command and specify "" as the storage profile name.
Step 4	UCS-A /org/service-profile* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to associate a storage profile with a service profile.

```
UCS-A# scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile # set storage-profile-name stp2
```

This example shows how to dissociate a service profile from a storage profile.

```
UCS-A# scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile # set storage-profile-name ""
```

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

	Command or Action	Purpose
Step 1	UCS-A /org/service-profile # show local-lun-ref	<p>Displays the following detailed information about all the local LUNs inherited by the specified service profile:</p> <ul style="list-style-type: none"> • Name—LUN name in the storage profile. • Admin State—Specifies whether a local LUN should be deployed or not. Admin state can be Online or Undeployed. <p>When the local LUN is being referenced by a service profile, if the auto-deploy status is no-auto-deploy then the admin state will be Undeployed, else it will be Online. After the local LUN is referenced by a service profile, any change made to this local LUN's auto-deploy status is not reflected in the admin state of the LUN inherited by the service profile.</p> <ul style="list-style-type: none"> • RAID Level—Summary of the RAID level of the disk group used. • Provisioned Size (GB)—Size, in GB, of the LUN specified in the storage profile. • Assigned Size (MB)—Size, in MB, assigned by UCSM. • Config State—State of LUN configuration. The states can be one of the following: <ul style="list-style-type: none"> • 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.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Not In Use—Service profile is using the virtual drive, but the virtual drive is not associated with a server. • Reference LUN—The preprovisioned virtual drive name, or UCSM-generated virtual drive name. • Deploy Name—The virtual drive name after deployment. • ID—Virtual drive ID. • Drive State—State of the virtual drive. The states are: <ul style="list-style-type: none"> • Unknown • Optimal • Degraded • Inoperable • Partially Degraded

Example

```
UCS-A /org/service-profile # show local-lun-ref
```

```
Local LUN Ref:
```

Profile Size (MB)	LUN Name Config	Admin State	RAID Level	Referenced Lun	RAID Level	Deploy Name	ID	Provisioned Size (GB)	Assigned Drive State
1024	luna Applied	Online	RAID 0	luna-1	Striped	luna-1	1003	1	Optimal
1024	lunb Applied	Online	RAID 0	lunb-1	Striped	lunb-1	1004	1	Optimal

```
UCS-A /org/service-profile #
```

```
Local LUN Ref:
```

Name	Admin State	RAID Level	Provisioned Size (GB)	Assigned
lun111	Online	RAID 0	30	30720

```

Applied      lun111-1      lun111-1      1001      Optimal
lun201      Online      Unspecified      1
Not Applied

```

Importing Foreign Configurations for a RAID Controller

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope raid-controller <i>raid-contr-id</i> {sas sata}	Enters RAID controller chassis mode.
Step 3	UCS-A /chassis/raid-controller # set admin-state import-foreign-configuration	Allows import of configurations from local disks that are in the Foreign Configuration state.

Example

This example shows how to import foreign configurations from local disks that are in the **Foreign Configuration** state:

```

UCS-A# scope chassis 1
UCS-A /chassis # scope raid-controller 1 sas
UCS-A /chassis/raid-controller # set admin-state import-foreign-configuration
UCS-A /chassis/raid-controller* #

```

Configuring Local Disk Operations

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope raid-controller <i>raid-contr-id</i> {sas sata}	Enters RAID controller chassis mode.
Step 3	UCS-A /chassis/raid-controller # scope local-disk <i>local-disk-id</i>	Enters local disk configuration mode.
Step 4	UCS-A /chassis/raid-controller/local-disk # set admin-state {clear-foreign-configuration dedicated-hot-spare [<i>admin-vd-id</i>] prepare-for-removal remove-hot-spare	Configures the local disk to one of the following states: <ul style="list-style-type: none"> • clear-foreign-configuration—Clears any foreign configuration that exists in a local

	Command or Action	Purpose
	unconfigured-good undo-prepare-for-removal}	<p>disk when it is introduced into a new configuration.</p> <ul style="list-style-type: none"> • dedicated-hot-spare—Specifies the local disk as a dedicated hot spare. The admin virtual drive ID that you can assign ranges from 0 to 4294967295. • prepare-for-removal—Specifies that the local disk is marked for removal from the chassis. • remove-hot-spare—Specifies that the local disk is no longer a hot spare. Use this only to clear any mismatch faults. • unconfigured-good—Specifies that the local disk can be configured. • undo-prepare-for-removal—Specifies that the local disk is no longer marked for removal from the chassis.

Example

This example shows how to clear any foreign configuration from a local disk:

```
UCS-A /chassis/raid-controller/local-disk # set admin-state clear-foreign-configuration
```

This example shows how to specify a local disk as a dedicated hot spare:

```
UCS-A /chassis/raid-controller/local-disk* # set admin-state dedicated-hot-spare 1001
```

This example shows how to specify that a local disk is marked for removal from the chassis:

```
UCS-A /chassis/raid-controller/local-disk* # set admin-state prepare-for-removal
```

This example shows how to specify that a local disk is marked for removal as a hot spare:

```
UCS-A /chassis/raid-controller/local-disk* # set admin-state remove-hot-spare
```

This example shows how to specify that a local disk is working, but is unconfigured for use:

```
UCS-A /chassis/raid-controller/local-disk* # set admin-state unconfigured-good
```

This example shows how to specify that a local disk is no longer marked for removal from the chassis:

```
UCS-A /chassis/raid-controller/local-disk* # set admin-state undo-prepare-for-removal
```

Configuring Virtual Drive Properties

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

If you try to associate to a server that does not support these properties, a configuration error will be generated.

Only the following storage controllers support these properties:

- LSI 6G MegaRAID SAS 9266-8i
- LSI 6G MegaRAID SAS 9271-8i
- LSI 6G MegaRAID 9265-8i
- LSI MegaRAID SAS 2208 ROMB
- LSI MegaRAID SAS 9361-8i

For the LSI MegaRAID SAS 2208 ROMB controller, these properties are supported only in the B420-M3 blade server. For the other controllers, these properties are supported in multiple rack servers.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org# scope disk-group-config-policy <i>disk-group-name</i>	Enters disk group configuration policy mode for the specified disk group name.
Step 3	UCS-A /org/disk-group-config-policy* # create virtual-drive-def	Creates a virtual drive definition and enters the virtual drive definition mode.
Step 4	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set access-policy <i>policy-type</i>	Specifies the access policy. This can be one of the following: <ul style="list-style-type: none"> • blocked • platform-default • read-only: • read-write
Step 5	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set drive-cache <i>state</i>	Specifies the state of the drive cache. This can be one of the following: <ul style="list-style-type: none"> • enable • disable • no-change • platform-default

	Command or Action	Purpose
		Important In Cisco UCS Release 2.5, the drive cache state cannot be changed. It will remain as platform-default , irrespective of the drive cache state that you select.
Step 6	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set io-policy <i>policy-type</i>	Specifies the I/O policy. This can be one of the following: <ul style="list-style-type: none"> • cached • direct • platform-default
Step 7	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set read-policy <i>policy-type</i>	Specifies the read policy. This can be one of the following: <ul style="list-style-type: none"> • normal • platform-default • read-ahead
Step 8	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set strip-size <i>strip-size</i>	Specifies the strip size. This can be one of the following: <ul style="list-style-type: none"> • 64 KB • 128 KB • 256 KB • 512 KB • 1024 KB • platform-default
Step 9	UCS-A /org/disk-group-config-policy/virtual-drive-def* # set write-cache-policy <i>policy-type</i>	Specifies the write-cache-policy. This can be one of the following: <ul style="list-style-type: none"> • always-write-back • platform-default • write-back-good-bbu • write-through
Step 10	UCS-A /org/disk-group-config-policy/virtual-drive-def* # commit-buffer	Commits the transaction to the system configuration.

	Command or Action	Purpose
Step 11	UCS-A /org/disk-group-config-policy/virtual-drive-def* # show	Displays the configured virtual drive properties.

Example

This example shows how to configure virtual disk properties:

```
UCS-A# scope org
UCS-A /org # scope disk-group-config-policy raid0policy
UCS-A /org/disk-group-config-policy # create virtual-drive-def
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set access-policy read-write
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set drive-cache enable
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set io-policy cached
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set read-policy normal
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set strip-size 1024
UCS-A /org/disk-group-config-policy/virtual-drive-def* # set write-cache-policy write-through
UCS-A /org/disk-group-config-policy/virtual-drive-def* # commit-buffer
UCS-A /org/disk-group-config-policy/virtual-drive-def # show

Virtual Drive Def:
  Strip Size (KB): 1024KB
  Access Policy: Read Write
  Read Policy: Normal
  Configured Write Cache Policy: Write Through
  IO Policy: Cached
  Drive Cache: Enable
UCS-A /org/disk-group-config-policy/virtual-drive-def #
```

What to do next

Create a Storage Profile

Deleting an Orphaned Virtual Drive

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope raid-controller <i>raid-contr-id</i> {sas sata}	Enters RAID controller chassis mode.
Step 3	(Optional) UCS-A /chassis/raid-controller # delete virtual-drive id <i>virtual-drive-id</i>	Deletes the orphaned virtual drive with the specified virtual drive ID.
Step 4	(Optional) UCS-A /chassis/raid-controller # delete virtual-drive name <i>virtual-drive-name</i>	Deletes the orphaned virtual drive with the specified virtual drive name.
Step 5	(Optional) UCS-A /chassis/raid-controller # scope virtual-drive <i>virtual-drive-id</i>	Enters virtual drive mode for the specified orphaned virtual drive.

