Sun Datacenter InfiniBand Switch 36

Administration Guide for Firmware Version 2.1

Part No.: E36269-01
March 2013, Revision A

Copyright © 2013, 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 END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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.

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. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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 © 2013, 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 END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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.

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. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

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.

Contents

Using This Documentation ix

Troubleshooting the Switch 1

Switch Hardware Problems 1

InfiniBand Fabric Problems 3

Network Management Troubleshooting Guidelines 6

Understanding Routing Through the Switch 7

Switch Chip Port to QSFP Connectors and Link LED Routes 7

QSFP Connectors and Link LEDs to Switch Chip Port Routes 8

Signal Route Through the Switch 9

Switch GUIDs Overview 9

Understanding Administrative Commands 13

Configuration Overview 13

Oracle ILOM Command Overview 15

Hardware Command Overview 16

InfiniBand Command Overview 16

Administering the Chassis 17

Monitoring the Chassis 17

- Display Switch General Health 18
- Display Power Supply Status 18
- Check Board-Level Voltages 19

- Display Internal Temperatures 20
- Display Fan Status 20
- Display Switch Environmental and Operational Data 21
- Display Flash Drive Information 22
- Display Free Memory and File System Space 23
- Display Chassis FRU ID 23
- Display Power Supply FRU ID 24
- Display Switch Firmware Versions 25

Oracle Auto Service Request 25

Controlling the Chassis 26

- Restart the Management Controller 26
- Change the `root` User Password 27

Administering the I4 Switch Chip 29

Monitoring the I4 Switch Chip 29

- Display the Switch Chip Port to QSFP Connector Mapping 30
- Locate an InfiniBand Device and Connector From the GUID and Port 31
- Locate a Switch Chip or Connector From the GUID and Port 31
- Display Switch Chip Boot Status 32
- Display Link Status 33
- Display Connectors Configured With Autodisable Functionality 33
- Display Switch Chip Port Status 34
- Display Switch Chip Port Counters 35

Controlling the I4 Switch Chip 36

- Disable a Switch Chip Port 36
- Enable a Switch Chip Port 37

Autodisabling Ports and Connectors 39

Autodisable Functionality 39

- Add Connectors to the Autodisable List 40
- Reenable Autodisabled Connectors 41
- Delete Connectors From the Autodisable List 42

Administering the Subnet Manager 45

Monitoring the Subnet Manager 45

- Display Subnet Manager Status 46
- Display Subnet Manager Location History 46
- Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm 47
- Display the Secret M_Key Configuration 48
- Verify M_Key Persistence 49
- Verify the Subnet Manager Configuration Integrity 50
- Display the Subnet Manager Log 51

Controlling the Subnet Manager 52

- Set the Subnet Manager Priority 52
- Set the Subnet Manager Prefix 53
- Enable Subnet Manager Controlled Handover 54
- Set the Subnet Manager Readable Management Key 55
- Clear the Subnet Manager Readable Management Key 56
- Set the Subnet Manager Routing Algorithm 57
- Enable the Subnet Manager 58
- Disable the Subnet Manager 59

Configuring Secure Fabric Management 60

Secret M_Key Overview 60

`smsubnetprotection` Command 61

- Prepare for Secret M_Key Functionality 62

Configuring Secret M_Keys 63

- Create the Secret M_Key Configuration 63

- Add Secret M_Keys to the Configuration 65
- Set the Secret M_Key to Use 68
- Delete Secret M_Keys From the Configuration 69
- Configure a Local Secret M_Key 71

- Set the Replication Password 72
- Enable or Disable Secret M_Key Persistence 72
- Enable or Disable Secret M_Key Functionality 73

Considering Situations In Regard to Secret M_Key 74

- Add a New Subnet Manager to the Secret M_Key Fabric 74
- Merge Subnets Into One Secret M_Key Fabric 75

Secret M_Key and Firmware Upgrade and Downgrade Guidelines 76

Administering the InfiniBand Fabric 79

Monitoring the InfiniBand Fabric 79

Displaying Information About Nodes 80

- Identify All Switches in the Fabric 80
- Identify All CAs in the Fabric 81
- Display the InfiniBand Fabric Topology (Simple) 82
- Display the InfiniBand Fabric Topology (Detailed) 82
- Display a Route Through the Fabric 83
- Display the Link Status of a Node 84
- Display Counters for a Node 85
- Display Low-Level Detailed Information About a Node 86
- Display Low-Level Detailed Information About a Port 87

Displaying Fabric Functionality Configuration Information 89

- Display the InfiniBand Fabric Partition Configuration 89
- Display the InfiniBand Fabric Configuration 90
- Display the `smnodes` List 91

- Display the Fabric Element Configuration 92
- Display the IPoIB Interface 92
- Display Configured Secret M_Keys 93

Controlling the InfiniBand Fabric 93

Investigating Nodes 94

- Perform Comprehensive Diagnostics for the Entire Fabric 94
- Determine Changes to the InfiniBand Fabric Topology 95
- Find 1x, SDR, or DDR Links in the Fabric 96
- Determine Which Links Are Experiencing Significant Errors 96

Controlling Nodes 97

- Clear Data and Error Counters 97
- Reset a Port 98
- Set Port Speed 99
- Disable or Enable an External Port 100

