



Cisco UCS Manager Network Management Guide Using the CLI, Release 3.1

First Published: 2016-01-20

Last Modified: 2017-04-27

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)

© 2016–2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

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

CHAPTER 1

| | |
|------------------------------------|----------|
| New and Changed Information | 1 |
| New and Changed Information | 1 |

CHAPTER 2

| | |
|--|----------|
| Overview | 3 |
| Overview | 3 |
| Cisco UCS Manager User CLI Documentation | 3 |
| Multilayer Network Design | 4 |

CHAPTER 3

| | |
|---|----------|
| LAN Connectivity | 7 |
| Fabric Interconnect Overview | 7 |
| Uplink Connectivity | 7 |
| Downlink Connectivity | 8 |
| Configuring the Fabric Interconnects | 8 |
| Fabric Interconnect Information Policy | 8 |
| Enabling the Information Policy on the Fabric Interconnect | 8 |
| Disabling the Information Policy on the Fabric Interconnect | 9 |
| Viewing the LAN Neighbors of the Fabric Interconnect | 10 |
| Viewing the SAN Neighbors of the Fabric Interconnect | 10 |
| Viewing the LLDP Neighbors of the Fabric Interconnect | 11 |
| Fabric Evacuation | 12 |

| | |
|--|---------------------------------------|
| Stopping Traffic on a Fabric Interconnect | 12 |
| Displaying the Status of Evacuation for a Fabric Interconnect | 13 |
| Displaying the Status of Evacuation for an IOM | 14 |
| Verifying Fabric Evacuation | 15 |
| Restarting Traffic on a Fabric Interconnect | 16 |
| Fabric Interconnect Port Types | 17 |
| Fabric Interconnect Switching Modes | 17 |
| Ethernet Switching Mode | 17 |
| Configuring Ethernet Switching Mode | 18 |
| Fibre Channel Switching Mode | 19 |
| Configuring Fibre Channel Switching Mode | 20 |
| <hr/> | |
| CHAPTER 4 | LAN Ports and Port Channels 21 |
| Unified Ports on the Cisco UCS 6200 Series and 6324 Fabric Interconnects | 21 |
| Port Modes | 22 |
| Port Types | 22 |
| Data Traffic Interruption from Port Mode Changing | 23 |
| Guidelines for Configuring Unified Ports | 23 |
| Cautions and Guidelines for Configuring Unified Uplink Ports and Unified Storage Ports | 25 |
| Configuring the Port Mode | 26 |
| Configuring Breakout Ports | 28 |
| Cisco UCS 6300 Fabric Interconnect Ethernet Breakout Ports | 28 |
| Configuring a 10G port with QSA Adapter on Cisco UCS FI 6332 | 29 |
| Configuring Multiple Breakout Ports | 30 |
| Configuring a Breakout Ethernet Uplink Port | 31 |
| Configuring a Breakout Ethernet Uplink Port Channel Member | 31 |
| Configuring Ethernet Uplink Breakout Port as a Pin Group Target | 32 |
| Configuring Breakout Appliance Ports | 33 |
| Configuring a Breakout Appliance Port Channel Member | 34 |
| Configuring Breakout FCoE Storage Ports | 34 |
| Configuring a Breakout FCoE Uplink Port | 35 |
| Configuring an FCoE Port Channel Member | 36 |
| Configuring a Breakout VLAN Member Port | 37 |
| Modifying a Breakout Port | 37 |

| | |
|---|----|
| Un-configuring Breakout Ports | 42 |
| Deleting Breakout Ports | 43 |
| Cisco UCS Mini Scalability Ports | 44 |
| Configuring Scalability Ports | 45 |
| Beacon LEDs for Unified Ports | 46 |
| Configuring the Beacon LEDs for Unified Ports | 46 |
| Physical and Backplane Ports | 47 |
| Displaying VIF Port Statistics Obtained From the Adaptor | 47 |
| Displaying VIF Port Statistics Obtained From the ASIC | 48 |
| Displaying VIF Ports That Correspond to NIV Ports | 48 |
| Verifying Status of Backplane Ports | 49 |
| Server Ports | 51 |
| Automatic Configuration of Fabric Interconnect Server Ports | 51 |
| Automatically Configuring Server Ports | 51 |
| Configuring a Server Port | 52 |
| Unconfiguring a Server Port | 53 |
| Uplink Ethernet Ports | 53 |
| Configuring an Uplink Ethernet Port | 53 |
| Unconfiguring an Uplink Ethernet Port | 54 |
| Appliance Ports | 55 |
| Configuring an Appliance Port | 55 |
| Assigning a Target MAC Address to an Appliance Port or Appliance Port Channel | 57 |
| Creating an Appliance Port | 58 |
| Mapping an Appliance Port to a Community VLAN | 59 |
| Unconfiguring an Appliance Port | 60 |
| FCoE Uplink Ports | 60 |
| Configuring a FCoE Uplink Port | 61 |
| Unconfiguring a FCoE Uplink Port | 61 |
| Viewing FCoE Uplink Ports | 62 |
| Unified Storage Ports | 63 |
| Configuring a Unified Storage Port | 63 |
| Unified Uplink Ports | 64 |
| Configuring a Unified Uplink Port | 64 |
| FCoE and Fibre Channel Storage Ports | 65 |

| | |
|---|----|
| Configuring a Fibre Channel Storage or FCoE Port | 65 |
| Unconfiguring a Fibre Channel Storage or FCoE Port | 66 |
| Restoring a Fibre Channel Storage Port Back to an Uplink Fibre Channel Port | 66 |
| Uplink Ethernet Port Channels | 67 |
| Configuring an Uplink Ethernet Port Channel | 67 |
| Unconfiguring an Uplink Ethernet Port Channel | 68 |
| Adding a Member Port to an Uplink Ethernet Port Channel | 69 |
| Deleting a Member Port from an Uplink Ethernet Port Channel | 69 |
| Appliance Port Channels | 70 |
| Configuring an Appliance Port Channel | 70 |
| Unconfiguring an Appliance Port Channel | 72 |
| Enabling or Disabling an Appliance Port Channel | 72 |
| Adding a Member Port to an Appliance Port Channel | 73 |
| Deleting a Member Port from an Appliance Port Channel | 74 |
| Fibre Channel Port Channels | 74 |
| Configuring a Fibre Channel Port Channel | 75 |
| Configuring a FCoE Port Channel | 76 |
| Adding Channel Mode Active To The Upstream NPIV Fibre Channel Port Channel | 76 |
| Enabling or Disabling a Fibre Channel Port Channel | 77 |
| Adding a Member Port to a Fibre Channel Port Channel | 78 |
| Deleting a Member Port from a Fibre Channel Port Channel | 79 |
| FCoE Port Channels | 79 |
| Configuring a FCoE Port Channel | 80 |
| Adding a Member Port to a FCoE Uplink Port Channel | 80 |
| Unified Uplink Port Channel | 81 |
| Configuring a Unified Uplink Port Channel | 82 |
| Event Detection and Action | 82 |
| Policy-Based Port Error Handling | 83 |
| Creating Threshold Definition | 83 |
| Configuring Error Disable on a Fabric Interconnect Port | 84 |
| Configuring Auto Recovery on a Fabric Interconnect Port | 85 |
| Viewing the Network Interface Port Error Counters | 86 |
| Adapter Port Channels | 87 |
| Viewing Adapter Port Channels | 87 |

| | |
|--|-----------------|
| Fabric Port Channels | 88 |
| Load Balancing Over Ports | 88 |
| Cabling Considerations for Fabric Port Channels | 89 |
| Configuring a Fabric Port Channel | 89 |
| Viewing Fabric Port Channels | 90 |
| Enabling or Disabling a Fabric Port Channel Member Port | 90 |
| <hr/> | |
| CHAPTER 5 | VLANs 93 |
| Named VLANs | 93 |
| Private VLANs | 94 |
| VLAN Port Limitations | 95 |
| Configuring Named VLANs | 97 |
| Creating a Named VLAN Accessible to Both Fabric Interconnects (Uplink Ethernet Mode) | 97 |
| Creating a Named VLAN Accessible to Both Fabric Interconnects (Ethernet Storage Mode) | 98 |
| Creating a Named VLAN Accessible to One Fabric Interconnect (Uplink Ethernet Mode) | 99 |
| Creating a Secondary VLAN for a Private VLAN (Accessible to One Fabric Interconnect) | 100 |
| Deleting a Named VLAN | 101 |
| Configuring Private VLANs | 102 |
| Creating a Primary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects) | 102 |
| Creating a Primary VLAN for a Private VLAN (Accessible to One Fabric Interconnect) | 103 |
| Creating a Secondary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects) | 104 |
| Creating a Secondary VLAN for a Private VLAN (Accessible to One Fabric Interconnect) | 105 |
| Allowing PVLANS on vNICs | 106 |
| Creating a Primary VLAN for a Private VLAN on an Appliance Cloud | 107 |
| Creating a Secondary VLAN for a Private VLAN on an Appliance Cloud | 108 |
| Community VLANs | 109 |
| Creating a Community VLAN | 109 |
| Viewing Community VLANs | 109 |
| Allowing Community VLANs on vNICs | 110 |
| Allowing PVLAN on Promiscuous Access or Trunk Port | 111 |
| Deleting a Community VLAN | 111 |
| Viewing the VLAN Port Count | 112 |
| VLAN Port Count Optimization | 113 |
| Enabling Port VLAN Count Optimization | 113 |

| | |
|---|------------|
| Disabling Port VLAN Count Optimization | 114 |
| Viewing the Port VLAN Count Optimization Groups | 114 |
| VLAN Groups | 115 |
| Creating a VLAN Group | 115 |
| Creating an Inband VLAN Group | 116 |
| Viewing VLAN Groups | 117 |
| Deleting a VLAN Group | 117 |
| VLAN Permissions | 118 |
| Creating VLAN Permissions | 118 |
| Viewing VLAN Permissions | 119 |
| Deleting a VLAN Permission | 119 |

CHAPTER 6

| | |
|-----------------------------|------------|
| LAN Pin Groups | 121 |
| LAN Pin Groups | 121 |
| Configuring a LAN Pin Group | 121 |

CHAPTER 7

| | |
|---------------------|------------|
| MAC Pools | 123 |
| MAC Pools | 123 |
| Creating a MAC Pool | 123 |
| Deleting a MAC Pool | 125 |

CHAPTER 8

| | |
|---|------------|
| Quality of Service | 127 |
| Quality of Service | 127 |
| Configuring System Classes | 128 |
| System Classes | 128 |
| Configuring a System Class | 129 |
| Disabling a System Class | 131 |
| Configuring Quality of Service Policies | 132 |
| Quality of Service Policy | 132 |
| Configuring a QoS Policy | 132 |
| Deleting a QoS Policy | 134 |
| Configuring Flow Control Policies | 135 |
| Flow Control Policy | 135 |
| Configuring a Flow Control Policy | 135 |

Deleting a Flow Control Policy 136

CHAPTER 9

Upstream Disjoint Layer-2 Networks 139

Upstream Disjoint Layer-2 Networks 139

Guidelines for Configuring Upstream Disjoint L2 Networks 140

Upstream Disjoint L2 Networks Pinning Considerations 141

Configuring Cisco UCS for Upstream Disjoint L2 Networks 143

Assigning Ports and Port Channels to VLANs 144

Removing Ports and Port Channels from VLANs 144

Viewing Ports and Port Channels Assigned to VLANs 145

CHAPTER 10

Network-Related Policies 147

vNIC Template 147

Creating vNIC Template Pairs 148

Undo vNIC Template Pairs 150

Configuring a vNIC Template 151

Deleting a vNIC Template 154

Ethernet Adapter Policies 154

Configuring an Ethernet Adapter Policy 154

Deleting an Ethernet Adapter Policy 156

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE 156

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with VXLAN 158

Ethernet and Fibre Channel Adapter Policies 160

Accelerated Receive Flow Steering 161

Guidelines and Limitations for Accelerated Receive Flow Steering 161

Interrupt Coalescing 162

Adaptive Interrupt Coalescing 162

Guidelines and Limitations for Adaptive Interrupt Coalescing 162

RDMA Over Converged Ethernet for SMB Direct 163

Guidelines and Limitations for SMB Direct with RoCE 163

Configuring a Default vNIC Behavior Policy 163

Deleting a vNIC from a LAN Connectivity Policy 164

Creating a LAN Connectivity Policy 165

Deleting a LAN Connectivity Policy 166

| | |
|--|-----|
| About the LAN and SAN Connectivity Policies | 166 |
| Privileges Required for LAN and SAN Connectivity Policies | 167 |
| Interactions between Service Profiles and Connectivity Policies | 167 |
| Creating a LAN Connectivity Policy | 167 |
| Creating a vNIC for a LAN Connectivity Policy | 168 |
| Deleting a vNIC from a LAN Connectivity Policy | 171 |
| Creating an iSCSI vNIC for a LAN Connectivity Policy | 172 |
| Deleting an iSCSI vNIC from a LAN Connectivity Policy | 174 |
| Network Control Policy | 174 |
| Configuring a Network Control Policy | 175 |
| Configuring Link Layer Discovery Protocol for Fabric Interconnect vEthernet Interfaces | 177 |
| Displaying Network Control Policy Details | 177 |
| Deleting a Network Control Policy | 178 |
| Creating a Multicast Policy | 178 |
| Deleting a Multicast Policy | 179 |
| Entering Multicast Policy Mode | 180 |
| Enter a Multicast Policy | 180 |
| Assigning a Global VLAN Multicast Policy | 181 |
| Disassociating a Global VLAN Multicast Policy | 181 |
| Disassociating a VLAN Multicast Policy | 182 |
| Configuring Ethernet Adapter Policies | 183 |
| Configuring an Ethernet Adapter Policy | 183 |
| Deleting an Ethernet Adapter Policy | 184 |
| Configuring the Default vNIC Behavior Policy | 185 |
| Default vNIC Behavior Policy | 185 |
| Configuring a Default vNIC Behavior Policy | 185 |
| Configuring a Network Control Policy | 186 |
| Deleting a Network Control Policy | 188 |
| Configuring Multicast Policies | 188 |
| Multicast Policy | 188 |
| Creating a Multicast Policy | 189 |
| Configuring IGMP Snooping Parameters | 190 |
| Modifying Multicast Policy Parameters | 191 |
| Assigning a VLAN Multicast Policy | 192 |

| | |
|---|-----|
| Deleting a Multicast Policy | 192 |
| LACP Policy | 193 |
| Creating a LACP Policy | 193 |
| Editing a LACP Policy | 194 |
| Assigning LACP Policy to Port-Channels | 194 |
| Configuring UDLD Link Policies | 195 |
| Understanding UDLD | 195 |
| UDLD Configuration Guidelines | 197 |
| Configuring a UDLD Link Policy | 197 |
| Modifying the UDLD System Settings | 198 |
| Configuring a Link Profile | 199 |
| Assigning a Link Profile to a Port Channel Ethernet Interface | 200 |
| Assigning a Link Profile to a Port Channel FCoE Interface | 200 |
| Assigning a Link Profile to an Uplink Ethernet Interface | 201 |
| Assigning a Link Profile to an Uplink FCoE Interface | 202 |
| VMQ Connection Policy | 202 |
| Creating a VMQ Connection Policy | 203 |



Preface

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

Audience

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

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

Conventions

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

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



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



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



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



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



Warning IMPORTANT SAFETY INSTRUCTIONS

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

SAVE THESE INSTRUCTIONS

Related Cisco UCS Documentation

Documentation Roadmaps

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

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

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

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](#), on page 1

New and Changed Information



CHAPTER 2

Overview

- [Overview, on page 3](#)
- [Cisco UCS Manager User CLI Documentation, on page 3](#)
- [Multilayer Network Design, on page 4](#)

Overview

This guide includes the following information:

- Configure/Enable Server Ports; Configure/Enable Uplink Ports; Configure/Enable FC Ports.
- Create LAN Pin Groups
- Create VLANs and VLAN groups
- Create Server Links
- Configure QoS System Class
- Configure Global Policies
- Monitor Network Health
- Traffic Monitoring

Cisco UCS Manager User CLI Documentation

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

| Guide | Description |
|---|---|
| Cisco UCS Manager Getting Started Guide | Discusses Cisco UCS architecture and Day 0 operations, including Cisco UCS Manager initial configuration, and configuration best practices. |

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

Multilayer Network Design

When you design a data center using a modular approach, the network is divided into three functional layers: Core, Aggregation, and Access. These layers can be physical or logical, and you can add and remove them without redesigning the entire data center network.

Because of the hierarchical topology of a modular design, the addressing is also simplified within the data center network. Modularity implies isolation of building blocks, which are separated from each other and communicate through specific network connections between the blocks. Modular design provides easy control of traffic flow and improved security. In other words, these blocks are independent from each other; a change in one block does not affect other blocks. Modularity also enables faster moves, adds, and changes (MACs) and incremental changes in the network.

Modular networks are scalable. Scalability allows for the considerable growth or reduction in the size of a network without making drastic changes. Scalable data center network design is achieved by using the principle of hierarchy and modularity.

Keep a network as simple as possible. Modular designs are simple to design, configure, and troubleshoot.

- **Access Layer**—The access layer is the first point of entry into the network for edge devices, end stations, and servers. The Access layer grants user access to network devices and provides connectivity to servers. The switches in the access layer are connected to two separate distribution layer switches for redundancy. The data center access layer provides Layer 2, Layer 3, and mainframe connectivity. The design of the access layer varies, depending on whether you use Layer 2 or Layer 3 access. The access layer in the data center is typically built at Layer 2, which allows better sharing of service devices across multiple servers. This design also enables the use of Layer 2 clustering, which requires the servers to be Layer 2 adjacent. With Layer 2 access, the default gateway, you can configure for the servers at the aggregation layer.
- **Aggregation Layer**—The aggregation (or distribution) layer aggregates the uplinks from the access layer to the data center core. This layer is the critical point for control and application services. Security and application service devices (such as load-balancing devices, SSL offloading devices, firewalls, and IPS devices) are often deployed as modules in the aggregation layer. The aggregation layer provides policy-based connectivity.
- **Core Layer**—Also known as backbone, the core layer provides high-speed packet switching, scalability and high availability, and fast convergence. Implementing a data center core is a best practice for large data centers. When you implement the core in an initial data center design, it eases network expansion and avoids disruption to the data center environment.

Use the following criteria to determine whether a core solution is appropriate: The data center typically connects to the campus core using Layer 3 links. The data center network is summarized, and the core injects a default route into the data center network.

- Ethernet bandwidth requirements
- Port density
- Administrative domains
- Anticipated future development



CHAPTER 3

LAN Connectivity

- [Fabric Interconnect Overview, on page 7](#)
- [Uplink Connectivity, on page 7](#)
- [Downlink Connectivity, on page 8](#)
- [Configuring the Fabric Interconnects, on page 8](#)
- [Fabric Evacuation, on page 12](#)
- [Fabric Interconnect Port Types, on page 17](#)
- [Fabric Interconnect Switching Modes, on page 17](#)

Fabric Interconnect Overview

The fabric interconnect is the core component of Cisco UCS. The Cisco UCS Fabric Interconnects provide uplink access to LAN, SAN, and out-of-band management segment. Cisco UCS infrastructure management is through the embedded management software, Cisco UCS Manager, for both hardware and software management. The Cisco UCS Fabric Interconnects are Top-of-Rack devices, and provide unified access to the Cisco UCS domain.

The Cisco UCS FIs provide network connectivity and management for the connected servers. The Cisco UCS Fabric Interconnects run the Cisco UCS Manager control software and consist of expansion modules for the Cisco UCS Manager software.

For more information about Cisco UCS Fabric Interconnects, see the *Cisco UCS Manager Getting Started Guide*.

Uplink Connectivity

Use fabric interconnect ports configured as uplink ports to connect to uplink upstream network switches. Connect these uplink ports to upstream switch ports as individual links, or as links configured as port channels. Port channel configurations provide bandwidth aggregation as well as link redundancy.

You can achieve northbound connectivity from the fabric interconnect through a standard uplink, a port channel, or a virtual port channel configuration. The port channel name and ID configured on fabric interconnect should match the name and ID configuration on the upstream Ethernet switch.

It is also possible to configure a port channel as a vPC, where port channel uplink ports from a fabric interconnect are connected to different upstream switches. After all uplink ports are configured, create a port channel for these ports.

Downlink Connectivity

Each fabric interconnect is connected to IOMs in the UCS chassis, which provides connectivity to each blade server. Internal connectivity from blade servers to IOMs is transparently provided by Cisco UCS Manager using 10BASE-KR Ethernet standard for backplane implementations, and no additional configuration is required. You must configure the connectivity between the fabric interconnect server ports and IOMs. Each IOM, when connected with the fabric interconnect server port, behaves as a line card to fabric interconnect, hence IOMs should never be cross-connected to the fabric interconnect. Each IOM is connected directly to a single fabric interconnect.

The Fabric Extender (also referred to as the IOM, or FEX) logically extends the fabric interconnects to the blade server. The best analogy is to think of it as a remote line card that's embedded in the blade server chassis, allowing connectivity to the external world. IOM settings are pushed via Cisco UCS Manager and are not managed directly. The primary functions of this module are to facilitate blade server I/O connectivity (internal and external), multiplex all I/O traffic up to the fabric interconnects, and help monitor and manage the Cisco UCS infrastructure.

Configure Fabric interconnect ports that should be connected to downlink IOM cards as server ports. Make sure there is physical connectivity between the fabric interconnect and IOMs. You must also configure the IOM ports and the global chassis discovery policy.

**Note**

For UCS 2200 I/O modules, you can also select the Port Channel option and all I/O module-connected server ports will be automatically added to a port channel.

Configuring the Fabric Interconnects

Fabric Interconnect Information Policy

You must configure the information policy to display the uplink switches that are connected to Cisco UCS.

**Important**

You must enable the information policy on the fabric interconnect to view the SAN, LAN, and LLDP neighbors of the fabric interconnect.

Enabling the Information Policy on the Fabric Interconnect

**Note**

By default, the information policy is disabled on the fabric interconnect.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A # scope system | Enters system mode. |
| Step 2 | UCS-A/system # scope info-policy | Enters the information policy state. |
| Step 3 | (Optional) UCS-A/system/info-policy # show | Shows if the information policy is enabled or disabled. |
| Step 4 | UCS-A/system/info-policy # enable | Enables the information policy on the fabric interconnect. |
| Step 5 | UCS-A/system/info-policy* # commit-buffer | Enables the information policy on the fabric interconnect. |

Example

The following example shows how to enable the information policy on the fabric interconnect:

```
UCS-A# scope system
UCS-A/system # scope info-policy
UCS-A/system/info-policy # show
Info Policy:
State: Disabled
UCS-A/system/info-policy # enable
UCS-A/system/info-policy* # commit-buffer
UCS-A/system/info-policy #
```

Disabling the Information Policy on the Fabric Interconnect

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A # scope system | Enters system mode. |
| Step 2 | UCS-A/system # scope info-policy | Enters the information policy state. |
| Step 3 | (Optional) UCS-A/system/info-policy # show | Shows if the information policy is enabled or disabled. |
| Step 4 | UCS-A/system/info-policy # disable | Disables the information policy on the fabric interconnect. |
| Step 5 | UCS-A/system/info-policy* # commit-buffer | Disables information policy on the fabric interconnect. |

Example

The following example shows how to disable the information policy on the fabric interconnect:

```

UCS-A# scope system
UCS-A/system # scope info-policy
UCS-A/system/info-policy # show
Info Policy:
State: Enabled
UCS-A/system/info-policy # disable
UCS-A/system/info-policy* # commit-buffer
UCS-A/system/info-policy #

```

Viewing the LAN Neighbors of the Fabric Interconnect

You must enable the information policy on the fabric interconnect to view the LAN neighbors.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope fabric-interconnect {a b} | Enters fabric interconnect mode for the specified fabric interconnect. |
| Step 2 | UCS-A/fabric-interconnect # show lan-neighbors | Displays the fabric interconnect LAN neighbors. |

Example

The following example shows how to display the LAN neighbors of the fabric interconnect:

```

UCS-A # scope fabric-interconnect a
UCS-Afabric-interconnect # show lan-neighbors
Info Policy:Enabled
Lan Neighbors:
Local Interface: Ethernet1/2
Device Id: bgl-samc02-B(SSI140305YK)
IPv4 Address: 10.105.214.105
FI Port DN: sys/switch-A/slot-1/switch-ether/port-2

```

Viewing the SAN Neighbors of the Fabric Interconnect

You must enable the information policy on the fabric interconnect to view the SAN neighbors.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope fabric-interconnect {a b} | Enters fabric interconnect mode for the specified fabric interconnect. |
| Step 2 | UCS-A/fabric-interconnect # show san-neighbors | Displays the fabric interconnect SAN neighbors. |

Example

The following example shows how to display the SAN neighbors of the fabric interconnect :

```
UCS-A # scope fabric-interconnect a
UCS-A/fabric-interconnect # show san-neighbors
Info Policy: Enabled
San neighbors:
Local Interface: fc2/1
Port VSAN: 100
Fabric Mgmt Addr: 10.65.124.252
Fabric pwnn: 20:02:00:05:9b:22:ad:c0
Fabric nwnn: 20:64:00:05:9b:22:ad:c1
My pwnn: 20:41:00:0d:ec:ee:dd:00
My nwnn: 20:64:00:0d:ec:ee:dd:01
FI Port DN: sys/switch-A/slot-2/switch-fc/port-1
```

Viewing the LLDP Neighbors of the Fabric Interconnect

You must enable the information policy on the fabric interconnect to view the LLDP neighbors.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope fabric-interconnect {a b} | Enters fabric interconnect mode for the specified fabric interconnect. |
| Step 2 | UCS-A/fabric-interconnect # show lldp-neighbors | Displays the fabric interconnect LLDP neighbors. |

Example

The following example shows how to display the LLDP neighbors of the fabric interconnect :

```
UCS-A # scope fabric-interconnect a
UCS-A/fabric-interconnect # show lldp-neighbors
Info Policy: Enabled

Lldp Neighbors:

Local Interface: Eth1/5
Chassis Id: 000d.ecff.5e90
Remote Interface: Eth1/9
Remote Port Description: Ethernet1/9
System Name: bgl-samc02-B
System Description: Cisco Nexus Operating System (NX-OS) Software TAC support:
http://www.cisco.com/tac Copyright (c) 2002-2011, Cisco Systems, Inc
System Capabilities: B
Enabled Capabilities: B
Native VLAN: 1
IPv4 Mgmt Address: 10.105.214.105
FI Port DN: sys/switch-A/slot-1/switch-ether/port-5
```

Fabric Evacuation

Cisco UCS Manager introduces fabric evacuation, which is the ability to evacuate all traffic that flows through a fabric interconnect from all servers attached to it through an IOM or FEX while upgrading a system. Fabric evacuation is not supported on direct-attached rack servers.

Upgrading the secondary fabric interconnect in a system disrupts active traffic on the fabric interconnect. This traffic fails over to the primary fabric interconnect. You can use fabric evacuation during the upgrade process as follows:

1. Stop all the traffic that is active through a fabric interconnect.
2. For vNICs configured with failover, verify that the traffic has failed over by using Cisco UCS Manager, or tools such as vCenter.
3. Upgrade the secondary fabric interconnect.
4. Restart all the stopped traffic flows.
5. Change the cluster lead to the secondary fabric interconnect.
6. Repeat steps 1 to 4 and upgrade the primary fabric interconnect.



Note

- Fabric interconnect traffic evacuation is supported only in a cluster configuration.
- You can evacuate traffic only from the subordinate fabric interconnect.
- The IOM or FEX backplane ports of the fabric interconnect on which evacuation is configured will go down, and their state will appear as **Admin down**. During the manual upgrade process, to move these backplane ports back to the Up state and resume traffic flow, you must explicitly configure **Admin Evac Mode** as **Off**.
- Starting with Cisco UCS Manager Release 3.1(3), you can use fabric evacuation during Auto Install.
- If you use fabric evacuation outside of the upgrade process, you must re-acknowledge the FEX to get the VIFs back to the online state.

Stopping Traffic on a Fabric Interconnect

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A # scope fabric-interconnect {a b} | Enters the fabric interconnect mode. |
| Step 2 | UCS-A /fabric-interconnect # stop server traffic [force] | Stops all the traffic that is active through the specified Fabric Interconnect. Use the force option to evacuate a fabric interconnect regardless of its current evacuation state. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 3 | UCS-A /fabric-interconnect # commit-buffer | Commits the transaction to the system configuration. |

Example

This example shows how to stop all traffic that is active through Fabric Interconnect B:

```
UCS-A# scope fabric-interconnect b
UCS-A /fabric-interconnect # stop server traffic
Warning: Enabling fabric evacuation will stop all traffic through this Fabric Interconnect
         from servers attached through IOM/FEX. The traffic will fail over to the Primary Fabric
         Interconnect for fail over vnics.
UCS-A /fabric-interconnect # commit-buffer
```

Displaying the Status of Evacuation for a Fabric Interconnect

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A # scope fabric-interconnect {a b} | Enters fabric interconnect mode for the specified fabric interconnect. |
| Step 2 | UCS-A /fabric-interconnect # show detail | Displays details about the specified fabric interconnect. |

Example

This example shows how to display the status of a fabric interconnect.