	Command or Action	Purpose
Step 6	UCS-A /chassis/raid-controller/virtual-drive # set admin-state delete	Deletes the orphaned virtual drive.
Step 7	UCS-A /chassis/raid-controller/virtual-drive # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to delete an orphan virtual drive by specifying the virtual drive ID.

```
UCS-A# scope chassis 1
UCS-A /chassis # scope raid-controller 1 sas
UCS-A /chassis/raid-controller # show virtual-drive
```

```
Virtual Drive:
  ID: 1001
  Name: lun111-1
  Block Size: 512
  Blocks: 62914560
  Size (MB): 30720
  Operability: Operable
  Presence: Equipped
  Oper Device ID: 0
  Change Qualifier: No Change
  Config State: Applied
  Deploy Action: No Action

  ID: 1002
  Name: luna-1
  Block Size: 512
  Blocks: 2097152
  Size (MB): 1024
  Operability: Operable
  Presence: Equipped
  Oper Device ID: 1
  Change Qualifier: No Change
  Config State: Orphaned
  Deploy Action: No Action

  ID: 1003
  Name: lunb-1
  Block Size: 512
  Blocks: 2097152
  Size (MB): 1024
  Operability: Operable
  Presence: Equipped
  Oper Device ID: 2
  Change Qualifier: No Change
  Config State: Orphaned
  Deploy Action: No Action

  ID: 1004
  Name: lunb-2
  Block Size: 512
  Blocks: 2097152
  Size (MB): 1024
  Operability: Operable
  Presence: Equipped
  Oper Device ID: 3
```

```
Change Qualifier: No Change
Config State: Orphaned
Deploy Action: No Action
```

```
ID: 1005
Name: luna-2
Block Size: 512
Blocks: 2097152
Size (MB): 1024
Operability: Operable
Presence: Equipped
Oper Device ID: 4
Change Qualifier: No Change
Config State: Orphaned
Deploy Action: No Action
```

```
...
```

```
UCS-A /chassis/raid-controller # delete virtual-drive id 1002
Warning: When committed, the virtual drive will be deleted, which may result in data loss.
```

```
UCS-A /chassis/raid-controller # commit-buffer
```

This example shows how to delete an orphan virtual drive by specifying the virtual drive name.

```
UCS-A# scope chassis 1
UCS-A /chassis # scope raid-controller 1 sas
UCS-A /chassis/raid-controller # show virtual-drive
```

```
Virtual Drive:
  ID: 1001
  Name: lun111-1
  Block Size: 512
  Blocks: 62914560
  Size (MB): 30720
  Operability: Operable
  Presence: Equipped
  Oper Device ID: 0
  Change Qualifier: No Change
  Config State: Applied
  Deploy Action: No Action

  ID: 1003
  Name: lunb-1
  Block Size: 512
  Blocks: 2097152
  Size (MB): 1024
  Operability: Operable
  Presence: Equipped
  Oper Device ID: 2
  Change Qualifier: No Change
  Config State: Orphaned
  Deploy Action: No Action

  ID: 1004
  Name: lunb-2
  Block Size: 512
  Blocks: 2097152
  Size (MB): 1024
  Operability: Operable
  Presence: Equipped
  Oper Device ID: 3
  Change Qualifier: No Change
  Config State: Orphaned
  Deploy Action: No Action
```

```

ID: 1005
Name: luna-2
Block Size: 512
Blocks: 2097152
Size (MB): 1024
Operability: Operable
Presence: Equipped
Oper Device ID: 4
Change Qualifier: No Change
Config State: Orphaned
Deploy Action: No Action

```

...

```

UCS-A /chassis/raid-controller # delete virtual-drive name lunb-1
Warning: When committed, the virtual drive will be deleted, which may result in data loss.

```

```

UCS-A /chassis/raid-controller # commit-buffer

```

This example shows how to delete an orphan virtual drive by setting the admin-state.

```

UCS-A# scope chassis 1
UCS-A /chassis # scope raid-controller 1 sas
UCS-A /chassis/raid-controller # scope virtual-drive 1004
UCS-A /chassis/raid-controller/virtual-drive # set admin-state delete

```

Warning: When committed, the virtual drive will be deleted, which may result in data loss.

```

UCS-A /chassis/raid-controller/virtual-drive # commit-buffer

```

Renaming an Orphaned Virtual Drive

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope raid-controller <i>raid-contr-id</i> {sas sata}	Enters RAID controller chassis mode.
Step 3	UCS-A /chassis/raid-controller # scope virtual-drive <i>virtual-drive-id</i>	Enters virtual drive mode for the specified virtual drive.
Step 4	UCS-A /chassis/raid-controller/virtual-drive # set name <i>virtual-drive-name</i>	Specifies a name for the orphan virtual drive.
Step 5	UCS-A /chassis/raid-controller/virtual-drive # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to specify a name for an orphan virtual drive.

```

UCS-A /chassis # scope raid-controller 1 sas
UCS-A /chassis/raid-controller # scope virtual-drive 1060
UCS-A /chassis/raid-controller/virtual-drive* # set name vdl

```

```
UCS-A /chassis/raid-controller/virtual-drive* # commit-buffer
```

Configuring the Boot Policy for a Local LUN



Note In Cisco UCS Manager Release 2.5, you cannot configure JBOD as a boot device.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope boot-policy <i>policy-name</i>	Enters organization boot policy mode for the specified boot policy.
Step 3	UCS-A /org/boot-policy # create storage	Creates a storage boot for the boot policy and enters organization boot policy storage mode.
Step 4	UCS-A /org/boot-policy/storage # create local	Creates a local storage location and enters the boot policy local storage mode.
Step 5	UCS-A /org/boot-policy/storage/local/ # create local-lun	Specifies a local hard disk drive as the local storage.
Step 6	UCS-A /org/boot-policy/storage/local/local-lun # create local-lun-image-path { primary secondary }	Specifies the boot order for the LUN that you specify. Important Cisco UCS Manager Release 2.2(4) does not support secondary boot order.
Step 7	UCS-A /org/boot-policy/storage/local/local-lun/local-lun-image-path # set lunname <i>lun_name</i>	Specifies the name of the LUN that you want to boot from.
Step 8	UCS-A /org/boot-policy/storage/local/local-storage-device # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create a boot policy named lab1-boot-policy, create a local hard disk drive boot for the policy, specify a boot order and a LUN to boot from, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope boot-policy lab1-boot-policy
```

```

UCS-A /org/boot-policy* # create storage
UCS-A /org/boot-policy/storage* # create local
UCS-A /org/boot-policy/storage/local* # create local-lun
UCS-A /org/boot-policy/storage/local/local-lun # create local-lun-image-path primary
UCS-A /org/boot-policy/storage/local/local-lun/local-lun-image-path # set lunname luna
UCS-A /org/boot-policy/storage/local/local-lun/local-lun-image-path # commit-buffer
UCS-A /org/boot-policy/storage/local/local-lun/local-lun-image-path #

```

What to do next

Include the boot policy in a service profile and template.

Local LUN Operations in a Service Profile

Although a service profile is derived from a service profile template, the following operations can be performed for each local LUN at the individual service profile level:

- [Preprovisioning a LUN Name or Claiming an Orphan LUN, on page 141](#)
- [Deploying and Undeploying a LUN, on page 142](#)
- [Renaming a Service Profile Referenced LUN, on page 143](#)



Note Preprovisioning a LUN name, claiming an orphan LUN, and deploying or undeploying a LUN result in server reboot.

Preprovisioning a LUN Name or Claiming an Orphan LUN

You can preprovision a LUN name or claim an orphan LUN by using the **set ref-name** command. Preprovisioning a LUN name or claiming an orphan LUN can be done only when the admin state of the LUN is **Undeployed**. You can also manually change the admin state of the LUN to **Undeployed** and claim an orphan LUN.



Important This operation will reboot the server.

If the LUN name is empty, set a LUN name before claiming it.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org# scope service-profile <i>service-profile-name</i>	Enters the specified service profile mode.

	Command or Action	Purpose
Step 3	UCS-A /org/service-profile# enter local-lun-ref <i>lun-name</i>	Enters the specified LUN.
Step 4	UCS-A /org/service-profile/local-lun-ref# set ref-name <i>ref-lun-name</i>	Sets the referenced LUN name. 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.

- If the LUN exists and is not orphaned, a configuration failure occurs.
- If a LUN is already referred to and the ref-name is changed, it will release the old LUN and will claim or create a LUN with the ref-name. The old LUN is marked as an orphan after the LUN reference is removed from the server.

Example

This examples shows how to preprovision a LUN name.

