

Sun Datacenter InfiniBand Switch 648

Administration Guide



Part No. 820-7739-12
May 2010, Revision A

Copyright © 2009, 2010 Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2009, 2010, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.



Adobe PostScript

Contents

Using This Documentation xi

Administering the Switch 1

Understanding Administrative Commands 1

Understanding the Component Addressing Scheme 2

Power Supply Addressing 3

CMC Addressing 4

Fabric Card Addressing 5

Line Card Addressing 6

Line Card CXP Connector Addressing 7

CLIA Command Overview 8

Switch-Specific Command Overview 9

ILOM Command Overview 9

InfiniBand Software Command Overview 10

Monitoring the Switch 10

Performing General Tasks 11

▼ Display the General Health of the Switch 11

▼ Display Components Present 11

▼ Display the Active Sessions 12

Checking Power Supplies 13

▼ Display Power Supplies Present 13

▼ Check the Status of a Power Supply 14

- ▼ Display the Firmware Version of a Power Supply 14
- Checking CMCs 15
 - ▼ Check Internal Power and Temperature of a CMC 15
 - ▼ Check the Status of a CMC 16
 - ▼ Check the Status LEDs of a CMC 17
 - ▼ Display the Firmware Version of a CMC 17
 - ▼ Display the Firmware Version of the Switch Chassis Manager 18
- Checking Fabric Cards 18
 - ▼ Check Fan Speed and Status 19
 - ▼ Display Fabric Cards Present 20
 - ▼ Check Fabric Card Power Faults 21
 - ▼ Check the Internal Power and Temperature of a Fabric Card 22
 - ▼ Check Fabric Card Internal Temperatures 23
 - ▼ Check Fabric Card Internal Voltages 23
 - ▼ Display the Base GUIDs of a Fabric Card 25
 - ▼ Check Fabric Card Link Status (Simple) 26
 - ▼ Check Fabric Card Link Status (Detailed) 26
 - ▼ Check Fabric Card Switch Chip Health 27
 - ▼ Check the IPMB State of a Fabric Card 28
 - ▼ Check the Status LEDs of a Fabric Card 29
 - ▼ Display the Firmware Versions of a Fabric Card 30
- Checking Line Cards 30
 - ▼ Display Line Cards Present 31
 - ▼ Check Line Card Power Faults 32
 - ▼ Check the Internal Power and Temperature of a Line Card 32
 - ▼ Check Line Card Internal Temperatures 33
 - ▼ Check Line Card Internal Voltages 33
 - ▼ Display the Base GUIDs of a Line Card 35

- ▼ Check Line Card Link Status (Simple) 36
- ▼ Check Line Card Link Status (Detailed) 37
- ▼ Check Line Card Switch Chip Health 38
- ▼ Check the IPMB State of a Line Card 38
- ▼ Check the Status LEDs of a Line Card 39
- ▼ Display the Firmware Versions of a Line Card 40

Checking Other Switch Characteristics 40

- ▼ Display the Date 41
- ▼ Display the User Accounts 41
- ▼ Display the Network Management Configuration 41
- ▼ Locate a Switch Chip or Connector From the GUID 43

Monitoring the InfiniBand Fabric 43

- ▼ Display Information About the Local HCA 44
- ▼ Identify All HCAs in the Fabric 45
- ▼ Identify All Switches in the Fabric 46
- ▼ Display the InfiniBand Fabric Topology 47
- ▼ Display a Route Through the Fabric 47
- ▼ Display the Link Status of a Node 49
- ▼ Display Counters for a Node 50
- ▼ Display Data Counters for a Node 51
- ▼ Display Low-Level Detailed Information About a Node 52
- ▼ Display Low-Level Detailed Information About a Port 53
- ▼ Map LIDs to GUIDs 54
- ▼ Determine the GUID and LID for a Node Within the Switch 55
- ▼ Display Subnet Manager Status 56
- ▼ Display OFED Software Version Information 57

Controlling the Switch 57

- Setting Up the Hardware 58

- ▼ Change the root Password 58
- ▼ Get Help on a CLIA Command 59
- ▼ Set the Date 59
- ▼ Change the CMC Host Name 60
- ▼ Reconfigure the Network Management Parameters 61
- ▼ Reconfigure the CMCs for Identical Addresses 63

Setting Up the Users 64

- ▼ Add a User 65
- ▼ Change a User's Password 65
- ▼ Delete a User 66

Managing Power Supplies 66

- ▼ Restart a Power Supply 67
- ▼ Enable and Activate a Power Supply 67
- ▼ Deactivate and Disable a Power Supply 68

Managing CMCs 69

- ▼ Restart a CMC 69
- ▼ Activate a CMC 70
- ▼ Deactivate a CMC 70
- ▼ Switch Over to Another CMC 71

Managing Fabric Cards 72

- ▼ Restart a Fabric Card or Filler 72
- ▼ Enable a Fabric Card or Filler 73
- ▼ Disable a Fabric Card or Filler 74
- ▼ Enable Standby Power for a Fabric Card or Filler 74
- ▼ Disable Standby Power for a Fabric Card or Filler 75
- ▼ Enable a Fabric Card Slot for Hot-Insertion 75
- ▼ Activate a Fabric Card or Filler 75
- ▼ Deactivate a Fabric Card or Filler 76

- ▼ Turn On a Fabric Card or Filler Locator LED 77
- ▼ Turn Off a Fabric Card or Filler Locator LED 78
- ▼ Enable Downed Fabric Card Links 78
- ▼ Enable a Fabric Card Switch Chip Port 79
- ▼ Disable a Fabric Card Switch Chip Port 80
- ▼ Reset a Fabric Card Switch Chip 80

Managing Line Cards 81

- ▼ Restart a Line Card 82
- ▼ Enable a Line Card 83
- ▼ Disable a Line Card 83
- ▼ Enable Standby Power for a Line Card 84
- ▼ Disable Standby Power for a Line Card 84
- ▼ Enable a Line Card Slot for Hot-Insertion 85
- ▼ Activate a Line Card 85
- ▼ Deactivate a Line Card 86
- ▼ Turn On a Line Card Locator LED 87
- ▼ Turn Off a Line Card Locator LED 88
- ▼ Enable Downed Line Card Links 88
- ▼ Enable a Line Card Switch Chip Port 89
- ▼ Disable a Line Card Switch Chip Port 89
- ▼ Reset a Line Card Switch Chip 90

Controlling the InfiniBand Fabric 91

- ▼ Perform Comprehensive Diagnostics for the Entire Fabric 91
- ▼ Perform Comprehensive Diagnostics for a Route 92
- ▼ Determine Changes to the InfiniBand Fabric Topology 92
- ▼ Find 1x or SDR or DDR Links in the Fabric 94
- ▼ Determine Which Links Are Experiencing Significant Errors 95
- ▼ Clear Error Counters 96

- ▼ Clear Data Counters 96
- ▼ Check All Ports 96
- ▼ Reset a Port 97
- ▼ Set the Speed of a Port 98
- ▼ Disable a Port 100
- ▼ Enable a Port 101
- Controlling the Subnet Manager 102
 - ▼ Start the Subnet Manager With Min Hop Routing 102
 - ▼ Start the Subnet Manager With Fat Tree Routing 103
 - ▼ Create the `guid.txt` File 104
 - ▼ Start the Subnet Manager With the `opensmd` Daemon 106
 - ▼ Stop the Subnet Manager With the `opensmd` Daemon 106
 - ▼ Create the `opensm.conf` File 106

Diagnosing a Problem 109

Identifying the Problem 109

Checking Logs and Error Messages 110

Switch Event Message Overview 110

- ▼ Display the Switch Message Log for Power Supplies 110
- ▼ Display the Switch Message Log for Fabric Cards 111
- ▼ Display the Switch Message Log for Line Cards 111
- ▼ Display the System Event Log 112

Checking Status LEDs 112

- ▼ Check Power Supply LEDs 112
- ▼ Check CMC Status LEDs 113
- ▼ Check the NET MGT LEDs 114
- ▼ Check Fabric Card LEDs 115
- ▼ Check Fan LEDs 115
- ▼ Check Line Card LEDs 116

▼ Check Link LEDs	117
Troubleshooting the Switch	117
Switch Hardware Problems	118
InfiniBand Fabric Problems	121
Understanding Signal Routing Through the Switch	124
CXP Connectors and LEDs to Line Card Switch Chip and Port Routes	124
Line Card Switch Chip to Line Card XBOW Connector Routes	126
Fabric Card XBOW Connector to Fabric Card Switch Chip Routes	127
Fabric Card Switch Chip to Fabric Card XBOW Connector Routes	128
Line Card XBOW Connector to Line Card Switch Chip Routes	129
Line Card Switch Chip and Port to CXP Connector and LED Routes	130
Signal Route Through the Switch	131
Understanding Switch Startup and Component Addition	132
Switch Startup Process	133
Fabric Card Addition Process	134
Line Card Addition Process	135
Switch GUIDs Overview	136
Index	139

Using This Documentation

This administration guide provides detailed procedures that describe monitoring and controlling the Sun Datacenter InfiniBand Switch 648 from Oracle®. This document is written for technicians, system administrators, authorized service providers, and users who have advanced experience administrating InfiniBand fabric hardware.

- “Related Documentation” on page xi
- “Documentation, Support, and Training” on page xii
- “Documentation Feedback” on page xii

Related Documentation

The documents listed as online are available at:

(<http://docs.sun.com/app/docs/prod/ib.switch.648>)

Application	Title	Part Number	Format	Location
Getting started	<i>Sun Datacenter InfiniBand Switch 648 Getting Started Guide</i>	820-7745	Printed PDF	Shipping kit Online
Last-minute information	<i>Sun Datacenter InfiniBand Switch 648 Product Notes</i>	820-7743	PDF	Online
Preparation and installation	<i>Sun Datacenter InfiniBand Switch 648 Installation Guide</i>	820-7738	PDF HTML	Online
Administration	<i>Sun Datacenter InfiniBand Switch 648 Administration Guide</i>	820-7739	PDF HTML	Online
Service	<i>Sun Datacenter InfiniBand Switch 648 Service Manual</i>	820-7740	PDF HTML	Online

Application	Title	Part Number	Format	Location
Command reference	<i>Sun Datacenter InfiniBand Switch 648 Command Reference</i>	820-7741	PDF HTML	Online
Compliance	<i>Sun Datacenter InfiniBand Switch 648 Safety and Compliance Guide</i>	820-7744	PDF	Online
ILOM information	<i>Oracle Integrated Lights Out Manager (ILOM) 3.0 Supplement for the Sun Datacenter InfiniBand Switch 648</i>	821-0896	PDF HTML	Online

Documentation, Support, and Training

These web sites provide additional resources:

- Documentation (<http://docs.sun.com>)
- Support (<http://www.sun.com/support>)
- Training (<http://www.sun.com/training>)

Documentation Feedback

Submit comments about this document by clicking the Feedback[+] link at (<http://docs.sun.com>). Include the title and part number of your document with your feedback:

Sun Datacenter InfiniBand Switch 648 Administration Guide, part number 820-7739-12.

Administering the Switch

The following topics describe the administration of the Sun Datacenter InfiniBand Switch 648.

- [“Understanding Administrative Commands” on page 1](#)
- [“Monitoring the Switch” on page 10](#)
- [“Monitoring the InfiniBand Fabric” on page 43](#)
- [“Controlling the Switch” on page 57](#)
- [“Controlling the InfiniBand Fabric” on page 91](#)

Related Information

- [“Diagnosing a Problem” on page 109](#)
- [Switch Installation](#)
- [Switch Remote Management](#)
- [Switch Service](#)
- [Switch Reference](#)

Understanding Administrative Commands

Administration of the switch (powering, configuration, reporting, and operation) is accomplished through the execution of various commands from a management interface. There are no manual controls on the switch. This design enables complete remote administration of the switch with the exception of service procedures. The administrative commands are described in the following topics:

- [“Understanding the Component Addressing Scheme” on page 2](#)
- [“CLIA Command Overview” on page 8](#)
- [“Switch-Specific Command Overview” on page 9](#)

- [“ILOM Command Overview” on page 9](#)
- [“InfiniBand Software Command Overview” on page 10](#)

Related Information

- [Switch Installation](#), understanding the installation
- [Switch Service](#), understanding service procedures
- [“Monitoring the Switch” on page 10](#)
- [“Monitoring the InfiniBand Fabric” on page 43](#)
- [“Controlling the Switch” on page 57](#)
- [“Controlling the InfiniBand Fabric” on page 91](#)

Understanding the Component Addressing Scheme

When you monitor or control the switch with administrative commands, you are often required to provide either a slot number or an Integrated Peripherals Management Bus (IPMB) address to the command. These values identify which component you want the command to affect.

Slot numbers for the components begin with 0 and are decimal. IPMB addresses are hexadecimal. Typically, the switch-specific commands use slot numbers, while the CLIA commands use IPMB addresses.

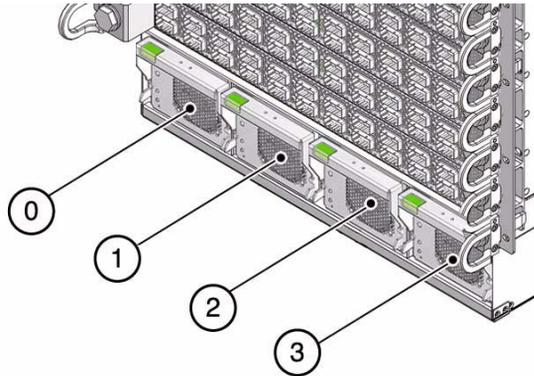
- [“Power Supply Addressing” on page 3](#)
- [“CMC Addressing” on page 4](#)
- [“Fabric Card Addressing” on page 5](#)
- [“Line Card Addressing” on page 6](#)
- [“Line Card CXP Connector Addressing” on page 7](#)

Related Information

- [“CLIA Command Overview” on page 8](#)
- [“Switch-Specific Command Overview” on page 9](#)

Power Supply Addressing

The power supplies mount at the lower front of the switch chassis and occupy slots numbered 0 through 3, going from left to right. That is, slot 0 is on the left side of the chassis and slot 3 is on the right side of the chassis, when looking at the slots.



Note – Power cords and power supplies are mated in a pass-through, one-to-one relationship. The power cord for power supply 0 is on the right side of the switch chassis when looking at the receptacles. The power cord for power supply 3 is on the left side of the switch chassis when looking at the receptacles.

The following table lists the slot numbers, associated IPMB addresses (in hexadecimal), and FRU numbers for the power supplies.

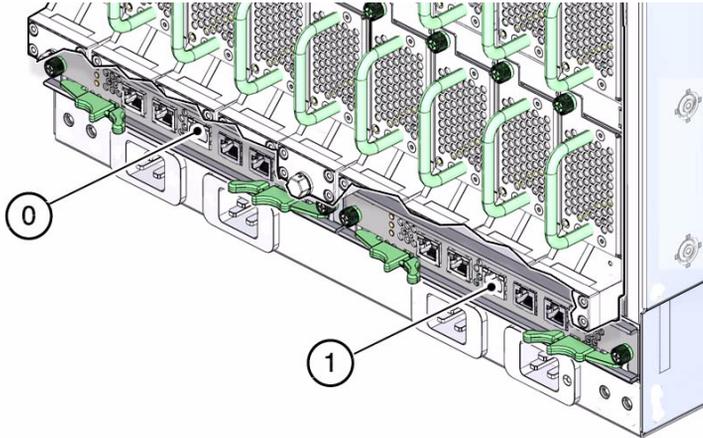
Slot Number	IPMB Address (in hexadecimal)	FRU Number
0	20	3
1	20	4
2	20	5
3	20	6

Related Information

- *Switch Reference*, `shelf` command
- [“CMC Addressing” on page 4](#)
- [“Fabric Card Addressing” on page 5](#)
- [“Line Card Addressing” on page 6](#)
- [“Line Card CXP Connector Addressing” on page 7](#)

CMC Addressing

The Chassis Management Controllers (CMCs) mount at the lower rear of the switch chassis and occupy slots numbered 0 and 1, going from left to right, when looking at the slots.



The following table lists the slot numbers and associated IPMB addresses (in hexadecimal) for the CMCs.

Slot Number	IPMB Address (in hexadecimal)
0	10
1	12

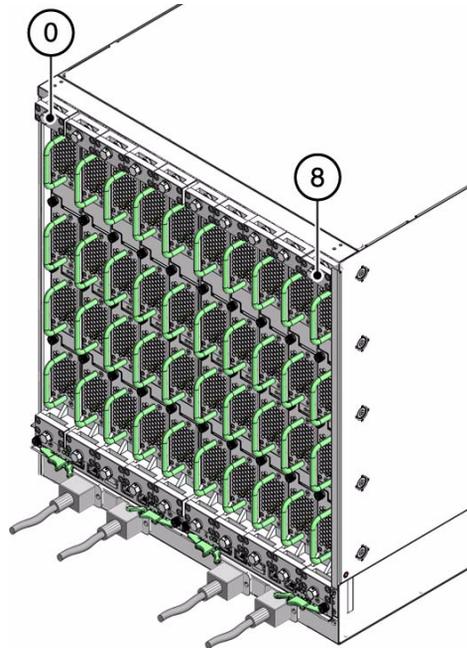
Related Information

- *Switch Reference*, `shelf` command
- [“Power Supply Addressing” on page 3](#)
- [“Fabric Card Addressing” on page 5](#)
- [“Line Card Addressing” on page 6](#)
- [“Line Card CXP Connector Addressing” on page 7](#)

Fabric Card Addressing

The fabric cards mount at the rear of the switch chassis and occupy slots numbered 0 through 8, going from left to right. That is, slot 0 is on the left side of the chassis and slot 8 is on the right side of the chassis when looking at the slots.

Within each fabric card are 4 fans, occupying slots numbered 0 through 3, going from bottom to top. The fans are not IPMB addressable.



The following table lists the slot numbers and associated IPMB addresses (in hexadecimal) for the fabric cards.

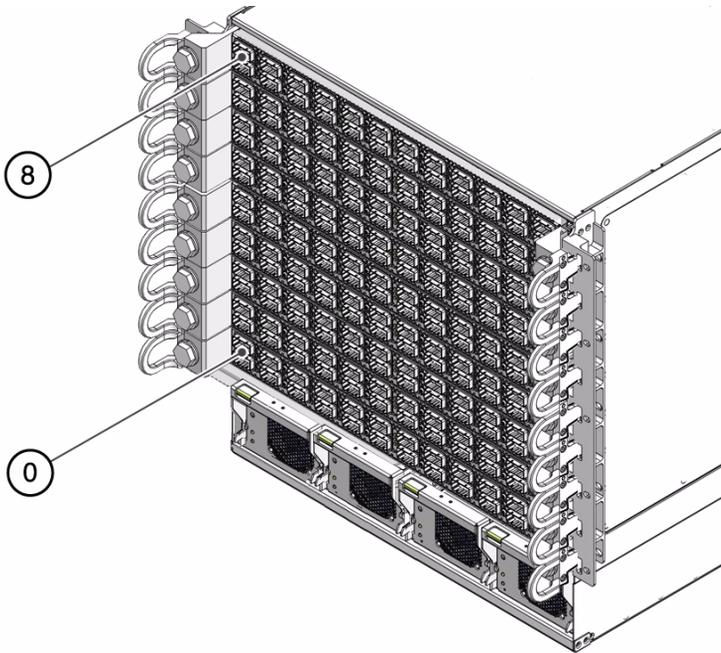
Slot Number	IPMB Address (in hexadecimal)
0	94
1	96
2	98
3	9A
4	9C
5	9E
6	A0
7	A2
8	A4

Related Information

- *Switch Reference*, `shelf` command
- “Power Supply Addressing” on page 3
- “CMC Addressing” on page 4
- “Line Card Addressing” on page 6
- “Line Card CXP Connector Addressing” on page 7

Line Card Addressing

The line cards mount at the front of the switch chassis and occupy slots numbered 0 through 8, going from bottom to top. That is, slot 0 is on the bottom, just above the power supply slots, and slot 8 is at the top of the chassis.



The following table lists the slot numbers and associated IPMB addresses (in hexadecimal) for the line cards.

Slot Number	IPMB Address (in hexadecimal)
0	82
1	84

Slot Number	IPMB Address (in hexadecimal)
2	86
3	88
4	8A
5	8C
6	8E
7	90
8	92

Related Information

- *Switch Reference*, `shelf` command
- [“Power Supply Addressing” on page 3](#)
- [“CMC Addressing” on page 4](#)
- [“Fabric Card Addressing” on page 5](#)
- [“Line Card CXP Connector Addressing” on page 7](#)

Line Card CXP Connector Addressing

Each line card panel has 24 InfiniBand CXP connectors, labeled from left to right, 0A, 1A, to 11A in the upper row, and 0B, 1B, to 11B in the lower row. That is, connector 0A is above 0B on the left side of the panel, and 11A is above 11B on the right side of the panel. This numbering scheme is the same for every line card.

The following table lists the connector numbers and associated IPMB address FRU numbers for the line cards.

Connector Number	IPMB Address FRU Number	Connector Number	IPMB Address FRU Number	Connector Number	IPMB Address FRU Number
0A	2	4A	10	8A	18
0B	1	4B	9	8B	17
1A	4	5A	12	9A	20
1B	3	5B	11	9B	19
2A	6	6A	14	10A	22
2B	5	6B	13	10B	21
3A	8	7A	16	11A	24

Connector Number	IPMB Address FRU Number	Connector Number	IPMB Address FRU Number	Connector Number	IPMB Address FRU Number
3B	7	7B	15	11B	23

Related Information

- *Switch Reference*, `shelf` command
- “Power Supply Addressing” on page 3
- “CMC Addressing” on page 4
- “Fabric Card Addressing” on page 5
- “Line Card Addressing” on page 6

CLIA Command Overview

The CMCs use a command-shell interface called CLIA for issuing commands and performing tasks. With these commands, you can administer the switch hardware. CLIA commands require you to prefix the command with the string `clia`. The switch documentation set describes 30 CLIA commands. More commands are available, but their use is not practical for switch applications and executing these commands results in errors or nonoperation.

Typically, CLIA commands use an IPMB address to identify the components to act upon. In some instances, the string `board slot`, where *slot* is the slot number, can be substituted for the IPMB address.

Related Information

- *Switch Reference*, `help` command
- *Switch Reference*, understanding CLIA commands
- “Switch-Specific Command Overview” on page 9
- “ILOM Command Overview” on page 9
- “InfiniBand Software Command Overview” on page 10

Switch-Specific Command Overview

The CMCs use a simplified Linux OS and file system. From the # prompt, you can type switch-specific commands to perform some administrative and management tasks. Some switch-specific commands are wrappers to the CLIA API. Switch-specific commands are more user friendly and can perform testing upon the switch chips themselves, enabling greater control of the switch and its operation.

Switch-specific commands use the slot number to identify a particular component to act upon.

Related Information

- *Switch Reference*, understanding switch-specific commands
- “CLIA Command Overview” on page 8
- “ILOM Command Overview” on page 9
- “InfiniBand Software Command Overview” on page 10

ILOM Command Overview

When you upgrade the CMC firmware to version 1.0.2 or higher, ILOM 3.0 support is also installed. The ILOM command line and web interfaces enable additional administration features and capabilities. The ILOM command line interface uses 13 commands that affect many ILOM targets. Alternatively, the ILOM web interface permits point-and-click administration of the ILOM components and services.

Information about ILOM support of the switch is available in the *Sun Datacenter InfiniBand Switch 648 Integrated Lights Out Manager (ILOM) 3.0 Supplement*, available online at:

(<http://docs.sun.com/app/docs/prod/ib.switch.648>)

Overall information about ILOM 3.0 is available online at:

(<http://docs.sun.com/app/docs/prod/int.lights.mgr30>)

Related Information

- *Switch Remote Access*, understanding the ILOM commands
- “CLIA Command Overview” on page 8
- “Switch-Specific Command Overview” on page 9
- “InfiniBand Software Command Overview” on page 10

InfiniBand Software Command Overview

The OFED software package contains the OpenIB software suite. The suite is comprised of many InfiniBand software commands that are a means of monitoring and controlling aspects of the InfiniBand fabric. These commands are installed on and run from the Linux InfiniBand host, which is also the host of the Subnet Manager. Use of these commands requires thorough knowledge of InfiniBand architecture and technology.

Related Information

- *Switch Reference*, understanding InfiniBand software commands
- [“CLIA Command Overview” on page 8](#)
- [“Switch-Specific Command Overview” on page 9](#)
- [“ILOM Command Overview” on page 9](#)

Monitoring the Switch

The following topics enable you to display and check the operation and status of the switch components.

Note – Commands described in these topics are issued on the CMC as the `root` user.

- [“Performing General Tasks” on page 11](#)
- [“Checking Power Supplies” on page 13](#)
- [“Checking CMCs” on page 15](#)
- [“Checking Fabric Cards” on page 18](#)
- [“Checking Line Cards” on page 30](#)
- [“Checking Other Switch Characteristics” on page 40](#)

Related Information

- [“Understanding Administrative Commands” on page 1](#)
- [“Monitoring the InfiniBand Fabric” on page 43](#)
- [“Controlling the Switch” on page 57](#)
- [“Controlling the InfiniBand Fabric” on page 91](#)

Performing General Tasks

These tasks provide you are quick overview of the switch's condition.