Note Admin Evacuation and Oper Evacuation and show the status of evacuation at the fabric interconnect.

```
UCS-A /fabric-interconnect # show detail
```

```
Fabric Interconnect:
  ID: B
  Product Name: Cisco UCS 6248UP
  PID: UCS-FI-6248UP
  VID: V01
  Vendor: Cisco Systems, Inc.
  Serial (SN): SSI171400HG
  HW Revision: 0
  Total Memory (MB): 16165
  OOB IP Addr: 10.193.32.172
  OOB Gateway: 10.193.32.1
  OOB Netmask: 255.255.255.0
  OOB IPv6 Address: ::
```

```

OOB IPv6 Gateway: ::
Prefix: 64
Operability: Operable
Thermal Status: Ok
Admin Evacuation: On
Oper Evacuation: On
Current Task 1:
Current Task 2:
Current Task 3:

```

Displaying the Status of Evacuation for an IOM

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 <i>iom-id</i> | Enters chassis IOM mode for the specified IOM. |
| Step 3 | UCS-A /chassis/iom # show detail | Displays evacuation status details for the specified IOM. |

Example

This example shows how to display the evacuation status details for an IOM.



Note **Oper Evacuation** shows the operational status of evacuation for the IOM.

```

UCS-A# scope chassis 1
UCS-A /chassis # scope iom 1
UCS-A /chassis/iom # show detail

IOM:
  ID: 1
  Side: Left
  Fabric ID: A
  User Label:
  Overall Status: Fabric Conn Problem
  Oper qualifier: Server Port Problem
  Operability: Operable
  Presence: Equipped
  Thermal Status: OK
  Discovery: Online
  Config State: Ok
  Peer Comm Status: Connected
  Product Name: Cisco UCS 2204XP
  PID: UCS-IOM-2204XP
  VID: V02
  Part Number: 73-14488-02
  Vendor: Cisco Systems Inc
  Serial (SN): FCH1718J9FT

```

```

HW Revision: 0
Mfg Date: 2013-05-12T00:00:00.000
Controller Subject: Iocard
Fabric Port Aggregation Capability: Port Channel
Oper Evacuation: On
Current Task 1:
Current Task 2:

```

Verifying Fabric Evacuation

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# show service-profile circuit server <i>server-id</i> | Shows the network circuit information for the service profile associated with the specified server. |

Example

The following example shows the VIF (Virtual NIC) paths before fabric evacuation.



Note

- VIF at Fabric Interconnect A shows that traffic is initially active through the fabric interconnect.
- VIF at Fabric Interconnect B is passive before evacuation.

```

UCS-A# show service-profile circuit server 1/6
Service Profile: test1
Server: 1/6
  Fabric ID: A
    Path ID: 1
      VIF      vNIC      Link State  Oper State  Prot State  Prot Role  Admin
Pin  Oper Pin  Transport
-----
      692 eth0      Up          Active      Active      Primary    0/0
      1/15      Ether
  Fabric ID: B
    Path ID: 1
      VIF      vNIC      Link State  Oper State  Prot State  Prot Role  Admin
Pin  Oper Pin  Transport
-----
      693 eth0      Up          Active      Passive     Backup     0/0
      1/15      Ether
UCS-A#

```

The following example shows the VIF paths after Fabric Interconnect A is evacuated.

**Note**

- After failover, the VIF state at Fabric Interconnect A goes into error.
- VIF at Fabric Interconnect B takes over as active.

```
UCS-A# show service-profile circuit server 1/6
Service Profile: test1
Server: 1/6
  Fabric ID: A
    Path ID: 1
      VIF      vNIC      Link State Oper State Prot State  Prot Role  Admin
Pin Oper Pin  Transport
-----
0/0      692 eth0      Error      Error      Active      Primary    0/0
      Ether
  Fabric ID: B
    Path ID: 1
      VIF      vNIC      Link State Oper State Prot State  Prot Role  Admin
Pin Oper Pin  Transport
-----
1/15     693 eth0      Up          Active     Passive     Backup     0/0
      Ether
UCS-A#
```

Restarting Traffic on a Fabric Interconnect

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A # scope fabric-interconnect {a b} | Enters the fabric interconnect mode. |
| Step 2 | UCS-A /fabric-interconnect # start server traffic | Restarts traffic through the specified fabric interconnect. |
| Step 3 | UCS-A /fabric-interconnect # commit-buffer | Commits the transaction to the system configuration. |

Example

This example shows how to restart traffic through Fabric Interconnect B:

```
UCS-A# scope fabric-interconnect b
UCS-A /fabric-interconnect # start server traffic
Warning: Resetting fabric evacuation will cause server traffic that failed over to the
Primary Fabric Interconnect to fail back to this Fabric Interconnect.
UCS-A /fabric-interconnect # commit-buffer
```


Fabric Interconnect Port Types

By default, all fabric interconnect ports are unconfigured. For Ethernet LAN connectivity, fabric interconnect ports can be in the following states:

- **Unconfigured**—Port is not configured and cannot be used.
- **Server Port**—Port is configured for downlink connection to an IOM Fabric Extender (FEX) module in a blade chassis.
- **Uplink Port**—Port is configured for uplink connection to the upstream Ethernet switch. Uplink ports are always configured as trunk ports.
- **Disabled**—Port is configured either as an uplink or server port and is currently disabled by the administrator.

For 6200 series fabric interconnects, all ports are unified ports; therefore you also configure all the ports as 1/10 Gigabit Ethernet, Fibre Channel (FC), FC uplink, appliance port, or FCoE port.

For 6300 series fabric interconnects, see the *UCS Manager Getting Started Guide*.

Fabric Interconnect Switching Modes

The Cisco UCS Fabric Interconnects are operated in two main switching modes: Ethernet or Fibre Channel. These modes are independent of each other. They determine how the fabric interconnect behaves as a device between the server and network/server and storage device.

Ethernet Switching Mode

The Ethernet switching mode determines how the fabric interconnect behaves as a switching device between the servers and the network. The fabric interconnect operates in either of the following Ethernet switching modes:

End-Host Mode

End-host mode allows the fabric interconnect to act as an end host to the network, representing all servers (hosts) connected to it through vNICs. This behavior is achieved by pinning (either dynamically pinning or hard pinning) vNICs to uplink ports, which provides redundancy to the network, and makes the uplink ports appear as server ports to the rest of the fabric.

In end-host mode, the fabric interconnect does not run the Spanning Tree Protocol (STP), but it avoids loops by denying uplink ports from forwarding traffic to each other and by denying egress server traffic on more than one uplink port at a time. End-host mode is the default Ethernet switching mode and should be used if either of the following is used upstream:

- Layer 2 switching for Layer 2 aggregation
- Virtual Switching System (VSS) aggregation layer



Note When you enable end-host mode, if a vNIC is hard pinned to an uplink port and this uplink port goes down, the system cannot repin the vNIC, and the vNIC remains down.

Switch Mode

Switch mode is the traditional Ethernet switching mode. The fabric interconnect runs STP to avoid loops, and broadcast and multicast packets are handled in the traditional way. Use the switch mode only if the fabric interconnect is directly connected to a router, or if either of the following is used upstream:

- Layer 3 aggregation
- VLAN in a box



Note For both Ethernet switching modes, even when vNICs are hard-pinned to uplink ports, all server-to-server unicast traffic in the server array is sent only through the fabric interconnect and is never sent through uplink ports. Server-to-server multicast and broadcast traffic is sent through all uplink ports in the same VLAN.

Cisco UCS Fabric Interconnect in Switch Mode with Cisco MDS 9000 Family Fibre Channel Switching Modules

While creating a port channel between a Cisco MDS 9000 family FC switching module and a Cisco UCS Fabric Interconnect in switch mode, use the following order:

1. Create the port channel on the MDS side.
2. Add the port channel member ports.
3. Create the port channel on the Fabric Interconnect side.
4. Add the port channel member ports.

If you create the port channel on the Fabric Interconnect side first, the ports will go into a suspended state.

When the Cisco UCS Fabric Interconnect is in switch mode, the port channel mode can only be in **ON** mode and not **Active**. However, to get the peer wwn information for the Fabric Interconnect, the port channel must be in **Active** mode.

Configuring Ethernet Switching Mode



Important

When you change the Ethernet switching mode, Cisco UCS Manager logs you out and restarts the fabric interconnect. For a cluster configuration, Cisco UCS Manager restarts both fabric interconnects. The subordinate fabric interconnect reboots first as a result of the change in switching mode. The primary fabric interconnect reboots only after you acknowledge it in **Pending Activities**. The primary fabric interconnect can take several minutes to complete the change in Ethernet switching mode and become system ready. The existing configuration is retained.

While the fabric interconnects are rebooting, all blade servers lose LAN and SAN connectivity, causing a complete outage of all services on the blades. This might cause the operating system to fail.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # set mode {end-host switch} | Sets the fabric interconnect to the specified switching mode. |
| Step 3 | UCS-A /eth-uplink # commit-buffer | Commits the transaction to the system configuration. Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager CLI. |

Example

The following example sets the fabric interconnect to end-host mode and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # set mode end-host
Warning: When committed, this change will cause the switch to reboot
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```

Fibre Channel Switching Mode

The Fibre Channel switching mode determines how the fabric interconnect behaves as a switching device between the servers and storage devices. The fabric interconnect operates in either of the following Fibre Channel switching modes:

End-Host Mode

End-host mode is synonymous with N Port Virtualization (NPV) mode. This mode is the default Fibre Channel Switching mode. End-host mode allows the fabric interconnect to act as an end host to the connected fibre channel networks, representing all servers (hosts) connected to it through virtual host bus adapters (vHBAs). This behavior is achieved by pinning (either dynamically pinning or hard-pinning) vHBAs to Fibre Channel uplink ports, which makes the Fibre Channel ports appear as server ports (N-ports) to the rest of the fabric. When in end-host mode, the fabric interconnect avoids loops by preventing uplink ports from receiving traffic from one another.



Note When you enable end-host mode, if a vHBA is hard-pinned to an uplink Fibre Channel port and this uplink port goes down, the system cannot repin the vHBA, and the vHBA remains down.

Switch Mode

Switch mode is not the default Fibre Channel switching mode. Switch mode allows the fabric interconnect to connect directly to a storage device. Enabling Fibre Channel switch mode is useful in Pod models where there is no SAN (for example, a single Cisco UCS domain that is connected directly to storage), or where a

SAN exists (with an upstream MDS). In Fibre Channel switch mode, SAN pin groups are irrelevant. Any existing SAN pin groups are ignored.

Configuring Fibre Channel Switching Mode



Note When the Fibre Channel switching mode is changed, both Cisco UCS fabric interconnects reload simultaneously. Reloading the fabric interconnects will cause a system-wide downtime for approximately 10 to 15 minutes.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope fc-uplink | Enters Fibre Channel uplink mode. |
| Step 2 | UCS-A /fc-uplink # set mode {end-host switch} | Sets the fabric interconnect to the specified switching mode. |
| Step 3 | UCS-A /fc-uplink # commit-buffer | Commits the transaction to the system configuration. Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager CLI. |

Example

The following example shows how to set the fabric interconnect to end-host mode and commit the transaction:

```
UCS-A # scope fc-uplink
UCS-A /fc-uplink # set mode end-host
UCS-A /fc-uplink* # commit-buffer
UCS-A /fc-uplink #
```



CHAPTER 4

LAN Ports and Port Channels

- [Unified Ports on the Cisco UCS 6200 Series and 6324 Fabric Interconnects, on page 21](#)
- [Physical and Backplane Ports, on page 47](#)
- [Server Ports, on page 51](#)
- [Uplink Ethernet Ports, on page 53](#)
- [Appliance Ports, on page 55](#)
- [FCoE Uplink Ports, on page 60](#)
- [Unified Storage Ports, on page 63](#)
- [Unified Uplink Ports, on page 64](#)
- [FCoE and Fibre Channel Storage Ports, on page 65](#)
- [Uplink Ethernet Port Channels, on page 67](#)
- [Appliance Port Channels, on page 70](#)
- [Fibre Channel Port Channels, on page 74](#)
- [FCoE Port Channels, on page 79](#)
- [Unified Uplink Port Channel, on page 81](#)
- [Event Detection and Action, on page 82](#)
- [Adapter Port Channels, on page 87](#)
- [Fabric Port Channels, on page 88](#)

Unified Ports on the Cisco UCS 6200 Series and 6324 Fabric Interconnects

Unified ports are ports on the Cisco UCS 6200 Series and 6324 Fabric Interconnects that you can configure to carry either Ethernet or Fibre Channel traffic. A Cisco UCS domain cannot use these un-reserved 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.

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.

Data Traffic Interruption from Port Mode Changing

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



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

Impact of Port Mode on an Expansion Module

After you make port mode changes on an expansion module, the module reboots. All traffic through port on the expansion module is interrupted for approximately 1 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.

When 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.

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 6200 series fabric interconnect with Cisco UCS Manager, version 2.0.

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.

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 a member port for an Ethernet or FCoE port channel, Cisco UCS Manager automatically makes this port 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.

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.

Expansions modules are not supported with Cisco UCS Mini.

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> <p>fc-uplink For configuring Fibre Channel uplink ports and Fibre Channel uplink port channels.</p> |

| | Command or Action | Purpose |
|---------------|--|---|
| 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 <i>Guidelines and Recommendations for Configuring Unified Ports</i> 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

Configuring Breakout Ports

Cisco UCS 6300 Fabric Interconnect Ethernet Breakout Ports

You can configure Cisco UCS 6300 Fabric Interconnects with 40 GB Ethernet ports as four breakout 10 GB ports using a supported breakout cable. The configuration requires a Small Form-Factor Pluggable adapter (SPF) that has one 40GB QSFP+ on one end to connect to the Fabric Interconnect and four 10 GB ports to connect to different end points supporting 10 GB connectivity. For more detailed information about the Cisco UCS 6300 Fabric Interconnects, see the *UCS Manager Getting Started Guide*.



Caution

Configuring breakout ports requires rebooting the Fabric Interconnect. Any existing configuration on a port is erased. It is recommended to break out all required ports in a single transaction.

Once you configure a breakout port, you can configure each 10 GB sub-port as server, uplink, FCoE uplink, FCoE storage or appliance as required.

The following table summarizes the constraints for breakout functionality for Cisco UCS 6300 series fabric interconnects:

| Fabric Interconnect | Breakout Configurable Ports | Normal Ports with no Breakout Support |
|---------------------|-----------------------------|---|
| UCS-FI-6332 | 1-12,15-26 | 13-14,27-32 Note <ul style="list-style-type: none"> • Auto-negotiate behavior is not supported on ports 27–32. • A maximum of four ports are allowed as breakout ports if using QoS jumbo frames. |

| Fabric Interconnect | Breakout Configurable Ports | Normal Ports with no Breakout Support |
|---------------------|-----------------------------|--|
| UCS-FI-6332-16UP | 17-34 | 1-16,35-40 Note <ul style="list-style-type: none"> • Auto-negotiate behavior is not supported on ports 35-40. • A maximum of four ports are allowed as breakout ports if using QoS jumbo frames. |



Caution Configuring breakout ports requires rebooting the Fabric Interconnect. Any existing configuration on a port is erased. It is recommended to break out all required ports in a single transaction.

Configuring a 10G port with QSA Adapter on Cisco UCS FI 6332

When a port on UCS FI-6332 is operating at the default 40G port speed, UCS Manager does not let you choose port speeds of 1GB or 10GB. To use a 40G port on UCS FI-6332 as a 10 GB port with a QSFP+Adapter (QSA) transceiver on the other end, you must configure it in the breakout mode.



Note When you try to change port speeds to 1GB or 10GB, UCS Manager displays a prompt to configure the port in breakout mode. Once you configure a breakout port, you can configure each 10GB sub-port as server, uplink, FCoE uplink, FCoE storage or appliance as required.

When you break out the port, only the first lane becomes usable as a 10G interface. If you use a breakout cable to split a single port into four 10G ports, and configure the ports in breakout mode, you can use all lanes as 10 GB ports.

Procedure

- Step 1** Configure breakout feature on the port that you want to use as a 10GB port on the Cisco UCS FI 6332. For more information about configuring the break out feature, see Configuring Fabric Interconnect Ethernet Breakout Ports.
- Caution** Configuring breakout ports requires rebooting the Fabric Interconnect. Any existing configuration on a port is erased. It is recommended to break out all required ports in a single transaction.
- Step 2** In Cisco UCS Manager, the first tuple interface is enabled after the QSA transceiver is plugged into the FI port. You can configure this interface based on your requirements.

The resulting ports after a break out of the 40G port are numbered using a 3-tuple naming convention. For example, the supported breakout ports are numbered Br-Ethernet 1/25/1, Br-Ethernet 1/25/2, Br-Ethernet 1/25/3, and Br-Ethernet 1/25/4, and only the first port becomes usable as a 10 GB port.

Configuring Multiple Breakout Ports

On a UCS 6300 Fabric Interconnect, you can specify a 40 Gigabit Ethernet port and create four 10 Gigabit Ethernet unconfigured breakout ports. Because configuring breakout on a port causes the reboot of the Fabric Interconnect, we recommend that you breakout all required ports in a single transaction.

Before you begin

Before configuring a breakout port, view the port status using the **show port** command.

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A # scope cabling | Enters the cabling mode. |
| Step 2 | UCS-A /cabling # scope fabric {a b} | Enters cabling fabric mode for the specified fabric. |
| Step 3 | UCS-A /cabling/fabric # create breakout slot-id port-id | Creates the breakout port on the selected slot and port. |
| Step 4 | UCS-A /cabling/fabric/breakout* # up | Returns you to fabric mode. Repeat steps 3 and 4 for each breakout port. |
| Step 5 | UCS-A /cabling/fabric/breakout* # commit-buffer | Commits the transaction to the server. |

Example

The following example creates breakout ports 1/1 through 1/4 on a UCS 6300 Fabric Interconnect and commits the transaction:

```
UCS-A# scope cabling
UCS-A /cabling # scope fabric a
UCS-A /cabling/fabric # create breakout 1 1
Warning: Port breakout create action reboots FI and any existing configurations on 40G port
will be erased.!
UCS-A /cabling/fabric/breakout* # up
UCS-A /cabling/fabric* # create breakout 1 2
Warning: Port breakout create action reboots FI and any existing configurations on 40G port
will be erased.!
UCS-A /cabling/fabric/breakout* # up
UCS-A /cabling/fabric* # create breakout 1 3
Warning: Port breakout create action reboots FI and any existing configurations on 40G port
will be erased.!
UCSM--A /cabling/fabric/breakout* # up
UCSM-shiva-a-A /cabling/fabric* # create breakout 1 4
Warning: Port breakout create action reboots FI and any existing configurations on 40G port
```

```
will be erased.!\nUCSM--A /cabling/fabric/breakout* # commit-buffer
```

What to do next

Verify that you created breakout ports on the fabric interconnect and on the NXOS switch. On the fabric interconnect use the **show breakout** command in cabling fabric mode for the specified fabric. In NXOS, use the **show interface brief** command.

Configuring a Breakout Ethernet Uplink 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. |
| Step 3 | UCS-A /eth-uplink/fabric # create aggr-interface slot-numaggregate port-num | Creates the interface for the specified aggregate (main) Ethernet uplink port. |
| Step 4 | UCS-A /eth-uplink/fabric/aggr-interface* # create br-interface breakout-port-num | Creates an interface for the specified breakout Ethernet uplink port. |
| Step 5 | UCS-A /eth-uplink/fabric/aggr-interface/br-interface # commit-buffer | Commits the transaction to the server. |

Example

The following example shows how to create an interface for breakout Ethernet uplink port 1 of the aggregate port 21 on slot 1 of fabric A:

```
UCS-A# scope eth-uplink\nUCS-A /eth-uplink # scope fabric a\nUCS-A /eth-uplink/fabric # enter aggr-interface 1 21\nUCS-A /eth-uplink/fabric/aggr-interface # create br-interface 1\nUCS-A /eth-uplink/fabric/aggr-interface/br-interface*# commit-buffer
```

Configuring a Breakout Ethernet Uplink Port Channel Member

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 mode for the specified fabric. |
| Step 3 | UCS-A# /eth-uplink/fabric # scope fcoe-port-channel fcoe-port-channel | Enters port channel for the specified FCoE uplink port. |
| Step 4 | UCS-A /eth-uplink/fabric/port-channel/fcoe-port-channel # enter aggr-interface slot-id port-id | Enters the interface for the specified aggregate(main) FCoE uplink port. |
| Step 5 | UCS-A /eth-uplink/fabric/port-channel/member-aggr-port # create br-member-port breakout-port-num | Creates the FCoE uplink port channel member. |
| Step 6 | UCS-A /eth-uplink/fabric/port-channel/member-aggr-port/br-member-port # commit-buffer Example: The following example creates an Ethernet uplink port channel member for an Ethernet port on port 2, and commits the transaction: UCS-A# scope eth-storage UCS-A /eth-uplink # scope fabric a UCS-A /eth-uplink/fabric # scope fcoe-port-channel 51 UCS-A /eth-uplink/fabric/port-channel/member-aggr-port # create br-member-port 2 UCS-A /eth-uplink/fabric/port-channel/member-aggr-port/br-member-port* # commit-buffer | Commits the transaction to the server. |

Configuring Ethernet Uplink Breakout Port as a Pin Group Target

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A# /eth-uplink/pin-group # enter pin-group pin-group-name | Enters the pin group with the specified name. |
| Step 3 | UCS-A# /eth-uplink/pin-group # set target {a b} breakout-ports slot-num aggregate-port-num breakout-port-num | Sets the selected target as the breakout port. |
| Step 4 | UCS-A # /eth-uplink/pin-group # commit-buffer Example: | Commits the transaction to the server. |

| | Command or Action | Purpose |
|--|---|---------|
| | <p>The following example sets the pin group target to breakout port 2 of the aggregate port 1 on slot 1, on fabric A , and commits the transaction:</p> <pre>UCS-A# scope eth-uplink UCS-A /eth-uplink # enter pin-group test UCS-A /eth-uplink/pin-group # set target a breakout-port 1 1 2 UCS-A /eth-uplink/pin-group* # commit-buffer</pre> | |

Configuring Breakout Appliance Ports

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 # enter aggr-interface slot-numaggregate-port-num | Enters the interface for the specified aggregate(main) appliance port. |
| Step 4 | UCS-A# /eth-storage/fabric/port-channel/member-aggr-port # create br -interface breakout-port-num | Creates an interface for the specified breakout appliance port. |
| Step 5 | UCS-A# /eth-storage/fabric/port-channel/member-aggr-port/member-port # commit-buffer | Commits the transaction to the server. |
| | <p>Example:</p> <p>The following example creates an interface for an appliance port 1 of the aggregate port 20 on slot 1 of fabric B, and commits the transaction:</p> <pre>UCS-A# scope eth-storage UCS-A /eth-storage # scope fabric a UCS-A /eth-storage/fabric # enter aggr-interface 1 20 UCS-A /eth-storage/fabric/aggr-interface # create br-interface 1 UCS-A /eth-storage/fabric/aggr-interface/br-interface* # commit-buffer</pre> | |

Configuring a Breakout Appliance Port Channel Member

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 # scope port-channel <i>port-channel-num</i> | Enters Ethernet storage mode for the specified port-channel. |
| Step 4 | UCS-A# /eth-storage/fabric # enter aggr-interface <i>slot-num</i> aggr-interface <i>slot-num</i> aggregate-port-num | Enters the interface for the specified aggregate(main) appliance port. |
| Step 5 | UCS-A /eth-storage/fabric/port-channel # enter member-aggr-port <i>slot-id</i> port-id | Enters the appliance port channel member port. |
| Step 6 | UCS-A /eth-storage/fabric/port-channel/member-aggr-port # create br-member-port <i>breakout-port-num</i> | Creates the appliance port channel member. |
| Step 7 | UCS-A /eth-storage/fabric/port-channel/member-aggr-port/ br-member-port # commit-buffer Example: The following example creates an appliance port channel member for an appliance port 2, and commits the transaction: UCS-A# scope eth-storage UCS-A /eth-storage # scope fabric a UCS-A /eth-storage/fabric # scope port-channel 21 UCS-A /eth-storage/fabric/port-channel # enter member-aggr-port 1 2 UCS-A /eth-storage/fabric/port-channel/member-aggr-port # create br-member-port 2 UCS-A /eth-storage/fabric/port-channel/member-aggr-port/ br-member-port* # commit-buffer | Commits the transaction to the server. |

Configuring Breakout FCoE Storage Ports

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. |

| | Command or Action | Purpose |
|---------------|---|---|
| Step 3 | UCS-A# /fc-storage/fabric enter aggr-interface slot-numaggregate port-num | Enter the interface for the specified aggregate(main) Fibre Channel storage port. |
| Step 4 | UCS-A# /fc-storage/fabric/aggr-interface # create br-interface br-fcoe breakout-port-num | Creates an interface for the specified breakout Fibre Channel storage port. |
| Step 5 | UCS-A# /fc-storage/fabric/aggr-interface/br-interface/br-fcoe # commit-buffer Example: The following example creates an interface for a breakout Fibre Channel storage port 1 of the aggregate port 21 on slot 1 of fabric a, and commits the transaction: UCS-A# scope fc-storage UCS-A /fc-storage # scope fabric a UCS-A /fc-storage/fabric # enter aggr-interface 1 21 UCS-A /fc-storage/fabric/aggr-interface # create br-interface 1 UCS-A /eth-uplink/fabric/aggr-interface/br-interface/br-fcoe # commit-buffer | Commits the transaction to the server. |

Configuring a Breakout 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 enter aggr-interface slot-numaggregate port-num | Enters interface for the specified aggregate(main) FCoE uplink port. |
| Step 4 | UCS-A# /fc-uplink/fabric/aggr-interface # create br-fcoeinterface breakout-port-num | Creates an interface for the specified breakout FCoE uplink port. |
| Step 5 | UCS-A# /fc-uplink/fabric/aggr-interface/br-fcoeinterface # commit-buffer Example: The following example shows how to create an interface for breakout FCoE uplink port 1 of the aggregate port 20 on slot 1 of fabric A: UCS-A# scope eth-uplink | Commits the transaction to the server. |

| | Command or Action | Purpose |
|--|---|---------|
| | <pre>UCS-A /fc-uplink # scope fabric a UCS-A /fc-uplink/fabric # enter aggr-interface 1 20 UCS-A /fc-uplink/fabric/aggr-interface # create br-fcoeinterface 1 UCS-A /fc-uplink/fabric/aggr-interface/br-fcoeinterface # commit-buffer</pre> | |

Configuring an FCoE Port Channel Member

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope fc-uplink | Enters Ethernet storage mode. |
| Step 2 | UCS-A# /fc-uplink # scope fabric {a b} | |
| Step 3 | UCS-A# /fc-uplink/fabric # scope fcoe-port-channel <i>fcoe-port-num</i> | |
| Step 4 | UCS-A /fc-uplink/fabric/port-channel # enter aggr-interface <i>slot-num</i> <i>port-num</i> aggregate-port-num | Enters the FCoE port channel member port. |
| Step 5 | UCS-A /fc-uplink/fabric/port-channel/member-aggr-port # create br-member-port <i>breakout-port-num</i> | Creates the FCoE port channel member for the specified breakout port. |
| Step 6 | UCS-A /fc-uplink/fabric/port-channel/member-aggr-port/br-member-port # commit-buffer Example: The following example creates a breakout FCoE port channel member port 4 on aggregate port 21, and commits the transaction: <pre>UCS-A# scope eth-storage UCS-A /fc-uplink # scope fabric a UCS-A /fc-uplink/fabric # scope port-channel 51 UCS-A /fc-uplink/fabric/port-channel # enter member-aggr-port 1 21 UCS-A /fc-uplink/fabric/port-channel/member-aggr-port # create br-member-port 4 UCS-A /fc-uplink/fabric/port-channel/member-aggr-port/br-member-port* # commit-buffer</pre> | Commits the transaction to the server. |

Configuring a Breakout VLAN Member Port

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | USA-A# scope eth-uplink | Enters Ethernet uplink mode for the specified fabric. |
| Step 2 | USA-A /eth-uplink # scope vlan id | Enters VLAN mode. |
| Step 3 | USA-A /eth-uplink/vlan # enter member-aggr-port {a b} slot-id port id | Enters an interface for the specified fabric, main aggregate port, and subport. breakout VLAN member port. |
| Step 4 | USA-A /eth-uplink/vlan/member-aggr-port # create br-member-port breakout-port-name | Creates an interface for the specified breakout VLAN member port. |
| Step 5 | USA-A /eth-uplink/vlan/member-aggr-port/br-member-port # commit-buffer Example: The following example creates an interface for a VLAN member on the aggregate port 4 on slot 1 of breakout Ethernet uplink port 1, and commits the transaction: <pre> USA-A# scope eth-uplink USA-A /eth-uplink # scope vlan id USA-A /eth-uplink/vlan # enter member-aggr-port a 1 1 USA-A /eth-uplink/vlan/member-aggr-port* # create br-member-port 4 USA-A /eth-uplink/vlan/member-aggr-port/br-member-port* # commit-buffer </pre> | Commits the transaction to the server. |

What to do next

Verify that you created the breakout VLAN Member port using the **show** command.