```
UCS-A# scope org
UCS-A /org # scope service-profile sp1
UCS-A /org/service-profile* # enter local-lun-ref lun1
UCS-A /org/service-profile/local-lun-ref* # set ref-name lun2
```

Deploying and Undeploying a LUN

You can deploy or undeploy a LUN by using the **admin-state** command. If the admin state of a local LUN is **Undeployed**, the reference of that LUN is removed and the LUN is not deployed.



Important

This operation will reboot the server.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org# scope service-profile <i>service-profile-name</i>	Enters the specified service profile mode.
Step 3	UCS-A /org/service-profile# enter local-lun-ref <i>lun-name</i>	Enters the specified LUN.
Step 4	UCS-A /org/service-profile/local-lun-ref# set admin-state { online undeployed }	Sets the admin state of the specified LUN to online or undeployed .

	Command or Action	Purpose
		If a LUN is already referred to and the admin state is set to undeployed , it will release the old LUN. The old LUN is marked as orphan after the LUN reference is removed from the server.

Example

This examples shows how to deploy a LUN.

```
UCS-A# scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile* # enter local-lun-ref lun1
UCS-A /org/service-profile/local-lun-ref* # set admin-state online
```

This examples shows how to undeploy a LUN.

```
UCS-A# scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile* # enter local-lun-ref lun1
UCS-A /org/service-profile/local-lun-ref* # set admin-state undeployed
```

Renaming a Service Profile Referenced LUN

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org# scope service-profile <i>service-profile-name</i>	Enters the specified service profile mode.
Step 3	UCS-A /org/service-profile# enter local-lun-ref <i>lun-name</i>	Enters the specified LUN.
Step 4	UCS-A /org/service-profile/local-lun-ref# set name	Renames the referenced LUN.

Example

This examples shows how to rename a LUN referenced by a service profile.

```
UCS-A# scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile* # enter local-lun-ref lun1
UCS-A /org/service-profile/local-lun-ref* # set name lun11
```




CHAPTER 11

Configuring SD Card Support

This chapter includes the following sections:

- [FlexFlash Secure Digital Card Support, on page 145](#)
- [FlexUtil Secure Digital Card Support, on page 148](#)

FlexFlash Secure Digital Card Support

Overview

Cisco UCS B-Series, C-Series M3 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.

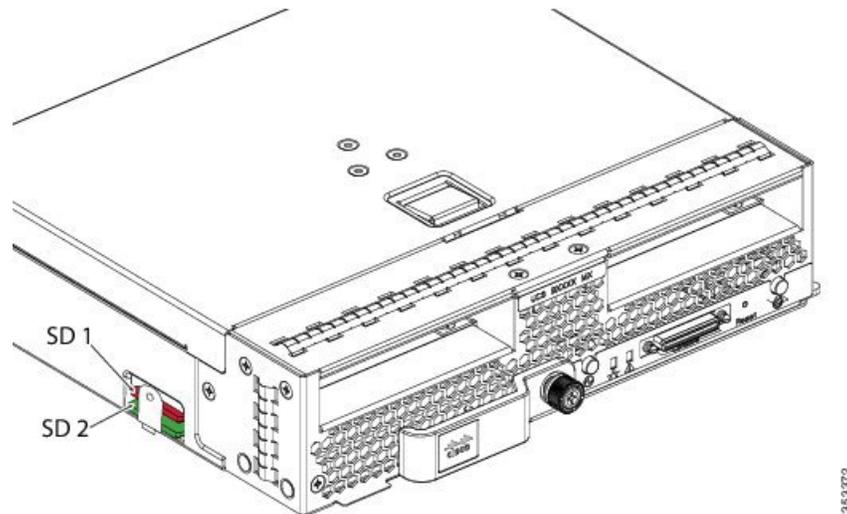
You can populate one or both the SD card slots that are provided. If two SD cards are populated, you can use them in a mirrored mode.



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 3: 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.
- 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 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.

Enabling or Disabling FlexFlash SD Card Support

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope local-disk-config-policy <i>policy-name</i>	Enters the specified local disk configuration policy mode.
Step 3	UCS-A /org/local-disk-config-policy # set flexflash-state {enable disable}	Specifies whether FlexFlash SD card support is enabled.
Step 4	UCS-A /org/local-disk-config-policy # set flexflash-raid-reporting-state {enable disable}	Specifies whether FlexFlash RAID reporting support is enabled.

	Command or Action	Purpose
		Note If only one SD card is installed, the FlexFlash inventory displays the RAID State as Disabled and the RAID Health as NA.
Step 5	UCS-A /org/local-disk-config-policy# set flexflash-removable-state { no no-change yes }	Specifies whether the FlexFlash removable state is disabled or enabled. The non-removable state is for hypervisors that require that option. All other remaining hypervisors can use the no-change option.
Step 6	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system.

Example

The following example shows how to enable FlexFlash SD card support and FlexFlash RAID reporting state on the local disk config policy default, and commits the transaction to the system:

```
UCS-A# scope org/
UCS-A /org # scope local-disk-config-policy default
UCS-A /org/local-disk-config-policy #set flexflash-state enable
UCS-A /org/local-disk-config-policy# #set flexflash-raid-reporting-state enable
UCS-A /org/local-disk-config-policy# #set flexflash-removable-state no
UCS-A /org/local-disk-config-policy* # commit-buffer
UCS-A /org/local-disk-config-policy #
```

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 12

Mini Storage

- [Mini Storage, on page 149](#)

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.

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

	Command or Action	Purpose
Step 1	UCS-A# scope server [<i>chassis-num/server-num</i> <i>dynamic-uuid</i>]	Enters server mode for the specified server.
Step 2	UCS-A /chassis/server # show mini-storage [<i>detail</i>]	Displays detailed information about the mini storage module for the specified server.

Example

This example displays detailed information about the mini storage module for server 6:

```

UCS-A# scope server 1/6
UCS-A /chassis/server # show mini-storage detail

Mini Storage Module:
  ID: 1
  Type: M2
  Model: UCS-MSTOR-M2
  Vendor: Cisco Systems Inc
  HW Rev: 0
  Serial: FCH2050JDHM
  VID: V00
  Part Number: 73-17926-04
  Product Name: Cisco UCS Mini-Storage Carrier for M.2
  Caption: Cisco UCS Mini-Storage Carrier for M.2 (holds up to 2)
  Description: Dual M.2 Mini-Storage Carrier (holds up to 2 M.2 modules)

```

Viewing the Storage Controller for the Mini Storage

Mini storage modules are supported only on M5 and higher servers.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server [<i>chassis-num/server-num</i> <i>dynamic-uuid</i>]	Enters server mode for the specified server.
Step 2	UCS-A /chassis/server # scope mini-storage <i>id m2</i> <i>sd</i>	Enters mini-storage mode for the specified server and mini storage card type.
Step 3	UCS-A /chassis/server/mini-storage # show referenced-controller [detail]	Displays information about the storage controller referenced by the mini storage module for the specified server.

Example

This example displays information about the storage controller for the M.2 mini storage card in server 6:

```

UCS-A# scope server 1/6
UCS-A /chassis/server # scope mini-storage 1 m2
UCS-A /chassis/server/mini-storage # show referenced-controller detail

```

```

Referenced Controller:
  ID: 1
  Type: PCH

```

This example displays information about the storage controller for the SD mini storage card in server 3:

```

UCS-A# scope server 1/3
UCS-A /chassis/server # scope mini-storage 1 sd
UCS-A /chassis/server/mini-storage # show referenced-controller detail

```