- "Display the General Health of the Switch" on page 11
- "Display Components Present" on page 11
- "Display the Active Sessions" on page 12

Related Information

- "Checking Power Supplies" on page 13
- "Checking CMCs" on page 15
- "Checking Fabric Cards" on page 18
- "Checking Line Cards" on page 30
- "Checking Other Switch Characteristics" on page 40

▼ Display the General Health of the Switch

- On the CMC, type:

```
# clia showunhealthy  
Pigeon Point Shelf Manager Command Line Interpreter  
There are no unhealthy components in the shelf.  
#
```

Related Information

- *Switch Reference*, showunhealthy command
- "Display Counters for a Node" on page 50

▼ Display Components Present

- On the CMC, type:

```
# showpresent  
PSU 00 present, state = ON  
PSU 01 present, state = ON  
PSU 03 present, state = ON  
FC 00 present and active (state = M4)  
FC 01 present but not active  
FCF 02 present and active (state = M4)  
.
```

```
.  
.
LC 07 present and active (state = M4)
LC 08 present and active (state = M4)
#
```

Note – The output in the example is just a portion of the full output.

In the example, power supply 2 is missing. Additionally, fabric card 1 is not active. Also, the string (state = M4), means the FRU is active.

Note – Fans and CMCs are not polled by the showpresent command.

Related Information

- *Switch Reference*, showpresent command
- [“Display Power Supplies Present” on page 13](#)
- [“Display Fabric Cards Present” on page 20](#)
- [“Display Line Cards Present” on page 31](#)

▼ Display the Active Sessions

- **On the CMC, type:**

```
# cli session
Pigeon Point Shelf Manager Command Line Interpreter
32 sessions possible, 2 sessions currently active
Session: 1
  User: ID 1, Name: ""; Privilege Level: "Administrator"
  Channel: 1 ("LAN_802_3"); Peer IP address: 172.16.2.203, Port:
1764
Session: 2
  User: ID 1, Name: ""; Privilege Level: "Administrator"
  Channel: 1 ("LAN_802_3"); Peer IP address: 172.16.2.203, Port:
1765
#
```

Related Information

- *Switch Reference*, session command

Checking Power Supplies

These tasks enable you to check the condition of the power supplies.

- [“Display Power Supplies Present”](#) on page 13
- [“Check the Status of a Power Supply”](#) on page 14
- [“Display the Firmware Version of a Power Supply”](#) on page 14

Related Information

- [“Managing Power Supplies”](#) on page 66
- [“Performing General Tasks”](#) on page 11
- [“Checking CMCs”](#) on page 15
- [“Checking Fabric Cards”](#) on page 18
- [“Checking Line Cards”](#) on page 30
- [“Checking Other Switch Characteristics”](#) on page 40

▼ Display Power Supplies Present

- On the CMC, type:

```
# showpresent | grep PSU
PSU 00 present, state = ON
PSU 01 present, state = ON
PSU 02 present, state = ON
PSU 03 present, state = ON
#
```

Related Information

- *Switch Reference*, `showpresent` command
- [“Display Fabric Cards Present”](#) on page 20
- [“Display Line Cards Present”](#) on page 31

▼ Check the Status of a Power Supply

- On the CMC, type:

```
# psustatus slot
```

where *slot* is 0–3. See “Power Supply Addressing” on page 3. For example:

```
# psustatus 0
PSU 0, 12 V on
#
```

Related Information

- *Switch Reference*, psustatus command
- “Check Power Supply LEDs” on page 112
- “Check the Status of a CMC” on page 16
- “Check Fan Speed and Status” on page 19
- “Check Fabric Card Power Faults” on page 21
- “Check Line Card Power Faults” on page 32

▼ Display the Firmware Version of a Power Supply

- On the CMC, type:

```
# getpsufwver slot
```

where *slot* is 0–3. See “Power Supply Addressing” on page 3. For example:

```
# getpsufwver 0
PSU 0 FW version 2.5
#
```

Related Information

- *Switch Reference*, mcmversion command
- “Display the Firmware Version of a CMC” on page 17
- “Display the Firmware Version of the Switch Chassis Manager” on page 18
- “Display the Firmware Versions of a Fabric Card” on page 30
- “Display the Firmware Versions of a Line Card” on page 40
- “Display OFED Software Version Information” on page 57

Checking CMCs

These tasks enable you to check the operation and status of the CMCs.

- “Check Internal Power and Temperature of a CMC” on page 15
- “Check the Status of a CMC” on page 16
- “Check the Status LEDs of a CMC” on page 17
- “Display the Firmware Version of a CMC” on page 17
- “Display the Firmware Version of the Switch Chassis Manager” on page 18

Related Information

- “Managing CMCs” on page 69
- “Performing General Tasks” on page 11
- “Checking Power Supplies” on page 13
- “Checking Fabric Cards” on page 18
- “Checking Line Cards” on page 30
- “Checking Other Switch Characteristics” on page 40

▼ Check Internal Power and Temperature of a CMC

- On the CMC, type:

```
# clia sensordata -v IPMB_address | grep -e LUN -e Processed
```

where *IPMB_address* is from “CMC Addressing” on page 4. For example:

```
# clia sensordata -v 10 | grep -e LUN -e Processed
10: LUN: 0, Sensor # 0 ("FRU 0 HOT_SWAP")
10: LUN: 0, Sensor # 1 ("IPMB LINK")
10: LUN: 0, Sensor # 2 ("Local Temp")
    Processed data: 27.000000 degrees C
10: LUN: 0, Sensor # 3 ("3.3STBY voltage")
    Processed data: 3.344000 Volts
10: LUN: 0, Sensor # 4 ("3.3MAIN voltage")
    Processed data: 3.361600 Volts
10: LUN: 0, Sensor # 5 ("VBAT")
    Processed data: 3.061500 Volts
10: LUN: 0, Sensor # 6 ("Analog 0")
    Processed data: 1.009400 Volts
```

```
10: LUN: 0, Sensor # 128 ("CPLD State")
10: LUN: 0, Sensor # 129 ("Reboot Reason")
#
```

Related Information

- *Switch Reference*, sensordata command
- “Check the Internal Power and Temperature of a Fabric Card” on page 22
- “Check the Internal Power and Temperature of a Line Card” on page 32

▼ Check the Status of a CMC

- On the CMC, type:

```
# clia shmstatus
Pigeon Point Shelf Manager Command Line Interpreter
Host: "Active"
#
```

Note – The CMC must have a Host: “Active” status to affect CMC and component administration and configuration changes.

Related Information

- *Switch Reference*, shmstatus command
- “Check the Status of a Power Supply” on page 14
- “Check Fan Speed and Status” on page 19
- “Check Fabric Card Power Faults” on page 21
- “Check Line Card Power Faults” on page 32

▼ Check the Status LEDs of a CMC

- On the CMC, type:

```
# clia getfruledstate -v IPMB_address |grep -e FRU -e State: -e supported
```

where *IPMB_address* is from “CMC Addressing” on page 4. For example:

```
# clia getfruledstate -v 10 |grep -e FRU -e State: -e supported
10: FRU # 0, Led # 0 ("BLUE LED"):
    Local Control LED State: LED OFF
    Colors supported(0x02): BLUE
10: FRU # 0, Led # 1 ("LED 1"):
    Local Control LED State: LED ON, color: GREEN
    Colors supported(0x08): GREEN
10: FRU # 0, Led # 2 ("LED 2"):
    Local Control LED State: LED OFF
    Colors supported(0x10): AMBER
#
```

Note – Only if an LED is stated LED ON or LED BLINKING, is the LED lit. Otherwise, the LED is off.

Related Information

- *Switch Reference*, getfruledstate command
- “Check CMC Status LEDs” on page 113
- “Check the Status LEDs of a Fabric Card” on page 29
- “Check the Status LEDs of a Line Card” on page 39

▼ Display the Firmware Version of a CMC

- On the CMC, type:

```
# clia version
Pigeon Point Shelf Manager Command Line Interpreter
Pigeon Point Shelf Manager ver. 2.6.3
Pigeon Point and the stylized lighthouse logo are trademarks of Pigeon Point
Systems.
Copyright (c) 2002-2008 Pigeon Point Systems
All rights reserved
Build date/time: Oct  6 2009 09:26:43
```

```
Carrier: SUN_M9
Carrier subtype: 0; subversion: 1
#
```

Related Information

- *Switch Reference*, version command
- “Display the Firmware Version of the Switch Chassis Manager” on page 18
- “Display the Firmware Version of a Power Supply” on page 14
- “Display the Firmware Versions of a Fabric Card” on page 30
- “Display the Firmware Versions of a Line Card” on page 40
- “Display OFED Software Version Information” on page 57

▼ Display the Firmware Version of the Switch Chassis Manager

- On the CMC, type:

```
# mcmversion
M9CM version 1.1.4
Build time: Oct 6 2009 09:18:56
#
```

Related Information

- *Switch Reference*, mcmversion command
- “Display the Firmware Version of a CMC” on page 17
- “Display the Firmware Version of a Power Supply” on page 14
- “Display the Firmware Versions of a Fabric Card” on page 30
- “Display the Firmware Versions of a Line Card” on page 40
- “Display OFED Software Version Information” on page 57

Checking Fabric Cards

The following tasks enable you to check and verify the fabric card operation and status.

Note – For the examples in this topic, fabric card 0 (IPMB address 94), is used. Additionally, fabric card fillers are installed in slots 2 and 6.

- “Check Fan Speed and Status” on page 19
- “Display Fabric Cards Present” on page 20
- “Check Fabric Card Power Faults” on page 21
- “Check the Internal Power and Temperature of a Fabric Card” on page 22
- “Check Fabric Card Internal Temperatures” on page 23
- “Check Fabric Card Internal Voltages” on page 23
- “Display the Base GUIDs of a Fabric Card” on page 25
- “Check Fabric Card Link Status (Simple)” on page 26
- “Check Fabric Card Link Status (Detailed)” on page 26
- “Check Fabric Card Switch Chip Health” on page 27
- “Check the IPMB State of a Fabric Card” on page 28
- “Check the Status LEDs of a Fabric Card” on page 29
- “Display the Firmware Versions of a Fabric Card” on page 30

Related Information

- “Managing Fabric Cards” on page 72
- “Performing General Tasks” on page 11
- “Checking Power Supplies” on page 13
- “Checking CMCs” on page 15
- “Checking Line Cards” on page 30
- “Checking Other Switch Characteristics” on page 40

▼ Check Fan Speed and Status

The fans in the fabric card are numbered 0 through 3, with fan 0 on the bottom and fan 3 on the top.

- On the CMC, type:

```
# checkfans | grep 'FC slot'
```

where *slot* is 0–8. See “Fabric Card Addressing” on page 5. For example:

```
# checkfans | grep 'FC 0'
FC 0 Fan 0 RPM = 20924.000000
FC 0 Fan 1 RPM = 21210.000000
FC 0 Fan 2 RPM = 20924.000000
Warning : FC 0 Fan 3 stopped
#
```

In this example, fan 3, the top fan in fabric card 0, has stopped.

Note – When checking fabric card filler fans, use `grep 'FCF slot'`.

Related Information

- *Switch Reference*, `checkfans` command
- “Check Fan LEDs” on page 115
- “Check the Status of a Power Supply” on page 14
- “Check the Status of a CMC” on page 16
- “Check Fabric Card Power Faults” on page 21
- “Check Line Card Power Faults” on page 32

▼ Display Fabric Cards Present

- On the CMC, type:

```
# showpresent | grep FC
FC 00 present and active (state = M4)
FC 01 present and active (state = M4)
FCF 02 present and active (state = M4)
FC 03 present and active (state = M4)
FC 04 present and active (state = M4)
FC 05 present and active (state = M4)
FCF 06 present and active (state = M4)
FC 07 present and active (state = M4)
FC 08 present and active (state = M4)
#
```

Note – In this example, fabric card fillers (FCF) are installed in slots 2 and 6.

In the output, (state = M4) means the FRU is active.

Related Information

- *Switch Reference*, showpresent command
- [“Display Power Supplies Present” on page 13](#)
- [“Display Line Cards Present” on page 31](#)

▼ Check Fabric Card Power Faults

- On the CMC, type:

```
# checkpwrfault | grep FC
FC 0 Power fault sensor = 0x00 OK
FC 1 Power fault sensor = 0x00 OK
FC 3 Power fault sensor = 0x00 OK
FC 4 Power fault sensor = 0x00 OK
FC 5 Power fault sensor = 0x00 OK
FC 7 Power fault sensor = 0x00 OK
FC 8 Power fault sensor = 0x00 OK
#
```

Note – Fabric card fillers are unable to report power faults, and hence are not displayed in the example

Related Information

- *Switch Reference*, checkpwrfault command
- [“Check the Status of a Power Supply” on page 14](#)
- [“Check the Status of a CMC” on page 16](#)
- [“Check Fan Speed and Status” on page 19](#)
- [“Check Line Card Power Faults” on page 32](#)

▼ Check the Internal Power and Temperature of a Fabric Card

- On the CMC, type:

```
# clia sensordata -v IPMB_address | grep -e LUN -e Processed
```

where *IPMB_address* is from “Fabric Card Addressing” on page 5. For example, for fabric card 0:

```
# clia sensordata -v 94 | grep -e LUN -e Processed
94: LUN: 0, Sensor # 0 ("Hot Swap")
94: LUN: 0, Sensor # 1 ("IPMB Physical")
94: LUN: 0, Sensor # 2 ("LM75_0")
    Processed data: 27.000000 degrees C
94: LUN: 0, Sensor # 3 ("LM75_1")
    Processed data: 26.000000 degrees C
94: LUN: 0, Sensor # 4 ("Ambient Temp")
    Processed data: 1.805400 Volts
94: LUN: 0, Sensor # 5 ("+1.8V STBY")
    Processed data: 1.805400 Volts
.
.
.
94: LUN: 0, Sensor # 57 ("I4B_temperature")
    Processed data: 49.000000 degrees C
#
```

Note – The output in the example is just a portion of the full output.

Note – Fabric card fillers report temperatures, external voltages, and fan speeds.

Related Information

- *Switch Reference*, sensordata command
- “Check Internal Power and Temperature of a CMC” on page 15
- “Check the Internal Power and Temperature of a Line Card” on page 32

▼ Check Fabric Card Internal Temperatures

- On the CMC, type:

```
# showtemps | grep FC
Temperature on FC 0, LM75 min = 26.00 C, LM75 max = 27.00 C, at adm1026 = 27.00 C
Temperature on FC 1, LM75 min = 22.00 C, LM75 max = 24.00 C, at adm1026 = 23.00 C
Temperature on FCF 2, LM75 min = 22.00 C, LM75 max = 27.00 C, at adm1026 = 25.00 C
Temperature on FC 3, LM75 min = 21.00 C, LM75 max = 26.00 C, at adm1026 = 25.00 C
Temperature on FC 4, LM75 min = 21.00 C, LM75 max = 27.00 C, at adm1026 = 24.00 C
Temperature on FC 5, LM75 min = 22.00 C, LM75 max = 26.00 C, at adm1026 = 24.00 C
Temperature on FCF 6, LM75 min = 21.00 C, LM75 max = 26.00 C, at adm1026 = 24.00 C
Temperature on FC 7, LM75 min = 22.00 C, LM75 max = 24.00 C, at adm1026 = 24.00 C
Temperature on FC 8, LM75 min = 24.00 C, LM75 max = 28.00 C, at adm1026 = 27.00 C
#
```

Note – Fabric card fillers also provide temperatures.

Related Information

- *Switch Reference*, showtemps command
- [“Check Line Card Internal Temperatures” on page 33](#)

▼ Check Fabric Card Internal Voltages

Note – The output of this procedure is for all fabric cards and line cards.

- Take one of the following actions:
 - On the CMC, type:

```
# showvoltages
Reading M9 voltages...
FC 0 readings
 1.8V      = 1.81
 2.5V_0   = 2.51
 2.5V_1   = 2.51
 1.2V_0   = 1.20
 1.2V_1   = 1.20
 3.3V STBY = 3.38
 3.3V     = 3.38
 2.5V STBY = 2.54
 5V       = 5.14
```

```
12V          = 12.10
1.8V STBY   = 1.79
1.2V STBY   = 1.22

FC 1 readings
1.8V        = 1.82
2.5V_0      = 2.50
.
.
.
FCF 2 readings
12V         = 12.10
3.3V STBY   = 3.37

FC 3 readings
1.8V        = 1.80
2.5V_0      = 2.52
.
.
.
#
```

Note – The output in the example is just a portion of the full output.

Note – Fabric card fillers report only supplied voltages, not internal voltages.

- On the CMC, type:

```
# checkvoltages
Reading M9 voltages...
Checking FC 0 ...
FC 0 OK

Checking FC 1 ...
FC 1 OK

Checking FCF 2 ...
FCF 2 OK
.
.
.
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `showvoltages` command
- *Switch Reference*, `checkvoltages` command
- [“Check Line Card Internal Voltages” on page 33](#)

▼ Display the Base GUIDs of a Fabric Card

- On the CMC, type:

```
# getbaseguid fc slot
```

where *slot* is 0–8. See [“Fabric Card Addressing” on page 5](#). For example:

```
# getbaseguid fc 0
Base GUID = 0x0021283a8311f000
#
```

Note – Fabric card fillers have no switch chips. Therefore, no reporting is provided from them.

The fabric card switch chip GUIDs are determined by replacing the trailing two zeros (00) of the fabric card base GUID with:

- 0xa2
- 0xb2

For example, 0x00bad0fc0033f100 becomes:

- 0x00bad0fc0033f1a2
- 0x00bad0fc0033f1b2

Related Information

- *Switch Reference*, `getbaseguid` command
- [“Identify All Switches in the Fabric” on page 46](#)
- [“Display the Base GUIDs of a Line Card” on page 35](#)

▼ Check Fabric Card Link Status (Simple)

- On the CMC, type:

```
# checklinks | grep FC
FC 0 Active, checking links....OK
FC 1 Active, checking links....OK
FC 3 Active, checking links....OK
FC 4 Active, checking links....OK
FC 5 Active, checking links....OK
FC 7 Active, checking links....OK
FC 8 Active, checking links....OK
#
```

Note – Fabric card fillers have no switch chips. Therefore, no reporting is provided from them.

If no faults are found, the output is simple, such as displayed in the example. If faults are found, the output is detailed. Such output can be confusing. Use the procedure [“Check Fabric Card Link Status \(Detailed\)”](#) on page 26, to present the output in a more organized way.

Related Information

- *Switch Reference*, checklinks command
- [“Enable Downed Fabric Card Links”](#) on page 78
- [“Check Line Card Link Status \(Simple\)”](#) on page 36

▼ Check Fabric Card Link Status (Detailed)

If the simple check (see [“Check Fabric Card Link Status \(Simple\)”](#) on page 26) showed a failure in a fabric card, use the detailed method to gather more information about that fabric card.

- On the CMC, type:

```
# checklinks | grep 'FC slot '
```

where *slot* is 00–08 in the form of two digits. See [“Fabric Card Addressing”](#) on page 5.

Note – There is a space following the second digit of the *slot* and preceding the quote (').

For example

```
# checklinks | grep 'FC 00 '  
Port 07 on I4 00 FC 00 is down  
Port 04 on I4 01 FC 00 is down  
Port 19 on I4 01 FC 00 is down  
Port 20 on I4 01 FC 00 is down  
.  
.  
.  
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `checklinks` command
- [“Enable a Fabric Card Switch Chip Port” on page 79](#)
- [“Reset a Fabric Card Switch Chip” on page 80](#)
- [“Check Line Card Link Status \(Detailed\)” on page 37](#)

▼ Check Fabric Card Switch Chip Health

- On the CMC, type:

```
# checkswitches  
Checking booted switches in M9...  
FC 0 Active, checking switches ....OK  
FC 1 Active, checking switches ....OK  
FC 3 Active, checking switches ....OK  
FC 4 Active, checking switches ....OK  
FC 5 Active, checking switches ....OK  
FC 7 Active, checking switches ....OK  
FC 8 Active, checking switches ....OK  
#
```

Note – Fabric card fillers have no switch chips. Therefore, no reporting is provided from them.

Related Information

- *Switch Reference*, checkswitches command
- “Reset a Fabric Card Switch Chip” on page 80
- “Check Line Card Switch Chip Health” on page 38

▼ Check the IPMB State of a Fabric Card

- On the CMC, type:

```
# clia getipmbstate IPMB_address
```

where *IPMB_address* is from “Fabric Card Addressing” on page 5. For example, for fabric card 0:

```
# clia getipmbstate 94
Pigeon Point Shelf Manager Command Line Interpreter
94: LUN: 0, Sensor # 1 ("IPMB Physical")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
#
```

Note – Fabric card fillers can also report their IPMB state.

Related Information

- *Switch Reference*, getipmbstate command
- “Check the IPMB State of a Line Card” on page 38

▼ Check the Status LEDs of a Fabric Card

- On the CMC, type:

```
# clia getfruledstate -v IPMB_address |grep -e FRU -e color: -e supported
```

where *IPMB_address* is from “Fabric Card Addressing” on page 5. For example:

```
# clia getfruledstate -v 94 |grep -e FRU -e color: -e supported
94: FRU # 0, Led # 0 ("BLUE LED"):
    Colors supported(0x02): BLUE
94: FRU # 0, Led # 1 ("LED 1"):
    Colors supported(0x10): AMBER
94: FRU # 0, Led # 2 ("LED 2"):
    Local Control LED State: LED ON, color: GREEN
    Colors supported(0x08): GREEN
94: FRU # 0, Led # 3 ("LED 3"):
    Colors supported(0x40): WHITE
#
```

In this example, the green OK LED is lit.

Note – The LED is lit only if an LED is listed as LED ON or LED BLINKING. Otherwise, the LED is off.

Related Information

- *Switch Reference*, getfruledstate command
- “Check Fabric Card LEDs” on page 115
- “Check the Status LEDs of a CMC” on page 17
- “Check the Status LEDs of a Line Card” on page 39

▼ Display the Firmware Versions of a Fabric Card

- On the CMC, type:

```
# getfwversion fc slot
```

where *slot* is 0–8. See “Fabric Card Addressing” on page 5. For example:

```
# getfwversion fc 0
FW versions for FC 0
H8 version           : 0.1.4
FC FPGA version      : 1.0.1
I4 FW image version  : 7.2.300
INI file version     : 1
#
```

Note – Fabric card fillers only report H8 firmware versions.

Related Information

- *Switch Reference*, `getfwversion` command
- “Display the Firmware Version of a Power Supply” on page 14
- “Display the Firmware Version of a CMC” on page 17
- “Display the Firmware Version of the Switch Chassis Manager” on page 18
- “Display the Firmware Versions of a Line Card” on page 40
- “Display OFED Software Version Information” on page 57

Checking Line Cards

The following tasks enable you to check and verify the line card operation and status.

Note – For the examples in this topic, line card 8 (IPMB address 92), is used.

- “Display Line Cards Present” on page 31
- “Check Line Card Power Faults” on page 32
- “Check the Internal Power and Temperature of a Line Card” on page 32
- “Check Line Card Internal Temperatures” on page 33
- “Check Line Card Internal Voltages” on page 33
- “Display the Base GUIDs of a Line Card” on page 35

- “Check Line Card Link Status (Simple)” on page 36
- “Check Line Card Link Status (Detailed)” on page 37
- “Check Line Card Switch Chip Health” on page 38
- “Check the IPMB State of a Line Card” on page 38
- “Check the Status LEDs of a Line Card” on page 39
- “Display the Firmware Versions of a Line Card” on page 40

Related Information

- “Managing Line Cards” on page 81
- “Performing General Tasks” on page 11
- “Checking Power Supplies” on page 13
- “Checking CMCs” on page 15
- “Checking Fabric Cards” on page 18
- “Checking Other Switch Characteristics” on page 40

▼ Display Line Cards Present

- On the CMC, type:

```
# showpresent | grep LC
LC 00 present and active (state = M4)
LC 01 present and active (state = M4)
LC 02 present and active (state = M4)
LC 03 present and active (state = M4)
LC 04 present and active (state = M4)
LC 05 present and active (state = M4)
LC 06 present and active (state = M4)
LC 07 present and active (state = M4)
LC 08 present and active (state = M4)
#
```

In the output, (state = M4) means the FRU is active.