Modifying a Breakout Port

The following table describes how to modify the supported breakout ports.

| Breakout Port Type | Scope | CLI Location From Which To Modify | Modify Options |
|-------------------------------------|------------|---|--|
| Ethernet Uplink | eth-uplink | UCS-A eth-uplink/fabric/agg-interface br-interface # create | mon-src — Creates a monitor source session. |
| | | UCS-A eth-uplink/fabric/agg-interface br-interface # set | eth-link-profile — Sets the Ethernet Link profile name. flow-control-policy — Sets the flow control policy that configures the receive and send flow control parameters for the LAN and Ethernet uplink ports. speed — Sets the speed for an Ethernet uplink port. user-label — Assigns an identifying label to the Ethernet Uplink port. |
| | | UCS-A eth-uplink/fabric/agg-interface br-interface # | disable — Disables the aggregate interface for the Ethernet Uplink breakout port. enable — Enables the aggregate interface for the Ethernet Uplink breakout port. |
| Ethernet Uplink port-channel member | fc-storage | UCS-A eth-uplink/port-channel/agg-interface member # set | eth-link-profile — Sets the Ethernet Link profile name. |
| | | UCS-A eth-uplink/port-channel/agg-interface member # | disable — Disables the aggregate interface for the breakout Ethernet Uplink port-channel member. enable — Enables the aggregate interface for the breakout Ethernet Uplink port-channel member. |

| Breakout Port Type | Scope | CLI Location From Which To Modify | Modify Options |
|---------------------------------|------------|---|--|
| FCoE Uplink | fc-uplink | UCS-A /uplink/fabric/agg-interface/fcoe-interface # create | mon-src — Creates a monitor source session. |
| | | UCS-A /uplink/fabric/agg-interface/fcoe-interface # set | eth-link-profile — Sets the Ethernet Link profile name. user-label — Assigns an identifying label to the FCoE uplink breakout port. |
| | | UCS-A /uplink/fabric/agg-interface/fcoe-interface # | disable — Disables the aggregate interface for the FCoE uplink breakout port. enable — Enables the aggregate interface for the FCoE uplink breakout port. |
| FCoE Uplink port-channel member | eth-uplink | UCS-A /eth/centralswitch/fabric/port # set | eth-link-profile — Sets the Ethernet Link profile name. |
| | | A /eth/centralswitch/fabric/port # | disable — Disables the aggregate interface for the breakout FCoE uplink port-channel member. enable — Enables the aggregate interface for the breakout FCoE uplink port-channel member. |
| FCoE Storage port | fc-storage | UCS-A /fc-storage/fabric/agg-interface/br-fcoe # create | mon-src — Creates a monitor source session. |
| | | UCS-A /fc-storage/fabric/agg-interface/br-fcoe # set | user-label — Assigns an identifying label to the server. |
| | | UCS-A /fc-storage/fabric/agg-interface/br-fcoe # | disable — Disables the aggregate interface for the breakout FCoE Storage port enable — Enables the aggregate interface for the breakout FCoE Storage port. |

| Breakout Port Type | Scope | CLI Location From Which To Modify | Modify Options |
|-------------------------------|--------------------|---|---|
| Appliance Port | eth-storage | UCS-A /eth-storage/fabric/agg-interface- # set | <p>adminsPEED— Sets the speed for a fabric interface.</p> <p>flowctrlpolicy—Sets the flow control policy that configures the receive and send flow control parameters for the appliance ports.</p> <p>nw-control-policy — Creates a network control policy for the appliance port.</p> <p>pinggroupname— Sets the pin group name for the fabric interface.</p> <p>portmode— Sets the appliance port mode.</p> <p>prio — Sets the QoS (Quality of Service) priority level.</p> <p>user-label— Assigns an identifying label to the appliance port.</p> |
| | | UCS-A /eth-storage/fabric/agg-interface- # create | <p>eth-target — Creates the Ethernet target endpoint.</p> <p>mon-src— Creates a monitor source session.</p> |
| | | UCS-A /eth-storage/fabric/agg-interface- # | <p>disable— Disables the aggregate interface for the appliance breakout port.</p> <p>enable—Enables the aggregate interface for the appliance breakout port.</p> |
| Appliance port-channel member | eth-storage | UCS-A /eth-storage/port-channel- # | <p>disable— Disables the aggregate interface for the breakout appliance port-channel member.</p> <p>enable—Enables the aggregate interface for the breakout appliance port-channel member.</p> |

| Breakout Port Type | Scope | CLI Location From Which To Modify | Modify Options |
|--|-----------------|---|---|
| VLAN Member | eth-uplink | A /eth-uplink/aggport/br- # set | isnative — Marks a member-port as a native VLAN. |
| Pin Group - Pin Target | eth-uplink | N/A | N/A |
| SPAN (Traffic Monitoring) Destination Port | eth-traffic-mon | A /eth-traffic-mon/aggport/br- # set | speed — Sets the speed for the SPAN (Traffic Monitoring) destination 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. |
| Step 3 | UCS-A /eth-uplink/fabric # scope aggr-interface port-number port-id . | Enters the interface for the specified aggregate(main) Ethernet uplink port. |
| Step 4 | UCS-A /eth-uplink/fabric/aggr-interface # scope br-interface port-id . | Enters the breakout Ethernet port for the specified port number. |
| Step 5 | UCS-A /eth-uplink/fabric/aggr-interface/br-interface # create mon-src . Example: The following example shows how to modify a Ethernet uplink port as a monitor source in breakout port 1 of the aggregate (main) interface in port 1 with an ID of 21. UCS-A# scope eth-uplink UCS-A /eth-uplink # scope fabric a UCS-A /eth-uplink/fabric # scope aggr-interface 1 21 UCS-A /eth-uplink/fabric/aggr-interface # scope br-interface 1 UCS-A /eth-uplink/fabric/aggr-interface/br-interface # create UCS-A /eth-uplink/fabric/aggr-interface/br-interface # create mon-src | Modifies the interface as a monitoring source. |

Modifying the Breakout Ethernet Uplink Port Speed and User Label

Enabling or Disabling a Breakout Ethernet Uplink Port

```
pranspat-3gfi-A /eth-uplink/fabric/aggr-interface/br-interface # set
eth-link-profile      Ethernet Link Profile name
flow-control-policy   flow control policy
speed                 Speed
user-label            User Label

pranspat-3gfi-A /eth-uplink/fabric/aggr-interface/br-interface #
disable               Disables services
enable                Enables services
```

Un-configuring Breakout Ports

If you have a breakout on port 2 in slot 1, you can un-configure the breakout port.

Before you begin

You can use the **show port** command to list the ports for the Fabric Interconnect (FI), and select the port that you want to breakout.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# / fabric-interconnect # show port Example: The following example lists the ports. <pre>Slot Aggr Port Port Oper State Mac Role Xcvr ----- 1 0 1 Link Down 84:B8:02:CA:37:56 Network 1000base T 1 2 1 Sfp Not Present 84:B8:02:CA:37:57 Unknown N/A 1 2 2 Sfp Not Present 84:B8:02:CA:37:57 Unknown N/A 1 2 3 Sfp Not Present 84:B8:02:CA:37:57 Unknown N/A 1 2 4 Sfp Not Present 84:B8:02:CA:37:57 Unknown N/A 1 0 3 Sfp Not Present 84:B8:02:CA:37:58 Unknown N/A</pre> | Displays the ports for the Fabric Interconnect. |
| Step 2 | UCS-A# scope cabling | Enters the cabling mode. |
| Step 3 | UCS-A# /cabling # scope fabric {a b} | Specifies fabric a or b. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 4 | UCS-A #/ cabling # delete breakout {1 2 | Warning The breakout port delete action reboots the FI, and any existing configurations on 10G ports are erased. |
| Step 5 | UCS-A /cabling/fabric/breakout* # commit . | Commits the transaction to the system configuration. The FI reboots. After the FI is back up, port 2 in slot one appears as a 40G port. |

What to do next

You can use the **show port** to view the unconfigured breakout ports.

Deleting Breakout Ports

You can delete 10 Gig Ethernet breakout ports. Use the **br-interface** or **br-member-port** scopes to select breakout sub-ports 1-4. You must provide the sub-port id for this scope. For example, **scope br-interface sub_port_id**.

The example described in this topic describes how to delete a breakout Ethernet uplink port. The following table describes how to delete the supported Ethernet breakout ports.

| Breakout Port Type | Scope | CLI Location From Which To Delete |
|-------------------------------------|-------------|--|
| Ethernet Uplink | eth-uplink | UCS-A /eth-uplink/fabric/aggr-interface # delete br-interface number |
| Ethernet Uplink port-channel member | eth-uplink | UCS-A /eth-uplink/fabric/port-channel/aggr-interface # delete br-member-port number |
| FCoE Uplink | fc-uplink | UCS-A /fc-uplink/fabric/aggr-interface # delete br-fcoeinterface number |
| FCoE Uplink port-channel member | eth-uplink | UCS-A /fc-uplink/fabric/fcoe-port-channel/aggr-interface # delete br-member-port number |
| FCoE Storage port | fc-storage | UCS-A /fc-storage/fabric/aggr-interface # delete br-interface br-fcoe number |
| Appliance Port | eth-storage | UCS--A /eth-storage/fabric/port-channel/member-aggr-port # delete br-member-port number |
| Appliance port-channel member | eth-storage | UCS-A /eth-storage/fabric/aggr-interface # delete br-interface number |
| VLAN Member | eth-uplink | UCS-A /eth-uplink/vlan/member-aggr-port # delete br-member-port number |
| Pin Group - Pin Target | eth-uplink | UCS-A /eth-uplink/pin-group # delete target number |

| Breakout Port Type | Scope | CLI Location From Which To Delete |
|---|-----------------|---|
| SPAN (Traffic Monitoring) Destination Port | eth-traffic-mon | UCS-A /eth-traffic-mon/fabric/eth-mon-session/dest-aggr-interface # delete br-dest-interface |

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A# scope eth-uplink | Enters the Ethernet uplink mode. |
| Step 2 | UCS-A# /eth-storage # scope fabric {a b} | Enters Ethernet storage mode for the specified fabric. |
| Step 3 | UCS-A /eth-uplink/fabric # scope port-channel number | Enters Ethernet uplink fabric port channel mode for the specified port channel. |
| Step 4 | UCS-A /eth-uplink/fabric/port-channel/aggr-interface # delete br-member-port number | Deletes the specified breakout port. |
| Step 5 | UCS-A /eth-uplink/fabric/port-channel/aggr-interface # commit-buffer Example: This example deletes an Ethernet Uplink port-channel member in breakout port 1 of the aggregate (main) interface port 1 slot 1. UCS-A# scope eth-uplink UCS-A /eth-uplink # scope fabric a UCS-A /eth-uplink/fabric # scope port-channel 1 UCS-A /eth-uplink/fabric/port-channel # enter aggr-interface 1 1 UCS-A /eth-uplink/fabric/port-channel/aggr-interface # delete br-member-port 1 UCS-A /eth-uplink/fabric/port-channel/aggr-interface* # commit-buffer | Commits the transaction to the server. |

What to do next

Verify that you deleted the specified breakout port using the **show** command.

Cisco UCS Mini Scalability Ports

The Cisco UCS 6324 Fabric Interconnect 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. |
| 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 #
```

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 #
```

Physical and Backplane Ports

Displaying VIF Port Statistics Obtained From the Adaptor

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 vethernet veth-id counters | Displays VIF port statistics that are obtained from the adaptor. |

Example

The following example shows how to display VIF port statistics that are obtained from the adaptor:

```
UCS-A /fabric-interconnect # connect nxos a
UCS-A(nxos)# show interface vethernet 684 counters
```

```
-----
Port                               InOctets                               InUcastPkts
-----
Veth684                             0                                       0

-----
Port                               InMcastPkts                             InBcastPkts
-----
Veth684                             0                                       0

-----
Port                               OutOctets                                OutUcastPkts
-----
Veth684                             0                                       0

-----
Port                               OutMcastPkts                             OutBcastPkts
-----
Veth684                             0                                       0
```

Displaying VIF Port Statistics Obtained From the ASIC

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 platform fwm info lif vethernet veth-id grep frame | Displays VIF-port RX and TX frame statistics obtained from the ASIC. RX statistics are for all type of frames. Tx statistics are only for known unicast frames. |

Example

The following example shows how to display VIF-port RX and TX frame statistics obtained from the ASIC:

```
UCS-A /fabric-interconnect # connect nxos a
UCS-A(nxos)# show platform fwm info lif vethernet 684 | grep frame

vif29 pd: rx frames: 0 tx frames: 0;

UCS-A(nxos)#
```

Displaying VIF Ports That Correspond to NIV Ports

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 platform fwm info lif vethernet veth-id grep niv | Displays VIF ports that correspond to NIV ports. |

Example

The following example shows how to display VIF ports that correspond to NIV ports:

```
UCS-A /fabric-interconnect # connect nxos a
UCS-A(nxos)# show platform fwm info lif vethernet 741 | grep niv

vif20 pd: niv_port_id 0x7000001f (the 0x1F or "31" is the Source/Dest-VP index)
```

Verifying Status of Backplane Ports

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 br | Displays the configuration of the interface, including the speed and status of the backplane ports. |

Example

The following example shows how to verify the status of backplane ports for fabric interconnect A:

```
UCS-A /fabric-interconnect # connect nxos a
UCS-A(nxos)# show interface br
```

```
-----
Ethernet      VLAN   Type Mode   Status Reason                Speed   Port
Interface                                           Ch #
-----
Eth1/1        1      eth  access down  SFP not inserted      40G(D) --
Eth1/2        1      eth  access down  SFP not inserted      40G(D) --
Br-Eth1/3/1   1      eth  access down  Administratively down  10G(D) --
Br-Eth1/3/2   1      eth  access down  Administratively down  10G(D) --
Br-Eth1/3/3   1      eth  access down  Administratively down  10G(D) --
Br-Eth1/3/4   1      eth  access down  Administratively down  10G(D) --
Eth1/4        1      eth  access down  SFP not inserted      40G(D) --
Br-Eth1/5/1   4044   eth  trunk  down  Link not connected     10G(D) --
Br-Eth1/5/2   4044   eth  trunk  down  Link not connected     10G(D) --
Br-Eth1/5/3   4044   eth  trunk  down  Link not connected     10G(D) --
Br-Eth1/5/4   4044   eth  trunk  down  Link not connected     10G(D) --
Eth1/6        1      eth  access down  SFP not inserted      40G(D) --
Eth1/7        1      eth  access down  SFP not inserted      40G(D) --
Eth1/8        1      eth  access down  SFP not inserted      40G(D) --
Eth1/9        1      eth  access down  SFP not inserted      40G(D) --
Eth1/10       1      eth  access down  SFP not inserted      40G(D) --
Eth1/11       1      eth  fabric up    none                   40G(D) --
```

Verifying Status of Backplane Ports

```

Eth1/12      1      eth access down  SFP not inserted      40G(D) --
Eth1/13      1      eth access down  SFP not inserted      40G(D) --
Eth1/14      1      eth access down  SFP not inserted      40G(D) --
Eth1/15      1      eth access down  SFP not inserted      40G(D) --
Eth1/16      1      eth access down  SFP not inserted      40G(D) --
Eth1/17      1      eth access down  SFP not inserted      40G(D) --
Eth1/18      1      eth access down  SFP not inserted      40G(D) --
Eth1/19      1      eth access down  SFP not inserted      40G(D) --
Eth1/20      1      eth access down  SFP not inserted      40G(D) --
Br-Eth1/21/1 1      eth trunk up     none                   10G(D) --
Br-Eth1/21/2 1      eth trunk up     none                   10G(D) --
Br-Eth1/21/3 1      eth trunk down   Link not connected    10G(D) --
Br-Eth1/21/4 1      eth trunk up     none                   10G(D) --
Eth1/22      1      eth access down  SFP not inserted      40G(D) --
Eth1/23      1      eth access down  SFP not inserted      40G(D) --
Eth1/24      1      eth access down  SFP not inserted      40G(D) --
Eth1/25      1      eth access down  SFP not inserted      40G(D) --
Eth1/26      1      eth access down  SFP not inserted      40G(D) --
Eth1/27      1      eth access down  SFP not inserted      40G(D) --
Eth1/28      1      eth access down  SFP not inserted      40G(D) --
Eth1/29      1      eth access down  SFP not inserted      40G(D) --
Eth1/30      1      eth access down  SFP not inserted      40G(D) --
Eth1/31      1      eth access down  SFP not inserted      40G(D) --
Eth1/32      1      eth access down  SFP not inserted      40G(D) --

```

```

-----
Port-channel VLAN      Type Mode      Status Reason              Speed      Protocol
Interface
-----
Po1285      1      eth vntag  up      none                a-10G(D)  none
Po1286      1      eth vntag  up      none                a-10G(D)  none
Po1287      1      eth vntag  up      none                a-10G(D)  none
Po1288      1      eth vntag  up      none                a-10G(D)  none
Po1289      1      eth vntag  up      none                a-10G(D)  none

```

```

-----
Port      VRF              Status IP Address              Speed      MTU
-----
mgmt0    --              down   10.197.157.252         --         1500

```

```

-----
Vethernet      VLAN      Type Mode      Status Reason              Speed
-----
Veth691      4047     virt trunk  down   nonParticipating    auto
Veth692      4047     virt trunk  up     none                 auto
Veth693      1        virt trunk  down   nonParticipating    auto
Veth695      1        virt trunk  up     none                 auto
Veth699      1        virt trunk  up     none                 auto

```

```

-----
Interface Secondary VLAN(Type)              Status Reason
-----
Vlan1      --              down   Administratively down

```

```

-----
Ethernet      VLAN      Type Mode      Status Reason              Speed      Port
Interface
-----
Eth1/1/1      1        eth vntag  up     none                10G(D) 1286
Eth1/1/2      1        eth access down  Administratively down 10G(D) --
Eth1/1/3      1        eth vntag  up     none                10G(D) 1286
Eth1/1/4      1        eth access down  Administratively down 10G(D) --
Eth1/1/5      1        eth vntag  up     none                10G(D) 1287
Eth1/1/6      1        eth access down  Administratively down 10G(D) --

```

| | | | | | | | |
|-----------|------|-----|--------|------|-----------------------|----------|------|
| Eth1/1/7 | 1 | eth | vntag | up | none | 10G (D) | 1287 |
| Eth1/1/8 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/9 | 1 | eth | vntag | up | none | 10G (D) | 1289 |
| Eth1/1/10 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/11 | 1 | eth | vntag | up | none | 10G (D) | 1289 |
| Eth1/1/12 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/13 | 1 | eth | vntag | up | none | 10G (D) | 1285 |
| Eth1/1/14 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/15 | 1 | eth | vntag | up | none | 10G (D) | 1285 |
| Eth1/1/16 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/17 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/18 | 1 | eth | vntag | up | none | 10G (D) | 1288 |
| Eth1/1/19 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/20 | 1 | eth | vntag | up | none | 10G (D) | 1288 |
| Eth1/1/21 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/22 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/23 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/24 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/25 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/26 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/27 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/28 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/29 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/30 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/31 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/32 | 1 | eth | access | down | Administratively down | 10G (D) | -- |
| Eth1/1/33 | 4044 | eth | trunk | up | none | 1000 (D) | -- |

Server Ports

Automatic Configuration of Fabric Interconnect Server Ports

Starting with Cisco UCS Manager release 3.1(3), you can automatically configure the fabric interconnect server ports. The server **Port Auto-Discovery Policy** determines how the system reacts when a new rack server, chassis, or FEX is added. By enabling this policy, Cisco UCS Manager automatically determines the type of device connected to the switch port and configures the switch port accordingly.



Note If you do not want a Cisco UCS C-Series appliance to be UCS Managed, pre-configure the appliance ports before connecting VIC ports to the Cisco UCS fabric interconnect.

Automatically Configuring Server Ports

Procedure

- Step 1** UCS-A# **scope org/**
Enters the root organization mode.

- Step 2** UCS-A / org# **scope por**
Enters organization port discovery policy mode.
- Step 3** UCS-A / org / port-disc-policy# **set descr**
Provides a description for the port discovery policy.
- Step 4** UCS-A / org / port-disc-policy# **set server-auto-disc**
Enables port auto-discovery.

Note By default `server-auto-disc` is disabled. Port auto-discovery is triggered by enabling `server-auto-disc`.

Example

The following example shows how to enable automatic configuration of fabric interconnect server ports:

```
UCS-A# scope org/
UCS-A /org# scope por
UCS-A / org / port-disc-policy # set descr
UCS-A / org / port-disc-policy # set server-auto-disc
```

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 #

```

Uplink Ethernet Ports

Configuring an Uplink Ethernet Port

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

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 <i>slot-num port-num</i> | Creates an interface for the specified Ethernet uplink port. |
| Step 4 | (Optional) UCS-A /eth-uplink/fabric # set speed {10gbps 1gbps} | Sets the speed for the specified Ethernet uplink port. Note For the 6100 series fabric interconnects, the admin speed is only configurable for the first eight ports on a 20-port fabric interconnect and the first 16 ports on a 40-port fabric interconnect. |
| Step 5 | UCS-A /eth-uplink/fabric # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create an interface for Ethernet uplink port 3 on slot 2 of fabric B, set the speed to 10 gbps, and commit the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric b
UCS-A /eth-uplink/fabric # create interface 2 3
UCS-A /eth-uplink/fabric # set speed 10gbps
UCS-A /eth-uplink/fabric* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Unconfiguring an Uplink Ethernet 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. |
| Step 3 | UCS-A /eth-uplink/fabric # delete interface <i>slot-num port-num</i> | Deletes the interface for the specified Ethernet uplink port. |
| Step 4 | UCS-A /eth-uplink/fabric # commit-buffer | Commits the transaction to the system configuration. |

Example

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

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

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} | Specifies whether the port mode is access or trunk. By default, the mode is set to trunk. |

| | Command or Action | Purpose |
|---------------|---|---|
| | | <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 <i>pin-group name</i> | 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 <i>sys-class-name</i> | <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 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. |
| Step 3 | UCS-A /eth-storage/fabric # scope interface slot-id port-id | 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 eth-target name | Specifies the name for the specified MAC address target. |
| Step 5 | UCS-A /eth-storage/fabric/interface/eth-target # set mac-address mac-address | 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. |
| 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 pubnname 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 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. 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 slot-num port-num | 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 8 on slot 1 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 #
```

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 8 on slot 1 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 #
```

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 #
```

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 a 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. |
| 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. |
| 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 #
```

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
```

Uplink Ethernet Port Channels

An uplink Ethernet port channel allows you to group several physical uplink Ethernet ports (link aggregation) to create one logical Ethernet link to provide fault-tolerance and high-speed connectivity. In Cisco UCS Manager, you create a port channel first and then add uplink Ethernet ports to the port channel. You can add up to 16 uplink Ethernet ports to a port channel.



Important

The state of a configured port changes to unconfigured in the following scenarios:

- The port is deleted or removed from a port channel. The port channel can be of any type, such as, uplink or storage.
- A port channel is deleted.



Note

Cisco UCS uses Link Aggregation Control Protocol (LACP), not Port Aggregation Protocol (PAgP), to group the uplink Ethernet ports into a port channel. If the ports on the upstream switch are not configured for LACP, the fabric interconnects treat all ports in an uplink Ethernet port channel as individual ports, and therefore forward packets.

Configuring an Uplink Ethernet Port Channel

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. |
| Step 3 | UCS-A /eth-uplink/fabric # create port-channel port-num | Creates a port channel on the specified Ethernet uplink port, and enters Ethernet uplink fabric port channel mode. |
| Step 4 | (Optional) UCS-A /eth-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 /eth-uplink/fabric/port-channel # set name port-chan-name | Specifies the name for the port channel. |
| Step 6 | (Optional) UCS-A /eth-uplink/fabric/port-channel # set flow-control-policy policy-name | Assigns the specified flow control policy to the port channel. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 7 | UCS-A /eth-uplink/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, sets the name to portchan13a, enables the administrative state, assigns the flow control policy named flow-con-pol432 to the port channel, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create port-channel 13
UCS-A /eth-uplink/fabric/port-channel* # enable
UCS-A /eth-uplink/fabric/port-channel* # set name portchan13a
UCS-A /eth-uplink/fabric/port-channel* # set flow-control-policy flow-con-pol432
UCS-A /eth-uplink/fabric/port-channel* # commit-buffer
UCS-A /eth-uplink/fabric/port-channel #
```

Unconfiguring an Uplink Ethernet Port Channel

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. |
| Step 3 | UCS-A /eth-uplink/fabric # delete port-channel port-num | Deletes the port channel on the specified Ethernet uplink port. |
| Step 4 | UCS-A /eth-uplink/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-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # delete port-channel 13
UCS-A /eth-uplink/fabric* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Adding a Member Port to an Uplink Ethernet Port Channel

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. |
| Step 3 | UCS-A /eth-uplink/fabric # scope port-channel port-num | Enters Ethernet uplink fabric port channel mode for the specified port channel. |
| Step 4 | UCS-A /eth-uplink/fabric/port-channel # create member-port slot-num port-num | Creates the specified member port from the port channel and enters Ethernet uplink fabric port channel member port mode. |
| Step 5 | UCS-A /eth-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 the port channel on port 13 of fabric A and commits the transaction.

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

Deleting a Member Port from an Uplink Ethernet Port Channel

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. |
| Step 3 | UCS-A /eth-uplink/fabric # scope port-channel port-num | Enters Ethernet uplink fabric port channel mode for the specified port channel. |
| Step 4 | UCS-A /eth-uplink/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-uplink/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-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope port-channel 13
UCS-A /eth-uplink/fabric/port-channel # delete member-port 1 7
UCS-A /eth-uplink/fabric/port-channel* # commit-buffer
UCS-A /eth-uplink/fabric/port-channel #
```

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. |

| | Command or Action | Purpose |
|----------------|--|--|
| 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. • 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 port-num | Deletes the port channel from the specified Ethernet storage port. |
| 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. |

| | 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 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. |
| 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 NPIV switch is configured with its channel mode as **active**. If both the member port(s) and peer port(s) do not have the same channel mode configured, the port

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

This example shows how to configure channel mode as active:

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

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 #

```

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 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. |

| | Command or Action | Purpose |
|---------------|---|--|
| 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 <i>port-chan-name</i> | Specifies the name for the port channel. |
| 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. |

| | Command or Action | Purpose |
|---------------|--|--|
| 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 #
```

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. |
| 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 #
```

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 | 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 Only active ports collect the network interface port statistics and send the information 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. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 6 | UCS-A/eth-server/stats-threshold-policy/class/property # create threshold-value { <i>above-normal</i> <i>below-normal</i> } { <i>cleared</i> <i>condition</i> <i>critical</i> <i>info</i> <i>major</i> <i>minor</i> <i>warning</i> } | Creates the specified threshold value for the class property and enters the organization statistics threshold policy class property threshold value mode. |
| 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 class-name | Enters the organization statistics threshold policy class mode for the specified statistics threshold policy class. |

| | Command or Action | Purpose |
|---------------|---|--|
| 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. 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 | UCSA/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 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. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 6 | UCS-A/eth-server/stats-threshold-policy/class/property* # set auto-recovery-time <i>time</i> | 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 { <i>a</i> <i>b</i> } | Enters chassis IOM mode for the specified IOM. |
| Step 3 | UCS-A/chassis/iom # scope port-group <i>fabric</i> | 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