```
Referenced Controller:  
  ID: 1  
  Type: PCH
```




CHAPTER 13

SED Security Policies

- [Security Policies for Self-Encrypting Drives, on page 153](#)
- [Security Flags of the Controller and Disk, on page 154](#)
- [Managing Local Security Policies, on page 154](#)
- [KMIP Client Certificate Policy, on page 159](#)
- [Managing Remote Security Policies, on page 162](#)
- [Securing an Existing Virtual Drive, on page 166](#)
- [Enabling Security on a Disk, on page 167](#)
- [Erasing a Secure Disk, on page 168](#)
- [Disabling Security on a Controller, on page 169](#)
- [Unlocking a Locked Disk , on page 169](#)
- [Erasing a Secure Foreign Configuration Disk, on page 170](#)
- [Displaying the Security Flags of a Controller , on page 171](#)
- [Displaying the Security Flags of a Local Disk , on page 172](#)
- [Displaying the Security Flags of a Virtual Drive , on page 173](#)

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 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.

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org	Enters the root organization mode.
Step 2	UCS-A /org # create storage-profile <i>storage-profile-name</i>	Creates a storage profile with the specified name at the organization level and enters the storage-profile configuration mode.
Step 3	UCS-A /org/storage-profile* # create security	Creates a security policy for the specified storage profile and enters the security policy mode.
Step 4	UCS-A /org/storage-profile/security* # create drive-security	Creates a drive security policy for the specified storage profile security and enters the drive security policy mode.

	Command or Action	Purpose
Step 5	UCS-A /org/storage-profile/security/drive-security* # create local	Creates a local security policy for the specified storage profile and enters the local policy mode.
Step 6	UCS-A /org/storage-profile/security/drive-security/local* # set security-key security-key	Sets the specified security key for the local policy. The security key must have 32 characters.
Step 7	UCS-A /org/storage-profile/security/drive-security/local* # commit-buffer	Commits the transaction to the system configuration.

Example

This example show how to create a local security policy with a security key:

```
UCS-A# scope org
UCS-A /org # create storage-profile stp-demo
UCS-A /org/storage-profile* # create security
UCS-A /org/storage-profile/security* # create drive-security
UCS-A /org/storage-profile/security/drive-security* # create local
UCS-A /org/storage-profile/security/drive-security/local* # set security-key
thereare32charactersinthisseckey
UCS-A /org/storage-profile/security/drive-security/local* # commit-buffer
UCS-A /org/storage-profile/security/drive-security/local #
```

Modifying the Security Key of a Local Security Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org	Enters the root organization mode.
Step 2	UCS-A /org # scope storage-profile <i>storage-profile-name</i>	Enters the storage-profile configuration mode for the specified storage profile.
Step 3	UCS-A /org/storage-profile # scope security	Enters the security policy mode for the specified storage profile.
Step 4	UCS-A /org/storage-profile/security # scope drive-security	Enters the drive security policy mode for the specified storage profile security.
Step 5	UCS-A /org/storage-profile/security/drive-security # scope local	Enters the local policy mode for the the specified storage profile.

	Command or Action	Purpose
Step 6	UCS-A /org/storage-profile/security/drive-security/local # set deployed-security-key <i>existing-security-key</i>	Specifies the existing key deployed on the server to configure a new key.
Step 7	UCS-A /org/storage-profile/security/drive-security/local* # set security-key <i>new-security-key</i>	Sets the new security key for the local policy.
Step 8	UCS-A /org/storage-profile/security/drive-security/local* # commit-buffer	Commits the transaction to the system configuration

Example

This example shows how to modify the security key of a local security policy:

```
UCS-A# scope org
UCS-A /org # scope storage-profile stp-demo
UCS-A /org/storage-profile # scope security
UCS-A /org/storage-profile/security # scope drive-security
UCS-A /org/storage-profile/security/drive-security # scope local
UCS-A /org/storage-profile/security/drive-security/local # set deployed-security-key
thereare32charactersinthisseckey
UCS-A /org/storage-profile/security/drive-security/local* # set security-key
thereare32charactersinthisnewkey
UCS-A /org/storage-profile/security/drive-security/local* # commit-buffer
UCS-A /org/storage-profile/security/drive-security/local #
```

Modifying the Security Policy from Local to Remote

Before you begin

Ensure that you have created a KMIP client certificate policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org	Enters the root organization mode.
Step 2	UCS-A /org # scope storage-profile <i>storage-profile-name</i>	Enters the storage-profile configuration mode for the selected storage profile.
Step 3	UCS-A /org/storage-profile # scope security	Enters the security policy mode for the specified storage profile.
Step 4	UCS-A /org/storage-profile/security # scope drive-security	Enters the drive security policy mode for the specified storage profile security.

	Command or Action	Purpose
Step 5	UCS-A /org/storage-profile/security/drive-security # create remote	Creates and enters the remote policy mode.
Step 6	UCS-A /org/storage-profile/security/drive-security/remote* # set deployed-security-key <i>existing-security-key</i>	Specifies the existing key deployed on the server.
Step 7	UCS-A /org/storage-profile/security/drive-security/remote* # set primary-server <i>primary-server-name</i>	Sets the primary server hostname or IP address.
Step 8	(Optional) UCS-A /org/storage-profile/security/drive-security/remote* # set secondary-server <i>secondary-server-name</i>	Sets the secondary server hostname or IP address.
Step 9	(Optional) UCS-A /org/storage-profile/security/drive-security/remote* # set port <i>kmip-server-port-number</i>	Sets the port number of the KMIP server. KMIP server port numbers can range from 1024 to 65535.
Step 10	UCS-A /org/storage-profile/security/drive-security/remote* # set server-certificate	Sets the KMIP certificate to the remote security policy.
Step 11	(Optional) UCS-A /org/storage-profile/security/drive-security/remote* # set timeout <i>timeout-seconds</i>	Sets the number of seconds in which communication between the storage and the KMIP server times out. Timeout can range from 5 seconds to 20 seconds.
Step 12	UCS-A /org/storage-profile/security/drive-security/remote* # commit-buffer	Commits the transaction to the system configuration.
Step 13	UCS-A /org/storage-profile/security/drive-security/remote # exit	Enters the drive security policy mode.
Step 14	UCS-A /org/storage-profile/security/drive-security # delete local	Deletes the existing local security policy.
Step 15	UCS-A /org/storage-profile/security/drive-security* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to modify a security policy from local to remote:

```
UCS-A # scope org
UCS-A /org # scope storage-profile stp-demo
```




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

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope security	Enters the security mode.
Step 2	UCS-A /security # create kmip-client-cert-policy	Creates the KMIP certificate policy and enters the KMIP client certificate policy mode.
Step 3	UCS-A /security/kmip-client-cert-policy* # set country <i>country-code</i>	Specifies the country code for the KMIP certificate policy. The country code must contain 2 letters in upper case.
Step 4	UCS-A /security/kmip-client-cert-policy* # set locality <i>locality-code</i>	Specifies the name of the locality or city for the KMIP certificate policy. Enter up to 32 characters for the locality name.
Step 5	UCS-A /security/kmip-client-cert-policy* # set org-name <i>org-name</i>	Specifies the organization name requesting the KMIP certificate policy. Enter up to 32 characters for the organization name.

	Command or Action	Purpose
Step 6	UCS-A /security/kmip-client-cert-policy* # set org-unit-name <i>unit-name</i>	Specifies the organizational unit name requesting the KMIP certificate policy. Enter up to 64 characters for the organizational unit name.
Step 7	UCS-A /security/kmip-client-cert-policy* # set state <i>state-code</i>	Specifies the name of the state, province, or county for the KMIP certificate policy. Enter up to 32 characters for the state name.
Step 8	(Optional) UCS-A /security/kmip-client-cert-policy* # set email <i>email-address</i>	Specifies the email address associated with the request.
Step 9	(Optional) UCS-A /security/kmip-client-cert-policy* # set validity <i>days</i>	Specifies the validity of the certificate in number of days. The validity can range between 365 days and 3650 days.
Step 10	UCS-A /security/kmip-client-cert-policy* # commit-buffer	Commits the transaction to the system configuration.
Step 11	UCS-A /security/kmip-client-cert-policy # show	Displays details of the KMIP certificate policy.

Example

This example shows how to create a KMIP certificate policy.

```
UCS-A# scope security
UCS-A /security # create kmip-client-cert-policy
UCS-A /security/kmip-client-cert-policy* # set country IN
UCS-A /security/kmip-client-cert-policy* # set locality BLR
UCS-A /security/kmip-client-cert-policy* # set org-name XYZ
UCS-A /security/kmip-client-cert-policy* # set org-unit-name Ops
UCS-A /security/kmip-client-cert-policy* # set state KA
UCS-A /security/kmip-client-cert-policy* # commit-buffer
UCS-A /security/kmip-client-cert-policy # show
```

```
KMIP Client certificate policy:
Certificate request country name: IN
State, province or county (full name): KA
Locality name (eg, city): BLR
Organisation name (eg, company): XYZ
Organisational Unit Name (eg, section): Ops
Certificate request e-mail name:
Validity of certificate in number of days: 1095
UCS-A /security/kmip-client-cert-policy #
```

Creating a KMIP Client Certificate 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

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-number</i>	Enters the server configuration mode for the specified server.
Step 2	UCS-A /server # create kmip-client-cert-policy	Creates the KMIP certificate policy and enters the KMIP client certificate policy mode.
Step 3	UCS-A /server/kmip-client-cert-policy* # set country <i>country-code</i>	Specifies the country code for the KMIP certificate policy. The country code must contain 2 letters in upper case.
Step 4	UCS-A /server/kmip-client-cert-policy* # set locality <i>locality-code</i>	Specifies the name of the locality or city for the KMIP certificate policy. Enter up to 32 characters for the locality name.
Step 5	UCS-A /server/kmip-client-cert-policy* # set org-name <i>org-name</i>	Specifies the organization name requesting the KMIP certificate policy. Enter up to 32 characters for the organization name.
Step 6	UCS-A /server/kmip-client-cert-policy* # set org-unit-name <i>unit-name</i>	Specifies the organizational unit name requesting the KMIP certificate policy. Enter up to 64 characters for the organizational unit name.
Step 7	UCS-A /server/kmip-client-cert-policy* # set state <i>state-code</i>	Specifies the name of the state, province, or county for the KMIP certificate policy. Enter up to 32 characters for the state name.
Step 8	(Optional) UCS-A /server/kmip-client-cert-policy* # set email <i>email-address</i>	Specifies the email address associated with the request.
Step 9	(Optional) UCS-A /server/kmip-client-cert-policy* # set validity <i>days</i>	Specifies the validity of the certificate in number of days. The validity can range between 365 days and 3650 days.
Step 10	UCS-A /server/kmip-client-cert-policy* # commit-buffer	Commits the transaction to the system configuration.
Step 11	UCS-A /server/kmip-client-cert-policy # show	Displays details of the KMIP certificate.

Example

This example shows how to create a KMIP certificate on a server.