Related Information

- *Switch Reference*, showpresent command
- “Display Power Supplies Present” on page 13
- “Display Fabric Cards Present” on page 20

▼ Check Line Card Power Faults

- On the CMC, type:

```
# checkpwrfault | grep LC
LC 0 Power fault sensor = 0x00      OK
LC 1 Power fault sensor = 0x00      OK
LC 2 Power fault sensor = 0x00      OK
LC 3 Power fault sensor = 0x00      OK
LC 4 Power fault sensor = 0x00      OK
LC 5 Power fault sensor = 0x00      OK
LC 6 Power fault sensor = 0x00      OK
LC 7 Power fault sensor = 0x00      OK
LC 8 Power fault sensor = 0x00      OK
#
```

Related Information

- *Switch Reference*, checkpwrfault command
- [“Check the Status of a Power Supply” on page 14](#)
- [“Check the Status of a CMC” on page 16](#)
- [“Check Fan Speed and Status” on page 19](#)
- [“Check Fabric Card Power Faults” on page 21](#)

▼ Check the Internal Power and Temperature of a Line Card

- On the CMC, type:

```
# clia sensordata IPMB_address | grep -e LUN -e Processed
```

where *IPMB_address* is from [“Line Card Addressing” on page 6](#). For example, for line card 8:

```
# clia sensordata 92 | grep -e LUN -e Processed
92: LUN: 0, Sensor # 0 ("Hot Swap")
92: LUN: 0, Sensor # 1 ("IPMB Physical")
92: LUN: 0, Sensor # 2 ("LM75_0")
    Processed data: 31.000000 degrees C
92: LUN: 0, Sensor # 3 ("LM75_1")
    Processed data: 37.000000 degrees C
92: LUN: 0, Sensor # 4 ("Ambient Temp")
    Processed data: 35.000000 degrees C
92: LUN: 0, Sensor # 5 (" +1.8V STBY")
```

```
Processed data: 1.805400 Volts
.
.
.
92: LUN: 0, Sensor # 96 ("I4D_temperature")
Processed data: 39.000000 degrees C
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `sensordata` command
- [“Check Internal Power and Temperature of a CMC” on page 15](#)
- [“Check the Internal Power and Temperature of a Fabric Card” on page 22](#)

▼ Check Line Card Internal Temperatures

- On the CMC, type:

```
# showtemps | grep LC
Temperature on LC 0, LM75 min = 36.00 C, LM75 max = 41.00 C, at adm1026 = 39.00 C
Temperature on LC 1, LM75 min = 31.00 C, LM75 max = 33.00 C, at adm1026 = 35.00 C
Temperature on LC 2, LM75 min = 25.00 C, LM75 max = 27.00 C, at adm1026 = 27.00 C
Temperature on LC 3, LM75 min = 28.00 C, LM75 max = 28.00 C, at adm1026 = 30.00 C
Temperature on LC 4, LM75 min = 30.00 C, LM75 max = 36.00 C, at adm1026 = 34.00 C
Temperature on LC 5, LM75 min = 36.00 C, LM75 max = 41.00 C, at adm1026 = 39.00 C
Temperature on LC 6, LM75 min = 25.00 C, LM75 max = 27.00 C, at adm1026 = 27.00 C
Temperature on LC 7, LM75 min = 32.00 C, LM75 max = 35.00 C, at adm1026 = 37.00 C
Temperature on LC 8, LM75 min = 31.00 C, LM75 max = 37.00 C, at adm1026 = 35.00 C
#
```

Related Information

- *Switch Reference*, `showtemps` command
- [“Check Fabric Card Internal Temperatures” on page 23](#)

▼ Check Line Card Internal Voltages

Note – The output of this procedure is for all fabric cards and line cards.

- Take one of the following actions:

- On the CMC, type:

```
# showvoltages
Reading M9 voltages...
FC 0 readings
 1.8V      = 1.81
 2.5V_0    = 2.51
 2.5V_1    = 2.51
 1.2V_0    = 1.20
.
.
.
LC 0 readings
 1.8V STBY = 1.79
 2.5V_0    = 2.54
 2.5V_1    = 2.55
 2.5V_2    = 2.54
 2.5V_3    = 2.54
 1.2V_0    = 1.19
 1.2V_1    = 1.19
 1.2V_2    = 1.20
 1.2V_3    = 1.19
 3.3V STBY = 3.34
 3.3V      = 3.31
 2.5V STBY = 2.56
 5V AB     = 5.12
 5V CD     = 5.17
 12V       = 12.10
 1.8V      = 1.81
 1.2V STBY = 1.20

LC 1 readings
 1.8V STBY = 1.81
.
.
.
#
```

Note – The output in the example is just a portion of the full output.

- On the CMC, type:

```
# checkvoltages
Reading M9 voltages...
Checking FC 0 ...
```

```
FC 0 OK
Checking FC 1 ...
FC 1 OK
.
.
.
Checking LC 0 ...
LC 0 OK
.
.
.
Checking LC 8 ...
LC 8 OK
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `showvoltages` command
- *Switch Reference*, `checkvoltages` command
- [“Check Fabric Card Internal Voltages” on page 23](#)

▼ Display the Base GUIDs of a Line Card

- On the CMC, type:

```
# getbaseguid lc slot
```

where *slot* is 0–8. See [“Line Card Addressing” on page 6](#). For example:

```
# getbaseguid lc 0
Base GUID = 0x0021283a83b71000
#
```

The line card switch chip GUIDs are determined by replacing the trailing two zeros (00) of the fabric card base GUID with:

- 0xa2
- 0xb2
- 0xc2

- 0xd2

For example, 0x00bad0fc30331800 becomes:

- 0x00bad0fc303318a2
- 0x00bad0fc303318b2
- 0x00bad0fc303318c2
- 0x00bad0fc303318d2

Related Information

- *Switch Reference*, getbaseguid command
- [“Identify All Switches in the Fabric” on page 46](#)
- [“Display the Base GUIDs of a Fabric Card” on page 25](#)

▼ Check Line Card Link Status (Simple)

- On the CMC, type:

```
# checklinks -c |grep LC
LC 0 Active, checking links.....OK
LC 1 Active, checking links.....OK
LC 2 Active, checking links.....OK
LC 3 Active, checking links.....OK
LC 4 Active, checking links.....OK
LC 5 Active, checking links.....OK
LC 6 Active, checking links.....OK
LC 7 Active, checking links.....OK
LC 8 Active, checking links.....OK
#
```

If no faults are found, the output is simple, such as displayed in the example. If faults are found, the output is detailed. Such output can be confusing. Use the procedure [“Check Line Card Link Status \(Detailed\)” on page 37](#), to present the output in a more organized way.

Related Information

- *Switch Reference*, checklinks command
- [“Enable Downed Line Card Links” on page 88](#)
- [“Check Fabric Card Link Status \(Simple\)” on page 26](#)

▼ Check Line Card Link Status (Detailed)

If the simple check (see [“Check Line Card Link Status \(Simple\)”](#) on page 36) showed a failure in a line card, use the detailed method to gather more information about that line card.

- On the CMC, type:

```
# checklinks -c | grep 'LC slot '
```

where *slot* is 00–08 in the form of two digits. See [“Line Card Addressing”](#) on page 6.

Note – There is a space following the second digit of the *slot* and preceding the quote (').

For example:

```
# checklinks -c | grep 'LC 08 '  
LC 8 Active, checking links.....  
Port 25 on I4 03 LC 08 is down      (Cable 10 A )  
Port 30 on I4 03 LC 08 is down      (Cable 10 A )  
.   
.   
.   
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- [Switch Reference](#), `checklinks` command
- [“Enable a Line Card Switch Chip Port”](#) on page 89
- [“Reset a Line Card Switch Chip”](#) on page 90
- [“Check Fabric Card Link Status \(Detailed\)”](#) on page 26

▼ Check Line Card Switch Chip Health

- On the CMC, type:

```
# checkswitches
Checking booted switches in M9...
.
.
.
LC 0 Active, checking switches ....OK
LC 1 Active, checking switches ....OK
LC 2 Active, checking switches ....OK
LC 3 Active, checking switches ....OK
LC 4 Active, checking switches ....OK
LC 5 Active, checking switches ....OK
LC 6 Active, checking switches ....OK
LC 7 Active, checking switches ....OK
LC 8 Active, checking switches ....OK
#
```

Related Information

- *Switch Reference*, checkswitches command
- [“Reset a Line Card Switch Chip” on page 90](#)
- [“Check Fabric Card Switch Chip Health” on page 27](#)

▼ Check the IPMB State of a Line Card

- On the CMC, type:

```
# clia getipmbstate IPMB_address
```

where *IPMB_address* is from [“Line Card Addressing” on page 6](#). For example, for line card 8:

```
# clia getipmbstate 92
Pigeon Point Shelf Manager Command Line Interpreter
92: LUN: 0, Sensor # 1 ("IPMB Physical")
  Bus Status: 0x08 (IPMB-A Enabled, IPMB-B Enabled)
  IPMB A State: 0x08 (LocalControl, No failure)
  IPMB B State: 0x08 (LocalControl, No failure)
#
```

Related Information

- *Switch Reference*, `getipmbstate` command
- [“Check the IPMB State of a Fabric Card” on page 28](#)

▼ Check the Status LEDs of a Line Card

- On the CMC, type:

```
# clia getfruledstate -v IPMB_address |grep -e FRU -e color: -e supported
```

where `IPMB_address` is from [“Line Card Addressing” on page 6](#). For example:

```
# clia getfruledstate -v 92 |grep -e FRU -e color: -e supported
92: FRU # 0, Led # 0 ("BLUE LED"):
    Colors supported(0x02): BLUE
92: FRU # 0, Led # 1 ("LED 1"):
    Colors supported(0x10): AMBER
92: FRU # 0, Led # 2 ("LED 2"):
    Local Control LED State: LED ON, color: GREEN
    Colors supported(0x08): GREEN
92: FRU # 0, Led # 3 ("LED 3"):
    Override LED State (current state): LED ON, color: WHITE
    Colors supported(0x40): WHITE
#
```

In this example, both the green OK LED and the white Locator LED are lit.

Note – The LED is lit only if an LED is listed as LED ON or LED BLINKING. Otherwise, the LED is off.

Related Information

- *Switch Reference*, `getfruledstate` command
- [“Check Line Card LEDs” on page 116](#)
- [“Check the Status LEDs of a CMC” on page 17](#)
- [“Check the Status LEDs of a Fabric Card” on page 29](#)

▼ Display the Firmware Versions of a Line Card

- On the CMC, type:

```
# getfwversion lc slot
```

where *slot* is 0–8. See “Line Card Addressing” on page 6. For example:

```
# getfwversion lc 8
FW versions for LC 8
H8 version           : 0.1.4
LC FPGA version     : 1.0.1
I4 FW image version  : 7.2.300
INI file version     : 1
#
```

Related Information

- *Switch Reference*, `getfwversion` command
- “Display the Firmware Version of a Power Supply” on page 14
- “Display the Firmware Version of a CMC” on page 17
- “Display the Firmware Version of the Switch Chassis Manager” on page 18
- “Display the Firmware Versions of a Fabric Card” on page 30
- “Display OFED Software Version Information” on page 57

Checking Other Switch Characteristics

The following tasks display various CMC information.

- “Display the Date” on page 41
- “Display the User Accounts” on page 41
- “Display the Network Management Configuration” on page 41
- “Locate a Switch Chip or Connector From the GUID” on page 43

Related Information

- “Performing General Tasks” on page 11
- “Checking Power Supplies” on page 13
- “Checking CMCs” on page 15
- “Checking Fabric Cards” on page 18
- “Checking Line Cards” on page 30

▼ Display the Date

- On the CMC, type:

```
# date
Fri Oct 30 01:52:21 UTC 2009
#
```

Related Information

- [“Set the Date” on page 59](#)

▼ Display the User Accounts

- On the CMC, type:

```
# clia user -v
Pigeon Point Shelf Manager Command Line Interpreter
1: ""
   Channels 0-15 Privilege level: "Administrator"
   Flags: "IPMI Messaging"
2: "openhpi"
   Channels 0-15 Privilege level: "OEM Proprietary"
   Flags: "IPMI Messaging"
#
```

Related Information

- *Switch Reference*, user command
- [“Setting Up the Users” on page 64](#)

▼ Display the Network Management Configuration

- On the CMC, type:

```
# clia getlanconfig 1
Pigeon Point Shelf Manager Command Line Interpreter
Authentication Type Support: 0x17 ( None MD2 MD5 Straight Password/Key )
Authentication Type Enables:
  Callback level: 0x00
  User level: 0x17 ( "None" "MD2" "MD5" "Straight Password/Key" )
  Operator level: 0x17 ( "None" "MD2" "MD5" "Straight Password/Key" )
  Administrator level: 0x17 ( "None" "MD2" "MD5" "Straight Password/Key" )
  OEM level: 0x00
```

```

IP Address: 10.60.34.20
IP Address Source: Static Address (Manually Configured) (0x01)
MAC Address: 00:18:49:00:86:32
Subnet Mask: 255.255.255.0
IPv4 Header Parameters: 0x40:0x40:0x10
Primary RMCP Port Number: 0x026f
Secondary RMCP Port Number: 0x0298
BMC-generated ARP Control: 0x02
    Enable BMC-generated ARP Response
Gratuitous ARP Interval: 2.0 seconds
Default Gateway Address: 10.60.34.254
Default Gateway MAC Address: 00:00:0c:07:ac:22
Backup Gateway Address: 0.0.0.0
Backup Gateway MAC Address: N/A
Community String: "public"
Number of Destinations: 16
Destination Type:
    N/A
Destination Address:
    N/A
802.1q VLAN ID: 0 (disabled)
VLAN priority: 0
Cipher Suite Entry count: 15
Supported Cipher Suite IDs: 0h, 1h, 2h, 3h, 4h, 5h, 6h, 7h, 8h, 9h, Ah, Bh, Ch,
Dh, Eh
Cipher Suite Privilege Levels:
    ID 00h, Priv.Level 'User' (2); ID 01h, Priv.Level 'User' (2);
    ID 02h, Priv.Level 'Administrator' (4); ID 03h, Priv.Level 'OEM Proprietary'
(5);
    ID 04h, Priv.Level 'OEM Proprietary' (5); ID 05h, Priv.Level 'OEM Proprietary'
(5);
    ID 06h, Priv.Level 'User' (2); ID 07h, Priv.Level 'Administrator'
(4);
    ID 08h, Priv.Level 'OEM Proprietary' (5); ID 09h, Priv.Level 'OEM Proprietary'
(5);
    ID 0Ah, Priv.Level 'OEM Proprietary' (5); ID 0Bh, Priv.Level 'Administrator'
(4);
    ID 0Ch, Priv.Level 'OEM Proprietary' (5); ID 0Dh, Priv.Level 'OEM Proprietary'
(5);
    ID 0Eh, Priv.Level 'OEM Proprietary' (5);
Destination Address VLAN TAGs:
    N/A
#

```

Related Information

- *Switch Reference*, getlanconfig command
- [“Reconfigure the Network Management Parameters” on page 61](#)

▼ Locate a Switch Chip or Connector From the GUID

The output of some InfiniBand software commands identify a node by its GUID. The `findport` command displays the location of the node within the switch and also indicates if the port is attached to a connector, and which one.

- On the CMC, type:

```
# findport -g guid port
```

where:

- *guid* is the global unit identifier with 0x truncated.
- *port* is the number of the port (1–36).

For example:

```
# findport -g 0021283a83ae11d2 19
Port 19 on switch 3 on LC 1 cable 9 B-3
#
```

In the example output, the node with GUID 0021283a83ae11d2 is switch chip 3 in line card 1. The link is to CXP cable connector 9B. The -3 means P3 carries the link, and link status is indicated with the left Link LED.

Related Information

- [“Switch GUIDs Overview”](#) on page 136
- [“Determine the GUID and LID for a Node Within the Switch”](#) on page 55

Monitoring the InfiniBand Fabric

The following tasks enable you to display and check the operation and status of the InfiniBand fabric and components.

Note – Commands in this topic are issued from the Linux InfiniBand host as the `root` user. You cannot issue the commands on the CMC.

- [“Display Information About the Local HCA”](#) on page 44
- [“Identify All HCAs in the Fabric”](#) on page 45
- [“Identify All Switches in the Fabric”](#) on page 46

- “Display the InfiniBand Fabric Topology” on page 47
- “Display a Route Through the Fabric” on page 47
- “Display the Link Status of a Node” on page 49
- “Display Counters for a Node” on page 50
- “Display Data Counters for a Node” on page 51
- “Display Low-Level Detailed Information About a Node” on page 52
- “Display Low-Level Detailed Information About a Port” on page 53
- “Map LIDs to GUIDs” on page 54
- “Determine the GUID and LID for a Node Within the Switch” on page 55
- “Display Subnet Manager Status” on page 56
- “Display OFED Software Version Information” on page 57

Related Information

- “Understanding Administrative Commands” on page 1
- “Monitoring the Switch” on page 10
- “Controlling the Switch” on page 57
- “Controlling the InfiniBand Fabric” on page 91
- “Determine Changes to the InfiniBand Fabric Topology” on page 92

▼ Display Information About the Local HCA

If you need to know the local HCA’s firmware version, vendor information, or port state information, you can use the `ibv_devinfo` command. If you add the `-v` option, the output is very detailed.

- **On the Linux InfiniBand host, type:**

```
# ibv_devinfo
hca_id:mlx4_0
  fw_ver:          2.5.9266
  node_guid:       0003:ba00:0100:c708
  sys_image_guid:  0003:ba00:0100:c70b
  vendor_id:       0x02c9
  vendor_part_id:  26428
  hw_ver:          0xA0
  board_id:        SUN0150000001
  phys_port_cnt:  2
    port:1
      state:       PORT_DOWN (1)
```

```
max_mtu:      2048 (4)
active_mtu:   2048 (4)
sm_lid:       0
port_lid:     0
port_lmc:     0x00
port:2
state:        PORT_ACTIVE (4)
max_mtu:      2048 (4)
active_mtu:   2048 (4)
sm_lid:       58
port_lid:     58
port_lmc:     0x00
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `ibv_devinfo` command

▼ Identify All HCAs in the Fabric

Similar to the `ibswitches` command, the `ibhosts` command displays identity information about the HCAs in the InfiniBand fabric. The output contains the GUID and name for each HCA.

Note – Unlike the `ibswitches` command, the `ibhosts` command does not display the LIDs.

- On the Linux InfiniBand host, type:

```
# ibhosts
Ca      : 0x5080020000911314 ports 1 "nnsn32-50 HCA-1"
Ca      : 0x5080020000911310 ports 1 "nnsn32-20 HCA-1"
Ca      : 0x50800200008e532c ports 1 "ib-71 HCA-1"
Ca      : 0x50800200008e5328 ports 1 "ib-70 HCA-1"
Ca      : 0x50800200008296a4 ports 2 "ib-90 HCA-1"
.
.
.
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `ibhosts` command

▼ Identify All Switches in the Fabric

If you need to know identity information about the switches in the InfiniBand fabric, you can use the `ibswitches` command. This command displays the GUID, name, LID, and LMC for each switch. The output of the command is a mapping of GUID to LID for switches in the fabric.

- **On the Linux InfiniBand host, type:**

```
# ibswitches
Switch : 0x0021283a83ae11d2 ports 36 "Sun DCS 648 shmm1500 LC slot 1 switch 3"
base port 0 lid 35 lmc 0
Switch : 0x0021283a83b112b2 ports 36 "Sun DCS 648 shmm1500 LC slot 2 switch 1"
base port 0 lid 17 lmc 0
Switch : 0x0021283a83b112a2 ports 36 "Sun DCS 648 shmm1500 LC slot 2 switch 0"
base port 0 lid 4 lmc 0
Switch : 0x0021283a83b112d2 ports 36 "Sun DCS 648 shmm1500 LC slot 2 switch 3"
base port 0 lid 36 lmc 0
Switch : 0x0021283a841513c2 ports 36 "Sun DCS 648 shmm1500 LC slot 3 switch 2"
base port 0 lid 31 lmc 0
.
.
.
#
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `ibswitches` command
- [“Display the Base GUIDs of a Fabric Card” on page 25](#)
- [“Display the Base GUIDs of a Line Card” on page 35](#)

▼ Display the InfiniBand Fabric Topology

To understand the routing that happens within your InfiniBand fabric, the `ibnetdiscover` command displays the node-to-node connectivity. The output of the command is dependent upon the size of your fabric.

- On the Linux InfiniBand host, type:

```
# ibnetdiscover
```

Note – You can use the `ibnetdiscover` command to determine the LIDs of the HCAs.

Related Information

- *Switch Reference*, `ibnetdiscover` command
- [“Determine Changes to the InfiniBand Fabric Topology”](#) on page 92
- [“Perform Comprehensive Diagnostics for the Entire Fabric”](#) on page 91

▼ Display a Route Through the Fabric

Often you need to know the route between two nodes. The `ibtracert` command can provide that information, displaying the GUIDs, ports, and LIDs of the nodes along the route.

- On the Linux InfiniBand host, type:

```
# ibtracert slid dlid
```

where:

- *slid* is the LID of the source node.

- *dlid* is the LID of the destination node.

For example:

```
# ibtracert 58 57
From ca {0x0003ba000100c708} portnum 2 lid 58-58 "nsn34-39 HCA-1"
[2] -> switch port {0x00bad0cc010016a2}[22] lid 1-1 "Sun DCS 648 shmm1500 LC slot
6 switch 0"
[1] -> switch port {0x00bad0fc0100f8b2}[7] lid 57-57 "Sun DCS 648 shmm1500 FC
slot 8 switch 1"
To switch {0x00bad0fc0100f8b2} portnum 0 lid 57-57 "Sun DCS 648 shmm1500 FC slot
8 switch 1"
#
```

For this example:

- The route starts at the local HCA hosting the Subnet Manager. The HCA has GUID 0x0003ba000100c708 and is using port 2. The HCA is LID 58 and in the description, the HCA host's name is nsn34-39.
- The route enters at port 2 of the switch with GUID 0x00bad0cc010016a2 and exits at port 22. The switch is LID 1, and from the description is inside line card 6 and is switch chip 0.
- The route enters at port 1 of the switch with GUID 0x00bad0fc0100f8b2 and exits at port 7. The switch is LID 57, and from the description is inside fabric card 8 and is switch chip 1.
- The route ends at switch with GUID 0x00bad0fc0100f8b2 at port 0. The switch is LID 57, and from the description is inside fabric card 8 and is switch chip 1.

Note – The ports that are identified in this route might *not* be the same ports as are identified by the switch-specific commands and CLIA commands.

Note – You can use the -G option to specify GUIDs instead of LIDs for the `ibtracert` command.

Related Information

- *Switch Reference*, `ibtracert` command
- [“Perform Comprehensive Diagnostics for a Route” on page 92](#)
- [“Understanding Signal Routing Through the Switch” on page 124](#)

▼ Display the Link Status of a Node

If you want to know the link status of a node in the InfiniBand fabric, the `ibportstate` command can tell you the state, width, and speed of that node.

- On the Linux InfiniBand host, type:

```
# ibportstate lid port
```

where:

- *lid* is the LID of the node.
- *port* is the port of the node.

For example:

```
# ibportstate 58 2  
PortInfo:  
# Port info: Lid 58 port 2  
LinkState:.....Active  
PhysLinkState:.....LinkUp  
LinkWidthSupported:.....1X or 4X  
LinkWidthEnabled:.....1X or 4X  
LinkWidthActive:.....4X  
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps  
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps  
LinkSpeedActive:.....10.0 Gbps  
#
```

In the output, the Active parameters are the current state of the port.

Note – You can use the `-G` option to specify a GUID and port instead of a LID and port.

Related Information

- *Switch Reference*, `ibportstate` command
- [“Check Fabric Card Link Status \(Detailed\)”](#) on page 26
- [“Check Line Card Link Status \(Detailed\)”](#) on page 37
- [“Find 1x or SDR or DDR Links in the Fabric”](#) on page 94
- [“Set the Speed of a Port”](#) on page 98

▼ Display Counters for a Node

To help ascertain the health of a node, the `perfquery` command is used to display the performance, error, and data counters for that node.

- **On the Linux InfiniBand host, type:**

```
# perfquery lid port
```

where:

- *lid* is the LID of the node.
- *port* is the port of the node.

Note – If a *port* value of 255 is specified for a switch node, the counters are the total for all switch ports.

For example:

```
# perfquery 57 255
# Port counters: Lid 57 port 255
PortSelect:.....255
CounterSelect:.....0x1b01
SymbolErrors:.....0
.
.
.
VL15Dropped:.....0
XmtData:.....2670336
RcvData:.....2672928
XmtPkts:.....37088
RcvPkts:.....37124
XmtWait:.....0
#
```

Note – The output in the example is just a portion of the full output.

Note – You can use the `-G` option to specify a GUID and port instead of a LID and port.

Related Information

- *Switch Reference*, `perfquery` command

- [“Clear Error Counters” on page 96](#)

▼ Display Data Counters for a Node

If you want to know the data counters for a node, the `ibdatacounts` command provides that subset of the `perfquery` command output.

- **On the Linux InfiniBand host, type:**

```
# ibdatacounts lid port
```

where:

- *lid* is the LID of the node.
- *port* is the port of the node.

Note – If no *port* is specified for a switch node, the counters are the total for all switch ports.