```

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

The 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 #

```

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:
  l2_da: 1 l2_sa: 1 l2_vlan: 0
  l3_da: 1 l3_sa: 1
  l4_da: 1 l4_sa: 1
  FCoE l2_da: 1 l2_sa: 1 l2_vlan: 0
  FCoE l3_did: 1 l3_sid: 1 l3_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 1: 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.

Configuring a Fabric Port Channel

Procedure

- Step 1** To include all links from the IOM to the fabric interconnect in a fabric port channel during chassis discovery, set the link grouping preference in the chassis discovery policy to port channel.
- Step 2** To include links from individual chassis in a fabric port channel during chassis discovery, set the link grouping preference in the chassis connectivity policy to port channel.

Step 3 After chassis discovery, enable or disable additional fabric port channel member ports.

What to do next

To add or remove chassis links from a fabric port channel after making a change to the chassis discovery policy or the chassis connectivity policy, reacknowledge the chassis. Chassis reacknowledgement is not required to enable or disable chassis member ports from a fabric port channel

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. |
| 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. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 3 | UCS-A /eth-server/fabric # scope fabric-port-channel <i>port-chan-id</i> | 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 <i>slot-id port-id</i> | 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 5

VLANs

- [Named VLANs, on page 93](#)
- [Private VLANs, on page 94](#)
- [VLAN Port Limitations, on page 95](#)
- [Configuring Named VLANs, on page 97](#)
- [Configuring Private VLANs, on page 102](#)
- [Community VLANs, on page 109](#)
- [Viewing the VLAN Port Count, on page 112](#)
- [VLAN Port Count Optimization, on page 113](#)
- [VLAN Groups, on page 115](#)
- [VLAN Permissions, on page 118](#)

Named VLANs

A named VLAN creates a connection to a specific external LAN. The VLAN isolates traffic to that external LAN, including broadcast traffic.

The name that you assign to a VLAN ID adds a layer of abstraction that allows you to globally update all servers associated with service profiles that use the named VLAN. You do not need to reconfigure the servers individually to maintain communication with the external LAN.

You can create more than one named VLAN with the same VLAN ID. For example, if servers that host business services for HR and Finance need to access the same external LAN, you can create VLANs named HR and Finance with the same VLAN ID. Then, if the network is reconfigured and Finance is assigned to a different LAN, you only have to change the VLAN ID for the named VLAN for Finance.

In a cluster configuration, you can configure a named VLAN to be accessible only to one fabric interconnect or to both fabric interconnects.

Guidelines for VLAN IDs



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

The VLAN name is case sensitive.

Private VLANs

A private VLAN (PVLAN) partitions the Ethernet broadcast domain of a VLAN into subdomains, and allows you to isolate some ports. Each subdomain in a PVLAN includes a primary VLAN and one or more secondary VLANs. All secondary VLANs in a PVLAN must share the same primary VLAN. The secondary VLAN ID differentiates one subdomain from another.

Isolated and Community VLANs

All secondary VLANs in a Cisco UCS domain can be Isolated or Community VLANs.



Note

You cannot configure an isolated VLAN to use with a regular VLAN.

Ports on Isolated VLANs

Communications on an isolated VLAN can only use the associated port in the primary VLAN. These ports are isolated ports and are not configurable in Cisco UCS Manager. A primary VLAN can have only one isolated VLAN, but multiple isolated ports on the same isolated VLAN are allowed. These isolated ports cannot communicate with each other. The isolated ports can communicate only with a regular trunk port or promiscuous port that allows the isolated VLAN.

An isolated port is a host port that belongs to an isolated secondary VLAN. This port has complete isolation from other ports within the same private VLAN domain. PVLANS block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports. You can have more than one isolated port in a specified isolated VLAN. Each port is completely isolated from all other ports in the isolated VLAN.

Guidelines for Uplink Ports

When you create PVLANS, use the following guidelines:

- The uplink Ethernet port channel cannot be in promiscuous mode.
- Each primary VLAN can have only one isolated VLAN.
- VIFs on VNTAG adapters can have only one isolated VLAN.

Guidelines for VLAN IDs



Note You cannot create VLANs with IDs from 3915 to 4042. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

The VLAN name is case sensitive.

VLAN Port Limitations

Cisco UCS Manager limits the number of VLAN port instances that you can configure under border and server domains on a fabric interconnect.

Types of Ports Included in the VLAN Port Count

The following types of ports are counted in the VLAN port calculation:

- Border uplink Ethernet ports
- Border uplink Ether-channel member ports
- FCoE ports in a SAN cloud
- Ethernet ports in a NAS cloud
- Static and dynamic vNICs created through service profiles
- VM vNICs created as part of a port profile in a hypervisor in hypervisor domain

Based on the number of VLANs configured for these ports, Cisco UCS Manager tracks the cumulative count of VLAN port instances and enforces the VLAN port limit during validation. Cisco UCS Manager reserves some pre-defined VLAN port resources for control traffic. These include management VLANs configured under HIF and NIF ports.

VLAN Port Limit Enforcement

Cisco UCS Manager validates VLAN port availability during the following operations:

- Configuring and unconfiguring border ports and border port channels
- Adding or removing VLANs from a cloud
- Configuring or unconfiguring SAN or NAS ports
- Associating or disassociating service profiles that contain configuration changes
- Configuring or unconfiguring VLANs under vNICs or vHBAs
- Receiving creation or deletion notifications from a VMWare vNIC and from an ESX hypervisor



Note This is outside the control of the Cisco UCS Manager.

- Fabric interconnect reboot
- Cisco UCS Manager upgrade or downgrade

Cisco UCS Manager strictly enforces the VLAN port limit on service profile operations. If Cisco UCS Manager detects that the VLAN port limit is exceeded, the service profile configuration fails during deployment.

Exceeding the VLAN port count in a border domain is less disruptive. When the VLAN port count is exceeded in a border domain Cisco UCS Manager changes the allocation status to Exceeded. To change the status back to **Available**, complete one of the following actions:

- Unconfigure one or more border ports
- Remove VLANs from the LAN cloud
- Unconfigure one or more vNICs or vHBAs

Configuring Named VLANs

Creating a Named VLAN Accessible to Both Fabric Interconnects (Uplink Ethernet Mode)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # create vlan <i>vlan-name</i> <i>vlan-id</i> | Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink VLAN mode. The VLAN name is case sensitive. |
| Step 3 | UCS-A /eth-uplink/fabric/vlan # set sharing { isolated none primary } | Sets the sharing for the specified VLAN. This can be one of the following: <ul style="list-style-type: none"> • isolated —This is a secondary VLAN associated with a primary VLAN. This VLAN is private. • none —This VLAN does not have any secondary or private VLANs. • primary —This VLAN can have one or more secondary VLANs. |
| Step 4 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, sets the sharing to none, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan accounting 2112
UCS-A /eth-uplink/vlan* # set sharing none
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Creating a Named VLAN Accessible to Both Fabric Interconnects (Ethernet Storage Mode)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

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. The VLAN name is case sensitive. |
| Step 3 | UCS-A /eth-storage/vlan # create member-port {a b} <i>slot-id port-id</i> | Creates a member port for the specified VLAN on the specified fabric. |
| Step 4 | UCS-A /eth-storage/vlan/member-port # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, creates a member port on slot 2, port 20, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # create vlan accounting 2112
UCS-A /eth-storage/vlan* # create member-port a 2 20
UCS-A /eth-storage/vlan/member-port* # commit-buffer
UCS-A /eth-storage/vlan/member-port #
```

Creating a Named VLAN Accessible to One Fabric Interconnect (Uplink Ethernet Mode)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

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 interconnect mode for the specified fabric interconnect (A or B). |
| Step 3 | UCS-A /eth-uplink/fabric # create vlan <i>vlan-name vlan-id</i> | Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink fabric interconnect VLAN mode. The VLAN name is case sensitive. |
| Step 4 | UCS-A /eth-uplink/fabric/vlan # set sharing { isolated none primary } | Sets the sharing for the specified VLAN. This can be one of the following: <ul style="list-style-type: none"> • isolated —This is a secondary VLAN associated with a primary VLAN. This VLAN is private. • none —This VLAN does not have any secondary or private VLANs. • primary —This VLAN can have one or more secondary VLANs. |
| Step 5 | UCS-A /eth-uplink/fabric/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, sets the sharing to none, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan finance 3955
UCS-A /eth-uplink/fabric/vlan* # set sharing none
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Creating a Secondary VLAN for a Private VLAN (Accessible to One Fabric Interconnect)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

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 interconnect mode for the specified fabric interconnect (A or B). |
| Step 3 | UCS-A /eth-uplink/fabric # create vlan <i>vlan-name vlan-id</i> | Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink fabric interconnect VLAN mode. The VLAN name is case sensitive. |
| Step 4 | UCS-A /eth-uplink/vlan # set sharing isolated | Sets the VLAN as the secondary VLAN. |
| Step 5 | UCS-A /eth-uplink/vlan # set pubnwnname <i>primary-vlan-name</i> | Specifies the primary VLAN to be associated with this secondary VLAN. |
| Step 6 | UCS-A /eth-uplink/fabric/vlan/member-port # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, makes this VLAN the secondary VLAN, associates the secondary VLAN with the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan finance 3955
UCS-A /eth-uplink/fabric/vlan* # set sharing isolated
UCS-A /eth-uplink/fabric/vlan* # set pubnwnname pvlan1000
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Deleting a Named VLAN

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

If you are deleting a private primary VLAN, ensure that you reassign the secondary VLANs to another working primary VLAN.

Before you begin

Before you delete a VLAN from a fabric interconnect, ensure that the VLAN was removed from all vNICs and vNIC templates.



Note If you delete a VLAN that is assigned to a vNIC or vNIC template, the vNIC might allow that VLAN to flap.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | (Optional) UCS-A /eth-uplink # scope fabric {a b} | Enters Ethernet uplink fabric mode. Use this command when you want to delete a named VLAN only from the specified fabric (a or b). |
| Step 3 | UCS-A /eth-uplink # delete vlan vlan-name | Deletes the specified named VLAN. |
| Step 4 | UCS-A /eth-uplink # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example deletes a named VLAN accessible to both fabric interconnects and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # delete vlan accounting
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```

The following example deletes a named VLAN accessible to one fabric interconnect and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # delete vlan finance
UCS-A /eth-uplink/fabric* # commit-buffer
UCS-A /eth-uplink/fabric #
```

Configuring Private VLANs

Creating a Primary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # create vlan <i>vlan-name</i> <i>vlan-id</i> | Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink VLAN mode. The VLAN name is case sensitive. |
| Step 3 | UCS-A /eth-uplink/vlan # set sharing primary | Sets the VLAN as the primary VLAN. |
| Step 4 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, makes this VLAN the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan accounting 2112
UCS-A /eth-uplink/vlan* # set sharing primary
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Creating a Primary VLAN for a Private VLAN (Accessible to One Fabric Interconnect)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

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 interconnect mode for the specified fabric interconnect. |
| Step 3 | UCS-A /eth-uplink/fabric # create vlan <i>vlan-name vlan-id</i> | Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink fabric interconnect VLAN mode. The VLAN name is case sensitive. |
| Step 4 | UCS-A /eth-uplink/fabric/vlan # set sharing primary | Sets the VLAN as the primary VLAN. |
| Step 5 | UCS-A /eth-uplink/fabric/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, makes this VLAN the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan finance 3955
UCS-A /eth-uplink/fabric/vlan* # set sharing primary
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Creating a Secondary VLAN for a Private VLAN (Accessible to Both Fabric Interconnects)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # create vlan <i>vlan-name</i> <i>vlan-id</i> | Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink VLAN mode. The VLAN name is case sensitive. |
| Step 3 | UCS-A /eth-uplink/vlan # set sharing isolated | Sets the VLAN as the secondary VLAN. |
| Step 4 | UCS-A /eth-uplink/vlan # set pubnwnname <i>primary-vlan-name</i> | Specifies the primary VLAN to be associated with this secondary VLAN. |
| Step 5 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for both fabric interconnects, names the VLAN accounting, assigns the VLAN ID 2112, makes this VLAN the secondary VLAN, associates the secondary VLAN with the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan accounting 2112
UCS-A /eth-uplink/vlan* # set sharing isolated
UCS-A /eth-uplink/vlan* # set pubnwnname pvlan1000
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Creating a Secondary VLAN for a Private VLAN (Accessible to One Fabric Interconnect)



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

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 interconnect mode for the specified fabric interconnect (A or B). |
| Step 3 | UCS-A /eth-uplink/fabric # create vlan <i>vlan-name vlan-id</i> | Creates a named VLAN, specifies the VLAN name and VLAN ID, and enters Ethernet uplink fabric interconnect VLAN mode. The VLAN name is case sensitive. |
| Step 4 | UCS-A /eth-uplink/vlan # set sharing isolated | Sets the VLAN as the secondary VLAN. |
| Step 5 | UCS-A /eth-uplink/vlan # set pubnwnname <i>primary-vlan-name</i> | Specifies the primary VLAN to be associated with this secondary VLAN. |
| Step 6 | UCS-A /eth-uplink/fabric/vlan/member-port # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for fabric interconnect A, names the VLAN finance, assigns the VLAN ID 3955, makes this VLAN the secondary VLAN, associates the secondary VLAN with the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # create vlan finance 3955
UCS-A /eth-uplink/fabric/vlan* # set sharing isolated
UCS-A /eth-uplink/fabric/vlan* # set pubnwnname pvlan1000
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Allowing PVLANS on vNICs

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org / | Enters root organization mode. |
| Step 2 | UCS-A /org # scope service-profile <i>profile-name</i> | Commits the transaction to the system configuration. |
| Step 3 | UCS-A /org/service-profile # scope vnic <i>vnic-name</i> | Enters command mode for the specified vNIC. |
| Step 4 | UCS-A /org/service-profile/vnic # create eth-if <i>community-vlan-name</i> | Allows the community VLAN to access the specified vNIC. |
| Step 5 | UCS-A /org/service-profile/vnic/eth-if* # exit | Exits the interface configuration mode for the specified vNIC. |
| Step 6 | UCS-A /org/service-profile/vnic* # create eth-if <i>primary-vlan-name</i> | Allows the primary VLAN to access the specified vNIC. |
| Step 7 | UCS-A /org/service-profile/vnic # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to assign the community VLAN cVLAN102 and the primary VLAN primaryVLAN100 to the vNIC vnic_1 and commits the transaction.

```
UCS-A# scope org /
UCS-A /org # scope service-profile GSP1
UCS-A /org/service-profile # scope vnic vnic_1
UCS-A /org/service-profile/vnic # create eth-if cVLAN102
UCS-A /org/service-profile/vnic/eth-if* # exit
UCS-A /org/service-profile/vnic # create eth-if primaryVLAN100
UCS-A /org/service-profile/vnic* # commit-buffer
```

Creating a Primary VLAN for a Private VLAN on an Appliance Cloud



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

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. The VLAN name is case sensitive. |
| Step 3 | UCS-A /eth-storage/vlan* # set sharing primary | Sets the VLAN as the primary VLAN. |
| Step 4 | UCS-A /eth-storage/vlan* # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for fabric interconnect A, names the VLAN, assigns the VLAN ID, makes this VLAN the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # create vlan primaryvlan500 500
UCS-A /eth-storage/vlan* # set sharing primary
UCS-A /eth-storage/vlan* # commit-buffer
UCS-A /eth-storage/vlan #
```

Creating a Secondary VLAN for a Private VLAN on an Appliance Cloud



Important

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

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

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. The VLAN name is case sensitive. |
| Step 3 | UCS-A /eth-storage/vlan* # set sharing isolated | Sets the VLAN as the secondary VLAN. |
| Step 4 | UCS-A /eth-storage/vlan* # set pubnwnname <i>primary-vlan-name</i> | Specifies the primary VLAN to be associated with this secondary VLAN. |
| Step 5 | UCS-A /eth-storage/vlan* # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a named VLAN for fabric interconnect A, names the VLAN, assigns the VLAN ID, makes this VLAN the secondary VLAN, associates the secondary VLAN with the primary VLAN, and commits the transaction:

```
UCS-A# scope eth-storage
UCS-A /eth-storage # create vlan isovlan501 501
UCS-A /eth-storage/vlan* # set sharing isolated
UCS-A /eth-storage/vlan* # set pubnwnname primaryvlan500
UCS-A /eth-storage/vlan* # commit-buffer
UCS-A /eth-storage/vlan # #
```

Community VLANs

Cisco UCS Manager supports Community VLANs in UCS Fabric Interconnects. Community ports communicate with each other and with promiscuous ports. Community ports have Layer 2 isolation from all other ports in other communities, or isolated ports within the PVLAN. Broadcasts are transmitted between the community ports associated with the PVLAN only and the other promiscuous ports. A promiscuous port can communicate with all interfaces, including the isolated and community ports within a PVLAN.

Creating a Community VLAN

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink . | Enters Ethernet uplink mode. |
| Step 2 | UCS-A# /eth-uplink/ # create vlan ID . | Create a VLAN with the specified VLAN ID. |
| Step 3 | UCS-A# /eth-uplink/ vlan # set sharing Type . | Specifies the vlan type. |
| Step 4 | UCS-A# /eth-uplink/ vlan # set pubnwnname Name . | Specifies the primary vlan association. |
| Step 5 | UCS-A# /eth-uplink/ vlan # commit-buffer . | Commits the transaction to the system configuration. |

Example

The following example shows how to create a Community VLAN:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan vlan203 203
UCS-A /eth-uplink/vlan* # set sharing community
UCS-A /eth-uplink/vlan* # set pubname vlan200
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan* # exit
UCS-A /vlan-group #
```

Viewing Community VLANs

Procedure

| | Command or Action | Purpose |
|---------------|-------------------------------|--|
| Step 1 | UCS-A# scope org | Enters Cisco UCS Manager organization. |
| Step 2 | UCS-A /org # show vlan | Displays the available groups in the organization. |

Example

The following example shows the available VLAN groups in the root org:

```
UCS-A# scope org
UCS-A# /org/# show vlan
VLAN Group:
```

| Name | VLAN ID | Fabric ID | Native VLAN | Sharing Type | Primary Vlan |
|---------|---------|-----------|-------------|--------------|--------------|
| vlan100 | 100 | Dual | No | Primary | vlan100 |
| vlan100 | 101 | Dual | No | Isolated | vlan100 |
| vlan100 | 203 | Dual | No | Community | vlan200 |

Allowing Community VLANs on vNICs

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>profile-name</i> | Commits the transaction to the system configuration. |
| Step 3 | UCS-A /org/service-profile # scope vnic <i>vnic-name</i> | Enters command mode for the specified vNIC. |
| Step 4 | UCS-A /org/service-profile/vnic # create eth-if <i>community-vlan-name</i> | Allows the community VLAN to access the specified vNIC. |
| Step 5 | UCS-A /org/service-profile/vnic # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to assign the community VLAN cVLAN101 to the vNIC vnic_1 and commits the transaction.

```
UCS-A# scope org /
UCS-A /org # scope service-profile GSP1
UCS-A /org/service-profile # scope vnic vnic_1
UCS-A /org/service-profile/vnic # create eth-if cVLAN101
UCS-A /org/service-profile/vnic* # commit-buffer
```

Allowing PVLAN on Promiscuous Access or Trunk Port

For a promiscuous access port, the isolated and community VLANs must be associated to the same primary VLAN.

For a promiscuous trunk port, isolated and community VLANs belonging to different primary VLANs are allowed, as well as regular VLANs.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A # scope eth-storage | Enters Ethernet storage mode. |
| Step 2 | UCS-A /eth-storage # scope vlan iso-vlan-name | Enters the specified isolated VLAN. |
| Step 3 | UCS-A /eth-storage/vlan # create member-port fabric slot- num port- num | Creates the member port for the specified fabric, assigns the slot number and port number, and enters member port configuration scope. |
| Step 4 | UCS-A /eth-storage/vlan/member-port # exit | Returns to VLAN mode. |
| Step 5 | UCS-A /eth-storage/vlan # exit | Returns to Ethernet storage mode. |
| Step 6 | UCS-A /eth-storage # scope vlan comm-vlan-name | Enters the specified community VLAN. |
| Step 7 | UCS-A /eth-storage/vlan # create member-port fabric slot- num port- num | Creates the member port for the specified fabric, assigns the slot number and port number, and enters member port configuration scope. |
| Step 8 | UCS-A /eth-storage/vlan/member-port # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to assign the isolated and community associated with the same primary VLAN to the same appliance port and commits the transaction.

```
UCS-A# scope eth-storage
UCS-A /eth-storage # scope vlan isovlan501
UCS-A /eth-storage/vlan # create member-port a 1 2
UCS-A /eth-storage/vlan/member-port* # exit
UCS-A /eth-storage/vlan* # exit
UCS-A /eth-storage* # scope vlan cvlan502
UCS-A /eth-storage/vlan* # create member-port a 1 2
UCS-A /eth-storage/vlan/member-port* # commit-buffer
UCS-A /eth-storage/vlan/member-port #
```

Deleting a Community VLAN

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

If you are deleting a private primary VLAN, ensure that you reassign the secondary VLANs to another working primary VLAN.

Before you begin

Before you delete a VLAN from a fabric interconnect, ensure that the VLAN was removed from all vNICs and vNIC templates.



Note If you delete a VLAN that is assigned to a vNIC or vNIC template, the vNIC might allow that VLAN to flap.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | (Optional) UCS-A /eth-uplink # scope fabric {a b} | Enters Ethernet uplink fabric mode. Use this command when you want to delete a named VLAN only from the specified fabric (a or b). |
| Step 3 | UCS-A /eth-uplink # delete community vlan <i>vlan-name</i> | Deletes the specified community VLAN. |
| Step 4 | UCS-A /eth-uplink # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example deletes a Community VLAN and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # delete community vlan vlan203
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```

Viewing the VLAN Port Count

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope fabric-interconnect {a b} | Enters fabric interconnect mode for the specified fabric interconnect. |
| Step 2 | UCS-A /fabric-interconnect # show vlan-port-count | Displays the VLAN port count. |

Example

The following example displays the VLAN port count for fabric interconnect A:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric-interconnect # show vlan-port-count

VLAN-Port Count:
VLAN-Port Limit      Access VLAN-Port Count      Border VLAN-Port Count      Alloc Status
-----
6000                  3                            0                            Available
```

VLAN Port Count Optimization

VLAN port count optimization enables mapping the state of multiple VLANs into a single internal state. When you enable the VLAN port count optimization, Cisco UCS Manager logically groups VLANs based on the port VLAN membership. This grouping increases the port VLAN count limit. VLAN port count optimization also compresses the VLAN state and reduces the CPU load on the fabric interconnect. This reduction in the CPU load enables you to deploy more VLANs over more vNICs. Optimizing VLAN port count does not change any of the existing VLAN configuration on the vNICs.

VLAN port count optimization is disabled by default. You can enable or disable the option based on your requirements.

**Important**

- Enabling VLAN port count optimization increases the number of available VLAN ports for use. If the port VLAN count exceeds the maximum number of VLANs in a non-optimized state, you cannot disable the VLAN port count optimization.
- VLAN port count optimization is not supported in Cisco UCS 6100 Series fabric interconnect.

Enabling Port VLAN Count Optimization

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink# set vlan-port-count-optimization enable | Enables the vlan for port VLAN count optimization. |
| Step 3 | UCS-A /eth-uplink* # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to enable VLAN port count optimization:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # set vlan-port-count-optimization enable
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink#
```

Disabling Port VLAN Count Optimization

If you have more Port VLAN count than that is allowed in the non port VLAN port count optimization state, you cannot disable the optimization.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink# set vlan-port-count-optimization disable | Disables the port VLAN count optimization. |
| Step 3 | UCS-A /eth-uplink # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to disable VLAN port count optimization:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # set vlan-port-count-optimization disable
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink#
```

Viewing the Port VLAN Count Optimization Groups

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink# show vlan-port-count-optimization group | Displays the vlan for port VLAN count optimization groups. |

Example

The following example shows port VLAN count optimization group in fabric a and b:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # show vlan-port-count-optimization group
VLAN Port Count Optimization Group:
  Fabric ID  Group ID  VLAN ID
  -----  -
  -----  -
  -----  -
```

| | | |
|---|----|-----|
| A | 5 | 6 |
| A | 5 | 7 |
| A | 5 | 8 |
| B | 10 | 100 |
| B | 10 | 101 |

VLAN Groups

VLAN groups allow you to group VLANs on Ethernet uplink ports, by function or by VLANs that belong to a specific network. You can define VLAN membership and apply the membership to multiple Ethernet uplink ports on the fabric interconnect.



Note Cisco UCS Manager supports a maximum of 200 VLAN Groups. If Cisco UCS Manager determines that you create more than 200 VLAN groups, the system disables VLAN compression.

You can configure inband and out-of-band (OOB) VLAN groups to use to access the Cisco Integrated Management Interface (CIMC) on blade and rack servers. Cisco UCS Manager supports OOB IPv4 and inband IPv4 and IPv6 VLAN groups for use with the uplink interfaces or uplink port channels.

After you assign a VLAN to a VLAN group, any changes to the VLAN group are applied to all Ethernet uplink ports that are configured with the VLAN group. The VLAN group also enables you to identify VLAN overlaps between disjoint VLANs.

You can configure uplink ports under a VLAN group. When you configure an uplink port for a VLAN group, that uplink port will support all the VLANs that are part of the associated VLAN groups and individual VLANs that are associated with the uplink using LAN Uplinks Manager, if any. Further, any uplink that is not selected for association with that VLAN group will stop supporting the VLANs that are part of that VLAN group.

You can create VLAN groups from the **LAN Cloud** or from the **LAN Uplinks Manager**.

Creating a VLAN Group

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope eth-uplink. | Enters Ethernet uplink mode. The VLAN Group name is case sensitive. |
| Step 2 | UCS-A# /eth-uplink/ #create vlan-group <i>Name</i> . | Create a VLAN group with the specified name. 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 | UCS-A# /eth-uplink/ vlan-group#create member-vlan <i>ID</i> . | Adds the specified VLANs to the created VLAN group. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 4 | UCS-A# /eth-uplink/vlan-group # create member-port [member-port-channel] . | Assigns the uplink Ethernet ports to the VLAN group. |
| Step 5 | UCS-A#/vlan-group* # commit-buffer . | Commits the transaction to the system configuration. |

Example

The following example shows how to create a VLAN group:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan-group eng
UCS-A /eth-uplink/vlan-group* # create member-vlan 3
UCS-A /eth-uplink/vlan-group* # commit-buffer
UCS-A /vlan-group #
```

Creating an Inband VLAN Group

Configure inband VLAN groups to provide access to remote users via an inband service profile.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth uplink | Enters Ethernet uplink configuration mode. |
| Step 2 | UCS-A /eth-uplink # create vlan-group inband-vlan-name | Creates a VLAN group with the specified name and enters VLAN group configuration mode. |
| Step 3 | UCS-A /eth-uplink/vlan-group # create member-vlan inband-vlan-name inband-vlan-id | Adds the specified VLAN to the VLAN group and enters VLAN group member configuration mode. |
| Step 4 | UCS-A /eth-uplink/vlan-group/member-vlan # exit | Exits VLAN group member configuration mode. |
| Step 5 | UCS-A /eth-uplink/vlan-group # create member-port fabricslot-num port-num | Creates the member port for the specified fabric, assigns the slot number, and port number and enters member port configuration. |
| Step 6 | UCS-A /eth-uplink/vlan-group/member-port # commit-buffer | Commits the transaction. |

Example

The example below creates a VLAN group named inband-vlan-group, creates a member of the group named Inband_VLAN and assigns VLAN ID 888, creates member ports for Fabric A and Fabric B, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create vlan-group inband-vlan-group
```

```

UCS-A /eth-uplink/vlan-group* # create member-vlan Inband_VLAN 888
UCS-A /eth-uplink/vlan-group/member-vlan* # exit
UCS-A /eth-uplink/vlan-group* # create member-port a 1 23
UCS-A /eth-uplink/vlan-group/member-port* # exit
UCS-A /eth-uplink/vlan-group* # create member-port b 1 23
UCS-A /eth-uplink/vlan-group/member-port* # commit-buffer
UCS-A /eth-uplink/vlan-group/member-port # exit
UCS-A /eth-uplink/vlan-group # exit

```

What to do next

Assign the inband VLAN group to an inband service profile.

Viewing VLAN Groups

Procedure

| | Command or Action | Purpose |
|---------------|-------------------------------------|--|
| Step 1 | UCS-A# scope org | Enters Cisco UCS Manager organization. |
| Step 2 | UCS-A /org # show vlan-group | Displays the available groups in the organization. |

Example

The following example shows the available VLAN groups in the root org:

```

UCS-A# scope org
UCS-A# /org/# show vlan-group
VLAN Group:
  Name
  ----
  eng
  hr
  finance

```

Deleting a VLAN Group

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink. | Enters Ethernet uplink mode. |
| Step 2 | UCS-A#/eth-uplink/ #delete vlan-groupName . | Deletes the specified VLAN group. |
| Step 3 | UCS-A#/eth-uplink* # commit-buffer. | Commits the transaction to the system configuration. |

Example

The following example shows how to delete a VLAN group:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # delete vlan-group eng
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```

VLAN Permissions

VLAN permissions restrict access to VLANs based on specified organizations and on the service profile organizations to which the VLANs belong. VLAN permissions also restrict the set of VLANs that you can assign to service profile vNICs. VLAN permissions is an optional feature and is disabled by default. You can enable or disable the feature based on your requirements. If you disable the feature, all of the VLANs are globally accessible to all organizations.



Note

If you enable the org permission in LAN > LAN Cloud > Global Policies > Org Permissions, when you create a VLAN, the **Permitted Orgs for VLAN(s)** option displays in the **Create VLANs** dialog box. If you do not enable the **Org Permissions**, the **Permitted Orgs for VLAN(s)** option does not display.

Enabling the org permission allows you to specify the organizations for the VLAN. When you specify the organizations, the VLAN becomes available to that specific organization and all of the sub organizations below the structure. Users from other organizations cannot access this VLAN. You can also modify the VLAN permission anytime based on changes to your VLAN access requirements.



Caution

When you assign the VLAN org permission to an organization at the root level, all sub organizations can access the VLANs. After assigning the org permission at the root level, and you change the permission for a VLAN that belongs to a sub organization, that VLAN becomes unavailable to the root level organization.

Creating VLAN Permissions

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope org. | Enters the Cisco UCS Manager VLAN organization. |
| Step 2 | UCS-A# /org/ # create vlan-permit <i>VLAN permission name.</i> | Creates the specified VLAN permission and assigns VLAN access permission to the organization. |
| Step 3 | UCS-A#/org* # commit-buffer. | Commits the transaction to the system configuration. |

Example

The following example shows how to create a VLAN permission for an organization:

```
UCS-A# scope org
UCS-A /org # create vlan-permit dev
UCS-A /org* # commit-buffer
UCS-A /org #
```

Viewing VLAN Permissions

Procedure

| | Command or Action | Purpose |
|---------------|--------------------------------------|---|
| Step 1 | UCS-A# scope org | Enters Cisco UCS Manager organization. |
| Step 2 | UCS-A /org # show vlan-permit | Displays the available permissions in the organization. |

Example

The following example shows the VLAN groups that have permission to access this VLAN:

```
UCS-A# scope org
UCS-A# /org/# show vlan-permit
VLAN Group:
  Name
  ----
  eng
  hr
  finance
```

Deleting a VLAN Permission

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope org. | Enters the Cisco UCS Manager VLAN organization. |
| Step 2 | UCS-A# /org/ # delete vlan-permit <i>VLAN permission name.</i> | Deletes the access permission to the VLAN. |
| Step 3 | UCS-A#/org* # commit-buffer. | Commits the transaction to the system configuration. |

Example

The following example shows how to delete a VLAN permission from an organization:

```
UCS-A# scope org
UCS-A /org # delete vlan-permit dev
UCS-A /org* # commit-buffer
UCS-A /org #
```




CHAPTER 6

LAN Pin Groups

- [LAN Pin Groups, on page 121](#)
- [Configuring a LAN Pin Group, on page 121](#)

LAN Pin Groups

Cisco UCS uses LAN pin groups to pin Ethernet traffic from a vNIC on a server to an uplink Ethernet port or port channel on the fabric interconnect. You can use this pinning to manage the distribution of traffic from the servers.

To configure pinning for a server, you must include the LAN pin group in a vNIC policy. The vNIC policy is then included in the service profile assigned to that server. All traffic from the vNIC travels through the I/O module to the specified uplink Ethernet port.



Note If you do not assign a pin group to a server interface through a vNIC policy, Cisco UCS Manager chooses an uplink Ethernet port or port channel for traffic from that server interface dynamically. This choice is not permanent. A different uplink Ethernet port or port channel may be used for traffic from that server interface after an interface flap or a server reboot.

If an uplink is part of a LAN pin group, the uplink is not necessarily reserved for only that LAN pin group. Other vNIC's policies that do not specify a LAN pin group can use the uplink as a dynamic uplink.

Configuring a LAN 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.

Before you begin

Configure the ports and port channels with which you want to configure the pin group. You can only include ports and port channels configured as uplink ports in a LAN pin group.

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # create pin-group <i>pin-group-name</i> | Creates an Ethernet (LAN) pin group with the specified name, and enters Ethernet uplink pin group mode. |
| Step 3 | (Optional) UCS-A /eth-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 /eth-uplink/pin-group # set target {a b dual} {port slot-num / port-num port-channel port-num} | Sets the Ethernet pin target to the specified fabric and port, or fabric and port channel. |
| Step 5 | UCS-A /eth-uplink/pin-group # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a LAN pin group named pingroup54 on fabric A, provides a description for the pin group, sets the pin group target to port channel 28, and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # create pin-group pingroup54
UCS-A /eth-uplink/pin-group* # set descr "This is my pin group #54"
UCS-A /eth-uplink/pin-group* # set target a port-channel 28
UCS-A /eth-uplink/pin-group* # commit-buffer
UCS-A /eth-uplink/pin-group #
```

What to do next

Include the pin group in a vNIC template.