```
UCS-A# scope server 5
UCS-A /server # create kmip-client-cert-policy
UCS-A /server/kmip-client-cert-policy* # set country IN
UCS-A /server/kmip-client-cert-policy* # set locality BLR
```

```
UCS-A /server/kmip-client-cert-policy* # set org-name XYZ
UCS-A /server/kmip-client-cert-policy* # set org-unit-name Ops
UCS-A /server/kmip-client-cert-policy* # set state KA
UCS-A /server/kmip-client-cert-policy* # commit-buffer
UCS-A /server/kmip-client-cert-policy* # show
```

```
KMIP Client certificate policy:
Certificate request country name: IN
State, province or county (full name): KA
Locality name (eg, city): BLR
Organisation name (eg, company): XYZ
Organisational Unit Name (eg, section): Ops
Certificate request e-mail name:
Validity of certificate in number of days: 1095
UCS-A /server/kmip-client-cert-policy #
```

Managing Remote Security Policies

Creating a Remote Security Policy

Before you begin

Ensure that you have created a KMIP client certificate policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org	Enters the root organization mode.
Step 2	UCS-A /org # scope storage-profile <i>storage-profile-name</i>	Enters the storage-profile configuration mode for the selected storage profile.
Step 3	UCS-A /org/storage-profile # create security	Creates and enters the security mode.
Step 4	UCS-A /org/storage-profile/security* # create drive-security	Creates and enters the drive-security mode.
Step 5	UCS-A /org/storage-profile/security/drive-security* # create remote	Creates and enters the remote policy mode.
Step 6	UCS-A /org/storage-profile/security/drive-security/remote* # set primary-server <i>primary-server-name</i>	Sets the primary server hostname or IP address.
Step 7	(Optional) UCS-A /org/storage-profile/security/drive-security/remote* # set secondary-server <i>secondary-server-name</i>	Sets the secondary server hostname or IP address.

	Command or Action	Purpose
Step 8	(Optional) UCS-A /org/storage-profile/security/drive-security/remote* # set port <i>kmip-server-port-number</i>	Sets the port number of the KMIP server. KMIP server port numbers can range from 1024 to 65535.
Step 9	UCS-A /org/storage-profile/security/drive-security/remote* # set server-certificate	Sets the KMIP certificate to the remote security policy.
Step 10	(Optional) UCS-A /org/storage-profile/security/drive-security/remote* # set timeout <i>timeout-seconds</i>	Sets the number of seconds in which communication between the storage and the KMIP server times out. Timeout can range from 5 seconds to 20 seconds.
Step 11	(Optional) UCS-A /org/storage-profile/security/drive-security/remote* # create login	Creates the login details for the KMIP server and enters the login mode.
Step 12	(Optional) UCS-A /org/storage-profile/security/drive-security/remote/login* # set username <i>username</i>	Sets the username to log into the KMIP server.
Step 13	(Optional) UCS-A /org/storage-profile/security/drive-security/remote/login* # set password <i>password</i>	Sets the password to log into the KMIP server.
Step 14	UCS-A /org/storage-profile/security/drive-security/remote/login* # commit-buffer	Commits the transaction to the system configuration.

Example

```
UCS-A # scope org
UCS-A /org # scope storage-profile stp-demo
UCS-A /org/storage-profile # create security
UCS-A /org/storage-profile/security* # create drive-security
UCS-A /org/storage-profile/security/drive-security* # create remote
UCS-A /org/storage-profile/security/drive-security/remote* # set primary-server 10.10.10.1
UCS-A /org/storage-profile/security/drive-security/remote* # set secondary-server 10.10.10.2
UCS-A /org/storage-profile/security/drive-security/remote* # set port 5696
UCS-A /org/storage-profile/security/drive-security/remote* # set server-certificate
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Prompt Certificate:
>-----BEGIN CERTIFICATE-----
MIIEEDCCAvigAwIBAgIGALofZVDsMA0GCSqGSIb3DQEBCwUAMIGQMSowKAYDVQQD
EyFDRyBDQSBTIG9uIHZvcmlldHJpY2RzbS5jaXNjby5jb20xFTATBgNVBAsTDFNh
dmJlU3RvcmlldjEwMBQGA1UEChMNQ21lZ28uU31zdGVtczERMA8GA1UEBxMIU2Fu
IEpvc2UxEzARBgNVBAgTCkNhbgG1mb3JuaWEeXzAxBG9NVBAYTA1VTMB4XDTE2MDkw
NzE5MzZmMVoXDTI2MDkwOTE5MzZmMVoVowgZAxKjAoBgNVBAMTlUNHIENBIkFmgb24g
dm9ybWV0cmlldjZHNtLmNpc2NvLmNvbTEVMBMGALUECMMU2F2YnVtdG9yZGV2MRYw
FAYDVQQKEw1DaXNjbyBTeXN0ZW1zMREwDwYDVQQHEWhTYW4gSm9zZTETMB>EGA1UE
CBMKQ2FsaWZvcmlldjEwMBQGA1UEBhMCMVVMwggEiMA0GCSqGSIb3DQEBAQUAA4IB
DwAwggEKAoIBAQDhX2UdIV>TQTchGo1FjAc5u1W9zAo/YkjD22ANpbEPiAmgWL97c
Xwj7yzArflrZ2kVwQcm4f6AdLOFUWzbuo+Fxd3rurd>w6BhJXdLj8Piq8094PqCLp
qdUF83SsRVVbCXHxOqdk9jsSQRvTcV4PloNrelMLq/mOqsaODs+us4ng7sMDTGXv
LeKFC8DUEm0G1GQACwiJ3s9O4+P2CI/d4P/>EyWwqABf3YJmAI1EQYUnoTwrg6EgY
```

```

ZvcPmHsmjXnbBzrL+ON7FBcbrTanvjyJxE6tFf5cRPghymfna7Fd3lfVwZCcGIoR+
> EOIAwgetzIRM6FzMiV2/tDT8STo/oo5Tg3dDagMBAAGjbj>BsMBIGA1UdEwEB/wQI
MAYBAf8CAQAwDgYDVR0PAQH/>>>>>>>>>>BAQDAgEGMB0GA1UdDgQWBBrnYyFiAK2lEDZJNC0Y
VlIqMgiUJdAnBgNVHSMEIDAegBRnYyFiAK2lEDZJNC0YVlIqMgiUJIIGALofZVds
MA0GCSqGSIB3DQEBChwAA4IBAQAfhB2+Ft8V2ELAFa7PcG/rU09ux7LYcCjt3STa
mzKdZ7Rn5COvknKrJX+EefT7xl03CQXT9aesAddQUOCy8fhiPoaMFr1Tgs1hdS0p
NJvfxV6QCun2UMRSuxWfG>0QFfofnXeIGkAmEYOpUdArSOTbtt4v6LjalA+KESvWW
5KaVemo2nsd+iD0IPCOhpShAgaAwpnYUq9mLfVgvV07Z+hmkuOIQTZ2+h+pJQtE0
+U5qaTts4pMXpqQPjlid0NMuaPug1SpSD7KBSjwR1SzehzPdnl6uprmvWa3VBk3
OK6y55FoIu+Wg9i/8kmfkghyGwTfo6weEKbleuVwupvpriMF>
-----END CERTIFICATE-----

UCS-A /org/storage-profile/security/drive-security/remote* # create login
UCS-A /org/storage-profile/security/drive-security/remote/login* # set username user1
UCS-A /org/storage-profile/security/drive-security/remote/login* # set password Password
UCS-A /org/storage-profile/security/drive-security/remote/login* # exit
UCS-A /org/storage-profile/security/drive-security/remote # exit
UCS-A /org/storage-profile/security/drive-security # show detail expand

```

Drive Security:

Remote:

```

Primary Server Name: 10.10.10.1
Secondary Server Name:10.10.10.2
KMIP Server Port: 5696
Deployed Security Key:

```

KMIP Server Certificate: -----BEGIN CERTIFICATE-----

```

MIIEEDCCAvigAwIBAgIGALofZVdsMA0GCSqGSIB3DQEBChwAA4IBAQAfhB2+Ft8V2ELAFa7PcG/rU09ux7LYcCjt3STa
mzKdZ7Rn5COvknKrJX+EefT7xl03CQXT9aesAddQUOCy8fhiPoaMFr1Tgs1hdS0p
NJvfxV6QCun2UMRSuxWfG0QFfofnXeIGkAmEYOpUdArSOTbtt4v6LjalA+KESvWW
5KaVemo2nsd+iD0IPCOhpShAgaAwpnYUq9mLfVgvV07Z+hmkuOIQTZ2+h+pJQtE0
+U5qaTts4pMXpqQPjlid0NMuaPug1SpSD7KBSjwR1SzehzPdnl6uprmvWa3VBk3
OK6y55FoIu+Wg9i/8kmfkghyGwTfo6weEKbleuVwupvpriMF
-----END CERTIFICATE-----

```

Modifying a Remote Security Key

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters the server mode for the specified server.

	Command or Action	Purpose
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # set admin-state modify-remote-key	Modifies the security key of a remote security policy.
Step 4	UCS-A /server/raid-controller # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to modify the remote security key on the controller:

```
UCS-A# scope server 3
UCS-A /server # scope raid-controller 1 sas
UCS-A /server/raid-controller # set admin-state modify-remote-key
UCS-A /server/raid-controller* # commit-buffer
UCS-A /server/raid-controller #
```

Modifying the Security Policy from Remote to Local

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org	Enters the root organization mode.
Step 2	UCS-A /org # scope storage-profile <i>storage-profile-name</i>	Enters the storage-profile configuration mode for the specified storage profile.
Step 3	UCS-A /org/storage-profile # scope security	Enters the security policy mode for the specified storage profile.
Step 4	UCS-A /org/storage-profile/security # scope drive-security	Enters the drive security policy mode for the specified storage profile security.
Step 5	UCS-A /org/storage-profile/security/drive-security # delete remote	Deletes the existing remote security policy.
Step 6	UCS-A /org/storage-profile/security/drive-security* # commit-buffer	Commits the transaction to the system configuration.
Step 7	UCS-A /org/storage-profile/security/drive-security # create local	Creates and enters the local policy mode.

	Command or Action	Purpose
Step 8	UCS-A /org/storage-profile/security/drive-security/local* # set security-key security-key	Sets the security key for the local policy.
Step 9	UCS-A /org/storage-profile/security/drive-security/local* # commit-buffer	Commits the transaction to the system configuration
Step 10		

Example

This example shows how to modify a security policy from remote to local:

```
UCS-A# scope org
UCS-A /org # scope storage-profile stp-demo
UCS-A /org/storage-profile # scope security
UCS-A /org/storage-profile/security # scope drive-security
UCS-A /org/storage-profile/security/drive-security # delete remote
UCS-A /org/storage-profile/security/drive-security* # commit-buffer
UCS-A /org/storage-profile/security/drive-security # create local
UCS-A /org/storage-profile/security/drive-security/local* # set security-key
thereare32charactersinthisseckey
UCS-A /org/storage-profile/security/drive-security/local* # commit-buffer
UCS-A /org/storage-profile/security/drive-security/local #
```

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.

Securing an Existing Virtual Drive

Before you begin

- The controller must be secure.
- The virtual drive must be in the **Orphaned** state.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server# scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller# scope virtual-drive <i>virtual-drive-id</i>	Enters the virtual drive mode for the specified orphaned virtual drive.
Step 4	UCS-A /server/raid-controller/virtual-drive# set admin-state secure-drive-group	Secures the existing virtual drive.
Step 5	UCS-A /server/raid-controller/virtual-drive*# commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to secure an existing virtual drive:

```
UCS-A# scope server 1
UCS-A /server# scope raid-controller 3 sas
UCS-A /server/raid-controller# scope virtual-drive 1000
UCS-A /server/raid-controller/virtual-drive # set admin-state secure-drive-group
UCS-A /server/raid-controller/virtual-drive*# commit-buffer
UCS-A /server/raid-controller/virtual-drive#
```

Enabling Security on a Disk

Before you begin

Ensure that the disk is a JBOD.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # scope local-disk <i>local-disk-id</i>	Enters the local disk configuration mode

	Command or Action	Purpose
Step 4	UCS-A /server/raid-controller/local-disk # set admin-state enable-security	Enables security on a JBOD.
Step 5	UCS-A /server/raid-controller/local-disk* # commit-buffer	Commits the transaction to the system configuration.

Example

The example shows how to enable security on a JBOD:

```
UCS-A# scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # scope local-disk 2
UCS-A /server/raid-controller/local-disk # set admin-state enable-security
UCS-A /server/raid-controller/local-disk* # commit-buffer
UCS-A /server/raid-controller/local-disk #
```

Erasing a Secure Disk

Before you begin

Ensure that the disk is in the **Unconfigured Good** state.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id</i> {SAS SAT}	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # scope local-disk <i>local-disk-id</i>	Enters the local disk configuration mode.
Step 4	UCS-A /server/raid-controller/local-disk # set admin-state clear secure-drive	Erases the secured disk and clears the security on the disk.
Step 5	UCS-A /server/raid-controller/local-disk* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to erase a secure disk:

```
UCS-A # scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # scope local-disk 2
UCS-A /server/raid-controller/local-disk # set admin-state clear secure-drive
UCS-A /server/raid-controller/local-disk* # commit-buffer
UCS-A /server/raid-controller/local-disk #
```

Disabling Security on a Controller

Before you begin

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # set admin-state disable-security	Disables security key on the controller.
Step 4	UCS-A /server/raid-controller # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to disable security on the controller:

```
UCS-A# scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # set admin-state disable-security
UCS-A /server/raid-controller* # commit-buffer
UCS-A /server/raid-controller #
```

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.

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id</i> { <i>SAS</i> <i>SAT</i> }	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SEDs only on SAS controllers.
Step 3	UCS-A /server/raid-controller # set admin-state unlock-disk [<i>security-key</i>]	<p>Unlocks the locked disks.</p> <p>If the security-key is set, this key is used to unlock disks that are in the locked state.</p> <p>If the security-key is not set, Cisco UCS Manager tries to unlock the disks by using the KMIP server. Setting the security-key is optional only if remote security is configured on the server.</p>
Step 4	UCS-A /server/raid-controller* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to unlock a locked disk with a local security policy by using a security-key:

```
UCS-A # scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # set admin-state unlock-disk thisisastring
UCS-A /server/raid-controller* # commit-buffer
UCS-A /server/raid-controller #
```

This example shows how to unlock a locked disk with a remote security policy by using the KMIP server:

```
UCS-A # scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # set admin-state unlock-disk
UCS-A /server/raid-controller* # commit-buffer
UCS-A /server/raid-controller #
```

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

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # scope local-disk <i>local-disk-id</i>	Enters the local disk configuration mode.
Step 4	UCS-A /server/raid-controller/local-disk # set admin-state clear secure-foreign-config-drive	Clears the secure foreign configuration drive.
Step 5	UCS-A /server/raid-controller/local-disk* # commit-buffer	Commits the transaction to the system configuration.

Example

This example shows how to clear a foreign configuration disk:

```
UCS-A# scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # scope local-disk 2
UCS-A /server/raid-controller/local-disk # set admin-state clear secure-foreign-config-drive

UCS-A /server/raid-controller/local-disk* # commit-buffer
UCS-A /server/raid-controller/local-disk #
```

Displaying the Security Flags of a Controller

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # show detail	Displays details of the RAID controller.

Example

This example shows to how to check if the security flag of controller is enabled:

```
UCS-A # scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # show detail

RAID Controller:
  ID: 3
  Type: SAS
  PCI Addr: 03:00.0
  Vendor: LSI Corp.
  Model: LSI MegaRAID SAS 3108
  Serial: SV55346948
  HW Rev: C0
  Raid Support: RAID0, RAID1, RAID5, RAID6, RAID10, RAID50, RAID60
  OOB Interface Supported: Yes
  Mode: RAID
  Rebuild Rate: 30
  Controller Status: Optimal
  Config State: Applied
  Pinned Cache Status: Disabled
  Sub OEM ID: 0
  Supported Strip Sizes: 1MB,64KB,256KB,512KB,128KB
  Default Strip Size: 64KB
  PCI Slot: HBA
  Controller Flags: Drive Security Capable
```

This example shows to how to check if the security flag of controller is enabled on a blade server:

Displaying the Security Flags of a Local Disk

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # scope local-disk <i>local-disk-id</i>	Enters the local disk configuration mode.
Step 4	UCS-A /server/raid-controller/local-disk # show detail	Displays details of the local disk.

Example

This example shows to how to display the security flag of a local disk:

```
UCS-A # scope server 1
UCS-A /server # scope raid-controller 3 sas
```

```

UCS-A /server/raid-controller #scope local-disk 2
UCS-A /server/raid-controller/local-disk # show detail

Local Disk:
  ID: 4
  Block Size: 512
  Physical Block Size: 4096
  Blocks: 1560545280
  Raw Size: 763097
  Size: 761985
  Technology: SSD
  Operability: Operable
  Oper Qualifier Reason: N/A
  Presence: Equipped
  Connection Protocol: SAS
  Product Variant: default
  Product Name: 800GB Enterprise performance SAS SED SSD (10 FWPD) - MTFDJAK800MBS
  PID: UCS-SD800GBEK9
  VID: V01
  Vendor: MICRON
  Model: S650DC-800FIPS
  Vendor Description: Micron
  Serial: ZAZ090VD0000822150Z3
  HW Rev: 0
  Running-Vers: MB13
  Average Seek Time (R/W): N/A
  Track to Track Seek Time (R/W): 115ms
  Part Number: 16-100911-01
  SKU: UCS-SD800GBEK9
  Drive State: Online
  Power State: Active
  Link Speed: 12 Gbps
  Enclosure Association Type: Direct Attached
  Device Version: MB13
  Drive Security Flags: Secured,Security Enabled,Security Capable

```

Displaying the Security Flags of a Virtual Drive

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server-id</i>	Enters the server mode for the specified server.
Step 2	UCS-A /server # scope raid-controller <i>raid-controller-id {SAS SAT}</i>	Enters the RAID controller mode. Currently, Cisco UCS Manager supports SED only on SAS controllers.
Step 3	UCS-A /server/raid-controller # scope virtual-drive <i>virtual-drive-id</i>	Enters the virtual drive mode.
Step 4	UCS-A /server/raid-controller/virtual-drive # show detail	Displays details of the virtual drive.

Example

This example shows to how to display the security flag of a virtual disk:

```
UCS-A # scope server 1
UCS-A /server # scope raid-controller 3 sas
UCS-A /server/raid-controller # scope virtual-drive 1000
UCS-A /server/raid-controller/virtual-drive # show detail

Virtual Drive:
  ID: 1000
  Name: luna
  Block Size: 512
  Blocks: 20971520
  Size: 10240
  Operability: Operable
  Presence: Equipped
  Lifecycle: Allocated
  Drive State: Optimal
  Type: RAID 0 Striped
  Strip Size (KB): 64
  Access Policy: Read Write
  Read Policy: Normal
  Configured Write Cache Policy: Write Through
  Actual Write Cache Policy: Write Through
  IO Policy: Direct
  Drive Cache: No Change
  Bootable: False
  Oper Device ID: 0
  Change Qualifier: No Change
  Config State: Applied
  Deploy Action: No Action
  Service Profile Lun Reference: org-root/ls-spl/vdrive-ref-lun-1
  Assigned To Server: sys/rack-unit-1
  Available Size on Disk Group (MB): 751745
  Unique Identifier: 90ae6ea0-6a39-49e1-9c0d-0f3e2e9ecfce
  Vendor Unique Identifier: 678da6e7-15b2-9c20-2011-c4f60c40e57a
  Security Flags: Drive Security Enable,Drive Security Capable
```



CHAPTER 14

Cisco UCS S3260 System Storage Management

- [Storage Server Features and Components Overview](#), on page 175
- [Cisco UCS S3260 Storage Management Operations](#), on page 183
- [Disk Sharing for High Availability](#), on page 184
- [Storage Enclosure Operations](#), on page 189
- [SAS Expander Configuration Policy](#), on page 190

Storage Server Features and Components Overview

Storage Server Features

The following table summarizes the Cisco UCS S3260 system features:

Table 2: 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.

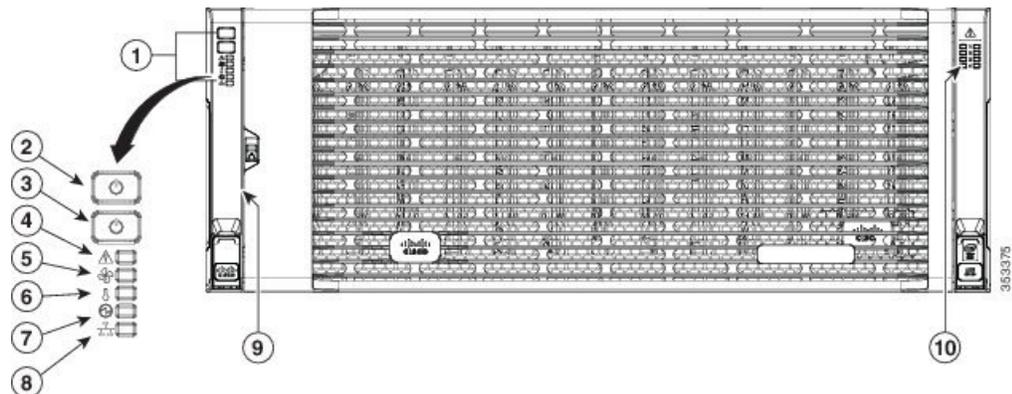
Feature	Description
Storage	<p>The system has the following storage options:</p> <ul style="list-style-type: none"> • Up to 56 top-loading 3.5-inch drives • Up to four 3.5-inch, rear-loading drives in the optional drive expander module • Up to four 2.5-inch, rear-loading SAS solid state drives (SSDs) • 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 (SIOC). 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 4: Front Panel Features



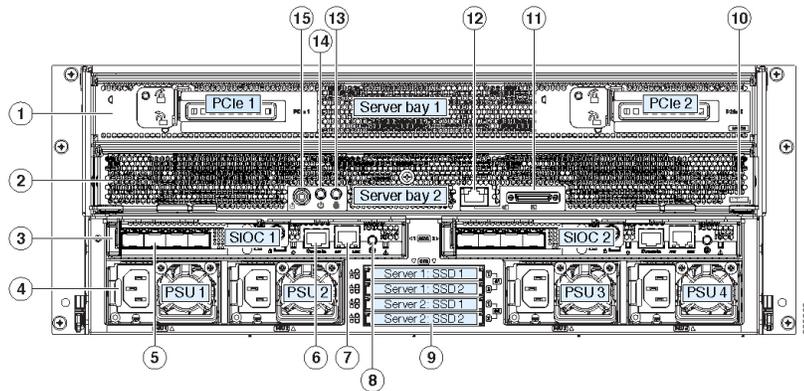
1	Operations panel	6	Temperature status LED
2	System Power button/LED	7	Power supply status LED
3	System unit identification button/LED	8	Network link activity LED

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 5: 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. 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	<p>Chassis Management Controller (CMS) Debug Firmware Utility port (one each SIOC)</p>	13	<p>Server node unit identification button/LED</p>

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

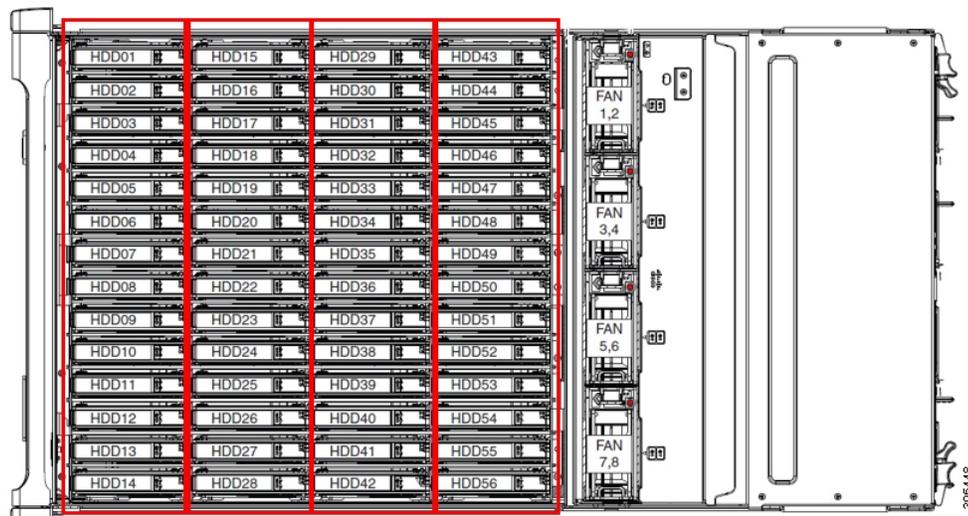
The Cisco UCS S3260 system consists of one or two server nodes, each with two CPUs, DIMM memory of 128, 256, or 512 GB, and a RAID card up to 4 GB cache or a pass-through controller. The server nodes can be one of the following:

- 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

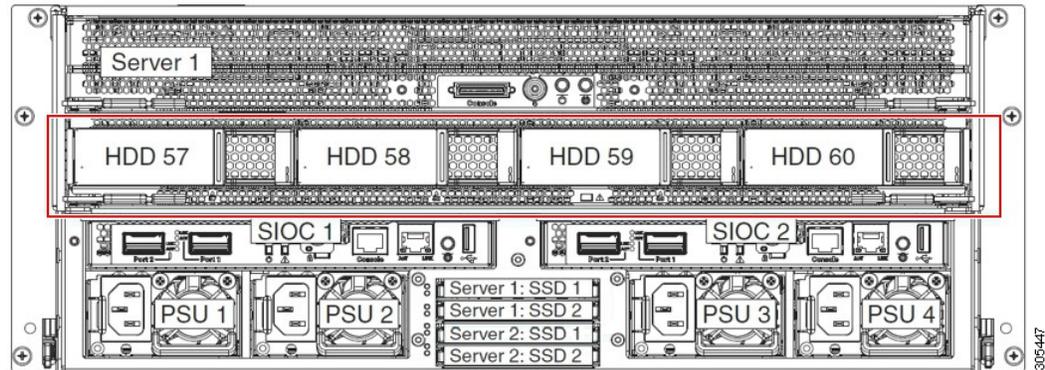
The Cisco UCS S3260 chassis has 4 rows of 14 disk slots on the HDD motherboard and 4 additional disk slots on the HDD expansion tray. The following image shows the disk arrangement for the 56 top-accessible, hot swappable 3.5-inch 6 TB or 4 TB 7200 rpm NL-SAS HDD drives. A disk slot has two SAS ports and each is connected a SAS expander in the chassis.

Figure 6: 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 7: 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 184](#)

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 3:

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 3.2</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**, **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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A org/ # create disk-zoning-policy <i>diskzoning policy-name</i>	Creates a disk zoning policy name with the specified disk zoning policy name.
Step 3	UCS-A /org/disk-zoning-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates the dzp1 disk zoning policy:

```
UCS-A# scope org
UCS-A /org # create disk-zoning-policy dzp1
UCS-A /org/disk-zoning-policy*# commit-buffer
UCS-A /org/disk-zoning-policy#
```

Creating Disk Slots and Assigning Ownership

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A org/ # disk-zoning-policy <i>disk-zoning-policy-name</i>	Enters the disk zoning policy.
Step 3	UCS-A org/disk-zoning-policy # create disk-slot <i>slot-id</i>	Creates disk slot with the specified slot number.
Step 4	UCS-A org/disk-zoning-policy/disk-slot* # set ownership <i>ownership-type</i> <i>{chassis-global-hot-spare dedicated shared unassigned}</i>	Specifies the disk ownership to be one of the following: <ul style="list-style-type: none"> • chassis-global-hot-spare—Chassis Global Hot Spare • dedicated—Dedicated Beginning with release 3.2(3a), Cisco UCS Manager can enable single path access to disk by configuring single DiskPort per

	Command or Action	Purpose
		<p>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. <p>Use the following command to set the drivepath:</p> <pre>set drivepath drivepath{path-0/path-1/path-both}</pre> <ul style="list-style-type: none"> • shared—Shared <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 4: Limitations for Shared Mode for Dual HBA Controller, on page 186.</p> <ul style="list-style-type: none"> • unassigned—Unassigned
Step 5	UCS-A org/disk-zoning-policy/disk-slot* # create controller-ref server-id sas controller-id	Creates controller reference for the specified server slot.
Step 6	UCS-A org/disk-zoning-policy/disk-slot # commit-buffer	Commits the transaction.

Table 4: 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

Example

The following example creates disk slot 1, sets the ownership as shared, creates a controller reference for the server slot 1, and commits the transaction:

```
UCS-A# scope org
UCS-A /org # scope disk-zoning-policy test
UCS-A /org/disk-zoning-policy* # create disk-slot 1
UCS-A /org/disk-zoning-policy/disk-slot* # set ownership shared
UCS-A /org/disk-zoning-policy/disk-slot* # create controller-ref 1 sas 1
UCS-A /org/disk-zoning-policy/disk-slot* # create controller-ref 2 sas 1
UCS-A /org/disk-zoning-policy/disk-slot* #commit-buffer
UCS-A /org/disk-zoning-policy/disk-slot #
```

Associating Disk Zoning Policies to Chassis Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A org/ # create chassis-profile <i>chassis-profile-name</i>	Creates a chassis profile with the specified name.
Step 3	UCS-A org/chassis-profile* # set disk-zoning-policy <i>disk-zoning-policy</i>	Sets the specified disk-zoning-policy.
Step 4	UCS-A org/chassis-profile* # commit-buffer	Commits the transaction.
Step 5	UCS-A org/chassis-profile# associate chassis <i>chassis-id</i>	Associates the disks in the disk zoning policy to the chassis with the specified chassis number.

Example

The following example creates the ch1 chassis profile, sets the disk zoning policy all56shared, commits the transaction and associates the disk in the all56shared policy with chassis 3:

```
UCS-A# scope org
UCS-A /org # create chassis-profile ch1
UCS-A /org/chassis-profile* # set disk-zoning-policy all56shared
UCS-A /org/chassis-profile* # commit-buffer
UCS-A /org/chassis-profile # associate chassis 3
UCS-A /org/fw-chassis-pack/pack-image #
```

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 3.2](#).

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis# scope virtual-drive-container <i>virtual-drive-container-num</i>	Enters the virtual drive container with the specified number.
Step 3	UCS-A /chassis/virtual-drive-container# scope virtual-drive <i>virtual-drive--num</i>	Enters the virtual drive for the specified virtual drive container.
Step 4	UCS-A /chassis/virtual-drive-container/virtual-drive# scope virtual-drive <i>virtual-drive--num</i> set admin-state <i>admin-state</i>	<p>Specifies one of the following admin states for the virtual drive:</p> <ul style="list-style-type: none"> • clear-transport-ready — Sets the state of the virtual drive to no longer be transport ready. • delete — Deletes the virtual drive. • hide— Choose this option for the safe migration of the virtual drive from one server to another. <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> <ul style="list-style-type: none"> • transport-ready — Choose this option for the safe migration of the virtual drive from one server to another. <p>Note When a virtual drive is marked as transport ready, the storage controller will disable all IO operations on the drive. In addition, after zoning the virtual drive and importing the foreign configuration, the virtual drive will be operational.</p>

	Command or Action	Purpose
Step 5	UCS-A /chassis/virtual-drive-container/virtual-drive# commit-buffer	Commits the transaction to the system configuration.

Example

The following example sets the state of the virtual drive 1001 in the virtual drive container 1 to transport ready:

```
UCS-A# scope chassis
UCS-A /chassis# scope virtual-drive-container 1
UCS-A /chassis/virtual-drive-container# scope virtual-drive 1001
UCS-A /chassis/virtual-drive-container/virtual-drive# set admin-state transport-ready
UCS-A /chassis/virtual-drive-container/virtual-drive# commit-buffer
```

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

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-id</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # remove storage-enclosure <i>storage-enclosure-name</i>	Removes the chassis level storage enclosure with the specified name.

Example

The following example removes storage enclosure 25 from chassis 2:

```
UCS-A# scope chassis 2
UCS-A /chassis# remove storage-enclosure 25
UCS-A /chassis#
```

SAS Expander Configuration Policy

Creating SAS Expander Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A org/ # create sas-expander-configuration-policy <i>sas-expander-configuration-policy-name</i>	Creates a SAS expander configuration policy with the specified policy name.
Step 3	(Optional) UCS-A /org/sas-expander-configuration-policy* # set descr <i>description</i>	Provides a description for the policy.
Step 4	(Optional) UCS-A /org/sas-expander-configuration-policy* # set 6g-12g-mixed-mode <i>disabled enabled no-change</i>	<p>Note Enabling or disabling 6G-12G Mixed Mode causes system reboot.</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. • No Change (Default) —Pre-existing configuration is retained.
Step 5	UCS-A /org/sas-expander-configuration-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates the secp1 SAS expander configuration policy:

```
UCS-A# scope org
UCS-A /org # create sas-expander-configuration-policy secp1
UCS-A /org/sas-expander-configuration-policy* # set 6g-12g-mixed-mode enabled
UCS-A /org/sas-expander-configuration-policy* # commit-buffer
UCS-A /org/sas-expander-configuration-policy#
```

Deleting a SAS Expander Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A org/ # delete sas-expander-configuration-policy <i>sas-expander-configuration-policy-name</i>	Deletes a SAS expander configuration policy with the specified policy name.
Step 3	UCS-A /org* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the secp1 SAS expander configuration policy:

```
UCS-A# scope org
UCS-A /org # delete create sas-expander-configuration-policy secp1
UCS-A /org*# commit-buffer
UCS-A /org/#
```