For example:

```
# ibdatacounts 57
XmtData:.....2675880
RcvData:.....2678472
XmtPkts:.....37165
RcvPkts:.....37201
#
```

Note – You can use the `-G` option to specify a GUID and port instead of a LID and port.

Related Information

- [Switch Reference, `ibdatacounts` command](#)
- [“Clear Data Counters” on page 96](#)

▼ Display Low-Level Detailed Information About a Node

If intensive troubleshooting is necessary to resolve a problem, the `smpquery` command can provide very detailed information about a node or one of its ports.

- On the Linux InfiniBand host, type:

```
# smpquery switchinfo lid
```

where *lid* is the LID of the node.

For example, to see detailed information about a switch with LID 200, type:

```
# smpquery switchinfo 57
# Switch info: Lid 57
LinearFdbCap:.....49152
RandomFdbCap:.....0
McastFdbCap:.....4096
LinearFdbTop:.....85
DefPort:.....0
DefMcastPrimPort:.....255
DefMcastNotPrimPort:.....255
LifeTime:.....18
StateChange:.....0
OptSLtoVLMapping:.....1
LidsPerPort:.....0
PartEnforceCap:.....32
InboundPartEnf:.....1
OutboundPartEnf:.....1
FilterRawInbound:.....1
FilterRawOutbound:.....1
EnhancedPort0:.....0
MulticastFDBTop:.....0x0000
#
```

Note – You can use the `-G` option to specify a GUID instead of a LID.

Related Information

- *Switch Reference*, `smpquery` command

▼ Display Low-Level Detailed Information About a Port

The `smpquery` command can also provide very detailed information about a port.

- On the Linux InfiniBand host, type:

```
# smpquery portinfo lid port
```

where:

- `lid` is the LID of the node.
- `port` is the port of the node.

For example, to see detailed information about port 1 on the switch with LID 200, type:

```
# smpquery portinfo 57 1
# Port info: Lid 57 port 1
Mkey:.....0x0000000000000000
GidPrefix:.....0x0000000000000000
Lid:.....0
SMLid:.....0
CapMask:.....0x0
DiagCode:.....0x0000
MkeyLeasePeriod:.....0
LocalPort:.....7
LinkWidthEnabled:.....1X or 4X
LinkWidthSupported:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkState:.....Active
PhysLinkState:.....LinkUp
LinkDownDefState:.....Polling
ProtectBits:.....0
LMC:.....0
.
.
.
SubnetTimeout:.....0
RespTimeVal:.....0
LocalPhysErr:.....8
OverrunErr:.....8
MaxCreditHint:.....85
RoundTrip:.....16777215
#
```

Note – The output in the example is just a portion of the full output.

Note – You can use the `-G` option to specify a GUID and port instead of a LID and port.

Related Information

- *Switch Reference*, `smpquery` command

▼ Map LIDs to GUIDs

In an InfiniBand fabric, the Subnet Manager and Subnet Administrator assign subnet specific LIDs to nodes. Often in the use of the InfiniBand software commands, you must provide a LID to issue a command to a particular InfiniBand device. Alternatively, the output of the command might identify InfiniBand devices by their LID.

You can create a file that is a mapping of node LIDs to node GUIDs, which can help with administrating your InfiniBand fabric. The following procedure creates a file that lists the LID in hexadecimal, the GUID in hexadecimal, and the node description.

Note – Creation of the mapping file is not a requirement for InfiniBand administration.

1. Create an inventory file:

```
# osmtest -f c -i inventory.txt
```

Note – The `inventory.txt` file can be used for other purposes too, besides this procedure.

2. Create a mapping file:

```
# cat inventory.txt | grep -e '^lid' -e 'port_guid' -e 'desc' | sed 's/^lid/\nlid/' > mapping.txt
```

3. Edit the latter half of the `mapping.txt` file to remove the nonessential information.

The content of the `mapping.txt` file looks similar to the following

```
lid                0x100
port_guid          0x00144f6c67be0002
# node_desc        nsn32-10 HCA-2

lid                0x101
port_guid          0x0144f6c666b50102
# node_desc        nsn32-10 HCA-1

lid                0x1
port_guid          0x0003ba000100c70a
# node_desc        nsn34-39 HCA-1
.
.
.
```

Note – The output in the example is just a portion of the full output.

Related Information

- *Switch Reference*, `osmtest` command

▼ Determine the GUID and LID for a Node Within the Switch

1. Create the search string literal to find the GUID.

```
LC slot slot switch switch_chip
```

```
FC slot slot switch switch_chip
```

where:

- *slot* the number of the line card (0–8) or fabric card (0–8).
- *switch_chip* is the number of the switch chip on the line card (0–3) or fabric card (0–1).

For example, for line card 3, switch chip 2:

```
LC slot 3 switch 2
```

2. On the Linux InfiniBand host, use the search string literal with the `ibswitches` command to determine the GUID and LID for the node.

For example:

```
# ibswitches |grep 'LC slot 3 switch 2'
Switch : 0x0021283a841513c2 ports 36 "Sun DCS 648 shmm1500 LC slot 3 switch 2"
base port 0 lid 31 lmc 0
#
```

3. Visually inspect the output.

For the example, the GUID is 0x0021283a841513c2 and the LID is 31.

Note – If you have more than one Sun Datacenter InfiniBand Switch 648 in your InfiniBand fabric, you might have more than one GUID displayed in the output. Verify that the NodeDescription displayed with the GUID is correct for the desired switch.

Related Information

- [“Switch GUIDs Overview” on page 136](#)
- [“Locate a Switch Chip or Connector From the GUID” on page 43](#)

▼ Display Subnet Manager Status

If you want to quickly determine your Subnet Manager’s priority and state, the `sminfo` command can also provide the LID and GUID of the hosting HCA.

- On the Linux InfiniBand host, type:

```
# sminfo
sminfo: sm lid 58 sm guid 0x3ba000100c70a, activity count 47808 priority 0 state
3 SMINFO_MAS
#
```

In the output, the Subnet Manager’s hosting HCA has LID 1 and GUID 0x3ba000100c70a. The Subnet Manager has a priority of 0 (lowest) and its state is 3 (master).

Related Information

- *Switch Reference*, `sminfo` command
- [“Start the Subnet Manager With Min Hop Routing” on page 102](#)
- [“Start the Subnet Manager With Fat Tree Routing” on page 103](#)

- [“Start the Subnet Manager With the opensmd Daemon”](#) on page 106

▼ Display OFED Software Version Information

If you are managing the software components of the OFED software suite in a piecemeal fashion, the `ofed_info` command can provide you the software versions of each component in the OFED suite.

- **On the Linux InfiniBand host, type:**

```
# ofed_info
```

Related Information

- *Switch Reference*, `ofed_info` command
- [“Display the Firmware Version of a Power Supply”](#) on page 14
- [“Display the Firmware Version of a CMC”](#) on page 17
- [“Display the Firmware Version of the Switch Chassis Manager”](#) on page 18
- [“Display the Firmware Versions of a Fabric Card”](#) on page 30
- [“Display the Firmware Versions of a Line Card”](#) on page 40

Controlling the Switch

The following topics describe how you can manage the switch and its components.

Note – Commands in this topic are issued on the CMC as the `root` user.

- [“Setting Up the Hardware”](#) on page 58
- [“Setting Up the Users”](#) on page 64
- [“Managing Power Supplies”](#) on page 66
- [“Managing CMCs”](#) on page 69
- [“Managing Fabric Cards”](#) on page 72
- [“Managing Line Cards”](#) on page 81

Related Information

- [“Understanding Administrative Commands”](#) on page 1

- “Monitoring the Switch” on page 10
- “Monitoring the InfiniBand Fabric” on page 43
- “Controlling the InfiniBand Fabric” on page 91

Setting Up the Hardware

The following commands enable you to set up basic CMC configuration information.

- “Change the root Password” on page 58
- “Get Help on a CLIA Command” on page 59
- “Set the Date” on page 59
- “Change the CMC Host Name” on page 60
- “Reconfigure the Network Management Parameters” on page 61
- “Reconfigure the CMCs for Identical Addresses” on page 63

Related Information

- “Setting Up the Users” on page 64
- “Managing Power Supplies” on page 66
- “Managing CMCs” on page 69
- “Managing Fabric Cards” on page 72
- “Managing Line Cards” on page 81

▼ Change the root Password

As shipped, the root password is not set and therefore not required. Press the Enter key when prompted for the password. To improve security, follow this procedure to set a new password.

- **On the CMC, type:**

```
# passwd
Changing password for root
Enter the new password (minimum of 5, maximum of 8 characters)
Please use a combination of upper and lower case letters and
numbers.
Enter new password: new-password
```

```
Re-enter new password: new-password
Password changed.
#
```

Related Information

- [“Change a User’s Password” on page 65](#)

▼ Get Help on a CLIA Command

- On the CMC, type:

```
# clia help command option
```

where:

- *command* – A clia command.
- *option* – A subordinate command or option to the command.

For example:

```
# clia help user add
Pigeon Point Shelf Manager Command Line Interpreter
Add user with the same setting on all channels
user add <user id> <user name> <flags> <privilege level> <password>
#
```

Related Information

- [Switch Reference, help command](#)

▼ Set the Date

- On the CMC, type:

```
# date [mmd][HHMM][yyyy] [.SS]
```

where:

- *mm* – Month
- *dd* – Date
- *HH* – Hour (24-hour format)
- *MM* – Minutes
- *yyyy* – Year

- SS – Seconds

For example:

```
# date 10300152
Fri Oct 30 01:52:01 UTC 2009
#
```

Related Information

- [“Display the Date” on page 41](#)

▼ Change the CMC Host Name

You can change the host name of the CMC to something which identifies the CMC or the switch where it is installed.

Note – The host name is propagated to the `NodeDescription` field seen in the output of many InfiniBand software commands.

1. **Access the CMC from its serial management port.**
2. **Log in, then invoke a reboot of the CMC:**

```
# reboot
```

The CMC reboots.

3. **When you see the following output, press any key to stop the boot process:**

```
hit any key to stop autoboot
```

The boot monitor prompt is displayed.

```
shmm1500
```

4. **Set the backup state IP address of the CMC:**

```
shmm1500 setenv hostname hostname
```

where *hostname* is the identifier to apply to the CMC. For example:

```
shmm1500 setenv hostname CMC-A
```

5. Save the configuration and continue booting the switch:

```
shmm1500 saveenv  
shmm1500 boot
```

The switch boots.

The boot monitor prompt is now the hostname. For example:

```
CMC-A boot
```

The hostname is also propagated to the NodeDescription field in the output of the m9diag command and InfiniBand software commands. For example:

```
# m9diag  
Oct 22 18:16:47 CMC-A user.warn kernel: POST memory PASSED  
Oct 22 18:16:47 CMC-A user.warn kernel: POST i2c PASSED  
Oct 22 18:16:47 CMC-A user.warn kernel: POST uart PASSED  
Oct 22 18:16:47 CMC-A user.warn kernel: POST ethernet PASSED  
.  
.  
.  
#
```

Related Information

- *Switch Reference*, `ibswitches` command
- [“Reconfigure the CMCs for Identical Addresses” on page 63](#)

▼ Reconfigure the Network Management Parameters

This task enables you to configure the NET MGT interface. You can set network parameters according to the specific details of your network using the `clia setlanconfig` command.

Note – The following procedure makes changes that do not require a reset or reboot.

1. **Log into the CMC from the serial management port.**
2. **Set the IP address for the CMC:**

```
# clia setlanconfig 1 3 IP_address
```

where *IP_address* is in the form of *xxx.xxx.xxx.xxx*.

3. Set the netmask for the CMC:

```
# clia setlanconfig 1 6 netmask
```

where *netmask* is the netmask. Typically, the netmask is 255.255.255.0. However, your network environment subnet might require a different netmask. Use a netmask number most appropriate to your environment.

4. Set the IP address for the CMC gateway:

```
# clia setlanconfig 1 12 gateway_IP_address
```

where *gateway_IP_address* is in the form of *xxx.xxx.xxx.xxx*.

5. Open a Telnet session and connect to the CMC by specifying the CMC's network address.

For example:

```
% telnet 123.45.67.89  
Trying 123.45.67.89...  
Connected to 123.45.67.89.  
Escape character is '^]'.  
Copyright 2003 Sun Microsystems, Inc. All rights reserved.  
Use is subject to license terms.  
Pigeon Point Shelf Manager Command Line Interpreter  
Please login:
```

6. Log in as root:

```
Please login: root  
Please Enter password: password  
#
```

The # prompt is displayed.

Related Information

- *Switch Reference*, setlanconfig command
- telnet man page
- “Display the Network Management Configuration” on page 41
- “Reconfigure the CMCs for Identical Addresses” on page 63
- “Change the CMC Host Name” on page 60

▼ Reconfigure the CMCs for Identical Addresses

Each CMC has two IP addresses, an active state IP address used when the CMC is active, and a backup state IP address used when the CMC is inactive. For ease of management console operation, the active state IP addresses of both CMCs should be the same. An advantage of identical active state IP addresses is that the network management console will always log into the active CMC.

This procedure assumes that both CMCs are currently configured.

1. Access the inactive CMC from its serial management port.
2. Log in, then invoke a reboot of the CMC:

```
# reboot
```

The CMC reboots.

3. When you see the following output, press any key to stop the boot process:

```
hit any key to stop autoboot
```

The boot monitor prompt is displayed.

```
shmm1500
```

4. Set the backup state IP address of the CMC:

```
shmm1500 setenv ipaddr backup_IP_address
```

where *backup_IP_address* is the backup state IP address common to both CMCs.

Note – Setting the backup state IP addresses of both CMCs to the same value causes a reallocation according to this rule: If the *backup_IP_address* is odd, then it is assigned to CMC1 and CMC0 is assigned *backup_IP_address+1*. If the *backup_IP_address* is even, then it is assigned to CMC0 and CMC1 is assigned *backup_IP_address +1*

5. Save the configuration and continue booting the switch:

```
shmm1500 saveenv  
shmm1500 boot
```

The switch boots.

6. Log into the CMC and set the active state IP address for the CMC:

```
# clia setlanconfig 1 3 active_IP_address
```

where *active_IP_address* is the active state IP address common to both CMCs.

7. Perform a switchover:

```
# clia switchover
```

8. If the common backup state IP address or common active state IP address is new to both CMCs, repeat Step 1 to Step 7 for the other CMC.

CMC0 and CMC1 now share the same active state IP address. Should one CMC become inactive, the second CMC is reachable at the same active state IP address.

Related Information

- *Switch Reference*, `setlanconfig` command
- *Switch Reference*, `switchover` command
- *Switch Installation*, powering on the switch
- “Reconfigure the Network Management Parameters” on page 61
- “Change the CMC Host Name” on page 60

Setting Up the Users

These tasks enable you to control user accessibility.

- “Add a User” on page 65
- “Change a User’s Password” on page 65
- “Delete a User” on page 66

Related Information

- “Setting Up the Hardware” on page 58
- “Managing Power Supplies” on page 66
- “Managing CMCs” on page 69
- “Managing Fabric Cards” on page 72
- “Managing Line Cards” on page 81

▼ Add a User

- On the CMC, type:

```
# cli user add userid "user_name" channel_access_flags privilege_level password
```

where:

- *userid* – A valid user identifying number.
- *user_name* – The user name (truncated to 16 characters).
- *channel-access_flag* – The first byte of the SetUserInfo command (only bits 4, 5, and 6 are meaningful).
 - Bit 6 – IPMI messaging enabled.
 - Bit 5 – Link authentication enabled.
 - Bit 4 – Restricted to callback.
- *privilege_level* – The user privilege level.
- *password* – The user's password (truncated to 16 characters).

For example:

```
# cli user add 3 "test_user" 0x40 4 "password"  
Pigeon Point Shelf Manager Command Line Interpreter  
User 3 added successfully  
#
```

Related Information

- *Switch Reference*, `user add` command
- [“Change a User’s Password” on page 65](#)
- [“Delete a User” on page 66](#)

▼ Change a User’s Password

- On the CMC, type:

```
# cli user passwd userid "password"
```

where:

- *userid* – A valid user identifying number.

- *password* – The user’s password (truncated to 16 characters).

For example:

```
# clia user passwd 3 "root"
Pigeon Point Shelf Manager Command Line Interpreter
  User 3, password changed successfully
#
```

Related Information

- *Switch Reference*, `user passwd` command
- [“Add a User” on page 65](#)
- [“Delete a User” on page 66](#)

▼ Delete a User

- On the CMC, type:

```
# clia user delete userid
```

where *userid* is a valid user identifying number. For example:

```
# clia user delete 3
Pigeon Point Shelf Manager Command Line Interpreter
  User 3 deleted successfully
#
```

Related Information

- *Switch Reference*, `user delete` command
- [“Add a User” on page 65](#)
- [“Change a User’s Password” on page 65](#)

Managing Power Supplies

These tasks enable you to power on and power off a power supply.

- [“Restart a Power Supply” on page 67](#)
- [“Enable and Activate a Power Supply” on page 67](#)
- [“Deactivate and Disable a Power Supply” on page 68](#)

Related Information

- “Checking Power Supplies” on page 13
- “Setting Up the Hardware” on page 58
- “Setting Up the Users” on page 64
- “Managing CMCs” on page 69
- “Managing Fabric Cards” on page 72
- “Managing Line Cards” on page 81

▼ Restart a Power Supply

Restarting a power supply is effectively disabling and re-enabling the power supply.

1. Disable the power supply.

See “Deactivate and Disable a Power Supply” on page 68.

2. Re-enable the power supply.

See “Enable and Activate a Power Supply” on page 67.

Related Information

- “Restart a CMC” on page 69
- “Restart a Fabric Card or Filler” on page 72
- “Restart a Line Card” on page 82

▼ Enable and Activate a Power Supply

- On the CMC, type:

```
# enablepsu slot
```

where *slot* is 0–3. See “Power Supply Addressing” on page 3. For example:

```
# enablepsu 0
Turning on 12V ...
PSU 0, 12 V on
#
```

Note – When a power supply is enabled, the OK LED is illuminated and the power supply goes to full power.

Related Information

- *Switch Reference*, enablepsu command
- “Restart a Power Supply” on page 67
- “Deactivate and Disable a Power Supply” on page 68
- “Activate a CMC” on page 70
- “Switch Over to Another CMC” on page 71
- “Activate a Fabric Card or Filler” on page 75
- “Activate a Line Card” on page 85

▼ Deactivate and Disable a Power Supply

- On the CMC, type:

```
# disablepsu slot
```

where *slot* is 0–3. See “Power Supply Addressing” on page 3.

Note – If you attempt to power off more power supplies than the minimum necessary, you are warned of such an action, and must commit to the action.

For example:

```
# disablepsu 1
PSU 1 can not be turned off. It is needed to power the active boards
Turn off PSU 1 anyway (y/n): y
PSU 1, 12 V is off
#
```

Note – When a power supply is disabled, the OK LED flashes.

Related Information

- *Switch Reference*, disablepsu command
- “Restart a Power Supply” on page 67
- “Enable and Activate a Power Supply” on page 67
- “Deactivate a CMC” on page 70
- “Switch Over to Another CMC” on page 71
- “Deactivate a Fabric Card or Filler” on page 76

- “Deactivate a Line Card” on page 86

Managing CMCs

The following tasks enable you to control CMC activity.

- “Restart a CMC” on page 69
- “Activate a CMC” on page 70
- “Deactivate a CMC” on page 70
- “Switch Over to Another CMC” on page 71

Related Information

- “Checking CMCs” on page 15
- “Setting Up the Hardware” on page 58
- “Setting Up the Users” on page 64
- “Managing Power Supplies” on page 66
- “Managing Fabric Cards” on page 72
- “Managing Line Cards” on page 81

▼ Restart a CMC



Caution – Restarting the CMC that is performing the administration effectively severs the link from the management console to the switch.

- On the CMC, type:

```
# reboot
Connection to 123.45.67.89 closed by foreign host.
%
```

Related Information

- “Activate a CMC” on page 70
- “Deactivate a CMC” on page 70
- “Switch Over to Another CMC” on page 71
- “Restart a Power Supply” on page 67
- “Restart a Fabric Card or Filler” on page 72

- [“Restart a Line Card” on page 82](#)

▼ Activate a CMC

Note – The CMC must have a Host: “Active” status to affect CMC and component administration and configuration changes.

- **On the CMC, type:**

```
# clia activate IPMB_address 0
```

where *IPMB_address* is from [“CMC Addressing” on page 4](#).

For example:

```
# clia activate 10 0  
Pigeon Point Shelf Manager Command Line Interpreter  
Command issued via IPMB, status = 0 (0x0)  
Command executed successfully  
#
```

Related Information

- [Switch Reference, activate command](#)
- [“Restart a CMC” on page 69](#)
- [“Deactivate a CMC” on page 70](#)
- [“Switch Over to Another CMC” on page 71](#)
- [“Enable and Activate a Power Supply” on page 67](#)
- [“Activate a Fabric Card or Filler” on page 75](#)
- [“Activate a Line Card” on page 85](#)

▼ Deactivate a CMC



Caution – Deactivating the CMC that is performing the administration effectively severs the link from the management console to the switch. Perform a switchover operation before deactivating the primary CMC.

- On the CMC, type:

```
# clia deactivate IPMB_address 0
```

where *IPMB_address* is from “CMC Addressing” on page 4.

For example:

```
# clia deactivate 10 0  
Pigeon Point Shelf Manager Command Line Interpreter  
Command issued via IPMB, status = 0 (0x0)  
Command executed successfully  
#
```

Related Information

- *Switch Reference*, deactivate command
- “Restart a CMC” on page 69
- “Activate a CMC” on page 70
- “Switch Over to Another CMC” on page 71
- “Deactivate and Disable a Power Supply” on page 68
- “Deactivate a Fabric Card or Filler” on page 76
- “Deactivate a Line Card” on page 86

▼ Switch Over to Another CMC

- On the CMC, type:

```
# clia switchover  
This Shelf Manager is now active, but is shutting down to trigger a switchover.  
#
```

Related Information

- *Switch Reference*, switchover command
- “Restart a CMC” on page 69
- “Activate a CMC” on page 70
- “Deactivate a CMC” on page 70

Managing Fabric Cards

The following tasks enable you to control a fabric card.

Note – For the examples in this topic, fabric card 0 (IPMB address 94), is used. Additionally, fabric card fillers are installed in slots 2 and 6.

- [“Restart a Fabric Card or Filler” on page 72](#)
- [“Enable a Fabric Card or Filler” on page 73](#)
- [“Disable a Fabric Card or Filler” on page 74](#)
- [“Enable Standby Power for a Fabric Card or Filler” on page 74](#)
- [“Disable Standby Power for a Fabric Card or Filler” on page 75](#)
- [“Enable a Fabric Card Slot for Hot-Insertion” on page 75](#)
- [“Activate a Fabric Card or Filler” on page 75](#)
- [“Deactivate a Fabric Card or Filler” on page 76](#)
- [“Turn On a Fabric Card or Filler Locator LED” on page 77](#)
- [“Turn Off a Fabric Card or Filler Locator LED” on page 78](#)
- [“Enable Downed Fabric Card Links” on page 78](#)
- [“Enable a Fabric Card Switch Chip Port” on page 79](#)
- [“Disable a Fabric Card Switch Chip Port” on page 80](#)
- [“Reset a Fabric Card Switch Chip” on page 80](#)

Related Information

- [“Checking Fabric Cards” on page 18](#)
- [“Setting Up the Hardware” on page 58](#)
- [“Setting Up the Users” on page 64](#)
- [“Managing Power Supplies” on page 66](#)
- [“Managing CMCs” on page 69](#)
- [“Managing Line Cards” on page 81](#)

▼ Restart a Fabric Card or Filler

Restarting a fabric card is effectively deactivating, disabling, re-enabling, and re-activating the fabric card.

1. Deactivate the fabric card.

See [“Deactivate a Fabric Card or Filler” on page 76](#).

2. **Disable standby voltage for the fabric card.**

See “Disable Standby Power for a Fabric Card or Filler” on page 75.

3. **Re-enable standby voltage for the fabric card.**

See “Enable Standby Power for a Fabric Card or Filler” on page 74.

4. **Re-activate the fabric card.**

See “Activate a Fabric Card or Filler” on page 75.

Related Information

- “Enable a Fabric Card or Filler” on page 73
- “Disable a Fabric Card or Filler” on page 74
- “Restart a Power Supply” on page 67
- “Restart a CMC” on page 69
- “Restart a Line Card” on page 82

▼ **Enable a Fabric Card or Filler**

- On the CMC, type:

```
# enableboard fc slot
```

where *slot* is 0–8. See “Fabric Card Addressing” on page 5. For example:

```
# enableboard fc 0
fc 0 is now enabled
#
```

Note – When a fabric card is enabled, the OK LED on that fabric card lights.

Related Information

- *Switch Reference*, enableboard command
- “Disable a Fabric Card or Filler” on page 74
- “Enable a Line Card” on page 83

▼ Disable a Fabric Card or Filler

- On the CMC, type:

```
# disableboard fc slot
```

where *slot* is from is 0–8. See [“Fabric Card Addressing” on page 5](#).