CHAPTER 7

MAC Pools

- [MAC Pools, on page 123](#)
- [Creating a MAC Pool, on page 123](#)
- [Deleting a MAC Pool, on page 125](#)

MAC Pools

A MAC pool is a collection of network identities, or MAC addresses, that are unique in their Layer 2 environment and are available to be assigned to vNICs on a server. If you use MAC pools in service profiles, you do not have to manually configure the MAC addresses to be used by the server associated with the service profile.

In a system that implements multitenancy, you can use the organizational hierarchy to ensure that MAC pools can be used only by specific applications or business services. Cisco UCS uses the name resolution policy to assign MAC addresses from the pool.

To assign a MAC address to a server, you must include the MAC pool in a vNIC policy. The vNIC policy is then included in the service profile assigned to that server.

You can specify your own MAC addresses or use a group of MAC addresses provided by Cisco.

Creating a MAC 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 # create mac-pool <i>mac-pool-name</i> | Creates a MAC pool with the specified name, and enters organization MAC pool mode. 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 |

| | Command or Action | Purpose |
|---------------|---|--|
| | | cannot change this name after the object is saved. |
| Step 3 | (Optional) UCS-A /org/mac-pool # set descr <i>description</i> | Provides a description for the MAC 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/mac-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/mac-pool # create block <i>first-mac-addr last-mac-addr</i> | Creates a block (range) of MAC addresses, and enters organization MAC pool block mode. You must specify the first and last MAC addresses in the address range using the form <i>nn:nn:nn:nn:nn:nn</i> , with the addresses separated by a space. Note A MAC pool can contain more than one MAC address block. To create multiple MAC address blocks, you must enter multiple create block commands from organization MAC pool mode. |
| Step 6 | UCS-A /org/mac-pool # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a MAC pool named pool37, provide a description for the pool, define a MAC address block by specifying the first and last MAC addresses in the block, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # create mac-pool pool37
UCS-A /org/mac-pool* # set descr "This is my MAC pool"
UCS-A /org/mac-pool* # create block 00:A0:D7:42:00:01 00:A0:D7:42:01:00
UCS-A /org/mac-pool/block* # commit-buffer
UCS-A /org/mac-pool/block #
```

What to do next

Include the MAC pool in a vNIC template.

Deleting a MAC 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 mac-pool <i>pool-name</i> | Deletes the specified MAC pool. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

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

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




CHAPTER 8

Quality of Service

- [Quality of Service, on page 127](#)
- [Configuring System Classes, on page 128](#)
- [Configuring Quality of Service Policies, on page 132](#)
- [Configuring Flow Control Policies, on page 135](#)

Quality of Service

Cisco UCS provides the following methods to implement quality of service:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames

Global QoS changes made to the QoS system class may result in brief data-plane interruptions for all traffic. Some examples of such changes are:

- Changing the MTU size for an enabled class
- Changing packet drop for an enabled class
- Changing the CoS value for an enabled class

Guidelines and Limitations for Quality of Service on Cisco UCS 6300 Series Fabric Interconnect

- Cisco UCS 6300 Series Fabric Interconnect uses a shared buffer for all system classes.
- Multicast optimization is not supported.
- When you change the QoS parameters for any class causes traffic disruption to all classes. The following table lists the changes in the QoS system class and the conditions that trigger a system reboot.

| QoS System class status | Condition | FI Reboot Status |
|-------------------------|-----------------------------------|------------------|
| Enabled | Change between drop and no drop | Yes |
| No-drop | Change between enable and disable | Yes |

| QoS System class status | Condition | FI Reboot Status |
|-------------------------|--------------------|------------------|
| Enable and no-drop | Change in MTU size | Yes |

- The subordinate FI reboots first as a result of the change in the QoS system class , followed by the primary FI.



Note Cisco UCS Manager displays a prompt that the fabric interconnect will reboot when the system policy is changed.

- **show queuing interface** command is not supported.

Guidelines and Limitations for Quality of Service on Cisco UCS Mini

- Cisco UCS Mini uses a shared buffer for all system classes.
- The bronze class shares the buffer with SPAN. We recommend using either SPAN or the bronze class.
- Multicast optimization is not supported.
- Changing the QoS parameters for any class causes traffic disruption to all classes.
- When mixing Ethernet and FC or FCoE traffic, the bandwidth distribution is not equal.
- Multiple streams of traffic from the same class may not be distributed equally.
- Use the same CoS values for all no-drop policies to avoid any FC or FCoE performance issues.
- Only the platinum and gold classes support no-drop policies.
- **show queuing interface** command is not supported.

Configuring System Classes

System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS domain. This industry standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. Two virtual lanes are reserved for internal system and management traffic. You can configure quality of service (QoS) for the other six virtual lanes. System classes determine how the DCE bandwidth in these six virtual lanes is allocated across the entire Cisco UCS domain.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic, which provides a level of traffic management, even in an oversubscribed system. For example, you can configure the **Fibre Channel Priority** system class to determine the percentage of DCE bandwidth allocated to FCoE traffic.

The following table describes the system classes that you can configure.

Table 1: System Classes

| System Class | Description |
|--------------------------------------|--|
| Platinum Gold Silver Bronze | <p>A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic.</p> <p>All properties of these system classes are available for you to assign custom settings and policies.</p> <p>For Cisco UCS Mini, packet drop can only be disabled on the platinum and gold classes. Only one platinum and one gold class can be configured as a no drop class at a time.</p> |
| Best Effort | <p>A system class that sets the quality of service for the lane reserved for basic Ethernet traffic.</p> <p>Some properties of this system class are preset and cannot be modified. For example, this class has a drop policy that allows it to drop data packets if required. You cannot disable this system class.</p> |
| Fibre Channel | <p>A system class that sets the quality of service for the lane reserved for Fibre Channel over Ethernet traffic.</p> <p>Some properties of this system class are preset and cannot be modified. For example, this class has a no-drop policy that ensures it never drops data packets. You cannot disable this system class.</p> <p>Note FCoE traffic has a reserved QoS system class that should not be used by any other type of traffic. If any other type of traffic has a CoS value that is used by FCoE, the value is remarked to 0.</p> |

Configuring a System Class

The type of adapter in a server might limit the maximum MTU supported. For example, network MTU above the maximums might cause the packet to be dropped for the following adapters:

- The Cisco UCS M71KR CNA adapter, which supports a maximum MTU of 9216.
- The Cisco UCS 82598KR-CI adapter, which supports a maximum MTU of 14000.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope eth-server | Enters Ethernet server mode. |
| Step 2 | UCS-A /eth-server # scope qos | Enters Ethernet server QoS mode. |
| Step 3 | UCS-A /eth-server/qos # scope eth-classified {bronze gold platinum silver} | Enters Ethernet server QoS Ethernet classified mode for the specified system class. |
| Step 4 | UCS-A /eth-server/qos/eth-classified # enable | Enables the specified system class. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 5 | UCS-A /eth-server/qos/eth-classified # set cos <i>cos-value</i> | <p>Specifies the class of service for the specified system class. Valid class of service values are 0 to 6.</p> <p>Important Use the same CoS values on UCS and N5K for all the no-drop policies. To insure that end-to-end PFC works correctly, have the same QoS policy configured on all intermediate switches.</p> <p>Note When the CoS value is set to 0 in any QoS class, this causes the adapter to use the same queue for best effort and the QoS class. When traffic congestion occurs, best effort and the QoS class will share the bandwidth equally instead of using the weight configured in the QoS class.</p> |
| Step 6 | UCS-A /eth-server/qos/eth-classified # set drop { drop no-drop } | <p>Specifies whether the channel can drop packets or not.</p> <p>Note Changes saved to the drop displays the following warning message: Warning: The operation will cause momentary disruption to traffic forwarding.</p> |
| Step 7 | UCS-A /eth-server/qos/eth-classified # set mtu { <i>mtu-value</i> fc normal } | <p>The maximum transmission unit, or packet size to be used. The maximum value for MTU is 9216.</p> <p>Note If the vNIC has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets might get dropped during data transmission.</p> <p>Changes saved to the MTU displays the following warning message: Warning: The operation will cause momentary disruption to traffic forwarding.</p> |
| Step 8 | UCS-A /eth-server/qos/eth-classified # set multicast-optimize { no yes } | <p>Specifies whether the class is optimized to for sending multicast packets.</p> |

| | Command or Action | Purpose |
|----------------|---|--|
| Step 9 | UCS-A /eth-server/qos/eth-classified # set weight { <i>weight-value</i> best-effort none } | Specifies the relative weight for the specified system class. Valid weight values are 0 to 10. |
| Step 10 | UCS-A /eth-server/qos/eth-classified # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to enable the platinum system class, allow the channel to drop packets, set the class of service to 6, set the MTU to normal, set the relative weight to 5, and commit the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope eth-classified platinum
UCS-A /eth-server/qos/eth-classified # enable
UCS-A /eth-server/qos/eth-classified* # set drop drop
Warning: The operation will cause momentary disruption to traffic forwarding
UCS-A /eth-server/qos/eth-classified* # set cos 6
UCS-A /eth-server/qos/eth-classified* # set mtu normal
Warning: The operation will cause momentary disruption to traffic forwarding
UCS-A /eth-server/qos/eth-classified* # set weight 5
UCS-A /eth-server/qos/eth-classified* # commit-buffer
UCS-A /eth-server/qos/eth-classified #
```

Disabling a System Class

If you disable a system class that is used in a QoS policy, Cisco UCS Manager uses the system class configured with CoS 0 for traffic on servers that are configured with the QoS policy. If no system class is configured as CoS 0, the Best Effort system class is used. You cannot disable the Best Effort or Fibre Channel system classes.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope eth-server | Enters Ethernet server mode. |
| Step 2 | UCS-A /eth-server # scope qos | Enters Ethernet server QoS mode. |
| Step 3 | UCS-A /eth-server/qos # scope eth-classified { bronze gold platinum silver } | Enters Ethernet server QoS Ethernet classified mode for the specified system class. |
| Step 4 | UCS-A /eth-server/qos/eth-classified # disable | Disables the specified system class. |
| Step 5 | UCS-A /eth-server/qos/eth-classified # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example disables the platinum system class and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope eth-classified platinum
UCS-A /eth-server/qos/eth-classified # disable
UCS-A /eth-server/qos/eth-classified* # commit-buffer
UCS-A /eth-server/qos/eth-classified #
```

Configuring Quality of Service Policies

Quality of Service Policy

A quality of service (QoS) policy assigns a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic. For certain adapters, you can also specify additional controls on the outgoing traffic, such as burst and rate.

You must include a QoS policy in a vNIC policy or vHBA policy and then include that policy in a service profile to configure the vNIC or vHBA.

Configuring a QoS Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | Switch-A# scope org <i>org-name</i> | Enters org mode for the specified organization. To enter the default org mode, type / as the <i>org-name</i> . |
| Step 2 | Switch-A /org # create qos-policy <i>policy-name</i> | Creates the specified QoS policy, and enters org QoS policy mode. |
| Step 3 | Switch-A /org/qos-policy # create egress-policy | Creates the egress policy (for both vNICs and vHBAs) to be used by the QoS policy, and enters org QoS policy egress policy mode. |
| Step 4 | Switch-A /org/qos-policy/egress-policy # set host-cos-control {full none} | (Optional) Specifies whether the host or Cisco UCS Manager controls the class of service (CoS) for a vNIC. This setting has no effect on a vHBA. Use the full keyword to have the host control the CoS. If the packet has a valid CoS value, the host uses that value. Otherwise, it uses the CoS value associated with the specified class priority. Use the none keyword to have Cisco |

| | Command or Action | Purpose |
|---------------|--|--|
| | | UCS Manager use the CoS value associated with the specified priority. |
| Step 5 | Switch-A /org/qos-policy/egress-policy # set prio <i>sys-class-name</i> | <p>Specifies the system class to be used for the egress policy. The <i>sys-class-name</i> 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 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 6 | Switch-A /org/qos-policy/egress-policy # set rate {<i>line-rate</i> <i>kbps</i>} burst <i>bytes</i> | <p>Specifies the expected average rate of traffic. Traffic that falls under this rate will always conform. The default is line-rate, which equals a value of 10,000,000. The minimum value is 8, and the maximum value is 40,000,000.</p> <p>Rate limiting is supported only on vNICs on the Cisco UCS VIC-1240 Virtual Interface Card and Cisco UCS VIC-1280 Virtual Interface Card. The Cisco UCS M81KR Virtual Interface Card supports rate limiting on both vNICs and vHBAs.</p> |
| Step 7 | Switch-A /org/qos-policy/egress-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates a QoS policy for vNIC traffic, assigns the platinum system class and sets the rate limit (traffic rate and burst size) for the egress policy, and commits the transaction:

```
Switch-A# scope org /
Switch-A /org # create qos-policy VnicPolicy34
Switch-A /org/qos-policy* # create egress-policy
Switch-A /org/qos-policy/egress-policy* # set prio platinum
Switch-A /org/qos-policy/egress-policy* # set rate 5000000 burst 65000
Switch-A /org/qos-policy/egress-policy* # commit-buffer
Switch-A /org/qos-policy/egress-policy #
```

The following example creates a QoS policy for vHBA traffic, assigns the fc (Fibre Channel) system class and sets the rate limit (traffic rate and burst size) for the egress policy, and commits the transaction:

```
Switch-A# scope org /
Switch-A /org # create qos-policy VhbaPolicy12
Switch-A /org/qos-policy* # create egress-policy
Switch-A /org/qos-policy/egress-policy* # set prio fc
Switch-A /org/qos-policy/egress-policy* # set rate 5000000 burst 65000
Switch-A /org/qos-policy/egress-policy* # commit-buffer
Switch-A /org/qos-policy/egress-policy #
```

What to do next

Include the QoS policy in a vNIC or vHBA template.

Deleting a QoS Policy

If you delete a QoS policy that is in use or you disable a system class that is used in a QoS policy, any vNIC or vHBA that uses that QoS policy is assigned to the Best Effort system class or to the system class with a CoS of 0. In a system that implements multitenancy, Cisco UCS Manager first attempts to find a matching QoS policy in the organization hierarchy.

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 # delete qos-policy <i>policy-name</i> | Deletes the specified QoS policy. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following deletes the QoS policy named QosPolicy34 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete qos-policy QosPolicy34
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Flow Control Policies

Flow Control Policy

Flow control policies determine whether the uplink Ethernet ports in a Cisco UCS domain send and receive IEEE 802.3x pause frames when the receive buffer for a port fills. These pause frames request that the transmitting port stop sending data for a few milliseconds until the buffer clears.

For flow control to work between a LAN port and an uplink Ethernet port, you must enable the corresponding receive and send flow control parameters for both ports. For Cisco UCS, the flow control policies configure these parameters.

When you enable the send function, the uplink Ethernet port sends a pause request to the network port if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels. If you enable the receive function, the uplink Ethernet port honors all pause requests from the network port. All traffic is halted on that uplink port until the network port cancels the pause request.

Because you assign the flow control policy to the port, changes to the policy have an immediate effect on how the port reacts to a pause frame or a full receive buffer.

Configuring a Flow Control Policy

Before you begin

Configure the network port with the corresponding setting for the flow control that you need. For example, if you enable the send setting for flow-control pause frames in the policy, ensure that the receive parameter in the network port is set to on or to desired. If you want the Cisco UCS port to receive flow-control frames, ensure that the send parameter is set to on or to desire on the network port. If you do not want to use flow control, you can set the send and receive parameters on the network port to off.

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # scope flow-control | Enters Ethernet uplink flow control mode. |
| Step 3 | UCS-A /eth-uplink/flow-control # create policy <i>policy-name</i> | Creates the specified flow control policy. |
| Step 4 | UCS-A /eth-uplink/flow-control/policy # set prio <i>prio-option</i> | Specifies one of the following flow control priority options: <ul style="list-style-type: none"> • auto —The Cisco UCS system and the network negotiate whether PPP will be used on this fabric interconnect. • on —PPP is enabled on this fabric interconnect. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 5 | UCS-A /eth-uplink/flow-control/policy # set receive <i>receive-option</i> | Specifies one of the following flow control receive options: <ul style="list-style-type: none"> • off—Pause requests from the network are ignored and traffic flow continues as normal. • on—Pause requests are honored and all traffic is halted on that uplink port until the network cancels the pause request. |
| Step 6 | UCS-A /eth-uplink/flow-control/policy # set send <i>send-option</i> | Specifies one of the following flow control send options: <ul style="list-style-type: none"> • off—Traffic on the port flows normally regardless of the packet load. • on—The Cisco UCS system sends a pause request to the network if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels. |
| Step 7 | UCS-A /eth-uplink/flow-control/policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following configures a flow control policy and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope flow-control
UCS-A /eth-uplink/flow-control # create policy FlowControlPolicy23
UCS-A /eth-uplink/flow-control/policy* # set prio auto
UCS-A /eth-uplink/flow-control/policy* # set receive on
UCS-A /eth-uplink/flow-control/policy* # set send on
UCS-A /eth-uplink/flow-control/policy* # commit-buffer
UCS-A /eth-uplink/flow-control/policy #
```

What to do next

Associate the flow control policy with an uplink Ethernet port or port channel.

Deleting a Flow Control Policy

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 flow-control | Enters Ethernet uplink flow control mode. |
| Step 3 | UCS-A /eth-uplink/flow-control # delete policy <i>policy-name</i> | Deletes the specified flow control policy. |
| Step 4 | UCS-A /eth-uplink/flow-control # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example deletes the flow control policy named FlowControlPolicy23 and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope flow-control
UCS-A /eth-uplink/flow-control # delete policy FlowControlPolicy23
UCS-A /eth-uplink/flow-control* # commit-buffer
UCS-A /eth-uplink/flow-control #
```




CHAPTER 9

Upstream Disjoint Layer-2 Networks

- [Upstream Disjoint Layer-2 Networks, on page 139](#)
- [Guidelines for Configuring Upstream Disjoint L2 Networks, on page 140](#)
- [Upstream Disjoint L2 Networks Pinning Considerations, on page 141](#)
- [Configuring Cisco UCS for Upstream Disjoint L2 Networks, on page 143](#)
- [Assigning Ports and Port Channels to VLANs, on page 144](#)
- [Removing Ports and Port Channels from VLANs, on page 144](#)
- [Viewing Ports and Port Channels Assigned to VLANs, on page 145](#)

Upstream Disjoint Layer-2 Networks

Upstream disjoint layer-2 networks (disjoint L2 networks) are required if you have two or more Ethernet clouds that never connect, but must be accessed by servers or virtual machines located in the same Cisco UCS domain. For example, you could configure disjoint L2 networks if you require one of the following:

- Servers or virtual machines to access a public network and a backup network
- Servers or virtual machines for more than one customer are located in the same Cisco UCS domain, and that need to access the L2 networks for both customers in a multi-tenant system



Note By default, data traffic in Cisco UCS works on a principle of mutual inclusion. All traffic for all VLANs and upstream networks travels along all uplink ports and port channels. If you have upgraded from a release that does not support upstream disjoint layer-2 networks, you must assign the appropriate uplink interfaces to your VLANs, or traffic for those VLANs continues to flow along all uplink ports and port channels.

The configuration for disjoint L2 networks works on a principle of selective exclusion. Traffic for a VLAN that is designated as part of a disjoint network can only travel along an uplink Ethernet port or port channel that is specifically assigned to that VLAN, and is selectively excluded from all other uplink ports and port channels. However, traffic for VLANs that are not specifically assigned to an uplink Ethernet port or port channel can still travel on all uplink ports or port channels, including those that carry traffic for the disjoint L2 networks.

In Cisco UCS, the VLAN represents the upstream disjoint L2 network. When you design your network topology for disjoint L2 networks, you must assign uplink interfaces to VLANs not the reverse.

For information about the maximum number of supported upstream disjoint L2 networks, see the appropriate *Cisco UCS Configuration Limits for Cisco UCS Manager Guide*.

Guidelines for Configuring Upstream Disjoint L2 Networks

When you plan your configuration for upstream disjoint L2 networks, consider the following:

Ethernet Switching Mode Must Be End-Host Mode

Cisco UCS only supports disjoint L2 networks when the Ethernet switching mode of the fabric interconnects is configured for end-host mode. You cannot connect to disjoint L2 networks if the Ethernet switching mode of the fabric interconnects is switch 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 VLANs.

VLAN Validity Criteria Are the Same for Uplink Ethernet Ports and Port Channels

The VLAN used for the disjoint L2 networks must be configured and assigned to an uplink Ethernet port or uplink Ethernet port channel. If the port or port channel does not include the VLAN, Cisco UCS Manager considers the VLAN invalid and does the following:

- Displays a configuration warning in the **Status Details** area for the server.
- Ignores the configuration for the port or port channel and drops all traffic for that VLAN.



Note

The validity criteria are the same for uplink Ethernet ports and uplink Ethernet port channels. Cisco UCS Manager does not differentiate between the two.

Overlapping VLANs Are Not Supported

Cisco UCS does not support overlapping VLANs in disjoint L2 networks. You must ensure that each VLAN only connects to one upstream disjoint L2 domain.

Each vNIC Can Only Communicate with One Disjoint L2 Network

A vNIC can only communicate with one disjoint L2 network. If a server needs to communicate with multiple disjoint L2 networks, you must configure a vNIC for each of those networks.

To communicate with more than two disjoint L2 networks, a server must have a Cisco VIC adapter that supports more than two vNICs.

Appliance Port Must Be Configured with the Same VLAN as Uplink Ethernet Port or Port Channel

For an appliance port to communicate with a disjoint L2 network, you must ensure that at least one uplink Ethernet port or port channel is in the same network and is therefore assigned to the same VLANs that are used by the appliance port. If Cisco UCS Manager cannot identify an uplink Ethernet port or port channel

that includes all VLANs that carry traffic for an appliance port, the appliance port experiences a pinning failure and goes down.

For example, a Cisco UCS domain includes a global VLAN named vlan500 with an ID of 500. vlan500 is created as a global VLAN on the uplink Ethernet port. However, Cisco UCS Manager does not propagate this VLAN to appliance ports. To configure an appliance port with vlan500, you must create another VLAN named vlan500 with an ID of 500 for the appliance port. You can create this duplicate VLAN in the **Appliances** node on the **LAN** tab of the Cisco UCS Manager GUI or the **eth-storage** scope in the Cisco UCS Manager CLI. If you are prompted to check for VLAN Overlap, accept the overlap and Cisco UCS Manager creates the duplicate VLAN for the appliance port.

Default VLAN 1 Cannot Be Configured Explicitly on an Uplink Ethernet Port or Port Channel

Cisco UCS Manager implicitly assigns default VLAN 1 to all uplink ports and port channels. Even if you do not configure any other VLANs, Cisco UCS uses default VLAN 1 to handle data traffic for all uplink ports and port channels.



Note After you configure VLANs in a Cisco UCS domain, default VLAN 1 remains implicitly on all uplink ports and port channels. You cannot explicitly assign default VLAN 1 to an uplink port or port channel, nor can you remove it from an uplink port or port channel.

If you attempt to assign default VLAN 1 to a specific port or port channel, Cisco UCS Manager raises an Update Failed fault.

Therefore, if you configure a Cisco UCS domain for disjoint L2 networks, do not configure any vNICs with default VLAN 1 unless you want all data traffic for that server to be carried on all uplink Ethernet ports and port channels and sent to all upstream networks.

VLANs for Both FIs Must be Concurrently Assigned

When you assign a port to a global VLAN, the VLAN is removed from all of the ports that are not explicitly assigned to the VLAN on both fabric interconnects. The ports on both FIs must be configured at the same time. If the ports are only configured on the first FI, traffic on the second FI will be disrupted.

Upstream Disjoint L2 Networks Pinning Considerations

Communication with an upstream disjoint L2 network requires that you ensure that the pinning is properly configured. Whether you implement soft-pinning or hard-pinning, a VLAN membership mismatch causes traffic for one or more VLANs to be dropped.

Soft-Pinning

Soft-pinning is the default behavior in Cisco UCS. If you plan to implement soft-pinning, you do not need to create LAN pin groups to specify a pin target for a vNIC. Instead, Cisco UCS Manager pins the vNIC to an uplink Ethernet port or port channel according to VLAN membership criteria.

With soft-pinning, Cisco UCS Manager validates data traffic from a vNIC against the VLAN membership of all uplink Ethernet ports and port channels. If you have configured disjoint L2 networks, Cisco UCS Manager must be able to find an uplink Ethernet port or port channel that is assigned to all VLANs on the vNIC. If no

uplink Ethernet port or port channel is configured with all VLANs on the vNIC, Cisco UCS Manager does the following:

- Brings the link down.
- Drops the traffic for all of the VLANs on the vNIC.
- Raises the following faults:
 - Link Down
 - VIF Down

Cisco UCS Manager does not raise a fault or warning about the VLAN configuration.

For example, a vNIC on a server is configured with VLANs 101, 102, and 103. Interface 1/3 is assigned only to VLAN 102. Interfaces 1/1 and 1/2 are not explicitly assigned to a VLAN, which makes them available for traffic on VLANs 101 and 103. As a result of this configuration, the Cisco UCS domain does not include a border port interface that can carry traffic for all three VLANs for which the vNIC is configured. As a result, Cisco UCS Manager brings down the vNIC, drops traffic for all three VLANs on the vNIC, and raises the Link Down and VIF Down faults.

hard-pinning

hard-pinning occurs when you use LAN pin groups to specify the pinning target for the traffic intended for the disjoint L2 networks. In turn, the uplink Ethernet port or port channel that is the pinning target must be configured to communicate with the appropriate disjoint L2 network.

With hard-pinning, Cisco UCS Manager validates data traffic from a vNIC against the VLAN membership of all uplink Ethernet ports and port channels, and validates the LAN pin group configuration to ensure it includes the VLAN and the uplink Ethernet port or port channel. If the validation fails at any point, Cisco UCS Manager does the following:

- Raises a Pinning VLAN Mismatch fault with a severity of Warning.
- Drops traffic for the VLAN.
- Does not bring the link down, so that traffic for other VLANs can continue to flow along it.

For example, if you want to configure hard-pinning for an upstream disjoint L2 network that uses VLAN 177, do the following:

- Create a LAN pin group with the uplink Ethernet port or port channel that carries the traffic for the disjoint L2 network.
- Configure at least one vNIC in the service profile with VLAN 177 and the LAN pin group.
- Assign VLAN 177 to an uplink Ethernet port or port channel included in the LAN pin group

If the configuration fails at any of these three points, then Cisco UCS Manager warns of a VLAN mismatch for VLAN 177 and drops the traffic for that VLAN only.



Note If changes are made to soft-pinning configurations resulting in vNIC VLANs not resolving with disjoint L2 uplink, a warning dialog box is displayed. The warning dialog box allows you to proceed with your configuration or cancel it. If you decide to proceed with the mis-configuration, you will experience a reduction in server traffic performance.

Configuring Cisco UCS for Upstream Disjoint L2 Networks

When you configure a Cisco UCS domain to connect with upstream disjoint L2 networks, you need to ensure that you complete all of the following steps.

Before you begin

Before you begin this configuration, ensure that the ports on the fabric interconnects are properly cabled to support your disjoint L2 networks configuration.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | Configure Ethernet switching mode for both fabric interconnects in Ethernet End-Host Mode. | The Ethernet switching mode must be in End-Host Mode for Cisco UCS to be able to communicate with upstream disjoint L2 networks. See LAN Ports and Port Channels, on page 21 . |
| Step 2 | Configure the ports and port channels that you require to carry traffic for the disjoint L2 networks. | |
| Step 3 | (Optional) Configure the LAN pin groups required to pin the traffic for the appropriate uplink Ethernet ports or port channels. | See Configuring a LAN Pin Group, on page 121 . |
| Step 4 | Create one or more VLANs. | These can be named VLANs or private VLANs. For a cluster configuration, we recommend that you create the VLANs accessible to both fabric interconnects. See VLANs, on page 93 and Upstream Disjointed Layer-2 Networks, on page 139 |
| Step 5 | Assign the desired ports or port channels to the VLANs for the disjoint L2 networks. | When this step is complete, traffic for these VLANs is sent through the trunks for the assigned ports and/or port channels. |
| Step 6 | Ensure that the service profiles for all servers that need to communicate with the disjoint L2 networks include the correct LAN connectivity configuration. This configuration ensures that | You can complete this configuration through one or more vNIC templates, or when you configure the networking options for the service profile. For more information about vNIC |

| | Command or Action | Purpose |
|--|---|---|
| | the vNICs direct the traffic to the appropriate VLAN. | templates and service profiles, see the <i>Cisco UCS Manager Storage Management Guide</i> . |

Assigning Ports and Port Channels to VLANs

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # scope vlan <i>vlan-name</i> | Enters Ethernet uplink VLAN mode for the specified VLAN. |
| Step 3 | UCS-A /eth-uplink/vlan # create member-port <i>fabric-interconnect slot-id port-id</i> | Assigns the specified VLAN to the specified uplink Ethernet port. |
| Step 4 | UCS-A /eth-uplink/vlan # create member-port-channel <i>fabric-interconnect member-port-chan-id</i> | Assigns the specified VLAN to the specified uplink Ethernet port channel. |
| Step 5 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. After a port or port channel is assigned to one or more VLANs, it is removed from all other VLANs. |

Example

The following example assigns uplink Ethernet ports to a named VLAN called VLAN100 on fabric interconnect A and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope vlan VLAN100
UCS-A /eth-uplink/vlan # create member-port a 2
UCS-A /eth-uplink/vlan # create member-port a 4
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Removing Ports and Port Channels from VLANs

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 vlan <i>vlan-name</i> | Enters Ethernet uplink VLAN mode for the specified VLAN. |
| Step 3 | UCS-A /eth-uplink/vlan # delete member-port <i>fabric-interconnect slot-id port-id</i> | Deletes the specified Uplink Ethernet member port assignment from the VLAN. |
| Step 4 | UCS-A /eth-uplink/vlan # delete member-port-channel <i>fabric-interconnect member-port-chan-id</i> | Deletes the specified Uplink Ethernet port channel assignment from the VLAN. |
| Step 5 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. Important If you remove all port or port channel interfaces from a VLAN, the VLAN returns to the default behavior and data traffic on that VLAN flows on all uplink ports and port channels. Based on the configuration in the Cisco UCS domain, this default behavior can cause Cisco UCS Manager to drop traffic for that VLAN. To avoid this occurrence, Cisco recommends that you assign at least one interface to the VLAN or delete the VLAN. |

Example

The following example deletes the association between uplink Ethernet port 2 on fabric interconnect A and the named VLAN called MyVLAN and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope vlan MyVLAN
UCS-A /eth-uplink/vlan # delete member-port a 2
UCS-A /eth-uplink/vlan* # commit-buffer
UCS-A /eth-uplink/vlan #
```

Viewing Ports and Port Channels Assigned to VLANs

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # scope vlan <i>vlan-name</i> | Enters Ethernet uplink VLAN mode for the specified VLAN. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 3 | UCS-A /eth-uplink/vlan # show member-port [detail expand] | Shows member ports assigned to the specified VLAN. |
| Step 4 | UCS-A /eth-uplink/vlan # show member-port-channel [detail expand] | Shows member port channels assigned to the specified VLAN. |
| Step 5 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example displays the full details for uplink Ethernet ports assigned to a named VLAN called MyVLAN:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope vlan MyVLAN
UCS-A /eth-uplink/vlan # show member-port detail
Member Port:
    Fabric ID: A
    Slot ID: 1
    Port ID: 2
    Mark Native Vlan: No
UCS-A /eth-uplink/vlan #
```



CHAPTER 10

Network-Related Policies

- [vNIC Template, on page 147](#)
- [Ethernet Adapter Policies, on page 154](#)
- [Ethernet and Fibre Channel Adapter Policies, on page 160](#)
- [Configuring a Default vNIC Behavior Policy, on page 163](#)
- [Deleting a vNIC from a LAN Connectivity Policy, on page 164](#)
- [Creating a LAN Connectivity Policy, on page 165](#)
- [Deleting a LAN Connectivity Policy, on page 166](#)
- [About the LAN and SAN Connectivity Policies, on page 166](#)
- [Network Control Policy, on page 174](#)
- [Creating a Multicast Policy, on page 178](#)
- [Deleting a Multicast Policy, on page 179](#)
- [Entering Multicast Policy Mode, on page 180](#)
- [Enter a Multicast Policy, on page 180](#)
- [Assigning a Global VLAN Multicast Policy, on page 181](#)
- [Disassociating a Global VLAN Multicast Policy, on page 181](#)
- [Disassociating a VLAN Multicast Policy, on page 182](#)
- [Configuring Ethernet Adapter Policies, on page 183](#)
- [Configuring the Default vNIC Behavior Policy, on page 185](#)
- [Configuring a Network Control Policy, on page 186](#)
- [Deleting a Network Control Policy, on page 188](#)
- [Configuring Multicast Policies, on page 188](#)
- [LACP Policy, on page 193](#)
- [Configuring UDLD Link Policies, on page 195](#)
- [VMQ Connection Policy, on page 202](#)

vNIC Template

The vNIC LAN connectivity policy defines how a vNIC on a server connects to the LAN.

Cisco UCS Manager does not automatically create a VM-FEX port profile with the correct settings when you create a vNIC template. If you want to create a VM-FEX port profile, you must configure the target of the vNIC template as a VM. You must include this policy in a service profile for it to take effect.

You can select VLAN groups in addition to any individual VLAN while creating a vNIC template.



Note If your server has two Emulex or QLogic NICs (Cisco UCS CNA M71KR-E or Cisco UCS CNA M71KR-Q), you must configure vNIC policies for both adapters in your service profile to get a user-defined MAC address for both NICs. If you do not configure policies for both NICs, Windows still detects both of them in the PCI bus. Then because the second eth is not part of your service profile, Windows assigns it a hardware MAC address. If you then move the service profile to a different server, Windows sees additional NICs because one NIC did not have a user-defined MAC address.

Creating vNIC Template Pairs

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A/ org # create vnic-templ <i>vnic-primary</i> . | Creates a Primary vNIC template. |
| Step 2 | UCS-A/ # org vnic-templ set type updating-template . | Set the template type to updating, which drives the configurations in the Primary vNIC template for shared configurations to the peer vNIC template. See the shared configurations listed below. |
| Step 3 | UCS-A/ # org vnic-templ [set fabric {a b}] . | Specifies the fabric for the Primary vNIC template. If you specify Fabric A for the Primary vNIC template, the Secondary vNIC template must be Fabric B or vice versa. |
| Step 4 | UCS-A/ # org vnic-templ set descr primaryinredundancypair . | Sets the template as the Primary vNIC template. |
| Step 5 | UCS-A/ # org vnic-templ set redundancy-type <i>primary</i> . | <p>Sets the redundancy template type as the Primary vNIC template.</p> <p>Following are descriptions of the Redundancy Types:</p> <p>Primary—Creates configurations that can be shared with the Secondary vNIC template. Any shared changes on the Primary vNIC template are automatically synchronized to the Secondary vNIC template.</p> <p>Secondary — All shared configurations are inherited from the Primary template.</p> <p>No Redundancy— Legacy vNIC template behavior.</p> <p>Following is a list of shared configurations:</p> <ul style="list-style-type: none"> • Network Control Policy |

| | Command or Action | Purpose |
|----------------|--|--|
| | | <ul style="list-style-type: none"> • QoS Policy • Stats Threshold Policy • Template Type • Connection Policies • VLANS • MTU <p>Following is a list of non-shared configurations:</p> <ul style="list-style-type: none"> • Fabric ID • CDN Source • MAC Pool • Description • Pin Group Policy |
| Step 6 | UCS-A/ # org vnic-templ exit . | Exits creating the redundancy template pairing. Note Ensure to commit the transaction after linking the Primary vNIC template to a peer Secondary vNIC template to create the redundancy pair. |
| Step 7 | UCS-A/ # org vnic-templ create vNIC-templ vNICsecondary . | Creates the Secondary vNIC template. |
| Step 8 | UCS-A/ # org vnic-templ set type updating-template . | Sets the template type to updating, which automatically inherits the configurations from the Primary vNIC template. |
| Step 9 | UCS-A/ org # vnic-templ [set fabric {a b}] . | Specifies the fabric for the Secondary vNIC template. If you specify Fabric A for the Primary vNIC template, the Secondary vNIC template must be Fabric B or vice versa. |
| Step 10 | UCS-A/ # org vnic-templ set descr secondaryredundancypair . | Sets the secondary vNIC template as a redundancy pair template. |
| Step 11 | UCS-A/ # org vnic-templ set redundancy-type secondary . | Sets the vNIC template type as Secondary. |
| Step 12 | UCS-A/ # org vnic-templ set peer-template-name vNIC-primary . | Sets the Primary vNIC template as the peer to the Secondary vNIC template. |

| | Command or Action | Purpose |
|----------------|---|--|
| Step 13 | UCS-A / # org vnic-templ commit-buffer . | Commits the transaction to the system configuration. |

Example

The following example configures a vNIC redundancy template pair and commits the transaction:

```
UCS-A /org* # create vnic-template vnic-primary
UCS-A /org/vnic-templ* # set type updating-template
UCS-A /org/vnic-templ* # set fabric a
UCS-A /org/vnic-templ* # set descr primaryinredundancypair
UCS-A /org/vnic-templ* # set redundancy-type primary
UCS-A /org/vnic-templ* # exit
UCS-A /org* # create vnic-templ vnicsecondary
UCS-A /org/vnic-templ* # set fabric b
UCS-A /org/vnic-templ* # set descr secondaryinredundancypair
UCS-A /org/vnic-templ* # set redundancy-type secondary
UCS-A /org/vnic-templ* # set peer-template-name vnic-primary
UCS-A /org/vnic-templ* # commit-buffer
UCS-A /org/vnic-templ #
```

What to do next

After you create the vNIC redundancy template pair, you can use the redundancy template pair to create redundancy vNIC pairs for any service profile in the same organization or sub- organization.

Undo vNIC Template Pairs

You can undo the vNIC template pair by changing the Peer Redundancy Template so that there is no peer template for the Primary or the Secondary template. When you undo a vNIC template pair, the corresponding vNIC pairs also becomes undone.

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A /org # scope vnic-templ <i>template1</i> . | Specifies the name of the vNIC template that you want to undo from the template pair. |
| Step 2 | UCS-A /org/ vnic-templ # set redundancy-type <i>no redundancy</i> . | Removes the pairing between the peer Primary or Secondary redundancy template used to perform the template pairing. |
| Step 3 | UCS-A /org/vnic-templ* # commit-buffer . | Commits the transaction to the system configuration. |

Example

The following example shows how to undo a template pairing:

```
UCS-A /org # scope vnic-templ template1
UCS-A /org/vnic-templ # set redundancy-type no-redundancy
UCS-A /org/vnic-templ* # commit buffer
```

Configuring a vNIC Template

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 vnic-templ <i>vnic-templ-name</i> [eth-if <i>vlan-name</i>] [fabric { a b }] [target [adapter vm]] | <p>Creates a vNIC template and enters organization vNIC template mode.</p> <p>The target you choose determines whether or not Cisco UCS Manager automatically creates a VM-FEX port profile with the appropriate settings for the vNIC template. This can be one of the following:</p> <ul style="list-style-type: none"> • Adapter—The vNICs apply to all adapters. No VM-FEX port profile is created if you choose this option. • VM—The vNICs apply to all virtual machines. A VM-FEX port profile is created if you choose this option. |
| Step 3 | (Optional) UCS-A /org/vnic-templ # set descr <i>description</i> | Provides a description for the vNIC template. |
| Step 4 | (Optional) UCS-A /org/vnic-templ # set fabric { a a-b b b-a } | <p>Specifies the fabric to use for the vNIC. If you did not specify the fabric when creating the vNIC template in Step 2, you have the option to specify it with this command.</p> <p>If you want this vNIC to be able to access the second fabric interconnect if the default one is unavailable, choose a-b (A is the primary) or b-a (B is the primary) .</p> |

| | Command or Action | Purpose |
|---------------|---|---|
| | | <p>Note Do not enable fabric failover for the vNIC under the following circumstances:</p> <ul style="list-style-type: none"> • If the Cisco UCS domain is running in Ethernet Switch Mode. vNIC fabric failover is not supported in that mode. If all Ethernet uplinks on one fabric interconnect fail, the vNICs do not fail over to the other. • If you plan to associate this vNIC to a server with an adapter that does not support fabric failover, such as the Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter. If you do so, Cisco UCS Manager generates a configuration fault when you associate the service profile with the server. |
| Step 5 | UCS-A /org/vnic-templ # set mac-pool <i>mac-pool-name</i> | The MAC address pool that vNICs created from this vNIC template should use. |
| Step 6 | UCS-A /org/vnic-templ # set mtu <i>mtu-value</i> | The maximum transmission unit, or packet size, that vNICs created from this vNIC template should use. Enter an integer between 1500 and 9000. |

| | Command or Action | Purpose |
|----------------|--|---|
| | | <p>Note If the vNIC template has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets may be dropped during data transmission.</p> <p>For VIC 14xx adapters, you can change the MTU size of the vNIC from the host interface settings. When the Overlay network is configured, make sure that the new value is equal to or less than the MTU specified in the associated QoS system class or packets could be dropped during data transmission.</p> |
| Step 7 | UCS-A /org/vnic-templ # set nw-control-policy <i>policy-name</i> | The network control policy that vNICs created from this vNIC template should use. |
| Step 8 | UCS-A /org/vnic-templ # set pin-group <i>group-name</i> | The LAN pin group that vNICs created from this vNIC template should use. |
| Step 9 | UCS-A /org/vnic-templ # set qos-policy <i>policy-name</i> | The quality of service policy that vNICs created from this vNIC template should use. |
| Step 10 | UCS-A /org/vnic-templ # set stats-policy <i>policy-name</i> | The statistics collection policy that vNICs created from this vNIC template should use. |
| Step 11 | UCS-A /org/vnic-templ # set type { initial-template updating-template } | Specifies the vNIC template update type. If you do not want vNIC 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 vNIC instances are updated when the vNIC template is updated. |
| Step 12 | UCS-A /org/vnic-templ # commit-buffer | Commits the transaction to the system configuration. |

Example

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

```
UCS-A# scope org /
UCS-A /org* # create vnic template VnicTempFoo
```

```

UCS-A /org/vnic-templ* # set descr "This is a vNIC template example."
UCS-A /org/vnic-templ* # set fabric a
UCS-A /org/vnic-templ* # set mac-pool pool137
UCS-A /org/vnic-templ* # set mtu 8900
UCS-A /org/vnic-templ* # set nw-control-policy ncp5
UCS-A /org/vnic-templ* # set pin-group PinGroup54
UCS-A /org/vnic-templ* # set qos-policy QosPol5
UCS-A /org/vnic-templ* # set stats-policy ServStatsPolicy
UCS-A /org/vnic-templ* # set type updating-template
UCS-A /org/vnic-templ* # commit-buffer
UCS-A /org/vnic-templ #

```

Deleting a vNIC Template

Procedure

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

Example

The following example deletes the vNIC template named VnicTemp42 and commits the transaction:

```

UCS-A# scope org /
UCS-A /org # delete vnic template VnicTemp42
UCS-A /org* # commit-buffer
UCS-A /org #

```

Ethernet Adapter Policies

Configuring an Ethernet Adapter Policy

Procedure

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

| | Command or Action | Purpose |
|----------------|--|--|
| Step 2 | UCS-A /org # create eth-policy <i>policy-name</i> | Creates the specified Ethernet adapter policy and enters organization Ethernet policy mode. |
| Step 3 | (Optional) UCS-A /org/eth-policy # set arfs acceleratedrfs { enabled disabled } | Configures Accelerated RFS. |
| Step 4 | (Optional) UCS-A /org/eth-policy # set comp-queue count <i>count</i> | Configures the Ethernet completion queue. |
| Step 5 | (Optional) UCS-A /org/eth-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 6 | (Optional) UCS-A /org/eth-policy # set failover timeout <i>timeout-sec</i> | Configures the Ethernet failover. |
| Step 7 | (Optional) UCS-A /org/eth-policy # set interrupt { coalescing-time <i>sec</i> coalescing-type { idle min } count <i>count</i> mode { intx msi msi-x }} | Configures the Ethernet interrupt. |
| Step 8 | (Optional) UCS-A /org/eth-policy # set nvgre adminstate { disabled enabled } | Configures NVGRE. |
| Step 9 | (Optional) UCS-A /org/eth-policy # set offload { large-receive tcp-rx-checksum tcp-segment tcp-tx-checksum } { disabled enabled } | Configures the Ethernet offload. |
| Step 10 | (Optional) UCS-A /org/eth-policy # set policy-owner { local pending } | Specifies the owner for the Ethernet adapter policy. |
| Step 11 | (Optional) UCS-A /org/eth-policy # set recv-queue { count <i>count</i> ring-size <i>size-num</i> } | Configures the Ethernet receive queue. |
| Step 12 | (Optional) UCS-A /org/eth-policy # set rss receivesidescaling { disabled enabled } | Configures the RSS. |
| Step 13 | (Optional) UCS-A /org/eth-policy # set trans-queue { count <i>count</i> ring-size <i>size-num</i> } | Configures the Ethernet transmit queue. |
| Step 14 | (Optional) UCS-A /org/eth-policy # set vxlan adminstate { disabled enabled } | Configures VXLAN. |
| Step 15 | UCS-A /org/eth-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example configures an Ethernet adapter policy, and commits the transaction:

```
UCS-A# scope org
UCS-A /org* # create eth-policy EthPolicy19
UCS-A /org/eth-policy* # set comp-queue count 16
UCS-A /org/eth-policy* # set descr "This is an Ethernet adapter policy example."
UCS-A /org/eth-policy* # set failover timeout 300
UCS-A /org/eth-policy* # set interrupt count 64
UCS-A /org/eth-policy* # set offload large-receive disabled
UCS-A /org/eth-policy* # set recv-queue count 32
UCS-A /org/eth-policy* # set rss receivesidescaling enabled
UCS-A /org/eth-policy* # set trans-queue
UCS-A /org/eth-policy* # commit-buffer
UCS-A /org/eth-policy #
```

Deleting an Ethernet 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# delete eth-policy <i>policy-name</i> | Deletes the specified Ethernet adapter policy. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example deletes the Ethernet adapter policy named EthPolicy19 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete eth-policy EthPolicy19
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE

Cisco UCS Manager supports stateless offloads with NVGRE only with Cisco UCS VIC 1340, 1380, 1385, 1387 and/or Cisco UCS VIC 1380 adapters that are installed on servers running Windows Server 2012 R2 operating systems. You cannot use NVGREs stateless offloads with Netflow, usNIC, or VM-FEX.

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 eth-policy <i>policy-name</i> | Creates the specified Ethernet adapter policy and enters organization Ethernet policy mode. |
| Step 3 | To enable stateless offloads with NVGRE, set the following options: | <ul style="list-style-type: none"> • Transmit Queues = 1 • Receive Queues = n (up to 8) • Completion Queues = # of Transmit Queues + # of Receive Queues • Interrupts = # Completion Queues + 2 • Network Virtualization using Generic Routing Encapsulation = Enabled • Interrupt Mode = Msi-X <p>Note If you set Interrupt Mode as Msi-X, and if pci=noms parameter is enabled in /boot/grub/grub.conf on RHEL system, then pci=noms would block the eNIC/fNIC driver to run in the Msi-X mode, impacting system performance.</p> <p>For more information on creating an Ethernet adapter policy, see Configuring an Ethernet Adapter Policy, on page 154.</p> |
| Step 4 | UCS-A /org/eth-policy # commit-buffer | Commits the transaction to the system configuration. |
| Step 5 | Install an eNIC driver Version 3.0.0.8 or later. | For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/Windows_b_Cisco_VIC_Drivers_for_Windows_Installation_Guide.html . |
| Step 6 | Reboot the server. | |

Example

The following example shows how to configure an Ethernet adapter policy to enable stateless offloads with NVGRE and commit the transaction:

```

UCS-A# scope org /
UCS-A /org* # create eth-policy NVGRE
UCS-A /org/eth-policy* # set descr "Ethernet adapter policy with stateless offloads"
UCS-A /org/eth-policy* # set nvgre adminstate enabled
UCS-A /org/eth-policy* # set comp-queue count 16
UCS-A /org/eth-policy* # set interrupt count 64
UCS-A /org/eth-policy* # set recv-queue count 32
UCS-A /org/eth-policy* # set rss receivesidescaling enabled
UCS-A /org/eth-policy* # set trans-queue 1
UCS-A /org/eth-policy* # set interrupt mode mxi-x
UCS-A /org/eth-policy* # commit-buffer
UCS-A /org/eth-policy #

```

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with VXLAN

Cisco UCS Manager supports VXLAN TSO and checksum offloads only with Cisco UCS VIC 1340, 1380, 1385, 1387, adapters that are running on ESXi 5.5 and later releases. Stateless offloads with VXLAN cannot be used with NetFlow, usNIC, VM-FEX, Netqueue, or VMQ.

VXLAN with Receive Side-Scaling (RSS) support starts with the Cisco UCS Manager 3.1(2) release. RSS is supported with VXLAN stateless offload on VIC adapters 1340, 1380, 1385, 1387, and SIOC on Cisco UCS S3260 system for ESXi 5.5 and later releases.



- Note** VXLAN stateless hardware offloads are not supported with Guest OS TCP traffic over IPv6 on UCS VIC 13xx adapters. To run VXLAN encapsulated TCP traffic over IPV6, disable the VXLAN stateless offloads feature.
- To disable the VXLAN stateless offload feature in UCS Manager, disable the 'Virtual Extensible LAN' field in the Ethernet Adapter Policy.
 - To disable the VXLAN stateless offload feature in the CIMC of a Cisco C-Series UCS server, uncheck 'Enable VXLAN' field in the Ethernet Interfaces pane's vNIC properties area.

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 eth-policy <i>policy-name</i> | Creates the specified Ethernet adapter policy and enters organization Ethernet policy mode. |
| Step 3 | To enable stateless offloads with VXLAN, set the following options: | <ul style="list-style-type: none"> • Transmit Queues = 1 • Receive Queues = n (up to 8) • Completion Queues = # of Transmit Queues + # of Receive Queues |

| | Command or Action | Purpose |
|---------------|---|---|
| | | <ul style="list-style-type: none"> • Interrupts = # Completion Queues + 2 • Virtual Extensible LAN = Enabled • Interrupt Mode = Msi-X <p>Note If you set Interrupt Mode as Msi-X, and if pci=noms parameter is enabled in /boot/grub/grub.conf on RHEL system, then pci=noms would block the eNIC/fNIC driver to run in the Msi-X mode, impacting system performance.</p> <ul style="list-style-type: none"> • Receive Side Scaling = Enabled <p>For more information on creating an Ethernet adapter policy, see Configuring an Ethernet Adapter Policy, on page 154.</p> |
| Step 4 | UCS-A /org/eth-policy # commit-buffer | Commits the transaction to the system configuration. |
| Step 5 | Install an eNIC driver Version 2.3.0.10 or later. | For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/ESX/2-0/b_Cisco_VIC_Drivers_for_ESX_Installation_Guide.html . |
| Step 6 | Reboot the server. | |

Example

The following example shows how to configure an Ethernet adapter policy to enable stateless offloads with VXLAN and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # create eth-policy VXLAN
UCS-A /org/eth-policy* # set descr "Ethernet adapter policy with stateless offloads"
UCS-A /org/eth-policy* # set vxlan adminstate enabled
UCS-A /org/eth-policy* # set comp-queue count 16
UCS-A /org/eth-policy* # set interrupt count 32
UCS-A /org/eth-policy* # set recv-queue count 8
UCS-A /org/eth-policy* # set rss receivesidescaling enabled
UCS-A /org/eth-policy* # set trans-queue 1
UCS-A /org/eth-policy* # set interrupt mode mxi-x
UCS-A /org/eth-policy* # commit-buffer
UCS-A /org/eth-policy #
```

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 to 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 is not responding to an IO request within the specified timeout, the FC adapter will abort the pending command then resend the same IO after the timer expires. The FC adapter valid range for this value is 1 to 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:

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

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

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

Accelerated Receive Flow Steering

Accelerated Receive Flow Steering (ARFS) is hardware-assisted receive flow steering that can increase CPU data cache hit rate by steering kernel level processing of packets to the CPU where the application thread consuming the packet is running.

Using ARFS can improve CPU efficiency and reduce traffic latency. Each receive queue of a CPU has an interrupt associated with it. You can configure the Interrupt Service Routine (ISR) to run on a CPU. The ISR moves the packet from the receive queue to the backlog of one of the current CPUs, which processes the packet later. If the application is not running on this CPU, the CPU must copy the packet to non-local memory, which adds to latency. ARFS can reduce this latency by moving that particular stream to the receive queue of the CPU on which the application is running.

ARFS is disabled by default and can be enabled through Cisco UCS Manager. To configure ARFS, do the following:

1. Create an adapter policy with ARFS enabled.
2. Associate the adapter policy with a service profile.
3. Enable ARFS on a host.
 1. Turn off Interrupt Request Queue (IRQ) balance.
 2. Associate IRQ with different CPUs.
 3. Enable ntuple by using ethtool.

Guidelines and Limitations for Accelerated Receive Flow Steering

- ARFS supports 64 filters per vNIC
- ARFS is supported on the following adapters:
 - Cisco UCS VIC 12XX
 - Cisco UCS VIC 13XX
- ARFS is supported on the following Operating Systems:

- Red Hat Enterprise Linux 6.5 and higher versions
- Red Hat Enterprise Linux 7.0 - 7.5
- SUSE Linux Enterprise Server 11 SP2 and higher versions
- SUSE Linux Enterprise Server 12 SP1 - SP3
- SUSE Linux Enterprise Server 15
- Ubuntu 14.04.2 and higher versions

Interrupt Coalescing

Adapters typically generate a large number of interrupts that a host CPU must service. Interrupt coalescing reduces the number of interrupts serviced by the host CPU. This is done by interrupting the host only once for multiple occurrences of the same event over a configurable coalescing interval.

When interrupt coalescing is enabled for receive operations, the adapter continues to receive packets, but the host CPU does not immediately receive an interrupt for each packet. A coalescing timer starts when the first packet is received by the adapter. When the configured coalescing interval times out, the adapter generates one interrupt with the packets received during that interval. The NIC driver on the host then services the multiple packets that are received. Reduction in the number of interrupts generated reduces the time spent by the host CPU on context switches. This means that the CPU has more time to process packets, which results in better throughput and latency.

Adaptive Interrupt Coalescing

Due to the coalescing interval, the handling of received packets adds to latency. For small packets with a low packet rate, this latency increases. To avoid this increase in latency, the driver can adapt to the pattern of traffic flowing through it and adjust the interrupt coalescing interval for a better response from the server.

Adaptive interrupt coalescing (AIC) is most effective in connection-oriented low link utilization scenarios including email server, databases server, and LDAP server. It is not suited for line-rate traffic.

Guidelines and Limitations for Adaptive Interrupt Coalescing

- Adaptive Interrupt Coalescing (AIC) does not provide any reduction in latency when the link utilization is more than 80 percent.
- Enabling AIC disables static coalescing.
- AIC is supported on the following Operating Systems:
 - Red Hat Enterprise Linux 6.4 and higher versions
 - SUSE Linux Enterprise Server 11 SP2 and higher versions
 - XenServer 6.5 and higher versions
 - Ubuntu 14.04.2 and higher versions

RDMA Over Converged Ethernet for SMB Direct

RDMA over Converged Ethernet (RoCE) allows direct memory access over an Ethernet network. RoCE is a link layer protocol, and hence, it allows communication between any two hosts in the same Ethernet broadcast domain. RoCE delivers superior performance compared to traditional network socket implementations because of lower latency, lower CPU utilization and higher utilization of network bandwidth. Windows 2012 and later versions use RDMA for accelerating and improving the performance of SMB file sharing and Live Migration.

Cisco UCS Manager Release 2.2(4) supports RoCE for Microsoft SMB Direct. It sends additional configuration information to the adapter while creating or modifying an Ethernet adapter policy.

Guidelines and Limitations for SMB Direct with RoCE

- Microsoft SMB Direct with RoCE is supported on Microsoft Windows, Release 2012 R2 for Cisco UCS Manager release 2.2(4) and later releases.
- For Microsoft SMB Direct with RoCE support on Microsoft Windows 2016 for Cisco UCS Manager release, check [UCS Hardware and Software Compatibility](#).
- Microsoft SMB Direct with RoCE is supported only with third generation Cisco UCS VIC 1340, 1380, 1385, 1387 adapters. Second generation UCS VIC 1225 and 1227 adapters are not supported.
- RoCE configuration is supported between Cisco adapters. Interoperability between Cisco adapters and third party adapters is not supported.
- Cisco UCS Manager does not support more than 4 RoCE-enabled vNICs per adapter.
- Cisco UCS Manager does not support RoCE with NVGRE, VXLAN, NetFlow, VMQ, or usNIC.
- Maximum number of queue pairs per adapter is 8192.
- Maximum number of memory regions per adapter is 524288.
- If you do not disable RoCE before downgrading Cisco UCS Manager from Release 2.2(4), downgrade will fail.