Note – If you disable a fabric card that is active you will receive a message warning you of such action. You must commit to the action.

For example:

```
# disableboard fc 0
fc 0 is active. Do you want to continue disabling this board (y/n)? y
fc 0 is M1 state, turning of stby
STBY for fc 0 is off
#
```

Note – When a fabric card is disabled, the OK LED on that fabric card flashes.

Related Information

- *Switch Reference*, disableboard command
- [“Enable a Fabric Card or Filler” on page 73](#)
- [“Disable a Line Card” on page 83](#)

▼ Enable Standby Power for a Fabric Card or Filler

- On the CMC, type:

```
# enablestby fc slot
```

where *slot* is from is 0–8. See [“Fabric Card Addressing” on page 5](#). For example:

```
# enablestby fc 0
STBY for fc 0 OK
#
```

Related Information

- *Switch Reference*, enablestby command

- “Disable Standby Power for a Fabric Card or Filler” on page 75
- “Enable Standby Power for a Line Card” on page 84

▼ Disable Standby Power for a Fabric Card or Filler

- On the CMC, type:

```
# disablestby fc slot
```

where *slot* is from is 0–8. See “Fabric Card Addressing” on page 5. For example:

```
# disablestby fc 0
. . . . . STBY for fc 0 is off
#
```

Related Information

- *Switch Reference*, disablestby command
- “Enable Standby Power for a Fabric Card or Filler” on page 74
- “Disable Standby Power for a Line Card” on page 84

▼ Enable a Fabric Card Slot for Hot-Insertion

- On the CMC, type:

```
# enablehotinsert fc slot
```

where *slot* is 0–8. See “Fabric Card Addressing” on page 5. For example:

```
# enablehotinsert fc 0
fc 0 is now enabled
#
```

Related Information

- *Switch Reference*, enablehotinsert command
- “Enable a Line Card Slot for Hot-Insertion” on page 85

▼ Activate a Fabric Card or Filler

- To bring a fabric card to full power, take one of the following actions:

- Type on the CMC:

```
# activate fc slot
```

where *slot* is 0–8. See “Fabric Card Addressing” on page 5. For example:

```
# activate fc 0
#
```

- Type on the CMC:

```
# clia activate IPMB_address 0
```

where *IPMB_address* is from “Fabric Card Addressing” on page 5. For example:

```
# clia activate 94 0
Pigeon Point Shelf Manager Command Line Interpreter
Command issued via IPMB, status = 0 (0x0)
Command executed successfully
#
```

When a fabric card is activated, the OK LED is illuminated, the fans spin up, and the fabric card goes to full power.

Note – Activating a fabric card can take up to 1 minute for both of the switch chips to boot. During the switch chip boot process, the fabric card might be identified by the `showpresent` command as being in an active (M4) state.

Related Information

- *Switch Reference*, activate command
- “Deactivate a Fabric Card or Filler” on page 76
- “Activate a Line Card” on page 85

▼ Deactivate a Fabric Card or Filler

- To power down a fabric card to a standby state, take one of the following actions:
 - Type on the CMC:

```
# deactivate fc slot
```

where *slot* is 0–8. See “Fabric Card Addressing” on page 5. For example:

```
# deactivate fc 0
Deactivating FC 0
Pigeon Point Shelf Manager Command Line Interpreter
  Command issued via IPMB, status = 0 (0x0)
  Command executed successfully
#
```

- Type on the CMC:

```
# clia deactivate IPMB_address 0
```

where *IPMB_address* is from “Fabric Card Addressing” on page 5. For example:

```
# clia deactivate 94 0
Pigeon Point Shelf Manager Command Line Interpreter
  Command issued via IPMB, status = 0 (0x0)
  Command executed successfully
#
```

Note – When a fabric card is deactivated, the OK LED flashes and the fans spin down.

Related Information

- *Switch Reference*, deactivate command
- “Activate a Fabric Card or Filler” on page 75
- “Deactivate a Line Card” on page 86

▼ Turn On a Fabric Card or Filler Locator LED

- On the CMC, type:

```
# clia setfruledstate IPMB_address 0 3 ON
```

where *IPMB_address* is from “Fabric Card Addressing” on page 5. For example:

```
# clia setfruledstate 94 0 3 ON
Pigeon Point Shelf Manager Command Line Interpreter
  Setting FRU's led state completed successfully, status = 0x0
#
```

Related Information

- *Switch Reference*, `setfruledstate` command
- “Turn Off a Fabric Card or Filler Locator LED” on page 78
- “Turn On a Line Card Locator LED” on page 87

▼ Turn Off a Fabric Card or Filler Locator LED

- On the CMC, type:

```
# clia setfruledstate IPMB_address 0 3 OFF
```

where *IPMB_address* is from “Fabric Card Addressing” on page 5. For example:

```
# clia setfruledstate 94 0 3 OFF  
Pigeon Point Shelf Manager Command Line Interpreter  
Setting FRU's led state completed successfully, status = 0x0  
#
```

Related Information

- *Switch Reference*, `setfruledstate` command
- “Turn On a Fabric Card or Filler Locator LED” on page 77
- “Turn Off a Line Card Locator LED” on page 88

▼ Enable Downed Fabric Card Links

If the output of the `checklinks` command indicates that some links are down, use the `-e` option of the `checklinks` command in attempt to enable those downed links.

1. On the CMC, type:

```
# checklinks -e
```

2. Use the `checklinks` command again, to verify which downed links have been enabled:

```
# checklinks |grep FC  
FC 0 Active, checking links....OK  
FC 1 Active, checking links....OK  
FC 3 Active, checking links....OK  
FC 4 Active, checking links....OK  
FC 5 Active, checking links....OK
```

```
FC 7 Active, checking links.....OK
FC 8 Active, checking links.....OK
#
```

Note – Fabric card fillers have no switch chips. Therefore, no reporting is provided from them.

Related Information

- *Switch Reference*, `checklinks` command
- [“Check Fabric Card Link Status \(Simple\)”](#) on page 26
- [“Check Fabric Card Switch Chip Health”](#) on page 27
- [“Enable Downed Line Card Links”](#) on page 88

▼ Enable a Fabric Card Switch Chip Port

- On the CMC, type:

```
# enableswitchport fc slot switch_chip port
```

where:

- *slot* is number of the fabric card (0–8). See [“Fabric Card Addressing”](#) on page 5.
- *switch_chip* is the number of the chip (0–1).
- *port* is the number of the port (1–36).

For example:

```
# enableswitchport fc 1 0 1
Enabling port 1 on switch 0 on FC 1
#
```

Note – Fabric card fillers have no switch chips. Therefore, this task is not possible with fillers.

Related Information

- *Switch Reference*, `enableswitchport` command
- [“Disable a Fabric Card Switch Chip Port”](#) on page 80
- [“Check Fabric Card Link Status \(Detailed\)”](#) on page 26
- [“Enable a Line Card Switch Chip Port”](#) on page 89

▼ Disable a Fabric Card Switch Chip Port

- On the CMC, type:

```
# disableswitchport fc slot switch_chip port
```

where:

- *slot* is number of the fabric card (0–8). See [“Fabric Card Addressing”](#) on page 5.
- *switch_chip* is the number of the chip (0–1).
- *port* is the number of the port (1–36).

For example:

```
# disableswitchport fc 1 0 1
Disabling port 1 on switch 0 on FC 1
#
```

Note – Fabric card fillers have no switch chips. Therefore, this task is not possible with fillers.

Related Information

- *Switch Reference*, `disableswitchport` command
- [“Enable a Fabric Card Switch Chip Port”](#) on page 79
- [“Check Fabric Card Link Status \(Detailed\)”](#) on page 26
- [“Disable a Line Card Switch Chip Port”](#) on page 89

▼ Reset a Fabric Card Switch Chip

- On the CMC, type:

```
# resetswitch fc slot switch_chip state
```

where:

- *slot* is number of the fabric card (0–8). See [“Fabric Card Addressing”](#) on page 5.
- *switch_chip* is the number of the chip (0–1).

- *state* is 0 to reset once, and 1 to hold in reset.

For example:

```
# resetswitch fc 1 0 0
#
```

Note – Fabric card fillers have no switch chips. Therefore, this task is not possible with fillers.

Related Information

- *Switch Reference*, `resetswitch` command
- [“Check Fabric Card Switch Chip Health” on page 27](#)
- [“Reset a Line Card Switch Chip” on page 90](#)

Managing Line Cards

The following tasks enable you to control the line cards.

Note – For the examples in this topic, line card 8 (IPMB address 92), is used.

- [“Restart a Line Card” on page 82](#)
- [“Enable a Line Card” on page 83](#)
- [“Disable a Line Card” on page 83](#)
- [“Enable Standby Power for a Line Card” on page 84](#)
- [“Disable Standby Power for a Line Card” on page 84](#)
- [“Enable a Line Card Slot for Hot-Insertion” on page 85](#)
- [“Activate a Line Card” on page 85](#)
- [“Deactivate a Line Card” on page 86](#)
- [“Turn On a Line Card Locator LED” on page 87](#)
- [“Turn Off a Line Card Locator LED” on page 88](#)
- [“Enable Downed Line Card Links” on page 88](#)
- [“Enable a Line Card Switch Chip Port” on page 89](#)
- [“Disable a Line Card Switch Chip Port” on page 89](#)
- [“Reset a Line Card Switch Chip” on page 90](#)

Related Information

- [“Checking Line Cards” on page 30](#)
- [“Setting Up the Hardware” on page 58](#)
- [“Setting Up the Users” on page 64](#)
- [“Managing Power Supplies” on page 66](#)
- [“Managing CMCs” on page 69](#)
- [“Managing Fabric Cards” on page 72](#)

▼ **Restart a Line Card**

Restarting a line card is effectively deactivating, disabling, re-enabling, and re-activating the fabric card.

1. Deactivate the line card.

See [“Deactivate a Line Card” on page 86](#).

2. Disable standby voltage for the line card.

See [“Disable Standby Power for a Line Card” on page 84](#).

3. Re-enable standby voltage for the line card.

See [“Enable Standby Power for a Line Card” on page 84](#).

4. Re-activate the line card.

See [“Activate a Line Card” on page 85](#).

Related Information

- [“Enable a Line Card” on page 83](#)
- [“Disable a Line Card” on page 83](#)
- [“Restart a Power Supply” on page 67](#)
- [“Restart a CMC” on page 69](#)
- [“Restart a Fabric Card or Filler” on page 72](#)

▼ Enable a Line Card

- On the CMC, type:

```
# enableboard lc slot
```

where *slot* is 0–8. See [“Line Card Addressing”](#) on page 6. For example:

```
# enableboard lc 8
lc 8 is now enabled
#
```

Note – When a line card is enabled, the OK LED on that line card lights.

Related Information

- [Switch Reference](#), enableboard command
- [“Disable a Line Card”](#) on page 83
- [“Enable a Fabric Card or Filler”](#) on page 73

▼ Disable a Line Card

- On the CMC, type:

```
# disableboard lc slot
```

where *slot* is 0–8. See [“Line Card Addressing”](#) on page 6. For example:

```
# disableboard lc 8
lc 8 is M1 state, turning of stby
STBY for lc 8 is off
#
```

Note – When a line card is disabled, the OK LED on that line card flashes.

Related Information

- [Switch Reference](#), disableboard command
- [“Enable a Line Card”](#) on page 83
- [“Disable a Fabric Card or Filler”](#) on page 74

▼ Enable Standby Power for a Line Card

- On the CMC, type:

```
# enablestby lc slot
```

where *slot* is 0–8. See “Line Card Addressing” on page 6. For example:

```
# enablestby lc 8
STBY for lc 8 OK
#
```

Related Information

- *Switch Reference*, enablestby command
- “Disable Standby Power for a Line Card” on page 84
- “Enable Standby Power for a Fabric Card or Filler” on page 74

▼ Disable Standby Power for a Line Card

- On the CMC, type:

```
# disablestby lc slot
```

where *slot* is 0–8. See “Line Card Addressing” on page 6. For example:

```
# disablestby lc 8
. . . . . STBY for lc 8 is off
#
```

Related Information

- *Switch Reference*, disablestby command
- “Enable Standby Power for a Line Card” on page 84
- “Disable Standby Power for a Fabric Card or Filler” on page 75

▼ Enable a Line Card Slot for Hot-Insertion

- On the CMC, type:

```
# enablehotinsert lc slot
```

where *slot* is 0–8. See “Line Card Addressing” on page 6. For example:

```
# enablehotinsert lc 8
lc 8 is now enabled
#
```

Related Information

- *Switch Reference*, enablehotinsert command
- “Enable a Fabric Card Slot for Hot-Insertion” on page 75

▼ Activate a Line Card

- To bring a line card to full power, take one of the following actions:
 - Type on the CMC:

```
# activate lc slot
```

where *slot* is 0–8. See “Line Card Addressing” on page 6. For example:

```
# activate lc 8
#
```

- Type on the CMC:

```
# clia activate IPMB_address 0
```

where *IPMB_address* is from “Line Card Addressing” on page 6. For example:

```
# clia activate 92 0
Pigeon Point Shelf Manager Command Line Interpreter
Command issued via IPMB, status = 0 (0x0)
Command executed successfully
#
```

When a line card is activated, the OK LED is illuminated and the line card goes to full power.

Note – Activating a line card can take up to 2 minutes for all of the switch chips to boot. During the switch chip boot process, the line card might be identified by the `showpresent` command as being in an active (M4) state.

Related Information

- *Switch Reference*, activate command
- “Deactivate a Line Card” on page 86
- “Activate a Fabric Card or Filler” on page 75

▼ Deactivate a Line Card

- To power down a line card to a standby state, take one of the following actions:
 - Type on the CMC:

```
# deactivate lc slot
```

where *slot* is 0–8. See “Line Card Addressing” on page 6. For example:

```
# deactivate lc 8
Deactivating LC 8
Pigeon Point Shelf Manager Command Line Interpreter
Command issued via IPMB, status = 0 (0x0)
Command executed successfully
#
```

- Type on the CMC:

```
# clia deactivate IPMB_address 0
```

where *IPMB_address* is from “Line Card Addressing” on page 6. For example:

```
# clia deactivate 92 0
Pigeon Point Shelf Manager Command Line Interpreter
  Command issued via IPMB, status = 0 (0x0)
  Command executed successfully
#
```

Note – When a line card is deactivated, the OK LED flashes.

Related Information

- *Switch Reference*, deactivate command
- “Activate a Line Card” on page 85
- “Deactivate a Fabric Card or Filler” on page 76

▼ Turn On a Line Card Locator LED

- On the CMC, type:

```
# clia setfruledstate IPMB_address 0 3 ON
```

where *IPMB_address* is from “Line Card Addressing” on page 6. For example:

```
# clia setfruledstate 92 0 3 ON
Pigeon Point Shelf Manager Command Line Interpreter
  Setting FRU's led state completed successfully, status = 0x0
#
```

Related Information

- *Switch Reference*, setfruledstate command
- “Turn Off a Line Card Locator LED” on page 88
- “Turn On a Fabric Card or Filler Locator LED” on page 77

▼ Turn Off a Line Card Locator LED

- On the CMC, type:

```
# clia setfruledstate IPMB_address 0 3 ON
```

where *IPMB_address* is from “Line Card Addressing” on page 6. For example:

```
# clia setfruledstate 92 0 3 OFF
Pigeon Point Shelf Manager Command Line Interpreter
Setting FRU's led state completed successfully, status = 0x0
#
```

Related Information

- *Switch Reference*, `setfruledstate` command
- “Turn On a Line Card Locator LED” on page 87
- “Turn Off a Fabric Card or Filler Locator LED” on page 78

▼ Enable Downed Line Card Links

If the output of the `checklinks` command indicates that some links are down, use the `-e` option of the `checklinks` command in attempt to enable those downed links.

1. On the CMC, type:

```
# checklinks -e -c
```

2. Use the `checklinks` command again, to verify which downed links have been enabled:

```
# checklinks -c |grep LC
LC 0 Active, checking links....OK
LC 1 Active, checking links....OK
LC 2 Active, checking links....OK
LC 3 Active, checking links....OK
LC 4 Active, checking links....OK
LC 5 Active, checking links....OK
LC 6 Active, checking links....OK
LC 7 Active, checking links....OK
LC 8 Active, checking links....OK
#
```

Related Information

- *Switch Reference*, `checklinks` command
- [“Check Line Card Link Status \(Simple\)”](#) on page 36
- [“Check Line Card Switch Chip Health”](#) on page 38
- [“Enable Downed Fabric Card Links”](#) on page 78

▼ Enable a Line Card Switch Chip Port

- On the CMC, type:

```
# enableswitchport lc slot switch_chip port
```

where:

- *slot* is number of the line card (0–8). See [“Line Card Addressing”](#) on page 6.
- *switch_chip* is the number of the chip (0–3).
- *port* is the number of the port (1–36).

For example:

```
# enableswitchport lc 8 0 1
Enabling port 1 on switch 0 on LC 8
#
```

Related Information

- *Switch Reference*, `enableswitchport` command
- [“Disable a Line Card Switch Chip Port”](#) on page 89
- [“Check Line Card Link Status \(Detailed\)”](#) on page 37
- [“Enable a Fabric Card Switch Chip Port”](#) on page 79

▼ Disable a Line Card Switch Chip Port

- On the CMC, type:

```
# disableswitchport lc slot switch_chip port
```

where:

- *slot* is number of the line card (0–8). See [“Line Card Addressing”](#) on page 6.
- *switch_chip* is the number of the chip (0–3).

- *port* is the number of the port (1–36).

For example:

```
# disableswitchport lc 8 0 1
Disabling port 1 on switch 0 on LC 8
#
```

Related Information

- *Switch Reference*, `disableswitchport` command
- “Enable a Line Card Switch Chip Port” on page 89
- “Check Line Card Link Status (Detailed)” on page 37
- “Disable a Fabric Card Switch Chip Port” on page 80

▼ Reset a Line Card Switch Chip

- On the CMC, type:

```
# resetswitch lc slot switch_chip state
```

where:

- *slot* is number of the line card (0–8). See “Line Card Addressing” on page 6.
- *switch_chip* is the number of the chip (0–3).
- *state* is 0 to reset once, and 1 to hold in reset.

For example:

```
# resetswitch lc 8 0 0
#
```

Related Information

- *Switch Reference*, `resetswitch` command
- “Check Line Card Switch Chip Health” on page 38
- “Reset a Fabric Card Switch Chip” on page 80

Controlling the InfiniBand Fabric

The following tasks describe how you can manage the InfiniBand fabric and its components.

Note – Commands in this topic are issued from the Linux InfiniBand host as the `root` user. You cannot issue the commands on the CMC.

- “Perform Comprehensive Diagnostics for the Entire Fabric” on page 91
- “Perform Comprehensive Diagnostics for a Route” on page 92
- “Determine Changes to the InfiniBand Fabric Topology” on page 92
- “Find 1x or SDR or DDR Links in the Fabric” on page 94
- “Determine Which Links Are Experiencing Significant Errors” on page 95
- “Clear Error Counters” on page 96
- “Clear Data Counters” on page 96
- “Check All Ports” on page 96
- “Reset a Port” on page 97
- “Set the Speed of a Port” on page 98
- “Disable a Port” on page 100
- “Enable a Port” on page 101
- “Controlling the Subnet Manager” on page 102

Related Information

- “Understanding Administrative Commands” on page 1
- “Monitoring the Switch” on page 10
- “Monitoring the InfiniBand Fabric” on page 43
- “Controlling the Switch” on page 57

▼ Perform Comprehensive Diagnostics for the Entire Fabric

Should you require a full testing of your InfiniBand fabric, the `ibdiagnet` command can perform many tests with verbose results. The command is a useful tool to determine the general overall health of the InfiniBand fabric.

- On the Linux InfiniBand host, type:

```
# ibdiagnet -v -r
```

The diagnostic are displayed. Additionally, the `ibdiagnet.log` file contains the log of the testing.

Related Information

- *Switch Reference*, `ibdiagnet` command
- [“Display the InfiniBand Fabric Topology” on page 47](#)
- [“Perform Comprehensive Diagnostics for a Route” on page 92](#)

▼ Perform Comprehensive Diagnostics for a Route

Similar to the `ibdiagnet` command, the `ibdiagpath` command can perform some of the same tests for a particular route.

- On the Linux InfiniBand host, type:

```
# ibdiagpath -v -l slid, dlid
```

where:

- *slid* is the LID of the source node.
- *dlid* is the LID of the destination node.

The diagnostics are displayed. Additionally, the `ibdiagpath.log` file contains the log of the testing.

Related Information

- *Switch Reference*, `ibdiagpath` command
- [“Display a Route Through the Fabric” on page 47](#)
- [“Perform Comprehensive Diagnostics for the Entire Fabric” on page 91](#)

▼ Determine Changes to the InfiniBand Fabric Topology

If your fabric has a number of nodes that are suspect, the `osmtest` command enables you to take a “snapshot” (inventory file) of your fabric and at a later time compare that file to the present conditions.

Note – Though this procedure is most useful after initializing the Subnet Manager, it can be performed at anytime.

1. **Initiate the Subnet Manager.**
2. **Take a “snapshot” of the fabric topology.**

```
# osmtest -f c
Command Line Arguments
Done with args
  Flow = Create Inventory
May 21 15:51:37 220542 [4CC45180] 0x7f -> Setting log level to: 0x03
May 21 15:51:37 220928 [4CC45180] 0x02 -> osm_vendor_init: 1000 pending umads
specified
May 21 15:51:37 248149 [4CC45180] 0x02 -> osm_vendor_bind: Binding to port
0x3ba000100c70a
May 21 15:51:37 281955 [4CC45180] 0x02 -> osmtest_validate_sa_class_port_info:
-----
SA Class Port Info:
  base_ver:1
  class_ver:2
  cap_mask:0x2602
  cap_mask2:0x0
  resp_time_val:0x10
-----
OSMTEST: TEST "Create Inventory" PASS
#
```

3. **After an event, compare the present topology to that saved in the inventory file:**

```
# osmtest -f v
Command Line Arguments
Done with args
  Flow = Validate Inventory
May 21 15:58:42 797761 [235DE180] 0x7f -> Setting log level to: 0x03
May 21 15:58:42 798204 [235DE180] 0x02 -> osm_vendor_init: 1000 pending umads
specified
May 21 15:58:42 825740 [235DE180] 0x02 -> osm_vendor_bind: Binding to port
0x3ba000100c70a
May 21 15:58:42 860037 [235DE180] 0x02 -> osmtest_validate_sa_class_port_info:
-----
SA Class Port Info:
  base_ver:1
  class_ver:2
  cap_mask:0x2602
  cap_mask2:0x0
```

```
resp_time_val:0x10
-----
May 21 15:58:42 936899 [235DE180] 0x01 -> osmtest_validate_node_data: Checking
node 0x00144f6c67be0000, LID 0x100
May 21 15:58:42 936956 [235DE180] 0x01 -> osmtest_validate_node_data: Checking
node 0x0144f6c666b50100, LID 0x101
.
.
.
May 21 15:55:45 306391 [22B6D180] 0x01 ->
osmtest_validate_single_path_rec_guid_pair:
  Checking src 0x508002000082dc35 to dest 0x508002000082dc35
May 21 15:55:45 307222 [22B6D180] 0x01 -> osmtest_validate_path_data: Checking
path SLID 0xFE to DLID 0xFE
May 21 15:55:45 308458 [22B6D180] 0x02 -> osmtest_run:
***** ALL TESTS PASS *****
OSMTEST: TEST "Validate Inventory" PASS
#
```

Note – Depending on the size of your InfiniBand fabric, the output from the `osmtest` command could be tens of thousands of lines long.

Related Information

- *Switch Reference*, `osmtest` command
- [“Display the InfiniBand Fabric Topology” on page 47](#)

▼ Find 1x or SDR or DDR Links in the Fabric

You can use the `ibdiagnet` command to determine which links are at 1x bandwidth, 2.5 Gbps, or 5 Gbps data rate.

- **On the Linux InfiniBand host, type:**

```
# ibdiagnet -lw 4x -ls 10 -pc -pm -skip all
```

In this instance of the `ibdiagnet` command, there is a check for all links to be at 4x QDR (`-lw 4x -ls 10`), and if not, to report those which are not. Additionally, the `ibdiagnet.log` file contains the log of the testing.

Related Information

- *Switch Reference*, `ibdiagnet` command
- [“Display the Link Status of a Node” on page 49](#)

- [“Determine Which Links Are Experiencing Significant Errors” on page 95](#)

▼ Determine Which Links Are Experiencing Significant Errors

You can use the `ibdiagnet` command to determine which links are experiencing symbol errors and recovery errors by injecting packets.

1. **On the Linux InfiniBand host, type.**

```
# ibdiagnet -c 100 -P all=1
```

In this instance of the `ibdiagnet` command, 100 test packets are injected into each link and the `-P all=1` option returns all counters that increment during the test.