Configuring a Default vNIC Behavior Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org / | Enters the root organization mode. |
| Step 2 | UCS-A/org # scope vnic-beh-policy | Enters default vNIC behavior policy mode. |
| Step 3 | UCS-A/org/vnic-beh-policy # set action { hw-inherit [template_name name] none } | Specifies the default vNIC behavior policy. This can be one of the following: <ul style="list-style-type: none"> • hw-inherit—If a service profile requires vNICs and none have been explicitly defined, Cisco UCS Manager creates the required vNICs based on the adapter |

| | Command or Action | Purpose |
|---------------|--|---|
| | | <p>installed in the server associated with the service profile.</p> <p>If you specify hw-inherit, you can also specify a vNIC template to create the vNICs.</p> <ul style="list-style-type: none"> • none—Cisco UCS Manager does not create default vNICs for a service profile. All vNICs must be explicitly created. |
| Step 4 | UCS-A/org/vnic-beh-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

This example shows how to set the default vNIC behavior policy to **hw-inherit**:

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

Deleting a vNIC from a LAN 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 lan-connectivity-policy <i>policy-name</i> | Enters LAN connectivity policy mode for the specified LAN connectivity policy. |
| Step 3 | UCS-A /org/lan-connectivity-policy # delete vnic <i>vnic-name</i> | Deletes the specified vNIC from the LAN connectivity policy. |
| Step 4 | UCS-A /org/lan-connectivity-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to delete a vNIC named vnic3 from a LAN connectivity policy named LanConnect42 and commit the transaction:

```

UCS-A# scope org /
UCS-A /org # scope lan-connectivity-policy LanConnect42
UCS-A /org/lan-connectivity-policy # delete vnic vnic3
UCS-A /org/lan-connectivity-policy* # commit-buffer
UCS-A /org/lan-connectivity-policy #

```

Creating a LAN 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 # create lan-connectivity-policy <i>policy-name</i> | Creates the specified LAN connectivity policy, and enters organization LAN connectivity 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/lan-connectivity-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a LAN connectivity policy named LanConnect42 and commit the transaction:

```

UCS-A# scope org /
UCS-A /org* # create lan-connectivity-policy LanConnect42
UCS-A /org/lan-connectivity-policy* # set descr "LAN connectivity policy"
UCS-A /org/lan-connectivity-policy* # commit-buffer
UCS-A /org/lan-connectivity-policy #

```

What to do next

Add one or more vNICs and/or iSCSI vNICs to this LAN connectivity policy.

Deleting a LAN Connectivity Policy

If you delete a LAN connectivity policy that is included in a service profile, it also deletes all vNICs and iSCSI vNICs from that service profile, and disrupt LAN 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 lan-connectivity-policy <i>policy-name</i> | Deletes the specified LAN connectivity policy. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to delete the LAN connectivity policy named LanConnectiSCSI42 from the root organization and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # delete lan-connectivity-policy LanConnectiSCSI42
UCS-A /org* # commit-buffer
UCS-A /org #
```

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 LAN Connectivity Policy

Procedure

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

| | Command or Action | Purpose |
|---------------|---|---|
| Step 2 | UCS-A /org # create lan-connectivity-policy <i>policy-name</i> | Creates the specified LAN connectivity policy, and enters organization LAN connectivity 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/lan-connectivity-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a LAN connectivity policy named LanConnect42 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org* # create lan-connectivity-policy LanConnect42
UCS-A /org/lan-connectivity-policy* # set descr "LAN connectivity policy"
UCS-A /org/lan-connectivity-policy* # commit-buffer
UCS-A /org/lan-connectivity-policy #
```

What to do next

Add one or more vNICs and/or iSCSI vNICs to this LAN connectivity policy.

Creating a vNIC for a LAN Connectivity Policy

If you are continuing from [Creating a LAN Connectivity Policy, on page 165](#), 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> . |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 2 | UCS-A /org # scope lan-connectivity-policy <i>policy-name</i> | Enters LAN connectivity policy mode for the specified LAN connectivity policy. |
| Step 3 | UCS-A /org/lan-connectivity-policy # create vnic vnic-name [eth-if <i>eth-if-name</i>] [fabric { a b }] | Creates a vNIC for the specified LAN connectivity policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved. |
| Step 4 | UCS-A /org/lan-connectivity-policy/vnic # set fabric { a a-b b b-a } | Specifies the fabric to use for the vNIC. If you did not specify the fabric when you created the vNIC in Step 3, you have the option to specify it with this command. If you want this vNIC to be able to access the second fabric interconnect if the default one is unavailable, choose a-b (A is the primary) or b-a (B is the primary) . Note Do not enable fabric failover for the vNIC under the following circumstances: <ul style="list-style-type: none"> • If the Cisco UCS domain is running in Ethernet Switch Mode. vNIC fabric failover is not supported in that mode. If all Ethernet uplinks on one fabric interconnect fail, the vNICs do not fail over to the other. • If you plan to associate this vNIC to a server with an adapter that does not support fabric failover, such as the Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter. If you do so, Cisco UCS Manager generates a configuration fault when you associate the service profile with the server. |
| Step 5 | UCS-A /org/lan-connectivity-policy/vnic # set adapter-policy <i>policy-name</i> | Specifies the adapter policy to use for the vNIC. |

| | Command or Action | Purpose |
|----------------|--|--|
| Step 6 | UCS-A /org/lan-connectivity-policy/vnic # set identity { dynamic-mac { <i>mac-addr</i> derived } mac-pool <i>mac-pool-name</i> } | Specifies the identity (MAC address) for the vNIC. You can set the identity using one of the following options: <ul style="list-style-type: none"> • Create a unique MAC address in the form <i>nn:nn:nn:nn:nn:nn</i>. • Derive the MAC address from one burned into the hardware at manufacture. • Assign a MAC address from a MAC pool. |
| Step 7 | UCS-A /org/lan-connectivity-policy/vnic # set mtu <i>size-num</i> | Specifies the maximum transmission unit, or packet size, that this vNIC accepts. Enter an integer between 1500 and 9216. Note If the vNIC has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets might get dropped during data transmission. |
| Step 8 | UCS-A /org/lan-connectivity-policy/vnic # set nw-control-policy <i>policy-name</i> | Specifies the network control policy that the vNIC should use. |
| Step 9 | UCS-A /org/lan-connectivity-policy/vnic # set order { <i>order-num</i> unspecified } | Specifies the relative order for the vNIC. |
| Step 10 | UCS-A /org/lan-connectivity-policy/vnic # set pin-group <i>group-name</i> | Specifies the LAN pin group that the vNIC should use. |
| Step 11 | UCS-A /org/lan-connectivity-policy/vnic # set qos-policy <i>policy-name</i> | Specifies the quality of service policy that the vNIC should use. |
| Step 12 | UCS-A /org/lan-connectivity-policy/vnic # set stats-policy <i>policy-name</i> | Specifies the statistics collection policy that the vNIC should use. |
| Step 13 | UCS-A /org/lan-connectivity-policy/vnic # set template-name <i>policy-name</i> | Specifies the dynamic vNIC connectivity policy to use for the vNIC. |
| Step 14 | UCS-A /org/lan-connectivity-policy/vnic # set vcon { 1 2 3 4 any } | Assigns the vNIC to the specified vCon. Use the any keyword to have Cisco UCS Manager automatically assign the vNIC. |
| Step 15 | UCS-A /org/lan-connectivity-policy/vnic # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to configure a vNIC for a LAN connectivity policy named LanConnect42 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope lan-connectivity-policy LanConnect42
UCS-A /org/lan-connectivity-policy* # create vnic vnic3 fabric a
UCS-A /org/lan-connectivity-policy/vnic* # set fabric a-b
UCS-A /org/lan-connectivity-policy/vnic* # set adapter-policy AdaptPol2
UCS-A /org/lan-connectivity-policy/vnic* # set identity mac-pool MacPool3
UCS-A /org/lan-connectivity-policy/vnic* # set mtu 8900
UCS-A /org/lan-connectivity-policy/vnic* # set nw-control-policy ncp5
UCS-A /org/lan-connectivity-policy/vnic* # set order 0
UCS-A /org/lan-connectivity-policy/vnic* # set pin-group EthPinGroup12
UCS-A /org/lan-connectivity-policy/vnic* # set qos-policy QosPol5
UCS-A /org/lan-connectivity-policy/vnic* # set stats-policy StatsPol2
UCS-A /org/lan-connectivity-policy/vnic* # set template-name VnicConnPol3
UCS-A /org/lan-connectivity-policy/vnic* # set vcon any
UCS-A /org/lan-connectivity-policy/vnic* # commit-buffer
UCS-A /org/lan-connectivity-policy/vnic #
```

What to do next

If desired, add another vNIC or an iSCSI vNIC to the LAN connectivity policy. If not, include the policy in a service profile or service profile template.

Deleting a vNIC from a LAN 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 lan-connectivity-policy <i>policy-name</i> | Enters LAN connectivity policy mode for the specified LAN connectivity policy. |
| Step 3 | UCS-A /org/lan-connectivity-policy # delete vnic <i>vnic-name</i> | Deletes the specified vNIC from the LAN connectivity policy. |
| Step 4 | UCS-A /org/lan-connectivity-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to delete a vNIC named vnic3 from a LAN connectivity policy named LanConnect42 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope lan-connectivity-policy LanConnect42
```

```
UCS-A /org/lan-connectivity-policy # delete vnic vnic3
UCS-A /org/lan-connectivity-policy* # commit-buffer
UCS-A /org/lan-connectivity-policy #
```

Creating an iSCSI vNIC for a LAN Connectivity Policy

If you are continuing from [Creating a LAN Connectivity Policy, on page 165](#), begin this procedure at Step 3.

Before you begin

The LAN connectivity policy must include an Ethernet vNIC that can be used as the overlay vNIC for the iSCSI device.

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 lan-connectivity-policy <i>policy-name</i> | Enters LAN connectivity policy mode for the specified LAN connectivity policy. |
| Step 3 | UCS-A /org/lan-connectivity-policy # create vnic-iscsi <i>iscsi-vnic-name</i> . | Creates an iSCSI vNIC for the specified LAN connectivity policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved. |
| Step 4 | (Optional) UCS-A /org/lan-connectivity-policy/vnic-iscsi # set iscsi-adaptor-policy <i>iscsi-adaptor-name</i> | Specifies the iSCSI adaptor policy that you have created for this iSCSI vNIC. |
| Step 5 | (Optional) UCS-A /org/lan-connectivity-policy/vnic-iscsi # set auth-name <i>authentication-profile-name</i> | Sets the authentication profile to be used by the iSCSI vNIC. The authentication profile must already exist for it to be set. For more information, see <i>Creating an Authentication Profile</i> . |
| Step 6 | UCS-A /org/lan-connectivity-policy/vnic-iscsi # set identity { dynamic-mac { <i>dynamic-mac-address</i> derived } mac-pool <i>mac-pool-name</i> } | Specifies the MAC address for the iSCSI vNIC. Note The MAC address is set only for the Cisco UCS NIC M51KR-B Adapters. |
| Step 7 | UCS-A /org/lan-connectivity-policy/vnic-iscsi # set iscsi-identity { <i>initiator-name</i> | Specifies the name of the iSCSI initiator or the name of an IQN pool from which the iSCSI |

| | Command or Action | Purpose |
|----------------|--|---|
| | <i>initiator-name</i> initiator-pool-name <i>iqn-pool-name</i> } | initiator name will be provided. The iSCSI initiator name can be up to 223 characters. |
| Step 8 | UCS-A /org/lan-connectivity-policy/vnic-iscsi # set overlay-vnic-name <i>overlay-vnic-name</i> | Specifies the Ethernet vNIC that is used by the iSCSI device as the overlay vNIC. For more information, see <i>Configuring a vNIC for a Service Profile</i> . |
| Step 9 | UCS-A /org/lan-connectivity-policy/vnic-iscsi # create eth-if | Creates an Ethernet interface for a VLAN assigned to the iSCSI vNIC. |
| Step 10 | UCS-A /org/ex/vnic-iscsi/eth-if # set vlnaname <i>vlan-name</i> | Specifies the VLAN name. The default VLAN is default. For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS VIC-1240 Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC. For the Cisco UCS M51KR-B Broadcom BCM57711 Adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC. |
| Step 11 | UCS-A /org/lan-connectivity-policy/vnic-iscsi # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to configure an iSCSI vNIC for a LAN connectivity policy named LanConnect42 and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope lan-connectivity-policy LanConnect42
UCS-A /org/lan-connectivity-policy # create vnic-iscsi iSCSI1
UCS-A /org/lan-connectivity-policy/vnic-iscsi* # set iscsi-adaptor-policy iscsiboot
UCS-A /org/lan-connectivity-policy/vnic-iscsi* # set auth-name initauth
UCS-A /org/lan-connectivity-policy/vnic-iscsi* # set identity dynamic-mac derived
UCS-A /org/lan-connectivity-policy/vnic-iscsi* # set iscsi-identity initiator-name iSCSI1
UCS-A /org/lan-connectivity-policy/vnic-iscsi* # set overlay-vnic-name eth1
UCS-A /org/lan-connectivity-policy/vnic-iscsi* # create eth-if
UCS-A /org/lan-connectivity-policy/vnic-iscsi/eth-if* # set vlnaname default
UCS-A /org/lan-connectivity-policy/vnic-iscsi/eth-if* # commit buffer
UCS-A /org/lan-connectivity-policy/vnic-iscsi/eth-if
```

What to do next

If desired, add another iSCSI vNIC or a vNIC to the LAN connectivity policy. If not, include the policy in a service profile or service profile template.

Deleting an iSCSI vNIC from a LAN 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 lan-connectivity-policy <i>policy-name</i> | Enters LAN connectivity policy mode for the specified LAN connectivity policy. |
| Step 3 | UCS-A /org/lan-connectivity-policy # delete vnic-iscsi <i>iscsi-vnic-name</i> | Deletes the specified iSCSI vNIC from the LAN connectivity policy. |
| Step 4 | UCS-A /org/lan-connectivity-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to delete an iSCSI vNIC named `iscsivnic3` from a LAN connectivity policy named `LanConnect42` and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope lan-connectivity-policy LanConnect42
UCS-A /org/lan-connectivity-policy # delete vnic-iscsi iscsivnic3
UCS-A /org/lan-connectivity-policy* # commit-buffer
UCS-A /org/lan-connectivity-policy #
```

Network Control Policy

This policy configures the network control settings for the Cisco UCS domain, including the following:

- Whether the Cisco Discovery Protocol (CDP) is enabled or disabled
- How the virtual interface (VIF) behaves if no uplink port is available in end-host mode
- The action that Cisco UCS Manager takes on the remote Ethernet interface, vEthernet interface , or vFibre Channel interface when the associated border port fails
- Whether the server can use different MAC addresses when sending packets to the fabric interconnect
- Whether MAC registration occurs on a per-VNIC basis or for all VLANs

Action on Uplink Fail

By default, the **Action on Uplink Fail** property in the network control policy is configured with a value of `link-down`. For adapters such as the Cisco UCS M81KR Virtual Interface Card, this default behavior directs Cisco UCS Manager to bring the vEthernet or vFibre Channel interface down if the associated border port fails. For Cisco UCS systems using a non-VM-FEX capable converged network adapter that supports both Ethernet and FCoE traffic, such as Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, this

default behavior directs Cisco UCS Manager to bring the remote Ethernet interface down if the associated border port fails. In this scenario, any vFibre Channel interfaces that are bound to the remote Ethernet interface are brought down as well.



Note If your implementation includes non-VM-FEX capable converged network adapters mentioned in this section and the adapter is expected to handle both Ethernet and FCoE traffic, we recommend that you configure the **Action on Uplink Fail** property with a value of warning. This configuration might result in an Ethernet teaming driver being unable to detect a link failure when the border port goes down.

MAC Registration Mode

MAC addresses are installed only on the native VLAN by default, which maximizes the VLAN port count in most implementations.



Note If a trunking driver is being run on the host and the interface is in promiscuous mode, we recommend that you set the MAC Registration Mode to All VLANs.

Configuring a Network Control Policy

MAC address-based port security for Emulex converged Network Adapters (N20-AE0102) is not supported. When MAC address-based port security is enabled, the fabric interconnect restricts traffic to packets that contain the MAC address that it first learns. This is either the source MAC address used in the FCoE Initialization Protocol packet, or the MAC address in an ethernet packet, whichever is sent first by the adaptor. This configuration can result in either FCoE or Ethernet packets being dropped.

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 nw-ctrl-policy <i>policy-name</i> | Creates the specified network control policy, and enters organization network control policy mode. |
| Step 3 | UCS-A /org/nw-ctrl-policy # {disable enable} cdp | Disables or enables Cisco Discovery Protocol (CDP). |
| Step 4 | UCS-A /org/nw-ctrl-policy # {disable enable} lldp transmit | Disables or enables the transmission of LLDP packets on an interface. |
| Step 5 | UCS-A /org/nw-ctrl-policy # {disable enable} lldp receive | Disables or enables the reception of LLDP packets on an interface. |
| Step 6 | UCS-A /org/nw-ctrl-policy # set uplink-fail-action {link-down warning} | Specifies the action to be taken when no uplink port is available in end-host mode. |

| | Command or Action | Purpose |
|----------------|--|--|
| | | Use the link-down keyword to change the operational state of a vNIC to down when uplink connectivity is lost on the fabric interconnect, and facilitate fabric failover for vNICs. Use the warning keyword to maintain server-to-server connectivity even when no uplink port is available, and disable fabric failover when uplink connectivity is lost on the fabric interconnect. The default uplink failure action is link-down. |
| Step 7 | UCS-A /org/nw-ctrl-policy # set mac-registration-mode {all-host-vlans only-native-vlan} | Whether adapter-registered MAC addresses are added only to the native VLAN associated with the interface or added to all VLANs associated with the interface. This can be one of the following: <ul style="list-style-type: none"> • Only Native Vlan—MAC addresses are only added to the native VLAN. This option is the default, and it maximizes the port+VLAN count. • All Host Vlans—MAC addresses are added to all VLANs with which they are associated. Select this option if your VLANs are configured to use trunking but are <i>not</i> running in Promiscuous mode. |
| Step 8 | UCS-A /org/nw-ctrl-policy # create mac-security | Enters organization network control policy MAC security mode. |
| Step 9 | UCS-A /org/nw-ctrl-policy/mac-security # set forged-transmit {allow deny} | Allows or denies the forging of MAC addresses when sending traffic. MAC security is disabled when forged MAC addresses are allowed, and MAC security is enabled when forged MAC addresses are denied. By default, forged MAC addresses are allowed (MAC security is disabled). |
| Step 10 | UCS-A /org/nw-ctrl-policy/mac-security # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a network control policy named ncp5, enable CDP, enable LLDP transmit and LLDP receive, set the uplink fail action to link-down, deny forged MAC addresses (enable MAC security), and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # create nw-ctrl-policy ncp5
UCS-A /org/nw-ctrl-policy* # enable cdp
```



```

UCS-A /org/nw-ctrl-policy* # enable lldp transmit
UCS-A /org/nw-ctrl-policy* # enable lldp receive
UCS-A /org/nw-ctrl-policy* # set uplink-fail-action link-down
UCS-A /org/nw-ctrl-policy* # create mac-security
UCS-A /org/nw-ctrl-policy/mac-security* # set forged-transmit deny
UCS-A /org/nw-ctrl-policy/mac-security* # commit-buffer
UCS-A /org/nw-ctrl-policy/mac-security #

```

The following example shows how to create a network control policy named ncp5, enable CDP, set the uplink fail action to link-down, and commit the transaction:

```

UCS-A# scope org /
UCS-A /org # create nw-ctrl-policy ncp5
UCS-A /org/nw-ctrl-policy* # enable cdp
UCS-A /org/nw-ctrl-policy* # set uplink-fail-action link-down
UCS-A /org/nw-ctrl-policy* # commit-buffer
UCS-A /org/nw-ctrl-policy #

```

Configuring Link Layer Discovery Protocol for Fabric Interconnect vEthernet Interfaces

Cisco UCS Manager allows you to enable and disable LLDP on a vEthernet interface. You can also retrieve information about these LAN uplink neighbors. This information is useful while learning the topology of the LAN connected to the UCS system and while diagnosing any network connectivity issues from the Fabric Interconnect (FI). The FI of a UCS system is connected to LAN uplink switches for LAN connectivity and to SAN uplink switches for storage connectivity. When using Cisco UCS with Cisco Application Centric Infrastructure (ACI), LAN uplinks of the FI are connected to ACI leaf nodes. Enabling LLDP on a vEthernet interface will help the Application Policy Infrastructure Controller (APIC) to identify the servers connected to the FI by using vCenter.

To permit the discovery of devices in a network, support for Link Layer Discovery Protocol (LLDP), a vendor-neutral device discovery protocol that is defined in the IEEE 802.1ab standard, is introduced. LLDP is a one-way protocol that allows network devices to advertise information about themselves to other devices on the network. LLDP transmits information about the capabilities and current status of a device and its interfaces. LLDP devices use the protocol to solicit information only from other LLDP devices.

You can enable or disable LLDP on a vEthernet interface based on the Network Control Policy (NCP) that is applied on the vNIC in the service profile.

Displaying Network Control Policy Details

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 nw-ctrl-policy { default <i>policy-name</i> } | Enters organization network control policy mode for the specified network control policy. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 3 | UCS-A /org/nw-ctrl-policy # show detail | Displays details about the specified network control policy. |

Example

The following example shows how to display the details of a network control policy named ncp5:

```
UCS-A# scope org /
UCS-A /org # scope nw-ctrl-policy ncp5
UCS-A /org/nw-ctrl-policy* # show detail

Network Control Policy:
  Name: ncp5
  CDP: Enabled
  LLDP Transmit: Enabled
  LLDP Receive: Enabled
  Uplink fail action: Link Down
  Adapter MAC Address Registration: Only Native Vlan
  Policy Owner: Local
  Description:

UCS-A /org/nw-ctrl-policy #
```

Deleting a Network Control Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org / | Enters the root organization mode. |
| Step 2 | UCS-A /org # delete nwctrl-policy <i>policy-name</i> | Deletes the specified network control policy. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example deletes the network control policy named ncp5 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete nwctrl-policy ncp5
UCS-A /org* # commit-buffer
UCS-A /org #
```

Creating a Multicast Policy

A multicast policy can be created only in the root organization and not in a sub-organization.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # create mcast-policy <i>policy-name</i> | Creates a multicast policy with the specified policy name, and enters organization multicast policy mode. |
| Step 3 | UCS-A /org/mcast-policy* # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a multicast policy named policy1:

```
UCS-A# scope org /
UCS-A /org # create mcast-policy policy1
UCS-A /org/mcast-policy* # commit-buffer
UCS-A /org/mcast-policy #
```

Deleting a Multicast Policy



Note If you assigned a non-default (user-defined) multicast policy to a VLAN and then delete that multicast policy, the associated VLAN inherits the multicast policy settings from the default multicast policy until the deleted policy is re-created.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # delete mcast-policy <i>policy-name</i> | Deletes a multicast policy with the specified policy name. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to delete a multicast policy named policy1:

```
UCS-A # scope org /
UCS-A /org # delete mcast-policy policy1
```

```
UCS-A /org* # commit-buffer
UCS-A /org #
```

Entering Multicast Policy Mode

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # scope mcast-policy <i>policy-name</i> | Enters organization multicast policy mode. |

Example

The following example shows how to create a multicast policy named policy1:

```
UCS-A# scope org /
UCS-A /org # scope mcast-policy policy1
UCS-A /org/mcast-policy #
```

Enter a Multicast Policy

You can enter an existing multicast policy using the **enter mcast-policy** *policy-name* command.

Before you begin

Create a multicast policy.

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # enter mcast-policy <i>policy-name</i> | Creates a new multicast policy with the specified policy name, and enters organization multicast policy mode. |

Example

The following example shows how to create a multicast policy named policy1 and enter mcast-policy mode:

```
UCS-A# scope org /
UCS-A /org # enter mcast-policy policy1
UCS-A /org/mcast-policy #
```

Assigning a Global VLAN Multicast Policy

You can assign a multicast policy to a global VLAN in the Ethernet uplink fabric mode.

Before you begin

Create a VLAN.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # scope vlan default | Enters Ethernet uplink VLAN mode. |
| Step 3 | UCS-A /eth-uplink/vlan # set mcastpolicy <i>policy-name</i> | Assigns a multicast policy to a global VLAN. |
| Step 4 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. |

Disassociating a Global VLAN Multicast Policy

You can disassociate a multicast policy from global VLANs in the Ethernet uplink fabric mode.



Note If you assigned a non-default (user-defined) multicast policy to a VLAN and then delete that multicast policy, the associated VLAN inherits the multicast policy settings from the default multicast policy until the deleted policy is re-created.

Before you begin

Create a Global VLAN and associate a multicast policy.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # scope vlan default | Enters Ethernet uplink VLAN mode. |
| Step 3 | UCS-A /eth-uplink/vlan # set mcastpolicy "" | Disassociates any multicast policy from the global VLAN. If you configure set mcastpolicy |

| | Command or Action | Purpose |
|---------------|---|---|
| | | "" in a VLAN, the VLAN will inherit multicast settings from the default multicast policy. |
| Step 4 | UCS-A /eth-uplink/vlan # commit-buffer | Commits the transaction to the system configuration. |

Disassociating a VLAN Multicast Policy

You can disassociate a VLAN from any multicast policy in the Ethernet uplink fabric mode by entering an empty string ("") as the policy name.

Before you begin

Create a VLAN and associate a multicast policy to the VLAN.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | Required: UCS-A /eth-uplink # scope fabric {a b} | Enters Ethernet uplink fabric mode for the specified fabric interconnect. |
| Step 3 | UCS-A /eth-uplink/fabric # scope vlan vlan-name | Enters Ethernet uplink fabric VLAN mode. |
| Step 4 | UCS-A /eth-uplink/fabric/vlan # set mcastpolicy "" | Disassociates any multicast policy for the VLAN. If you configure set mcastpolicy "" in a VLAN, the VLAN will inherit multicast settings from the default multicast policy. |
| Step 5 | UCS-A /eth-uplink/fabric/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example disassociates any multicast policy from a VLAN named vlan1 and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope vlan vlan1
UCS-A /eth-uplink/fabric/vlan # set mcastpolicy policy1
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Configuring Ethernet Adapter Policies

Configuring an Ethernet 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 <i>/</i> as the <i>org-name</i> . |
| Step 2 | UCS-A /org # create eth-policy <i>policy-name</i> | Creates the specified Ethernet adapter policy and enters organization Ethernet policy mode. |
| Step 3 | (Optional) UCS-A /org/eth-policy # set arfs acceleratdrfs { enabled disabled } | Configures Accelerated RFS. |
| Step 4 | (Optional) UCS-A /org/eth-policy # set comp-queue count <i>count</i> | Configures the Ethernet completion queue. |
| Step 5 | (Optional) UCS-A /org/eth-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 6 | (Optional) UCS-A /org/eth-policy # set failover timeout <i>timeout-sec</i> | Configures the Ethernet failover. |
| Step 7 | (Optional) UCS-A /org/eth-policy # set interrupt { coalescing-time <i>sec</i> coalescing-type { idle min } count <i>count</i> mode { intx msi msi-x }} | Configures the Ethernet interrupt. |
| Step 8 | (Optional) UCS-A /org/eth-policy # set nvgre adminstate { disabled enabled } | Configures NVGRE. |
| Step 9 | (Optional) UCS-A /org/eth-policy # set offload { large-receive tcp-rx-checksum tcp-segment tcp-tx-checksum } { disabled enabled } | Configures the Ethernet offload. |
| Step 10 | (Optional) UCS-A /org/eth-policy # set policy-owner { local pending } | Specifies the owner for the Ethernet adapter policy. |
| Step 11 | (Optional) UCS-A /org/eth-policy # set recv-queue { count <i>count</i> ring-size <i>size-num</i> } | Configures the Ethernet receive queue. |

| | Command or Action | Purpose |
|----------------|---|--|
| Step 12 | (Optional) UCS-A /org/eth-policy # set rss receivesidescaling {disabled enabled} | Configures the RSS. |
| Step 13 | (Optional) UCS-A /org/eth-policy # set trans-queue {count <i>count</i> ring-size <i>size-num</i> } | Configures the Ethernet transmit queue. |
| Step 14 | (Optional) UCS-A /org/eth-policy # set vxlan adminstate {disabled enabled} | Configures VXLAN. |
| Step 15 | UCS-A /org/eth-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example configures an Ethernet adapter policy, and commits the transaction:

```
UCS-A# scope org
UCS-A /org* # create eth-policy EthPolicy19
UCS-A /org/eth-policy* # set comp-queue count 16
UCS-A /org/eth-policy* # set descr "This is an Ethernet adapter policy example."
UCS-A /org/eth-policy* # set failover timeout 300
UCS-A /org/eth-policy* # set interrupt count 64
UCS-A /org/eth-policy* # set offload large-receive disabled
UCS-A /org/eth-policy* # set recv-queue count 32
UCS-A /org/eth-policy* # set rss receivesidescaling enabled
UCS-A /org/eth-policy* # set trans-queue
UCS-A /org/eth-policy* # commit-buffer
UCS-A /org/eth-policy #
```

Deleting an Ethernet 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 # delete eth-policy <i>policy-name</i> | Deletes the specified Ethernet adapter policy. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example deletes the Ethernet adapter policy named EthPolicy19 and commits the transaction:


```
UCS-A# scope org /
UCS-A /org # delete eth-policy EthPolicy19
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring the Default vNIC Behavior Policy

Default vNIC Behavior Policy

Default vNIC behavior policy allows you to configure how vNICs are created for a service profile. You can choose to create vNICs manually, or you can create them automatically.

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

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



Note If you do not specify a default behavior policy for vNICs, **HW Inherit** is used by default.

Configuring a Default vNIC Behavior Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org / | Enters the root organization mode. |
| Step 2 | UCS-A/org # scope vnic-beh-policy | Enters default vNIC behavior policy mode. |
| Step 3 | UCS-A/org/vnic-beh-policy # set action { hw-inherit [template_name name] none } | Specifies the default vNIC behavior policy. This can be one of the following: <ul style="list-style-type: none"> • hw-inherit—If a service profile requires vNICs and none have been explicitly defined, Cisco UCS Manager creates the required vNICs based on the adapter installed in the server associated with the service profile. <p>If you specify hw-inherit, you can also specify a vNIC template to create the vNICs.</p> |

| | Command or Action | Purpose |
|---------------|--|--|
| | | <ul style="list-style-type: none"> none—Cisco UCS Manager does not create default vNICs for a service profile. All vNICs must be explicitly created. |
| Step 4 | UCS-A/org/vnic-beh-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

This example shows how to set the default vNIC behavior policy to **hw-inherit**:

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

Configuring a Network Control Policy

MAC address-based port security for Emulex converged Network Adapters (N20-AE0102) is not supported. When MAC address-based port security is enabled, the fabric interconnect restricts traffic to packets that contain the MAC address that it first learns. This is either the source MAC address used in the FCoE Initialization Protocol packet, or the MAC address in an ethernet packet, whichever is sent first by the adaptor. This configuration can result in either FCoE or Ethernet packets being dropped.

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 nw-ctrl-policy <i>policy-name</i> | Creates the specified network control policy, and enters organization network control policy mode. |
| Step 3 | UCS-A /org/nw-ctrl-policy # {disable enable} cdp | Disables or enables Cisco Discovery Protocol (CDP). |
| Step 4 | UCS-A /org/nw-ctrl-policy # {disable enable} lldp transmit | Disables or enables the transmission of LLDP packets on an interface. |
| Step 5 | UCS-A /org/nw-ctrl-policy # {disable enable} lldp receive | Disables or enables the reception of LLDP packets on an interface. |
| Step 6 | UCS-A /org/nw-ctrl-policy # set uplink-fail-action {link-down warning} | Specifies the action to be taken when no uplink port is available in end-host mode. |

| | Command or Action | Purpose |
|----------------|--|--|
| | | Use the link-down keyword to change the operational state of a vNIC to down when uplink connectivity is lost on the fabric interconnect, and facilitate fabric failover for vNICs. Use the warning keyword to maintain server-to-server connectivity even when no uplink port is available, and disable fabric failover when uplink connectivity is lost on the fabric interconnect. The default uplink failure action is link-down. |
| Step 7 | UCS-A /org/nw-ctrl-policy # set mac-registration-mode {all-host-vlans only-native-vlan} | Whether adapter-registered MAC addresses are added only to the native VLAN associated with the interface or added to all VLANs associated with the interface. This can be one of the following: <ul style="list-style-type: none"> • Only Native Vlan—MAC addresses are only added to the native VLAN. This option is the default, and it maximizes the port+VLAN count. • All Host Vlans—MAC addresses are added to all VLANs with which they are associated. Select this option if your VLANs are configured to use trunking but are <i>not</i> running in Promiscuous mode. |
| Step 8 | UCS-A /org/nw-ctrl-policy # create mac-security | Enters organization network control policy MAC security mode. |
| Step 9 | UCS-A /org/nw-ctrl-policy/mac-security # set forged-transmit {allow deny} | Allows or denies the forging of MAC addresses when sending traffic. MAC security is disabled when forged MAC addresses are allowed, and MAC security is enabled when forged MAC addresses are denied. By default, forged MAC addresses are allowed (MAC security is disabled). |
| Step 10 | UCS-A /org/nw-ctrl-policy/mac-security # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a network control policy named ncp5, enable CDP, enable LLDP transmit and LLDP receive, set the uplink fail action to link-down, deny forged MAC addresses (enable MAC security), and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # create nw-ctrl-policy ncp5
UCS-A /org/nw-ctrl-policy* # enable cdp
```

```

UCS-A /org/nw-ctrl-policy* # enable lldp transmit
UCS-A /org/nw-ctrl-policy* # enable lldp receive
UCS-A /org/nw-ctrl-policy* # set uplink-fail-action link-down
UCS-A /org/nw-ctrl-policy* # create mac-security
UCS-A /org/nw-ctrl-policy/mac-security* # set forged-transmit deny
UCS-A /org/nw-ctrl-policy/mac-security* # commit-buffer
UCS-A /org/nw-ctrl-policy/mac-security #

```

The following example shows how to create a network control policy named ncp5, enable CDP, set the uplink fail action to link-down, and commit the transaction:

```

UCS-A# scope org /
UCS-A /org # create nw-ctrl-policy ncp5
UCS-A /org/nw-ctrl-policy* # enable cdp
UCS-A /org/nw-ctrl-policy* # set uplink-fail-action link-down
UCS-A /org/nw-ctrl-policy* # commit-buffer
UCS-A /org/nw-ctrl-policy #

```

Deleting a Network Control Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org / | Enters the root organization mode. |
| Step 2 | UCS-A /org # delete nwctrl-policy <i>policy-name</i> | Deletes the specified network control policy. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example deletes the network control policy named ncp5 and commits the transaction:

```

UCS-A# scope org /
UCS-A /org # delete nwctrl-policy ncp5
UCS-A /org* # commit-buffer
UCS-A /org #

```

Configuring Multicast Policies

Multicast Policy

This policy is used to configure Internet Group Management Protocol (IGMP) snooping and IGMP querier. IGMP Snooping dynamically determines hosts in a VLAN that should be included in particular multicast transmissions. You can create, modify, and delete a multicast policy that can be associated to one or more VLANs. When a multicast policy is modified, all VLANs associated with that multicast policy are re-processed to apply the changes.

By default, IGMP snooping is enabled and IGMP querier is disabled. When IGMP snooping is enabled, the fabric interconnects send the IGMP queries only to the hosts. They do not send IGMP queries to the upstream network. To send IGMP queries to the upstream, do one of the following:

- Configure IGMP querier on the upstream fabric interconnect with IGMP snooping enabled
- Disable IGMP snooping on the upstream fabric interconnect
- Change the fabric interconnects to switch mode

The following limitations and guidelines apply to multicast policies:

- On a 6200 series fabric interconnect, user-defined multicast policies can also be assigned along with the default multicast policy.
- Only the default multicast policy is allowed for a global VLAN.
- If a Cisco UCS domain includes 6300 and 6200 series fabric interconnects, any multicast policy can be assigned.
- We highly recommend you use the same IGMP snooping state on the fabric interconnects and the associated LAN switches. For example, if IGMP snooping is disabled on the fabric interconnects, it should be disabled on any associated LAN switches as well.

Creating a Multicast Policy

A multicast policy can be created only in the root organization and not in a sub-organization.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # create mcast-policy <i>policy-name</i> | Creates a multicast policy with the specified policy name, and enters organization multicast policy mode. |
| Step 3 | UCS-A /org/mcast-policy* # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a multicast policy named policy1:

```
UCS-A# scope org /
UCS-A /org # create mcast-policy policy1
UCS-A /org/mcast-policy* # commit-buffer
UCS-A /org/mcast-policy #
```

Configuring IGMP Snooping Parameters

You can enable or disable IGMP snooping for a multicast policy. By default, the IGMP snooping state is enabled for a multicast policy. You can also set the IGMP snooping querier state and IPv4 address for the multicast policy.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # create mcast-policy <i>policy-name</i> | Creates a new multicast policy with the specified policy name, and enters organization multicast policy mode. |
| Step 3 | UCS-A /org/mcast-policy* # set querier {enabled disabled} | Enables or disables IGMP snooping querier. By default, IGMP snooping querier is disabled for a multicast policy. |
| Step 4 | UCS-A /org/mcast-policy* # set querierip <i>IGMP snooping querier IPv4 address</i> | Specifies the IPv4 address for the IGMP snooping querier. |
| Step 5 | UCS-A /org/mcast-policy* # set snooping {enabled disabled} | Enables or disables IGMP snooping. By default, IGMP snooping is enabled for a multicast policy. |
| Step 6 | UCS-A /org/mcast-policy* # commit-buffer | Commits the transaction to the system configuration. Note Follow these guidelines if you choose to set IGMP Snooping querier IP addresses for a multicast policy: <ol style="list-style-type: none"> 1. In the Ethernet Switch-Mode configuration, you must set the querier IP addresses for each FI in the domain. 2. In the Ethernet End-Host mode, you can set the querier IP address just for FI A, and optionally for FI B as well. If an IP address is not set explicitly for FI-B, it uses the same address set for FI A. |

Example

The following example shows how to create and enter a multicast policy named policy1:

```

UCS-A# scope org /
UCS-A /org # create mcast-policy policy1
UCS-A /org/mcast-policy* # set querier enabled
UCS-A /org/mcast-policy* # set querierip 1.2.3.4
UCS-A /org/mcast-policy* # set snooping enabled
UCS-A /org/mcast-policy* # commit-buffer
UCS-A /org/mcast-policy #

```

Modifying Multicast Policy Parameters

You can modify an existing multicast policy to change the state of IGMP snooping or IGMP snooping querier. When a multicast policy is modified, all VLANs associated with that multicast policy are re-processed to apply the changes.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # scope mcast-policy <i>policy-name</i> | Enters organization multicast policy mode. |
| Step 3 | UCS-A /org/mcast-policy* # set querier {enabled disabled} | Enables or disables IGMP snooping querier. By default, IGMP snooping querier is disabled for a multicast policy. |
| Step 4 | UCS-A /org/mcast-policy* # set querierip <i>IGMP snooping querier IPv4 address</i> | Specifies the IPv4 address for the IGMP snooping querier. |
| Step 5 | UCS-A /org/mcast-policy* # set snooping {enabled disabled} | Enables or disables IGMP snooping. By default, IGMP snooping is enabled for a multicast policy. |
| Step 6 | UCS-A /org/mcast-policy* # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a multicast policy named policy1:

```

UCS-A# scope org /
UCS-A /org # scope mcast-policy policy1
UCS-A /org/mcast-policy* # set querier enabled
UCS-A /org/mcast-policy* # set querierip 1.2.3.4
UCS-A /org/mcast-policy* # set snooping enabled
UCS-A /org/mcast-policy* # commit-buffer
UCS-A /org/mcast-policy #

```

Assigning a VLAN Multicast Policy

You can set a multicast policy for a VLAN in the Ethernet uplink fabric mode. You cannot set a multicast policy for an isolated VLAN.

Before you begin

Create a VLAN.

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | Required: UCS-A /eth-uplink # scope fabric {a b} | Enters Ethernet uplink fabric mode for the specified fabric interconnect. |
| Step 3 | UCS-A /eth-uplink/fabric # scope vlan vlan-name | Enters Ethernet uplink fabric VLAN mode. |
| Step 4 | UCS-A /eth-uplink/fabric/vlan # set mcastpolicy policy-name | Assigns a multicast policy for the VLAN. |
| Step 5 | UCS-A /eth-uplink/fabric/vlan # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example sets a named VLAN accessible to one fabric interconnect and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope vlan vlan1
UCS-A /eth-uplink/fabric/vlan # set mcastpolicy policy1
UCS-A /eth-uplink/fabric/vlan* # commit-buffer
UCS-A /eth-uplink/fabric/vlan #
```

Deleting a Multicast Policy



Note

If you assigned a non-default (user-defined) multicast policy to a VLAN and then delete that multicast policy, the associated VLAN inherits the multicast policy settings from the default multicast policy until the deleted policy is re-created.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | UCS-A# scope org | Enters organization mode for the specified organization. |
| Step 2 | UCS-A /org # delete mcast-policy <i>policy-name</i> | Deletes a multicast policy with the specified policy name. |
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to delete a multicast policy named policy1:

```
UCS-A # scope org /
UCS-A /org # delete mcast-policy policy1
UCS-A /org* # commit-buffer
UCS-A /org #
```

LACP Policy

Link Aggregation combines multiple network connections in parallel to increase throughput and to provide redundancy. Link aggregation control protocol (LACP) provides additional benefits for these link aggregation groups. Cisco UCS Manager enables you to configure LACP properties using LACP policy.

You can configure the following for a lacp policy:

- **Suspended-individual:** If you do not configure the ports on an upstream switch for lacp, the fabric interconnects treat all ports as uplink Ethernet ports to forward packets. You can place the lacp port in suspended state to avoid loops. When you set suspend-individual on a port-channel with lacp, if a port that is part of the port-channel does not receive PDUs from the peer port, it will go into suspended state.
- **Timer values:** You can configure rate-fast or rate-normal. In rate-fast configuration, the port is expected to receive 1 PDU every 1 second from the peer port. The time out for this is 3 seconds. In rate-normal configuration, the port is expected to receive 1 PDU every 30 seconds. The timeout for this is 90 seconds.

System creates a default lacp policy at system start up. You can modify this policy or create new. You can also apply one lacp policy to multiple port-channels.

Creating a LACP Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|------------------------------------|
| Step 1 | UCS-A# scope org | Enters the root organization mode. |
| Step 2 | UCS-A /org # create lacppolicy <i>policy nam.</i> | Creates the specified lacp policy. |

| | Command or Action | Purpose |
|---------------|-----------------------------------|--|
| Step 3 | UCS-A /org # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example creates the lacp policy and commits the transaction:

```
UCS-A # scope org
UCS-A /org # create lacppolicy lacpl
UCS-A /org* # commit-buffer
UCS-A /org #
```

Editing a LACP Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope org | Enters the root organization mode. |
| Step 2 | UCS-A /org # scope lacppolicy <i>policy-name</i> . | Enters the specified lacp policy. |
| Step 3 | UCS-A /org/lacp policy/ <i>policy-name</i> # set suspend-individual <i>true</i> . | Sets suspend individual for the policy. |
| Step 4 | UCS-A /org/lacp policy/ <i>policy-name</i> # set lacp-rate <i>fast</i> . | Sets LACP rate for the policy. |
| Step 5 | UCS-A /org/lacp policy/ <i>policy-name</i> # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example modifies the lacp policy and commits transaction:

```
UCS-A# scope org
UCS-A/org # scope lacppolicy policy-name
UCS-A /org/lacp policy policy-name # set suspend-individual true
UCS-A/prg/policy policy-name # set lacp-rate fast
UCS-A /org* # commit-buffer
UCS-A /org #
```

Assigning LACP Policy to Port-Channels

Default lacp policy is assigned to port channels by default. You can assign a different lacp policy to the port channel. If the assigned policy does not exist, system generates a fault. You can create the same policy to clear the fault.



Note You can assign lacp policy to port-channels, FCoE port-channels, and ethernet storage port-channels. This procedure describes assigning the lacp policy to port-channels.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | UCS-A# scope eth-uplink | Enters Ethernet uplink mode. |
| Step 2 | UCS-A /eth-uplink # scope fabric | Enters the fabric mode. |
| Step 3 | UCS-A /eth-uplink/fabric # scope port-channel | Enters the port-channel mode. |
| Step 4 | UCS-A /eth-uplink/fabric/port-channel # set lacp-policy-name <i>policy-name</i> | Specifies the lacp policy for this port-channel. |
| Step 5 | UCS-A /eth-uplink/ fabric/port-channel commit-buffer | Commits the transaction to the system. |

Example

The following example shows assigning a lacp policy to a port-channel:

```
UCS-A# scope eth-uplink
UCS-A UCS-A/eth-uplink # scope fabric
UCS-A UCS-A/eth-uplink/fabric # scope port-channel
UCS-A UCS-A/eth-uplink/port-channel # set lacp-policy-name
UCS-A UCS-A/eth-uplink/port-channel* # commit-buffer
UCS-A UCS-A/eth-uplink/port-channel #
```

Configuring UDLD Link Policies

Understanding UDLD

UniDirectional Link Detection (UDLD) is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it marks the link as unidirectional. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

UDLD works with the Layer 1 mechanisms to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected interfaces. When you enable both autonegotiation and UDLD, the Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic sent by a local device is received by its neighbor but traffic from the neighbor is not received by the local device.

Modes of Operation

UDLD supports two modes of operation: normal (the default) and aggressive. In normal mode, UDLD can detect unidirectional links due to misconnected interfaces on fiber-optic connections. In aggressive mode, UDLD can also detect unidirectional links due to one-way traffic on fiber-optic and twisted-pair links and to misconnected interfaces on fiber-optic links.

In normal mode, UDLD detects a unidirectional link when fiber strands in a fiber-optic interface are misconnected and the Layer 1 mechanisms do not detect this misconnection. If the interfaces are connected correctly but the traffic is one way, UDLD does not detect the unidirectional link because the Layer 1 mechanism, which is supposed to detect this condition, does not do so. In case, the logical link is considered undetermined, and UDLD does not disable the interface. When UDLD is in normal mode, if one of the fiber strands in a pair is disconnected and autonegotiation is active, the link does not stay up because the Layer 1 mechanisms did not detect a physical problem with the link. In this case, UDLD does not take any action, and the logical link is considered undetermined.

UDLD aggressive mode is disabled by default. Configure UDLD aggressive mode only on point-to-point links between network devices that support UDLD aggressive mode. With UDLD aggressive mode enabled, when a port on a bidirectional link that has a UDLD neighbor relationship established stops receiving UDLD packets, UDLD tries to reestablish the connection with the neighbor and administratively shuts down the affected port. UDLD in aggressive mode can also detect a unidirectional link on a point-to-point link on which no failure between the two devices is allowed. It can also detect a unidirectional link when one of the following problems exists:

- On fiber-optic or twisted-pair links, one of the interfaces cannot send or receive traffic.
- On fiber-optic or twisted-pair links, one of the interfaces is down while the other is up.
- One of the fiber strands in the cable is disconnected.

Methods to Detect Unidirectional Links

UDLD operates by using two mechanisms:

- Neighbor database maintenance

UDLD learns about other UDLD-capable neighbors by periodically sending a hello packet (also called an advertisement or probe) on every active interface to keep each device informed about its neighbors. When the switch receives a hello message, it caches the information until the age time (hold time or time-to-live) expires. If the switch receives a new hello message before an older cache entry ages, the switch replaces the older entry with the new one.

UDLD clears all existing cache entries for the interfaces affected by the configuration change whenever an interface is disabled and UDLD is running, whenever UDLD is disabled on an interface, or whenever the switch is reset. UDLD sends at least one message to inform the neighbors to flush the part of their caches affected by the status change. The message is intended to keep the caches synchronized.

- Event-driven detection and echoing

UDLD relies on echoing as its detection mechanism. Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-sync neighbor, it restarts the detection window on its side of the connection and sends echo messages in reply. Because this behavior is the same on all UDLD neighbors, the sender of the echoes expects to receive an echo in reply.

If the detection window ends and no valid reply message is received, the link might shut down, depending on the UDLD mode. When UDLD is in normal mode, the link might be considered undetermined and

might not be shut down. When UDLD is in aggressive mode, the link is considered unidirectional, and the interface is shut down.

If UDLD in normal mode is in the advertisement or in the detection phase and all the neighbor cache entries are aged out, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbors.

If you enable aggressive mode when all the neighbors of a port have aged out either in the advertisement or in the detection phase, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbor. UDLD shuts down the port if, after the fast train of messages, the link state is still undetermined.

UDLD Configuration Guidelines

The following guidelines and recommendations apply when you configure UDLD:

- A UDLD-capable interface also cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.
- When configuring the mode (normal or aggressive), make sure that the same mode is configured on both sides of the link.
- UDLD should be enabled only on interfaces that are connected to UDLD capable devices. The following interface types are supported:
 - Ethernet uplink
 - FCoE uplink
 - Ethernet uplink port channel member
 - FCoE uplink port channel member

Configuring a UDLD Link Policy

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | UCS-A# scope org / | Enters the root organization mode. |
| Step 2 | UCS-A /org # create udld-link-policy <i>link-policy-name</i> | Creates a UDLD link policy with the specified name, and enters UDLD link policy mode. |
| Step 3 | UCS-A /org/udld-link-policy # commit-buffer | Commits the transaction to the system configuration. |
| Step 4 | UCS-A /org/udld-link-policy # exit | Returns to the previous mode. |
| Step 5 | UCS-A /org # scope udld-link-policy <i>link-policy-name</i> | Enters UDLD link policy mode for the specified UDLD link policy. |
| Step 6 | UCS-A /org/udld-link-policy # set mode { aggressive normal } | Specifies the mode for the UDLD link policy. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 7 | UCS-A /org/udld-link-policy # set admin-state {disabled enabled} | Disables or enables UDLD on the interface. |
| Step 8 | UCS-A /org/udld-link-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a link profile called UDLDPol1, sets the mode to aggressive, and enables UDLD on the interface.

```
UCS-A# scope org /
UCS-A /chassis/org # create udld-link-policy UDLDPol1
UCS-A /chassis/org/udld-link-policy* # commit-buffer
UCS-A /chassis/org/udld-link-policy # exit
UCS-A /chassis/org # scope udld-link-policy UDLDPol1
UCS-A /chassis/org/udld-link-policy # set mode aggressive
UCS-A /chassis/org/udld-link-policy* # set admin-state enabled
UCS-A /chassis/org/udld-link-policy* # commit-buffer
UCS-A /chassis/org/udld-link-policy #
```

Modifying the UDLD System Settings

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope org / | Enters the root organization mode. |
| Step 2 | UCS-A /org # show udld-policy | Displays the current UDLD system settings. |
| Step 3 | UCS-A /org # scope udld-policy default | Enters UDLD policy mode for the global UDLD policy. |
| Step 4 | UCS-A /org/udld-policy # set message-interval <i>seconds</i> | Specifies the time interval (in seconds) between UDLD probe messages on ports that are in advertisement mode. Enter an integer between 7 and 60. The default is 15 seconds. |
| Step 5 | UCS-A /org/udld-policy # set recovery-action [reset none] | Specifies the action to be taken on any ports that are disabled when UDLD aggressive mode is enabled. The default is none. |
| Step 6 | UCS-A /org/udld-policy # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to update the default UDLD system settings for a 30 second time interval.

```
UCS-A# scope org /
UCS-A /chassis/org # show udld-policy

UDLD system settings:
  Name           Message interval (sec) Recovery action
  -----
  default        15                               None

UCS-A /chassis/org # scope udld-policy default
UCS-A /chassis/org/udld-policy # set message-interval 30
UCS-A /chassis/org/udld-policy* # commit-buffer
UCS-A /chassis/org/udld-policy #
```

Configuring a Link Profile

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | UCS-A# scope org / | Enters the root organization mode. |
| Step 2 | UCS-A /org # create eth-link-profile <i>link-profile-name</i> | Creates a link profile with the specified name, and enters link profile mode. |
| Step 3 | UCS-A /org/eth-link-profile # commit-buffer | Commits the transaction to the system configuration. |
| Step 4 | UCS-A /org/eth-link-profile # exit | Returns to the previous mode. |
| Step 5 | UCS-A /org # scope eth-link-profile <i>link-profile-name</i> | Enters link profile mode for the specified link profile. |
| Step 6 | UCS-A /org/eth-link-profile # set udld-link-policy <i>link-policy-name</i> | Assigns the specified UDLD link policy to the link profile. |
| Step 7 | UCS-A /org/eth-link-profile # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to create a link profile called LinkProfile1 and assign the default UDLD link policy.

```
UCS-A# scope org /
UCS-A /chassis/org # create eth-link-profile LinkProfile1
UCS-A /chassis/org/eth-link-profile* # commit-buffer
UCS-A /chassis/org/eth-link-profile # exit
UCS-A /chassis/org # scope eth-link-profile LinkProfile1
```

```
UCS-A /chassis/org/eth-link-profile # set udd-link-policy default
UCS-A /chassis/org/eth-link-profile* # commit-buffer
```

Assigning a Link Profile to a Port Channel Ethernet Interface

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. |
| Step 3 | UCS-A /eth-uplink/fabric # scope port-channel port-chan-id | Enters Ethernet uplink fabric port channel mode for the specified port channel. |
| Step 4 | UCS-A /eth-uplink/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-uplink/fabric/port-channel/member-port # set eth-link-profile link-profile-name | Assigns the specified link profile. |
| Step 6 | UCS-A /eth-uplink/fabric/port-channel/member-port # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to assign link profile LinkProfile1 to a port channel Ethernet interface:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope port-channel 88
UCS-A /eth-uplink/fabric/port-channel # scope member-port 1 31
UCS-A /eth-uplink/fabric/port-channel/member-port # set eth-link-profile LinkProfile1
UCS-A /eth-uplink/fabric/port-channel/member-port* # commit-buffer
UCS-A /eth-uplink/fabric/port-channel/member-port #
```

Assigning a Link Profile to a Port Channel FCoE Interface

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. |

| | Command or Action | Purpose |
|---------------|---|---|
| Step 3 | UCS-A /fc-uplink/fabric # scope fcoe-port-channel <i>port-chan-id</i> | Enters Fibre Channel uplink fabric port channel mode for the specified port channel. |
| Step 4 | UCS-A /fc-uplink/fabric/fcoe-port-channel # scope fcoe-member-port <i>slot-id port-id</i> | Enters Fibre Channel server fabric, fabric port channel mode for the specified member port. |
| Step 5 | UCS-A /fc-uplink/fabric/fcoe-port-channel/fcoe-member-port # set eth-link-profile <i>link-profile-name</i> | Assigns the specified link profile. |
| Step 6 | UCS-A /fc-uplink/fabric/fcoe-port-channel/fcoe-member-port # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to assign link profile LinkProfile1 to a port channel FCoE interface:

```
UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope fcoe-port-channel 192
UCS-A /fc-uplink/fabric/fcoe-port-channel # scope fcoe-member-port 1 20
UCS-A /fc-uplink/fabric/fcoe-port-channel/fcoe-member-port # set eth-link-profile LinkProfile1
UCS-A /fc-uplink/fabric/fcoe-port-channel/fcoe-member-port* # commit-buffer
UCS-A /fc-uplink/fabric/fcoe-port-channel/fcoe-member-port #
```

Assigning a Link Profile to an Uplink Ethernet Interface

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. |
| Step 3 | UCS-A /eth-uplink/fabric # scope interface <i>slot-num port num</i> | Enters the interface command mode for the specified uplink port. |
| Step 4 | UCS-A /eth-uplink/fabric/interface # set eth-link-profile <i>link-profile-name</i> | Assigns the specified link profile. |
| Step 5 | UCS-A /eth-uplink/fabric/interface # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to assign link profile LinkProfile1 to an uplink Ethernet interface:

```

UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope fabric a
UCS-A /eth-uplink/fabric # scope interface 2 2
UCS-A /eth-uplink/fabric/interface # set eth-link-profile LinkProfile1
UCS-A /eth-uplink/fabric/interface* # commit-buffer
UCS-A /eth-uplink/fabric/interface #

```

Assigning a Link Profile to an Uplink FCoE Interface

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 fcoeinterface slot-num port num | Enters the Fibre Channel interface command mode for the specified uplink port. |
| Step 4 | UCS-A /fc-uplink/fabric/fcoeinterface # set eth-link-profile link-profile-name | Assigns the specified link profile. |
| Step 5 | UCS-A /fc-uplink/fabric/fcoeinterface # commit-buffer | Commits the transaction to the system configuration. |

Example

The following example shows how to assign link profile LinkProfile1 to an uplink FCoE interface:

```

UCS-A# scope fc-uplink
UCS-A /fc-uplink # scope fabric a
UCS-A /fc-uplink/fabric # scope fcoeinterface 2 2
UCS-A /fc-uplink/fabric/fcoeinterface # set eth-link-profile LinkProfile1
UCS-A /fc-uplink/fabric/fcoeinterface* # commit-buffer
UCS-A /fc-uplink/fabric/fcoeinterface #

```

VMQ Connection Policy

Cisco UCS Manager enables you to configure VMQ connection policy for a vNIC. VMQ provides improved network performance to the entire management operating system. Configuring a VMQ vNIC connection policy involves the following:

- Create a VMQ connection policy
- Create a static vNIC in a service profile
- Apply the VMQ connection policy to the vNIC

If you want to configure the VMQ vNIC on a service profile for a server, at least one adapter in the server must support VMQ. Make sure the servers have at least one the following adapters installed:

- UCS-VIC-M82-8P
- UCSB-MLOM-40G-01
- UCSC-PCIE-CSC-02

The following are the supported Operating Systems for VMQ:

- Windows 2012
- Windows 2012R2

You can apply only any one of the vNIC connection policies on a service profile at any one time. Make sure to select one of the three options such as Dynamic, usNIC or VMQ connection policy for the vNIC. When a VMQ vNIC is configured on service profile, make sure you have the following settings:

- Select SRIOV in the BIOS policy.
- Select Windows in the Adapter policy.

Creating a VMQ Connection 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 vmq-conn-policy <i>policy-name</i> | Specifies the name for this VMQ connection policy. |
| Step 3 | UCS-A /org/vmq-conn-policy* # set queue-count <i>queue count</i> | Specifies the queue count for the VMQ connection policy. |
| Step 4 | UCS-A /org/vmq-conn-policy* # set interrupt-count <i>interrupt count</i> | Specifies the interrupt count for the VMQ connection policy. |
| Step 5 | UCS-A /org/vmq-conn-policy* # commit-buffer | Commits the transaction to the system. |

Example

The following example creates a VMQ connection policy:

```
UCS-A# scope org
UCS-A /org # create vmq-conn-policy policy name
UCS-A /org/vmq-conn-policy* # set queue-count queue count (number)
UCS-A /org/vmq-conn-policy* # set interrupt-count queue count (number)
UCS-A /org/vmq-conn-policy* # commit-buffer
UCS-A /org/vmq-conn-policy #
```