2. **In the output of the `ibdiagnet` command, search for the `symbol_error_counter` string.**

That line contains the symbol error count in hexadecimal. The preceding lines identify the node and port with the errors. Symbol errors are minor errors, and if there are relatively few during the diagnostic, they can be monitored.

Note – According to the InfiniBand specification 10E-12 BER, the maximum allowable symbol error rate is 120 errors per hour.

3. **Also in the output of the `ibdiagnet` command, search for the `link_error_recovery_counter` string.**

That line contains the recovery error count in hexadecimal. The preceding lines identify the node and port with the errors. Recovery errors are major errors and the respective links must be investigated for the cause of the rapid symbol error propagation.

Note – Additionally, the `ibdiagnet.log` file contains the log of the testing.

Related Information

- [Switch Reference, `ibdiagnet` command](#)
- [“Find 1x or SDR or DDR Links in the Fabric” on page 94](#)
- [“Display Counters for a Node” on page 50](#)

▼ Clear Error Counters

If you are troubleshooting a port, the `perfquery` command will provide counters of errors occurring at that port. To determine if the problem has been resolved, you can reset all of the error counters to 0 with the `ibclearerrors` command.

- On the Linux InfiniBand host, type:

```
# ibclearerrors
```

Related Information

- *Switch Reference*, `ibclearerrors` command
- [“Clear Data Counters” on page 96](#)
- [“Display Counters for a Node” on page 50](#)

▼ Clear Data Counters

When you are optimizing the InfiniBand fabric for performance, you might want to know how the throughput increases or decreases according to changes you are making to the fabric and Subnet Manager. The `ibclearcounters` command enables you to reset the data counters for all ports to 0.

- On the Linux InfiniBand host, type:

```
# ibclearcounters
```

Related Information

- *Switch Reference*, `ibclearcounters` command
- [“Clear Error Counters” on page 96](#)
- [“Display Data Counters for a Node” on page 51](#)

▼ Check All Ports

To perform a quick check of all ports of all nodes in your InfiniBand fabric, you can use the `ibcheckstate` command.

- On the Linux InfiniBand host, type:

```
# ibcheckstate -v
# Checking Switch: nodeguid 0x00066a00d80001dd
Node check lid 4: OK
Port check lid 4 port 24: OK
Port check lid 4 port 19: OK
.
.
.
# Checking Ca: nodeguid 0x0002c90200001818
Node check lid 5: OK
Port check lid 5 port 1: OK

## Summary: 5 nodes checked, 0 bad nodes found
##           10 ports checked, 0 ports with bad state found
#
```

Note – The `ibcheckstate` command requires time to complete, depending upon the size of your InfiniBand fabric. Without the `-v` option, the output contains only failed ports. The output in the example is only a small portion of the actual output.

Related Information

- *Switch Reference*, `ibcheckstate` command
- [“Check Fabric Card Link Status \(Simple\)” on page 26](#)
- [“Check Line Card Link Status \(Simple\)” on page 36](#)

▼ Reset a Port

You might need to reset a port to determine its functionality.

- On the Linux InfiniBand host, type:

```
# ibportstate lid port reset
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example:

```
# ibportstate 57 2 reset
Initial PortInfo:
# Port info: Lid 57 port 2
LinkState:.....Active
PhysLinkState:.....LinkUp
LinkWidthSupported:.....1X or 4X
LinkWidthEnabled:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.....10.0 Gbps

After PortInfo set:
# Port info: Lid 57 port 2
LinkState:.....Down
PhysLinkState:.....Disabled

After PortInfo set:
# Port info: Lid 57 port 2
LinkState:.....Down
PhysLinkState:.....Polling
#
```

Note – You can use the `-G` option to specify a GUID and port instead of a LID and port.

Related Information

- *Switch Reference*, `ibportstate` command
- [“Disable a Port” on page 100](#)
- [“Enable a Port” on page 101](#)
- [“Reset a Fabric Card Switch Chip” on page 80](#)
- [“Reset a Line Card Switch Chip” on page 90](#)

▼ Set the Speed of a Port

You can manually set the speed of a single port to help determine symbol error generation. The `ibportstate` command can set the speed to 2.5, 5.0, or 10.0 Gbyte/sec.

- On the Linux InfiniBand host, type:

```
# ibportstate lid port speed speed
```

where:

- *lid* is the LID of the node.
- *port* is the port of the node.
- *speed* is the speed of the port: 1 for 2.5 Gbyte/sec, 2 for 5.0 Gbyte/sec, and 4 for 10.0 Gbyte/sec.

Note – Adding speed values enable either speed. For example, speed 7 is 2.5, 5.0, and 10.0 Gbyte/sec.

For example:

```
# ibportstate 57 2 speed 1
Initial PortInfo:
# Port info: Lid 57 port 2
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps

After PortInfo set:
# Port info: Lid 57 port 2
LinkSpeedEnabled:.....2.5 Gbps
# ibportstate 57 2 speed 4
Initial PortInfo:
# Port info: Lid 57 port 2
LinkSpeedEnabled:.....2.5 Gbps
After PortInfo set:
# Port info: Lid 57 port 2
LinkSpeedEnabled:.....10.0 Gbps (IBA extension)
#
```

Note – You can use the -G option to specify a GUID, port, and speed instead of a LID, port, and speed.

Related Information

- *Switch Reference*, `ibportstate` command
- [“Display the Link Status of a Node” on page 49](#)

▼ Disable a Port

If a port is found to be problematic, you can disable it.

- **On the Linux InfiniBand host, type:**

```
# ibportstate lid port disable
```

where:

- *lid* is the LID of the node.
- *port* is the port of the node.

For example:

```
# ibportstate 57 2 disable
Initial PortInfo:
# Port info: Lid 57 port 2
LinkState:.....Active
PhysLinkState:.....LinkUp
LinkWidthSupported:.....1X or 4X
LinkWidthEnabled:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.....10.0 Gbps

After PortInfo set:
# Port info: Lid 57 port 2
LinkState:.....Down
PhysLinkState:.....Disabled
#
```

Note – You can use the `-G` option to specify a GUID and port instead of a LID and port.

Related Information

- *Switch Reference*, `ibportstate` command
- [“Enable a Port” on page 101](#)
- [“Disable a Fabric Card Switch Chip Port” on page 80](#)
- [“Disable a Line Card Switch Chip Port” on page 89](#)

▼ Enable a Port

After disabling a port, you can enable the port with the `ibportstate` command.

- On the Linux InfiniBand host, type:

```
# ibportstate lid port enable
```

where:

- *lid* is the LID of the node.
- *port* is the port of the node.

For example:

```
# ibportstate 57 2 enable
Initial PortInfo:
# Port info: Lid 57 port 2
LinkState:.....Down
PhysLinkState:.....Disabled
LinkWidthSupported:.....1X or 4X
LinkWidthEnabled:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:.....2.5 Gbps

After PortInfo set:
# Port info: Lid 57 port 2
LinkState:.....Down
PhysLinkState:.....PortConfigurationTraining
#
```

Note – You can use the `-G` option to specify a GUID and port instead of a LID and port.

Related Information

- *Switch Reference*, `ibportstate` command
- [“Disable a Port” on page 100](#)
- [“Enable a Fabric Card Switch Chip Port” on page 79](#)
- [“Enable a Line Card Switch Chip Port” on page 89](#)

Controlling the Subnet Manager

The OpenSM Subnet Manager is started with the `opensm` command or by the `opensmd` daemon and the `/etc/opensm/opensm.conf` file. The advantage of the daemon and file combination is that starting, stopping, and configuring the subnet manager is easier.

The following topics describe how to control the Subnet Manager:

- “Start the Subnet Manager With Min Hop Routing” on page 102
- “Start the Subnet Manager With Fat Tree Routing” on page 103
- “Create the `guid.txt` File” on page 104
- “Start the Subnet Manager With the `opensmd` Daemon” on page 106
- “Stop the Subnet Manager With the `opensmd` Daemon” on page 106
- “Create the `opensm.conf` File” on page 106

Related Information

- “Controlling the Switch” on page 57

▼ Start the Subnet Manager With Min Hop Routing

The simplest way to start the OpenSM Subnet Manager is with the Min Hop routing algorithm. Min Hop is not the most optimal algorithm.

- On the Linux InfiniBand host, type:

```
# opensm -R minhop -A -s 120 -B
-----
OpenSM 3.3.2_20090831_ce6f8dd
Command Line Arguments:
  Activate 'minhop' routing engine(s)
  Unicast routing cache option is on
  sweep interval = 120
  Log File: /var/log/opensm.log
-----
OpenSM 3.3.2_20090831_ce6f8dd
Entering DISCOVERING state
Using default GUID 0x3ba000100c708
Entering MASTER state
SUBNET UP
#
```

Related Information

- *Switch Reference*, opensm command
- “Start the Subnet Manager With Fat Tree Routing” on page 103
- “Start the Subnet Manager With the opensmd Daemon” on page 106
- “Display Subnet Manager Status” on page 56
- “Display Information About the Local HCA” on page 44

▼ Start the Subnet Manager With Fat Tree Routing

The Fat-Tree routing algorithm is the most optimal for the Sun Datacenter InfiniBand Switch 648. The Fat-Tree algorithm requires a root node GUID file (`guid.txt`).

- On the Linux InfiniBand host, type:

```
# opensm -R ftree -A -a /etc/opensm/guid.txt -s 120 -B
-----
OpenSM 3.3.2_20090831_ce6f8dd
Command Line Arguments:
  Activate 'ftree' routing engine(s)
  Unicast routing cache option is on
  Root Guid File: /etc/opensm/guid.txt
  sweep interval = 120
  Log File: /var/log/opensm.log
-----
OpenSM 3.3.2_20090831_ce6f8dd
Entering DISCOVERING state
Using default GUID 0x3ba000100c708
Entering MASTER state
SUBNET UP
#
```

Related Information

- *Switch Reference*, opensm command
- “Start the Subnet Manager With Min Hop Routing” on page 102
- “Start the Subnet Manager With the opensmd Daemon” on page 106
- “Create the `guid.txt` File” on page 104
- “Display Subnet Manager Status” on page 56
- “Display Information About the Local HCA” on page 44

▼ Create the `guid.txt` File

If your InfiniBand fabric uses the Fat-Tree routing algorithm, you must provide a modified listing of fabric card GUIDs, called the root node GUID file, for the Subnet Manager. The `/etc/opensm/guid.txt` file is such a file. This procedure describes how to create the `guid.txt` file.

Note – In this procedure, you issue commands to both the CMC *and* the Linux InfiniBand host.

1. Type the following command on the CMC to retrieve the fabric card base GUID:

```
# getbaseguid fc slot
```

where *slot* is the number of the fabric card (0–8). For example:

```
# getbaseguid fc 0
Base GUID = 0x00bad0fc0009f000
#
```

2. Record the GUID to a text editor on your Linux InfiniBand host.

3. Repeat Step 1 and Step 2 for each fabric card.

Once you have recorded each fabric card's GUID, you will have a list similar to the following on your Linux InfiniBand host:

```
0x00bad0fc0009f000
0x00bad0fc0033f100
0x00bad00017fcf200
0x00bad00005fcf300
0x00bad00003fcf400
0x00bad00007fcf500
0x00bad0fc0011f600
0x00bad0b00002f700
0x00bad0b00001f800
```

Note – If you have fewer than 9 fabric cards, your list will be smaller.

4. For each GUID, add the following values to make two new sums:

- 0xa2
- 0xb2

Note – For GUIDs that end with two zeros (00), you are effectively replacing the zeros with a2 and b2.

For example, 0x00bad0fc0009f000 becomes:

- 0x00bad0fc0009f0a2
- 0x00bad0fc0009f0b2

5. Repopulate your list with the modified GUIDs.

For example:

```
0x00bad0fc0009f0a2
0x00bad0fc0009f0b2
0x00bad0fc0033f1a2
0x00bad0fc0033f1b2
0x00bad00017fcf2a2
0x00bad00017fcf2b2
0x00bad00005fcf3a2
0x00bad00005fcf3b2
0x00bad00003fcf4a2
0x00bad00003fcf4b2
0x00bad00007fcf5a2
0x00bad00007fcf5b2
0x00bad0fc0011f6a2
0x00bad0fc0011f6b2
0x00bad0b00002f7a2
0x00bad0b00002f7b2
0x00bad0b00001f8a2
0x00bad0b00001f8b2
```

6. Save the list on the Linux InfiniBand host as the /etc/opensm/guid.txt file.

7. Set the permissions for the /etc/opensm/guid.txt file:

```
# chmod 0644 /etc/opensm/guid.txt
#
```

Related Information

- *Switch Command Reference*, getbaseguid command
- *Switch Command Reference*, opensm command
- [“Create the opensm.conf File” on page 106](#)
- [“Start the Subnet Manager With Fat Tree Routing” on page 103](#)
- [“Start the Subnet Manager With the opensmd Daemon” on page 106](#)

▼ Start the Subnet Manager With the opensmd Daemon

The opensmd daemon enables you to start the OpenSM Subnet Manager without providing command line configuration parameters.

- On the Linux InfiniBand host, type:

```
# /etc/init.d/opensmd start
Starting IB Subnet Manager.           [ OK ]
#
```

The Subnet Manager is started.

Related Information

- *Switch Reference*, opensmd command
- “Start the Subnet Manager With Min Hop Routing” on page 102
- “Start the Subnet Manager With Fat Tree Routing” on page 103
- “Create the opensm.conf File” on page 106
- “Stop the Subnet Manager With the opensmd Daemon” on page 106

▼ Stop the Subnet Manager With the opensmd Daemon

You can stop the OpenSM Subnet Manager with the opensmd daemon.

- On the Linux InfiniBand host, type:

```
# /etc/init.d/opensmd stop
Stopping IB Subnet Manager...        [ OK ]
#
```

The Subnet Manager is stopped.

Related Information

- *Switch Reference*, opensmd command
- “Start the Subnet Manager With the opensmd Daemon” on page 106

▼ Create the opensm.conf File

The opensmd daemon uses the opensm.conf file to configure the OpenSM Subnet Manager. You must create an opensm.conf file that is optimized for the Sun Datacenter InfiniBand Switch 648.

1. On the Linux InfiniBand host, stop the OpenSM Subnet Manager.
See [“Stop the Subnet Manager With the opensmd Daemon”](#) on page 106.
2. If a `/etc/opensm/opensm.conf` file does not exist, create one:

```
# opensm --create-config /etc/opensm/opensm.conf
-----
OpenSM 3.3.2_20090831_ce6f8dd
Command Line Arguments:
  Creating config file template '/etc/opensm/opensm.conf'.
  Log File: /var/log/opensm.log
-----
#
```

The `/etc/opensm/opensm.conf` file is created.

3. Open the `/etc/opensm/opensm.conf` file in a text editor.
4. Find the following parameters and change them to as listed.
 - `sweep_interval 120`
 - `routing_engine ftree`
 - `use_ucast_cache TRUE`
 - `root_guid_file /etc/opensm/guid.txt`

Note – If the `/etc/opensm/guid.txt` file does not exist, you must create it. See [“Create the guid.txt File”](#) on page 104.

5. Save the `/etc/opensm/opensm.conf` file.
6. Set the permissions for the `/etc/opensm/opensm.conf` file:

```
# chmod 0644 /etc/opensm/opensm.conf
#
```

Related Information

- *Switch Reference*, `opensm` command
- *Switch Reference*, `opensmd` daemon
- [“Create the guid.txt File”](#) on page 104
- [“Start the Subnet Manager With the opensmd Daemon”](#) on page 106

Diagnosing a Problem

The following topics describe tools that help you diagnose a problem.

- [“Identifying the Problem” on page 109](#)
- [“Troubleshooting the Switch” on page 117](#)
- [“Understanding Signal Routing Through the Switch” on page 124](#)
- [“Understanding Switch Startup and Component Addition” on page 132](#)
- [“Switch GUIDs Overview” on page 136](#)

Related Information

- [“Administering the Switch” on page 1](#)
- [Switch Installation](#)
- [Switch Remote Management](#)
- [Switch Service](#)
- [Switch Reference](#)

Identifying the Problem

Part of the troubleshooting process is identifying if a problem exists. These topics describe tools you can use to initially check for problems.

- [“Checking Logs and Error Messages” on page 110](#)
- [“Checking Status LEDs” on page 112](#)

Related Information

- [“Troubleshooting the Switch” on page 117](#)
- [“Understanding Signal Routing Through the Switch” on page 124](#)
- [“Understanding Switch Startup and Component Addition” on page 132](#)

Checking Logs and Error Messages

One of the first steps in troubleshooting a problem is to determine when and where a problem occurred. The system log and switch logs provide time-stamped entries of events and errors.

- [“Switch Event Message Overview”](#) on page 110
- [“Display the Switch Message Log for Power Supplies”](#) on page 110
- [“Display the Switch Message Log for Fabric Cards”](#) on page 111
- [“Display the Switch Message Log for Line Cards”](#) on page 111
- [“Display the System Event Log”](#) on page 112

Related Information

- [“Checking Status LEDs”](#) on page 112

Switch Event Message Overview

The switch message log contains events that happened to the switch hardware. This log is stored on the CMC in the `/var/log/syslog` file. There are two ways to commands that display the log.

The `showlogs` command provides a more user friendly means of understanding the system events. The output is succinct and events are system-wide.

The CLIA `sel` command is used to view the system event log (`sel`) on a specific IPM controller. The `sel` command gives greater flexibility to viewing the log and the detail of the log entries than the `showlogs` command.

Related Information

- *Switch Reference*, `showlogs` command
- *Switch Reference*, `sel` command

▼ Display the Switch Message Log for Power Supplies

- On the CMC, type:

```
# showlogs | grep PSU slot
```

where *slot* is 00–03.

Related Information

- *Switch Reference*, showlogs command
- “Check Power Supply LEDs” on page 112
- “Display the Switch Message Log for Fabric Cards” on page 111
- “Display the Switch Message Log for Line Cards” on page 111
- “Display the System Event Log” on page 112

▼ Display the Switch Message Log for Fabric Cards

- On the CMC, type:

```
# showlogs | grep FC slot
```

where *slot* is 00–08.

Related Information

- *Switch Reference*, showlogs command
- “Check Fabric Card LEDs” on page 115
- “Check Fan LEDs” on page 115
- “Display the Switch Message Log for Power Supplies” on page 110
- “Display the Switch Message Log for Line Cards” on page 111
- “Display the System Event Log” on page 112

▼ Display the Switch Message Log for Line Cards

- On the CMC, type:

```
# showlogs | grep LC slot
```

where *slot* is 00–08.

Related Information

- *Switch Reference*, showlogs command
- “Check Line Card LEDs” on page 116
- “Display the Switch Message Log for Power Supplies” on page 110
- “Display the Switch Message Log for Fabric Cards” on page 111
- “Display the System Event Log” on page 112

▼ Display the System Event Log

- To display the system event messages, type on the CMC:

```
# clia sel
```

Related Information

- *Switch Reference*, `sel` command
- [“Display the Switch Message Log for Power Supplies” on page 110](#)
- [“Display the Switch Message Log for Fabric Cards” on page 111](#)
- [“Display the Switch Message Log for Line Cards” on page 111](#)

Checking Status LEDs

A simple way to quickly determine the general health of a switch component is to observe the component's status LEDs.

- [“Check Power Supply LEDs” on page 112](#)
- [“Check CMC Status LEDs” on page 113](#)
- [“Check the NET MGT LEDs” on page 114](#)
- [“Check Fabric Card LEDs” on page 115](#)
- [“Check Fan LEDs” on page 115](#)
- [“Check Line Card LEDs” on page 116](#)
- [“Check Link LEDs” on page 117](#)

Related Information

- [“Checking Logs and Error Messages” on page 110](#)

▼ Check Power Supply LEDs

The power supply status LEDs are located on the power supply at the front of the chassis.

1. **Visually inspect the status LEDs.**

2. Compare what is displayed on the power supply to the following table.

Glyph	Name	Color	State and Meaning
	OK	Green	On – Power supply enabled, 12 VDC is supplied. Off – No input voltage. Flashing – Power supply disabled, 12 VDC is not supplied. Standby voltage only.
	Attention	Amber	On – Fault detected, 12 VDC shut down. Off – Normal operation. Flashing – No function.



Caution – If a power supply has shut down because of a thermal or overcurrent condition, signified by the amber Attention LED lighting, remove the respective power cord from the chassis. Allow the power supply to completely cool for at least 15 minutes. Any less cooling time might cause damage to the power supply when the power cord is reattached. If the amber Attention LED lights upon reattaching the power cord, replace the power supply.

Related Information

- [“Check the Status of a Power Supply” on page 14](#)
- [“Display Power Supplies Present” on page 13](#)
- [“Check Fan LEDs” on page 115](#)

▼ Check CMC Status LEDs

The CMC status LEDs are located on the left side of the CMC panel.

1. **Visually inspect the status LEDs.**

2. Compare what is displayed on the CMC to the following table.

Glyph	Name	Color	State and Meaning
	OK	Green	On – CMC is operating normally. Off – No power is being supplied to the CMC. Flashing – CMC is inactive status as the backup CMC.
	Attention	Amber	On – A fault or critical error has been detected. Off – Normal operation. Flashing – No function.
	Ready-to-Remove	Blue	On – CMC has been deactivated and is ready to be removed. Off – Do not remove Flashing – No function.

Related Information

- [“Check the Status LEDs of a CMC” on page 17](#)
- [“Check the Status of a CMC” on page 16](#)
- [“Check Internal Power and Temperature of a CMC” on page 15](#)
- [“Check the NET MGT LEDs” on page 114](#)
- [“Check Fabric Card LEDs” on page 115](#)
- [“Check Line Card LEDs” on page 116](#)

▼ **Check the NET MGT LEDs**

The network management status LEDs are located at the center of the CMC panel.

1. Visually inspect the status LEDs.
2. Compare what is displayed at the NET MGT connector to the following table.

Name	Color	State and Meaning
Link	Green	On – 10BASE-T link. Off – No link or link down. Flashing – 100BASE-T link.
Activity	Amber	On – No function. Off – No activity. Flashing – Packet activity.

Related Information

- [“Display the Network Management Configuration” on page 41](#)

- [“Reconfigure the Network Management Parameters” on page 61](#)
- [“Check CMC Status LEDs” on page 113](#)
- [“Check Link LEDs” on page 117](#)

▼ Check Fabric Card LEDs

The fabric card status LEDs are located at the center of the fabric card.

1. **Visually inspect the status LEDs.**
2. **Compare what is displayed on the fabric card to the following table.**

Glyph	Location	Name	Color	State and Meaning
	Left	Ok Ready-to-Remove	Green	On – Normal operation at full power. Off – No power at all applied. Flashing – Fabric card has been disabled and is ready to be removed, standby voltage available.
	Center	Attention	Amber	On – A fault or critical error has been detected. Off – Normal operation. Flashing – No function.
	Right	Locator	White	On – No function. Off – Normal operation. Flashing – The fabric card is identifying itself.

Related Information

- [“Check the Status LEDs of a Fabric Card” on page 29](#)
- [“Display Fabric Cards Present” on page 20](#)
- [“Check Fabric Card Power Faults” on page 21](#)
- [“Check the Internal Power and Temperature of a Fabric Card” on page 22](#)
- [“Check Fabric Card Switch Chip Health” on page 27](#)
- [“Check CMC Status LEDs” on page 113](#)
- [“Check Fan LEDs” on page 115](#)
- [“Check Line Card LEDs” on page 116](#)

▼ Check Fan LEDs

The fan status LED is located in the lower left corner of the fan.

1. Visually inspect the status LED.
2. If the LED is lit, there is a problem with that fan.

Related Information

- [“Check Fan Speed and Status” on page 19](#)
- [“Display Fabric Cards Present” on page 20](#)
- [“Check Power Supply LEDs” on page 112](#)
- [“Check Fabric Card LEDs” on page 115](#)

▼ Check Line Card LEDs

The line card status LEDs are located at the center of the line card.

1. Visually inspect the status LEDs.
2. Compare what is displayed on the fabric card to the following table.

Glyph	Location	Name	Color	State and Meaning
	Top	Locator	White	On – No function. Off – Normal operation. Flashing – The line card is identifying itself.
	Middle	Attention	Amber	On – A fault or critical error has been detected. Off – Normal operation. Flashing – No function.
	Bottom	Ok	Green	On – Normal operation at full power.
		Ready-to-Remove		Off – No power at all applied. Flashing – Fabric card has been disabled and is ready to be removed, standby voltage available.

Related Information

- [“Check the Status LEDs of a Line Card” on page 39](#)
- [“Display Line Cards Present” on page 31](#)
- [“Check Line Card Power Faults” on page 32](#)
- [“Check the Internal Power and Temperature of a Line Card” on page 32](#)
- [“Check Line Card Switch Chip Health” on page 38](#)
- [“Check CMC Status LEDs” on page 113](#)
- [“Check Fabric Card LEDs” on page 115](#)

▼ Check Link LEDs

The link LEDs are located along the bottom edge of the front of the line card.

1. Visually inspect the link LEDs.
2. Compare what is displayed for a particular link to the following table.

Name	Color	State and Meaning
Link	Green	On – Link established. Off – No link or link down. Flashing – Symbol errors.

Related Information

- [“Check Line Card Link Status \(Detailed\)” on page 37](#)
- [“Display the Link Status of a Node” on page 49](#)
- [“Display Counters for a Node” on page 50](#)
- [“Check the NET MGT LEDs” on page 114](#)

Troubleshooting the Switch

The following topics help you resolve some basic problems that might occur with your switch.

- [“Switch Hardware Problems” on page 118](#)
- [“InfiniBand Fabric Problems” on page 121](#)

Related Information

- [“Identifying the Problem” on page 109](#)
- [“Understanding Signal Routing Through the Switch” on page 124](#)
- [“Understanding Switch Startup and Component Addition” on page 132](#)

Switch Hardware Problems

The following table lists situations that might occur with switch hardware and corrective steps that can be taken to resolve the problem.

Situation	Corrective Steps
The Attention LED on a power supply is lit or the power supply seems dysfunctional.	<ol style="list-style-type: none">1. Check that the power supply is present. See “Display Power Supplies Present” on page 13.2. Check the power supply status. See “Check the Status of a Power Supply” on page 14.3. Unplug the respective power cord, wait 15 minutes, then reattach the power cord.4. If the previous steps do not rectify the situation, replace the power supply. See <i>Switch Service</i>, servicing a power supply.
The Attention LED on a CMC is lit or the CMC seems dysfunctional.	<ol style="list-style-type: none">1. If you are able to access the CMC, reboot the CMC. See “Restart a CMC” on page 69.2. If the previous step does not rectify the situation, perform a hot-swap of the CMC. See <i>Switch Service</i>, removing a CMC, installing a CMC.3. If you are unable to access the CMC, or the previous step does not rectify the situation, replace the CMC. See <i>Switch Service</i>, servicing a CMC.
The Attention LED on a fan is lit or the fan seems dysfunctional.	<ol style="list-style-type: none">1. Check that the respective fabric card is present. See “Display Fabric Cards Present” on page 20.2. Check the fan speed. See “Check Fan Speed and Status” on page 19.3. Check the logs for any indication of fan failure. See “Display the Switch Message Log for Fabric Cards” on page 111.4. If the previous steps do not rectify the situation, replace the fan. See the <i>Switch Service</i>, replacing a fan.5. If replacing the suspect fan with a known good fan does not rectify the situation, replace the fabric card. See <i>Switch Service</i>, servicing a fabric card.

Situation	Corrective Steps
The Attention LED on a fabric card is lit or the fabric card seems dysfunctional.	<ol style="list-style-type: none"> 1. Check that the fabric card is present. See “Display Fabric Cards Present” on page 20. 2. Ensure that the environmental conditions are within limits. See <i>Switch Installation</i>, environmental requirements. 3. Check overall switch health. See “Display the General Health of the Switch” on page 11 4. Check that there are no fabric card power faults. See “Check Fabric Card Power Faults” on page 21. 5. Check the logs for any indication of fabric card failure. See “Display the Switch Message Log for Fabric Cards” on page 111. 6. Restart the fabric card. See “Restart a Fabric Card or Filler” on page 72. 7. If the previous steps do not rectify the situation, replace the fabric card. See <i>Switch Service</i>, servicing a fabric card.
The Attention LED on a line card is lit or the line card seems dysfunctional.	<ol style="list-style-type: none"> 1. Check that the line card is present. See “Display Line Cards Present” on page 31. 2. Ensure that the environmental conditions are within limits. See <i>Switch Installation</i>, environmental requirements. 3. Check overall switch health. See “Display the General Health of the Switch” on page 11 4. Check that there are no line card power faults. See “Check Line Card Power Faults” on page 32. 5. Check the logs for any indication of line card failure. See “Display the Switch Message Log for Line Cards” on page 111. 6. Restart the line card. See “Restart a Line Card” on page 82. 7. If the previous steps do not rectify the situation, replace the line card. See <i>Switch Service</i>, servicing a line card.

Situation	Corrective Steps
After installation, not all internal links are operational.	<ol style="list-style-type: none"> 1. Determine if the non-operational links are localized to one fabric card or line card. See “Check Fabric Card Link Status (Simple)” on page 26 and “Check Line Card Link Status (Simple)” on page 36. 2. If the non-operational links are localized to one fabric card, reset the links for that card. See “Enable Downed Fabric Card Links” on page 78. 3. If the non-operational links are localized to one line card, verify that each IB cable connection has appropriate Link LEDs lit. See “Check Link LEDs” on page 117. 4. If the non-operational links are localized to one line card, reset the links for that card. See “Enable Downed Line Card Links” on page 88. 5. If both fabric cards and line cards have non-operational links, verify if there is commonality to the links. See “Understanding Signal Routing Through the Switch” on page 124 6. If there is commonality through an XBOW connector: Unscrew the retainer bolts of the affected fabric cards or line cards or both by 3/4 turn. Alternate between retainer bolts, turning each 1/4 turn counter-clockwise each time. 7. If the previous step did not rectify the situation, reseal the fabric card and line card. See <i>Switch Service</i>, servicing a fabric card, servicing a line card. 8. If the previous steps do not rectify the situation, remove the fabric card or line card and inspect the midplane XBOW connectors and the fabric card or line card XBOW connectors. See <i>Switch Service</i>, removing a fabric card, removing a line card. See <i>Switch Installation</i>, inspecting the midplane connectors, inspecting the XBOW connectors. 9. If the XBOW connectors are in good condition, exchange the slot positions of the suspect fabric card with a known good fabric card. See <i>Switch Service</i>, servicing a fabric card. 10. If the problem has moved to the other fabric card slot, replace the suspect fabric card with a new one. If the problem remains at the original fabric card slot, replace the suspect line card with a new one. See <i>Switch Service</i>, replacing a fabric card, replacing a line card.

Related Information

- [“InfiniBand Fabric Problems”](#) on page 121

InfiniBand Fabric Problems

The following table lists situations that might occur with the InfiniBand fabric and corrective steps that can be taken to resolve the problem.

Situation	Corrective Steps
An InfiniBand Link LED is blinking.	<ol style="list-style-type: none"><li data-bbox="394 371 1239 423">1. Disconnect and properly reconnect both ends of the respective InfiniBand cable. See <i>Switch Service</i>, servicing an InfiniBand cable.<li data-bbox="394 435 1310 522">2. If the LED is still blinking, determine the significance of the errors through use of the <code>ibdiagnet</code> command. See “Determine Which Links Are Experiencing Significant Errors” on page 95.<li data-bbox="394 534 972 586">3. Determine which connectors map to the affected link. See “Locate a Switch Chip or Connector From the GUID” on page 43.<li data-bbox="394 598 1310 651">4. If some of the links are running at 1x or SDR, use that situation elsewhere in this table to rectify the problem.<li data-bbox="394 663 1100 715">5. Disable and re-enable the respective ports. See “Disable a Port” on page 100 and “Enable a Port” on page 101.<li data-bbox="394 727 1310 779">6. If the errors are still significant, swap the cable with a known good one or reconnect the cable to a known good remote port, and repeat from 2.<li data-bbox="394 791 1293 881">7. Depending upon what does or does not rectify the problem, replace that component. See <i>Switch Service</i>, replacing a InfiniBand cable, replacing a line card. See remote port’s documentation for replacement procedures.
There are errors on some InfiniBand links.	<ol style="list-style-type: none"><li data-bbox="394 900 822 961">1. Clear the error counters. See “Clear Error Counters” on page 96.<li data-bbox="394 968 662 991">2. Start a fabric stress test.<li data-bbox="394 998 1310 1085">3. Identify the suspect links using the <code>ibdiagnet</code> command. See “Determine Which Links Are Experiencing Significant Errors” on page 95. Look for text like the following: <pre data-bbox="419 1097 1186 1222">-W- lid=0x0006 guid=0x0021283a8816c0a0 dev=48438 Port=34 Performance Monitor counter : Value link_recovery_error_counter : 0x1 symbol_error_counter : 0x25 (Increase by 3 during ibdiagnet)</pre><li data-bbox="394 1229 1279 1281">4. For links that are experiencing recovery errors or substantial symbol errors, refer to other parts of this table to help identify the cause and rectify the problem.

Situation	Corrective Steps
Output of InfiniBand software commands provides only GUID and port, not switch chip numbers or CXP connectors.	<ol style="list-style-type: none"> 1. The <code>findport</code> switch-specific command can translate GUID port combinations to the location in the switch. See “Locate a Switch Chip or Connector From the GUID” on page 43. 2. If the port immediately links to a CXP connector, the <code>findport</code> command identifies that connector. See <i>Switch Reference</i>, <code>findport</code> command.
Some InfiniBand links are running at 1x or SDR.	<p>For a temporary solution:</p> <ol style="list-style-type: none"> 1. Identify the suspect links using the <code>ibdiagnet</code> command. See “Find 1x or SDR or DDR Links in the Fabric” on page 94. Look for text like the following: <pre>-W- link with SPD=2.5 found at direct path "1,19" From: a Switch PortGUID=0x00066a00d80001dd Port=19 To: a Switch PortGUID=0x00066a00d80001dd Port=24</pre> 2. Determine which connectors map to the affected link. See “Locate a Switch Chip or Connector From the GUID” on page 43. 3. Verify the cable connection at both ends. See <i>Switch Service</i>, servicing an InfiniBand cable. 4. Disable and re-enable the respective ports. See “Disable a Port” on page 100 and “Enable a Port” on page 101. 5. If the previous steps do not rectify the problem, disable the port. See “Disable a Port” on page 100. <p>For a permanent solution:</p> <ol style="list-style-type: none"> 1. Perform the steps for a temporary solution, steps 1 to step 5. 2. Swap the cable with a known good one or reconnect the cable to a known good remote port, and repeat from 1. 3. Depending upon what does or does not rectify the problem, replace that component. See <i>Switch Service</i>, replacing a InfiniBand cable, replacing a fabric card, replacing a line card. See remote port’s documentation for replacement procedures.

Situation	Corrective Steps
Performance of the InfiniBand fabric seems diminished.	<ol style="list-style-type: none"> 1. Determine if there are errors or problems with the InfiniBand fabric. See: <ul style="list-style-type: none"> “Perform Comprehensive Diagnostics for the Entire Fabric” on page 91 “Find 1x or SDR or DDR Links in the Fabric” on page 94 “Determine Which Links Are Experiencing Significant Errors” on page 95 2. Locate the affected nodes by the GUID provided in the output of the <code>ibdiagnet</code> command. See “Locate a Switch Chip or Connector From the GUID” on page 43. 3. If the problem is at a cable connection, swap the suspect cable with a known good cable or reconnect the cable to a known good remote port and repeat Step 1. See <i>Switch Service</i>, servicing an InfiniBand cable. 4. If the problem still remains at the cable connection, disable and re-enable that port on the line card and repeat Step 1. See “Disable a Port” on page 100 and “Enable a Port” on page 101. 5. If the problem is within a line card or fabric card, disable and re-enable the respective port. See “Disable a Port” on page 100 and “Enable a Port” on page 101. 6. If the problem still remains within a line card or fabric card, reduce the local deflection of the midplane. Unscrew the retainer bolts of the affected fabric cards or line cards or both by 3/4 turn. Alternate between retainer bolts, turning each 1/4 turn counter-clockwise each time and then reseal the fabric card or line card. See <i>Switch Service</i>, servicing a fabric card, servicing a line card. <p>Temporary solution:</p> <ul style="list-style-type: none"> • If the problem still remains, disable the affected port. See “Disable a Port” on page 100. <p>Permanent solution:</p> <ul style="list-style-type: none"> • If the problem still remains, replace the affected component. See <i>Switch Service</i>, replacing a fabric card, replacing a line card, replacing an InfiniBand cable. See remote port’s documentation for replacement procedures.

Related Information

- “Switch Hardware Problems” on page 118

Understanding Signal Routing Through the Switch

Several tables in the following topics describe the routing through the line cards and fabric cards. The tables map the switch chip and port to an XBOW connector or CXP connector. The tables also provide a reverse map. When command output provides a switch chip and port, you can use these tables to determine the route that link is following.

Additionally, if a command describes a CXP connection, or you want to associate a line card link LED to its respective switch chip port, use these tables for that purpose too.

- [“CXP Connectors and LEDs to Line Card Switch Chip and Port Routes”](#) on page 124
- [“Line Card Switch Chip to Line Card XBOW Connector Routes”](#) on page 126
- [“Fabric Card XBOW Connector to Fabric Card Switch Chip Routes”](#) on page 127
- [“Fabric Card Switch Chip to Fabric Card XBOW Connector Routes”](#) on page 128
- [“Line Card XBOW Connector to Line Card Switch Chip Routes”](#) on page 129
- [“Line Card Switch Chip and Port to CXP Connector and LED Routes”](#) on page 130
- [“Signal Route Through the Switch”](#) on page 131

Related Information

- [“Display the InfiniBand Fabric Topology”](#) on page 47
- [“Display a Route Through the Fabric”](#) on page 47
- [“Perform Comprehensive Diagnostics for a Route”](#) on page 92
- [“Identifying the Problem”](#) on page 109
- [“Troubleshooting the Switch”](#) on page 117
- [“Understanding Switch Startup and Component Addition”](#) on page 132

CXP Connectors and LEDs to Line Card Switch Chip and Port Routes

The following table provides a mapping of the CXP connector and its link LEDs to the respective I4 switch chip and port. The I4 switch chip and port are given as *chip-port*, where:

- *chip* – The identifying number of the I4 switch chip (0–3).
- *port* – The identifying number of the port (19–36).

Connector Group	CXP Connector A			CXP Connector B		
	Left LED P3	Center LED P2	Right LED P1	Left LED P3	Center LED P2	Right LED P1
0	0-20	0-21	0-22	0-19	0-23	0-24
1	0-26	0-27	0-28	0-25	0-29	0-30
2	0-35	0-34	0-33	0-36	0-32	0-31
3	1-20	1-21	1-22	1-19	1-23	1-24
4	1-26	1-27	1-28	1-25	1-29	1-30
5	1-35	1-34	1-33	1-36	1-32	1-31
6	2-20	2-21	2-22	2-19	2-23	2-24
7	2-26	2-27	2-28	2-25	2-29	2-30
8	2-35	2-34	2-33	2-36	2-32	2-31
9	3-20	3-21	3-22	3-19	3-23	3-24
10	3-26	3-27	3-28	3-25	3-29	3-30
11	3-35	3-34	3-33	3-36	3-32	3-31

Note – When using CXP to QSFP splitter cables, the P1, P2, and P3 QSFP connectors are related to the right, center, and left link LEDs respectively.

Related Information

- [“Check Line Card Link Status \(Detailed\)” on page 37](#)
- [“Display a Route Through the Fabric” on page 47](#)
- [“Perform Comprehensive Diagnostics for a Route” on page 92](#)
- [“Line Card Switch Chip and Port to CXP Connector and LED Routes” on page 130](#)

Line Card Switch Chip to Line Card XBOW Connector Routes

The following table provides a mapping of the I4 switch chips and their ports to the XBOW connector and its ports, for the line cards.

I4	Port	→	XBOW	Port	I4	Port	→	XBOW	Port	I4	Port	→	XBOW	Port	I4	Port	→	XBOW	Port
0	1	→	8	2	1	1	→	8	4	2	1	→	8	6	3	1	→	8	8
0	2	→	8	1	1	2	→	8	3	2	2	→	8	5	3	2	→	8	7
0	3	→	7	2	1	3	→	7	4	2	3	→	7	6	3	3	→	7	8
0	4	→	7	1	1	4	→	7	3	2	4	→	7	5	3	4	→	7	7
0	5	→	6	2	1	5	→	6	4	2	5	→	6	6	3	5	→	6	8
0	6	→	6	1	1	6	→	6	3	2	6	→	6	5	3	6	→	6	7
0	7	→	5	2	1	7	→	5	4	2	7	→	5	6	3	7	→	5	8
0	8	→	5	1	1	8	→	5	3	2	8	→	5	5	3	8	→	5	7
0	9	→	4	2	1	9	→	4	4	2	9	→	4	6	3	9	→	4	8
0	10	→	4	1	1	10	→	4	3	2	10	→	4	5	3	10	→	4	7
0	11	→	3	2	1	11	→	3	4	2	11	→	3	6	3	11	→	3	8
0	12	→	3	1	1	12	→	3	3	2	12	→	3	5	3	12	→	3	7
0	13	→	0	1	1	13	→	0	3	2	13	→	0	5	3	13	→	0	7
0	14	→	0	2	1	14	→	0	4	2	14	→	0	6	3	14	→	0	8
0	15	→	1	1	1	15	→	1	3	2	15	→	1	5	3	15	→	1	7
0	16	→	1	2	1	16	→	1	4	2	16	→	1	6	3	16	→	1	8
0	17	→	2	1	1	17	→	2	3	2	17	→	2	5	3	17	→	2	7
0	18	→	2	2	1	18	→	2	4	2	18	→	2	6	3	18	→	2	8

Related Information

- [“Check Line Card Link Status \(Detailed\)” on page 37](#)
- [“Display a Route Through the Fabric” on page 47](#)
- [“Perform Comprehensive Diagnostics for a Route” on page 92](#)
- [“Line Card XBOW Connector to Line Card Switch Chip Routes” on page 129](#)

Fabric Card XBOW Connector to Fabric Card Switch Chip Routes

The following table provides a reverse-lookup mapping of the XBOW connector and its ports, to the I4 switch chip and its port, for the fabric cards.

XBOW	Port	→	I4	Port	XBOW	Port	→	I4	Port	XBOW	Port	→	I4	Port
0	1	→	0	1	3	1	→	0	36	6	1	→	0	12
0	2	→	1	24	3	2	→	1	13	6	2	→	1	7
0	3	→	1	25	3	3	→	1	31	6	3	→	1	6
0	4	→	0	2	3	4	→	0	28	6	4	→	0	17
0	5	→	0	3	3	5	→	0	29	6	5	→	0	16
0	6	→	1	23	3	6	→	1	32	6	6	→	1	9
0	7	→	1	27	3	7	→	1	14	6	7	→	1	5
0	8	→	0	4	3	8	→	0	26	6	8	→	0	15
1	1	→	0	22	4	1	→	0	33	7	1	→	0	8
1	2	→	1	28	4	2	→	1	17	7	2	→	1	2
1	3	→	1	29	4	3	→	1	16	7	3	→	1	3
1	4	→	0	21	4	4	→	0	34	7	4	→	0	11
1	5	→	0	20	4	5	→	0	35	7	5	→	0	10
1	6	→	1	26	4	6	→	1	15	7	6	→	1	4
1	7	→	1	36	4	7	→	1	12	7	7	→	1	1
1	8	→	0	19	4	8	→	0	30	7	8	→	0	18
2	1	→	0	27	5	1	→	0	14	8	1	→	0	5
2	2	→	1	34	5	2	→	1	11	8	2	→	1	21
2	3	→	1	35	5	3	→	1	10	8	3	→	1	20
2	4	→	0	24	5	4	→	0	13	8	4	→	0	7
2	5	→	0	25	5	5	→	0	31	8	5	→	0	6
2	6	→	1	30	5	6	→	1	18	8	6	→	1	19
2	7	→	1	33	5	7	→	1	8	8	7	→	1	22
2	8	→	0	23	5	8	→	0	32	8	8	→	0	9

Related Information

- [“Check Fabric Card Link Status \(Detailed\)” on page 26](#)
- [“Display a Route Through the Fabric” on page 47](#)
- [“Perform Comprehensive Diagnostics for a Route” on page 92](#)
- [“Fabric Card Switch Chip to Fabric Card XBOW Connector Routes” on page 128](#)

Fabric Card Switch Chip to Fabric Card XBOW Connector Routes

The following table provides a mapping of the I4 switch chips and their ports to the XBOW connector and its ports, for the fabric cards.

I4	Port	→	XBOW	Port	I4	Port	→	XBOW	Port	I4	Port	→	XBOW	Port	I4	Port	→	XBOW	Port
0	1	→	0	1	0	19	→	1	8	1	1	→	7	7	1	19	→	8	6
0	2	→	0	4	0	20	→	1	5	1	2	→	7	2	1	20	→	8	3
0	3	→	0	5	0	21	→	1	4	1	3	→	7	3	1	21	→	8	2
0	4	→	0	8	0	22	→	1	1	1	4	→	7	6	1	22	→	8	7
0	5	→	8	1	0	23	→	2	8	1	5	→	6	7	1	23	→	0	6
0	6	→	8	5	0	24	→	2	4	1	6	→	6	3	1	24	→	0	2
0	7	→	8	4	0	25	→	2	5	1	7	→	6	2	1	25	→	0	3
0	8	→	7	1	0	26	→	3	8	1	8	→	5	7	1	26	→	1	6
0	9	→	8	8	0	27	→	2	1	1	9	→	6	6	1	27	→	0	7
0	10	→	7	5	0	28	→	3	4	1	10	→	5	3	1	28	→	1	2
0	11	→	7	4	0	29	→	3	5	1	11	→	5	2	1	29	→	1	3
0	12	→	6	1	0	30	→	4	8	1	12	→	4	7	1	30	→	2	6
0	13	→	5	4	0	31	→	5	5	1	13	→	3	2	1	31	→	3	3
0	14	→	5	1	0	32	→	5	8	1	14	→	3	7	1	32	→	3	6
0	15	→	6	8	0	33	→	4	1	1	15	→	4	6	1	33	→	2	7
0	16	→	6	5	0	34	→	4	4	1	16	→	4	3	1	34	→	2	2
0	17	→	6	4	0	35	→	4	5	1	17	→	4	2	1	35	→	2	3
0	18	→	7	8	0	36	→	3	1	1	18	→	5	6	1	36	→	1	7

Related Information

- “Check Fabric Card Link Status (Detailed)” on page 26
- “Display a Route Through the Fabric” on page 47
- “Perform Comprehensive Diagnostics for a Route” on page 92
- “Fabric Card XBOW Connector to Fabric Card Switch Chip Routes” on page 127

Line Card XBOW Connector to Line Card Switch Chip Routes

The following table provides a reverse-lookup mapping of the XBOW connector and its ports, to the I4 switch chip and its port, for the line cards.

XBOW	Port	→	I4	Port	XBOW	Port	→	I4	Port	XBOW	Port	→	I4	Port
0	1	→	0	13	3	1	→	0	12	6	1	→	0	6
0	2	→	0	14	3	2	→	0	11	6	2	→	0	5
0	3	→	1	13	3	3	→	1	12	6	3	→	1	6
0	4	→	1	14	3	4	→	1	11	6	4	→	1	5
0	5	→	2	13	3	5	→	2	12	6	5	→	2	6
0	6	→	2	14	3	6	→	2	11	6	6	→	2	5
0	7	→	3	13	3	7	→	3	12	6	7	→	3	6
0	8	→	3	14	3	8	→	3	11	6	8	→	3	5
1	1	→	0	15	4	1	→	0	10	7	1	→	0	4
1	2	→	0	16	4	2	→	0	9	7	2	→	0	3
1	3	→	1	15	4	3	→	1	10	7	3	→	1	4
1	4	→	1	16	4	4	→	1	9	7	4	→	1	3
1	5	→	2	15	4	5	→	2	10	7	5	→	2	4
1	6	→	2	16	4	6	→	2	9	7	6	→	2	3
1	7	→	3	15	4	7	→	3	10	7	7	→	3	4
1	8	→	3	16	4	8	→	3	9	7	8	→	3	3
2	1	→	0	17	5	1	→	0	8	8	1	→	0	2
2	2	→	0	18	5	2	→	0	7	8	2	→	0	1
2	3	→	1	17	5	3	→	1	8	8	3	→	1	2
2	4	→	1	18	5	4	→	1	7	8	4	→	1	1

XBOW	Port	→	I4	Port	XBOW	Port	→	I4	Port	XBOW	Port	→	I4	Port
2	5	→	2	17	5	5	→	2	8	8	5	→	2	2
2	6	→	2	18	5	6	→	2	7	8	6	→	2	1
2	7	→	3	17	5	7	→	3	8	8	7	→	3	2
2	8	→	3	18	5	8	→	3	7	8	8	→	3	1

Related Information

- [“Check Line Card Link Status \(Detailed\)” on page 37](#)
- [“Display a Route Through the Fabric” on page 47](#)
- [“Perform Comprehensive Diagnostics for a Route” on page 92](#)
- [“Line Card Switch Chip to Line Card XBOW Connector Routes” on page 126](#)

Line Card Switch Chip and Port to CXP Connector and LED Routes

A reverse-lookup mapping of the I4 switch chip and port to CXP connector and its link LEDs is provided in the following table. The connectors and LEDs are given as *connectorletter-location*, where:

- *connector* – The identifying number of the connector group (0–11).
- *letter* – The letter identifying the upper (A) or lower (B) connector.
- *location* – The location of the LED, left, center, or right.

Port	I4 Switch 0	I4 Switch 1	I4 Switch 2	I4 Switch 3
19	0B-left-P3	3B-left-P3	6B-left-P3	9B-left-P3
20	0A-left-P3	3A-left-P3	6A-left-P3	9A-left-P3
21	0A-center-P2	3A-center-P2	6A-center-P2	9A-center-P2
22	0A-right-P1	3A-right-P1	6A-right-P1	9A-right-P1
23	0B-center-P2	3B-center-P2	6B-center-P2	9B-center-P2
24	0B-right-P1	3B-right-P1	6B-right-P1	9B-right-P1
25	1B-left-P3	4B-left-P3	7B-left-P3	10B-left-P3
26	1A-left-P3	4A-left-P3	7A-left-P3	10A-left-P3
27	1A-center-P2	4A-center-P2	7A-center-P2	10A-center-P2
28	1A-right-P1	4A-right-P1	7A-right-P1	10A-right-P1

Port	I4 Switch 0	I4 Switch 1	I4 Switch 2	I4 Switch 3
29	1B-center-P2	4B-center-P2	7B-center-P2	10B-center-P2
30	1B-right-P1	4B-right-P1	7B-right-P1	10B-right-P1
31	2B-right-P1	5B-right-P1	8B-right-P1	11B-right-P1
32	2B-center-P2	5B-center-P2	8B-center-P2	11B-center-P2
33	2A-right-P1	5A-right-P1	8A-right-P1	11A-right-P1
34	2A-center-P2	5A-center-P2	8A-center-P2	11A-center-P2
35	2A-left-P3	5A-left-P3	8A-left-P3	11A-left-P3
36	2B-left-P3	5B-left-P3	8B-left-P3	11B-left-P3

Note – When using CXP to QSCP splitter cables, the P1, P2, and P3 QSCP connectors are related to the right, center, and left link LEDs respectively.

Related Information

- [“Check Line Card Link Status \(Detailed\)” on page 37](#)
- [“Display a Route Through the Fabric” on page 47](#)
- [“Perform Comprehensive Diagnostics for a Route” on page 92](#)
- [“CXP Connectors and LEDs to Line Card Switch Chip and Port Routes” on page 124](#)

Signal Route Through the Switch

By combining the information from [“Understanding Signal Routing Through the Switch” on page 124](#), it is possible to determine a route through the switch. This topic describes a sample situation that might occur.

1. A route is initiated at line card LC 6, connector 2A. The left LED blinks.
2. Using [“CXP Connectors and LEDs to Line Card Switch Chip and Port Routes” on page 124](#), it is determined that the link routes to I4 switch chip 0, through port 35.
3. The Subnet Manager instructs I4 switch chip 0 to use port 4 to forward the link.
4. Using [“Line Card Switch Chip to Line Card XBOW Connector Routes” on page 126](#), it is determined that the link routes to XBOW connector 7, port 1.
5. The link leaves line card LC 6 and enters fabric card FC 7 at XBOW connector 6, port 1.

Note – At line card-midplane-fabric card XROW connector interface, the port remains the same.

6. Using [“Fabric Card XROW Connector to Fabric Card Switch Chip Routes”](#) on page 127, it is determined that the link routes to I4 switch chip 0, through port 12.
7. The Subnet Manager instructs I4 switch chip 0 to use port 22 to forward the link.
8. Using [“Fabric Card Switch Chip to Fabric Card XROW Connector Routes”](#) on page 128, it is determined that the link routes to XROW connector 1, port 1.
9. The link leaves fabric card FC 7 and enters line card LC 1 at XROW connector 7 port 1.

Note – At fabric card-midplane-line card XROW connector interface, the port remains the same.

10. Using [“Line Card XROW Connector to Line Card Switch Chip Routes”](#) on page 129, it is determined that the link routes to I4 switch chip 0, through port 4.
11. The Subnet Manager instructs I4 switch chip 0 to use port 24 to forward the link.
12. Using [“Line Card Switch Chip and Port to CXP Connector and LED Routes”](#) on page 130, it is determined that the link exits line card LC 1 at connector 0B. The right LED blinks.

Related Information

- [“Display a Route Through the Fabric”](#) on page 47
- [“Perform Comprehensive Diagnostics for a Route”](#) on page 92

Understanding Switch Startup and Component Addition

These topics describe the behavior of the switch and how it affects the InfiniBand fabric during startup and when InfiniBand fabric components are added.

- [“Switch Startup Process”](#) on page 133
- [“Fabric Card Addition Process”](#) on page 134
- [“Line Card Addition Process”](#) on page 135

Related Information

- “Identifying the Problem” on page 109
- “Troubleshooting the Switch” on page 117
- “Understanding Signal Routing Through the Switch” on page 124

Switch Startup Process

1. Switch startup commences when standby power is enabled for the Chassis Management Controllers (CMCs). Standby power is applied when the power supplies are energized.
2. The redundant CMCs perform self-tests and determine if the other instance is present and operational. The CMCs then negotiate active and standby roles.
3. The active CMC becomes available on the management network.
4. The CMC checks that the required number of power supplies are present and operational in each power domain, and then activates full power mode for each power supply.
5. If the power supply population is functional, commands to the CMC enable standby power for each fabric card and line card. Standby power is provided to each card sequentially to avoid overloading the power supplies.
6. The CMC performs further discovery and checks of the fabric cards and line cards using Intelligent Platform Management interface (IPMI) communication with the Intelligent Platform Management controllers (IPMCs) in the fabric cards and line cards.

Note – Conversion from standby to full-power mode can be an automatic or explicitly controlled operation.

7. The CMC determines if the available power is sufficient for the number of fabric cards and line cards in each power domain, and then sends commands to the CMC to initiate power-on sequencing. All the fabric cards are powered on first, followed by the line cards.
8. After power-on of each fabric card, the CMC checks that a sufficient number of fans on the fabric card are fully operational. If too few fans are operational, the fabric card is shut down.
9. Completing the fabric card power-on, the CMC checks that both the card and supplied power status is okay.

10. When all fabric cards have been powered on and checked, the CMC re-evaluates the cooling capacity to ensure that line cards are only powered on if sufficient cooling capacity is available.
11. The line cards are then powered on, as directed by commands to the CMC.
12. When a fabric card or line card is instructed to go to full power, the card's on-board IPMC verifies that the local power sensors indicate okay status, and then enables power for the switch chips.
13. The IPMC performs basic card self-test functions that include being able to read sensible data from the key I²C slave devices on the card. The switch chips also have I²C slave devices.
14. The IPMC constantly monitors the various sensors and signals on the card, including cable and link status.
15. The IPMC enables all InfiniBand ports and links on the fabric cards, and all internal ports and links on the line cards. The CXP connectors on the line cards are not enabled until the IPMC is instructed to do so by the CMC.
16. After power sequencing and basic tests have been completed, the IPMC reports the card as operational to the CMC.
17. When the CMC has received operational status confirmation from all cards, the CMC correlates InfiniBand link state information to verify that all relevant internal links have become operational.
18. Links that fail to become operational are recorded.
19. The CMC then instructs the IPMCs on the line cards to enable external links through the CXP connectors on the line cards.

Fabric Card Addition Process

Adding a fabric card is normally a repair action or expansion of a less than fully configured switch.

1. The administrator informs the CMC of the slot where the fabric card is to be installed. The CMC verifies that the slot is empty and writes a log event.
2. The CMC acknowledges the new card has been inserted by observing the corresponding presence sensors.
3. Through administrator commands, the CMC enables standby power for the new fabric card and verifies that the IPMC on the card is fully operational.

4. The CMC checks that the available power supplies in the corresponding power domain can accommodate another fabric card powering on. The administrator commands the CMC to power on the card and the CMC waits for the IPMC to report the card as fully operational.
5. If installing the fabric card compliments the chassis cooling budget, the CMC powers on any previously disabled components and informs the Fabric Manager.
6. The CMC checks that all InfiniBand links connected to operational line cards have trained correctly, and that no excessive error counts exist for the fabric card ports and corresponding line card ports.

Related Information

- *Switch Service*, installing a fabric card

Line Card Addition Process

Adding a line card is normally a repair action or expansion of a less than fully configured switch.

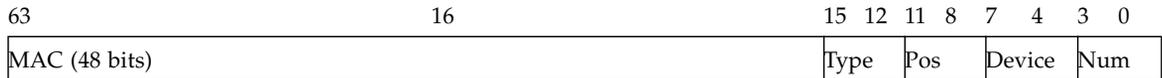
1. The administrator informs the CMC identifying the slot where the line card is to be installed. The CMC verifies that the slot is empty and writes a log event.
2. The CMC acknowledges that the new card has been inserted by observing the corresponding presence sensors.
3. Through administrator commands, the CMC enables standby power for the new line card and verifies that the IPMC on the card is fully operational.
4. The CMC checks that the available power supplies in the corresponding power domain can accommodate another line card powering on. The administrator commands the CMC to power on the card, and waits for the IPMC to report the card as fully operational.
5. The CMC checks that all InfiniBand links connected to operational fabric cards have trained correctly, and that no excessive error counts exists for the line card ports and corresponding fabric card ports.
6. The CMC records the current CXP cable status for the line card.
7. All the external links on the line card remain disabled.

Related Information

- *Switch Service*, installing a line card

Switch GUIDs Overview

Global unit identifiers (GUIDs) are unique 64-bit strings which identify nodes such as switches and channel adapters. For Oracle's Sun Datacenter InfiniBand Switches, GUIDs are modified to identify the node's role and location. The following table describes the GUID's structure.



The five fields of the GUID are described as follows:

- MAC – Bits 63 through 16 are the Machine Allocation Code (MAC) address. A standard for network components, the MAC address is typically provided by manufacturers in a 6-byte, colon delimited string. For example, 00:11:22:33:44:55.
- Type – Bits 15 through 12 is the type of board on which the node resides.
- Pos – Bits 11 through 8 identify the position of the board within the switch.
- Device – Bits 7 through 4 identify which device on the board has that node.
- Num – Bits 3 through 0 are numbers reserved for the programs which modifies the GUID. In most occurrences, the value is 0x2.

The following table provides values for Type, Pos, Device.

Board	Type	Position	Device
Fabric card	0xF	0x0 (Fabric card 0)–0x8 (Fabric card 8)	0xA (I4 chip 0)–0xB (I4 chip 1)
Line card	0x1	0x0 (Line card 0)–0x8 (Line card 8)	0xA (I4 chip 0)–0xD (I4 chip 3)
36-Port	0xA	0x0	0xA (I4 chip 0)
72-Port	0xB	0x0	0xA (I4 chip 0)–0xF (I4 chip 5)

Note – In the previous table, the I4 switch chips are identified with alphabetical characters (A, B, C, and D). For the input of switch-specific and CLIA commands, and in other tables, the I4 switch chips are identified by numeric characters (0, 1, 2, and 3), respectively.

For example, given the following output from the `ibswitches` command:

```
Switch : 0x0021283a83b112b2 ports 36 "Sun DCS 648 shmm1500 LC slot 2 switch 1"  
base port 0 lid 17 lmc 0
```

The GUID is 0x0021283a83b112b2 or 0021283a83b1 1 2 B 2. Using the information provided in this topic:

- The MAC address is 0x0021283a83b1 or 00:21:28:3A:83:B1.
- The type is 0x1, or a line card.
- The position is 0x2, or line card slot 2.
- The device is 0xB, or I4 chip 1.
- The number is 2.

Related Information

- *Switch Reference*, `ibnetdiscover` command
- *Switch Reference*, `ibnodes` command
- *Switch Reference*, `ibswitches` command
- *Switch Reference*, `ibhosts` command
- *Switch Reference*, `ibrouters` command
- [“Identify All Switches in the Fabric” on page 46](#)
- [“Display the Base GUIDs of a Fabric Card” on page 25](#)
- [“Display the Base GUIDs of a Line Card” on page 35](#)

Index

A

activate command, 70, 75, 85
activating
 CMC, 70
 fabric card, 75
 line card, 85
 power supply, 67
adding
 fabric card, 134
 line card, 135
 user, 65
addressing
 CMC, 4
 CXP connector, 7
 fabric card, 5
 line card, 6
 power supply, 3
administering the switch, 1

B

boot command, 60, 63

C

changing password
 root, 58
 user, 65
checkfans command, 19
checking
 CMC, 15
 LEDs, 17, 113
 power, 15
 status, 16
 temperature, 15
 error messages, 110
 fabric card, 18
 fault, 21
 IPMB state, 28

 LEDs, 29, 115
 link status, 26
 power, 22
 switch chip health, 27
 temperature, 22, 23
 voltage, 23

fan

 LEDs, 115
 speed, 19
 status, 19

LEDs, 112

line card, 30
 fault, 32
 IPMB state, 38
 LEDs, 39, 116
 link status, 36, 37
 power, 32
 switch chip health, 38
 temperature, 32, 33
 voltage, 33

logs, 110

port, 96

power supply, 13, 14
 LEDs, 112

switch characteristics, 40

checklinks command, 26, 36, 37, 78, 88

checkpwrfault command, 21, 32

checkswitches command, 27, 38

checkvoltages command, 23, 33

clear

 data counters, 96
 error counters, 96

CLIA command overview, 8

CMC

 activating, 70
 addressing, 4
 checking, 15

- LEDs, 17, 113
 - power, 15
 - status, 16
 - temperature, 15
- configuring identical addresses, 63
- deactivating, 70
- firmware
 - displaying, 17, 18
 - managing, 69
 - rebooting, 69
 - setting host name, 60
 - switchover, 71
- command
 - activate, 70, 75, 85
 - boot, 60, 63
 - checkfans, 19
 - checklinks, 26, 36, 37, 78, 88
 - checkpwrfault, 21, 32
 - checkswitches, 27, 38
 - checkvoltages, 23, 33
 - date, 41, 59
 - deactivate, 70, 76, 86
 - disableboard, 74, 83
 - disablepsu, 68
 - disablestby, 75, 84
 - disableswitchport, 80, 89
 - enableboard, 73, 83
 - enablehotinsert, 75, 85
 - enablepsu, 67
 - enablestby, 74, 84
 - enableswitchport, 79, 89
 - findport, 43
 - getbaseguid, 25, 35, 103, 104
 - getfruledstate, 17, 29, 39
 - getfwversion, 30, 40
 - getipmbstate, 28, 38
 - getlanconfig, 41
 - getpsufwver, 14
 - help, 59
 - ibcheckstate, 96
 - ibclearcounters, 96
 - ibclearerrors, 96
 - ibdatacounts, 51
 - ibdiagnet, 91, 94, 95
 - ibdiagpath, 92
 - ibhosts, 45
 - ibnetdiscover, 47
 - ibportstate, 49, 97, 98, 100, 101
 - ibswitches, 46, 55, 60
 - ibtracert, 47
 - ibv_devinfo, 44
 - mcmversion, 18
 - ofed_info, 57
 - opensm, 102, 103, 104
 - opensmd, 106
 - osmtest, 54, 92
 - passwd, 58
 - perfquery, 50
 - psustatus, 14
 - reboot, 63, 69
 - resetswitch, 80, 90
 - saveenv, 60, 63
 - sel, 112
 - sensordata, 15, 22, 32
 - session, 12
 - setenv, 60, 63
 - setfruledstate, 77, 78, 87, 88
 - setlanconfig, 61, 63
 - shmstatus, 16
 - showlogs, 110, 111
 - showpresent, 11, 13, 20, 31
 - showtemps, 23, 33
 - showunhealthy, 11
 - showvoltages, 23, 33
 - sminfo, 56
 - smpquery, 52, 53
 - switchover, 63, 71
 - telnet, 61
 - user, 41
 - user add, 65
 - user delete, 66
 - user passwd, 65
 - version, 17
- component
 - addressing, 2
- controlling
 - InfiniBand fabric, 91
 - Subnet Manager, 102
 - switch, 57
- creating
 - opensm.conf file, 106
- CXP connector
 - addressing, 7
 - mapping
 - from GUID, 43
 - to switch chip, 124

D

- data counters
 - clearing, 96
 - displaying, 51
- date
 - displaying, 41
 - setting, 59
- date command, 41, 59
- deactivate command, 70, 76, 86
- deactivating
 - CMC, 70
 - fabric card, 76
 - line card, 86
 - power supply, 68
- deleting
 - a user, 66
- diagnosing
 - a problem, 109
- diagnostics
 - fabric, 91
 - route, 92
- disableboard command, 74, 83
- disablepsu command, 68
- disablestby command, 75, 84
- disableswitchport command, 80, 89
- disabling
 - fabric card, 74, 75
 - line card, 83, 84
 - port, 80, 89, 100
- displaying
 - active sessions, 12
 - CMC firmware, 17, 18
 - components present, 11
 - date, 41
 - fabric card
 - base GUIDs, 25
 - firmware, 30
 - message log, 111
 - presence, 20
 - line card
 - base GUID, 35
 - firmware, 40
 - message log, 111
 - presence, 31
 - local HCA, 44
 - network management configuration, 41
 - node

- counters, 50, 51
- information, 52
- link status, 49
- OFED version, 57
- port information, 53
- power supply
 - firmware, 14
 - message log, 110
 - presence, 13
- route, 47
- Subnet Manager status, 56
- switch health, 11
- system event log, 112
- topology, 47
- user accounts, 41

E

- enableboard command, 73, 83
- enablehotinsert command, 75, 85
- enablepsu command, 67
- enablestby command, 74, 84
- enableswitchport command, 79, 89
- enabling
 - fabric card, 73, 74
 - links, 78
 - port, 79
 - line card, 83, 84
 - links, 88
 - port, 89
 - port, 101
- error counters
 - clearing, 96
 - displaying, 50
- event messages overview, 110
- example route through switch, 131

F

- fabric card
 - activating, 75
 - adding, 134
 - addressing, 5
 - checking, 18
 - fault, 21
 - IPMB state, 28
 - LEDs, 29, 115
 - link status, 26
 - power, 22

- switch chip health, 27
- temperature, 22, 23
- voltage, 23
- deactivating, 76
- disabling, 74, 75
- displaying
 - base GUIDs, 25
 - firmware, 30
 - message log, 111
 - presence, 20
- enabling, 73, 74
- hot insertion, 75
- LEDs
 - Locator off, 78
 - Locator on, 77
- link
 - enabling, 78
 - status, 26
- managing, 72
- port
 - disabling, 80
 - enabling, 79
- reset switch chip, 80
- restarting, 72

fan

- checking
 - LEDs, 115
 - speed, 19
 - status, 19

fault

- fabric card, 21
- line card, 32

findport command, 43

firmware

- CMC, 17, 18
- fabric card, 30
- line card, 40
- power supply, 14

G

- general tasks, 11
- getbaseguid command, 25, 35, 103, 104
- getfruledstate command, 17, 29, 39
- getfwversion command, 30, 40
- getipmbstate command, 28, 38
- getlanconfig command, 41
- getpsufwver command, 14

GUID

- fabric card, 25
- line card, 35
- mapping
 - to CXP connector, 43
 - to switch chip, 43
- nomenclature, 136

H

- hardware problems, 118
- HCA
 - identity, 45
 - information, 44
- help, 59
- help command, 59

I

- ibcheckstate command, 96
- ibclearcounters command, 96
- ibclearerrors command, 96
- ibdatacounts command, 51
- ibdiagnet command, 91, 94, 95
- ibdiagpath command, 92
- ibhosts command, 45
- ibnetdiscover command, 47
- ibportstate command, 49, 97, 98, 100, 101
- ibswitches command, 46, 55, 60
- ibtracert command, 47
- ibv_devinfo command, 44
- identifying
 - HCAs, 45
 - problem, 109
 - switches, 46

ILOM command overview, 9

- InfiniBand fabric
 - monitoring, 43
 - problems, 121
- InfiniBand software commands
 - overview, 10

IPMB

- state
 - fabric card, 28
 - line card, 38

L

- LEDs

- checking, 112
- CMC, 17, 113
- fabric card, 29, 115
 - Locator off, 78
 - Locator on, 77
- fan, 115
- line card, 39, 116
 - Locator off, 88
 - Locator on, 87
- link, 117
- NET MGT, 114
- power supply, 112
- line card
 - activating, 85
 - adding, 135
 - addressing, 6
 - checking, 30
 - fault, 32
 - IPMB state, 38
 - LEDs, 39, 116
 - link status, 36, 37
 - power, 32
 - switch chip health, 38
 - temperature, 32, 33
 - voltage, 33
 - deactivating, 86
 - disabling, 83, 84
 - displaying
 - base GUID, 35
 - firmware, 40
 - message log, 111
 - presence, 31
 - enabling, 83, 84
 - hot insertion, 85
 - LEDs
 - Locator off, 88
 - Locator on, 87
 - link
 - enabling, 88
 - status, 36, 37
 - managing, 81
 - port
 - disabling, 89
 - enabling, 89
 - reset switch chip, 90
 - restarting, 82
- link status
 - fabric card, 26
 - line card, 36, 37
- node, 49

M

- managing
 - CMC, 69
 - fabric card, 72
 - line card, 81
 - power supply, 66
- mapping
 - CXP to line card switch chip, 124
 - fabric card switch chip to XBOW, 128
 - GUID to CXP connector, 43
 - GUID to switch chip, 43
 - LED to line card switch chip, 124
 - LID to GUIDs, 54
 - line card switch chip to CXP, 130
 - line card switch chip to LEDs, 130
 - line card switch chip to XBOW, 126
 - switch chip to GUID, 55
 - switch chip to LID, 55
 - XBOW to fabric card switch chip, 127
 - XBOW to line card switch chip, 129
- mcmversion command, 18
- message log
 - fabric card, 111
 - line card, 111
 - power supply, 110
- monitoring
 - InfiniBand fabric, 43
 - switch, 10

N

- network management
 - configuration, 41
- node
 - counters, 50, 51
 - information, 52
 - link status, 49

O

- ofed_info command, 57
- opensm command, 102, 103, 104
- opensm.conf file, 106
- opensmd daemon, 106
- osmtest command, 54, 92
- overview
 - CLIA commands, 8

- event messages, 110
- ILOM commands, 9
- InfiniBand software commands, 10
- switch-specific commands, 9

P

- passwd command, 58
- password
 - root, 58
 - user, 65
- perfquery command, 50
- poor performing links, 94
- port
 - checking, 96
 - disabling, 100
 - enabling, 101
 - information, 53
 - resetting, 97
 - speed, 98
- power
 - CMC, 15
 - fabric card, 22
 - line card, 32
- power supply
 - activating, 67
 - addressing, 3
 - checking, 13
 - LEDs, 112
 - status, 14
 - deactivating, 68
 - displaying
 - firmware, 14
 - message log, 110
 - presence, 13
 - managing, 66
 - restarting, 67
- presence
 - components, 11
 - fabric card, 20
 - line card, 31
 - power supply, 13
- problem
 - diagnosing, 109
 - hardware, 118
 - identifying, 109
 - InfiniBand fabric, 121
- psustatus command, 14

R

- reboot command, 63, 69
- recovery errors, 95
- resetswitch command, 80, 90
- resetting
 - fabric card switch chip, 80
 - line card switch chip, 90
 - port, 97
- restarting
 - fabric card, 72
 - line card, 82
 - power supply, 67
- root node GUID file, 104
- route, 47
 - CXP to line card switch chip, 124
 - fabric card switch chip to XBOW, 128
 - line card switch chip to CXP, 130
 - line card switch chip to XBOW, 126
 - through switch, 124
 - XBOW to fabric card switch chip, 127
 - XBOW to line card switch chip, 129

S

- saveenv command, 60, 63
- sel command, 112
- sensordata command, 15, 22, 32
- session command, 12
- setenv command, 60, 63
- setfruledstate command, 77, 78, 87, 88
- setlanconfig command, 61, 63
- setting
 - CMC host name, 60
 - date, 59
 - hardware, 58
 - network management parameters, 61
 - port speed, 98
 - users, 64
- shmstatus command, 16
- showlogs command, 110, 111
- showpresent command, 11, 13, 20, 31
- showtemps command, 23, 33
- showunhealthy command, 11
- showvoltages command, 23, 33
- sminfo command, 56
- smpquery command, 52, 53

starting Subnet Manager, 102, 103, 106

status

CMC, 16

fan, 19

power supply, 14

Subnet Manager, 56

stopping Subnet Manager, 106

Subnet Manager

starting, 102, 103, 106

status, 56

stopping, 106

switch

administering, 1

controlling, 57

GUIDs, 136

identity, 46

monitoring, 10

startup, 133

troubleshooting, 117

switch chip

health

fabric card, 27

line card, 38

mapping

to GUID, 55

to LID, 55

switchover command, 63, 71

switch-specific commands

overview, 9

symbol errors, 95

T

telnet command, 61

temperature

CMC, 15

fabric card, 22, 23

line card, 32, 33

topology

change, 92

discovery, 47

troubleshooting the switch, 117

U

understanding

administrative commands, 1

component addition, 132

switch startup, 132

user

add, 65

deleting, 66

password, 65

user add command, 65

user command, 41

user delete command, 66

user passwd command, 65

V

version command, 17

voltage

fabric card, 23

line card, 33

X

XBOW connector mapping, 127, 129