Partitioning the InfiniBand Fabric 101

`smpartition` Command 102

`smpartition` Command Guidelines 103

Partitions and P_Keys 103

- Determine the Partitions and P_Keys 104
- Create a User Partition 105
- Add or Remove a Port From a Partition 107
- Modify a Partition or Port 109
- Delete a Partition 112
- Remove User Partitions for Firmware Downgrade 113

Correlating Fabric Elements 114

`createfabric` Command 115

- Create a Fabric Element Configuration 115
- Add or Delete an Element From the Configuration 117

▼ Modify an Element of the Configuration 118
Configuring the Fabric Director Node List 120
Fabric Director and Fabric Elements 120
`fdconfig` Command Overview 121
▼ Create a Fabric Configuration 122
▼ Add or Remove an Element From the Fabric Configuration 124
▼ Modify an Element of the Configuration 126
Performing Supportive Tasks 127
▼ Create the `smnodes` List 128
▼ Update the `smnodes` List 129
▼ Create the IPoIB Interface 130
▼ Delete the IPoIB Interface 131

Index 133

Using This Documentation

This administration guide provides detailed procedures that describe administration of the Sun Datacenter InfiniBand Switch 36 Firmware Version 2.1 from Oracle. This document is written for technicians, system administrators, and users who have advanced experience administrating InfiniBand fabric hardware.

- "Product Notes" on page ix
- "Related Documentation" on page x
- "Feedback" on page x
- "Access to Oracle Support" on page x

Product Notes

For late-breaking information and known issues about this product, refer to the product notes at:

`http://docs.oracle.com/cd/E36265_01`

Related Documentation

Documentation	Links
Sun Datacenter InfiniBand Switch 36 Firmware Version 2.1	http://docs.oracle.com/cd/E36265_01
Oracle Integrated Lights Out Manager (ILOM) 3.0	http://docs.oracle.com/cd/E19860-01
All Oracle products	http://www.oracle.com/documentation

Feedback

Provide feedback on this documentation at:

http://www.oracle.com/goto/docfeedback

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs visit if you are hearing impaired.

Troubleshooting the Switch

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

- "Switch Hardware Problems" on page 1
- "InfiniBand Fabric Problems" on page 3
- "Network Management Troubleshooting Guidelines" on page 6
- "Understanding Routing Through the Switch" on page 7
- "Switch GUIDs Overview" on page 9

Related Information

- "Understanding Administrative Commands" on page 13
- "Administering the Chassis" on page 17
- "Administering the I4 Switch Chip" on page 29
- "Administering the InfiniBand Fabric" on page 79
- "Administering the Subnet Manager" on page 45

Switch Hardware Problems

This table lists situations that might occur with switch hardware and corrective steps that you can take to resolve the problem.

Situation	Corrective Steps
The Attention LED on a power supply is lit or the power supply seems dysfunctional.	1. Check the power supply status. See "Display Power Supply Status" on page 18. 2. Unplug the respective power cord, wait 15 minutes, then reattach the power cord. 3. If the previous steps do not rectify the situation, replace the power supply. Refer to *Switch Service*, servicing power supplies.

Situation	Corrective Steps
There is no network communication with the management controller.	1. Verify that the management controller is powered without faults. Refer to *Switch Service*, checking the chassis status LEDs. 2. Verify that there is a LinkUp condition at the NET MGT port. Refer to *Switch Service*, checking the NET MGT port LEDs. 3. Verify that the DHCP server is providing the IP address you are using to access the management controller. See DHCP server documentation. 4. If you are able to access the management controller through the USB connector, restart the management controller. Refer to *Switch Installation*, accessing the management controller and "Restart the Management Controller" on page 26. 5. If you are unable to access the management controller through the USB connector, power cycle the switch. Refer to *Switch Service*, powering off power supplies and *Switch Service*, powering on power supplies. 6. If the previous steps do not rectify the situation, replace the switch. Refer to *Switch Service*, removing the switch from the rack. Refer to *Switch Installation*, installing the switch into the rack.
The Attention LED on the switch chassis is lit or the management controller seems dysfunctional.	1. If you are unable to access the management controller, power cycle the switch. Refer to *Switch Service*, powering off the power supply and *Switch Service*, powering on the power supply. 2. If you are able to access the management controller, restart the management controller. See "Restart the Management Controller" on page 26. 3. Check overall switch health. See "Display Switch General Health" on page 18 4. Verify that the switch is within operating temperatures and voltages. See "Display Switch Environmental and Operational Data" on page 21. 5. If the previous steps do not rectify the situation, replace the switch. Refer to *Switch Service*, removing the switch from the rack. Refer to *Switch Installation*, installing the switch into the rack.
The Attention LED on a fan is lit or the fan seems dysfunctional.	1. Check the fan speed. See "Display Fan Status" on page 20. 2. If the previous step does not rectify the situation, replace the fan. Refer to *Switch Service*, servicing fans. 3. If the previous step does not rectify the situation, install the fan into another available slot. Refer to *Switch Service*, servicing fans. 4. If no other slots are available, replace the switch. Refer to *Switch Service*, removing the switch from the rack. Refer to *Switch Installation*, installing the switch into the rack.

Related Information

- "InfiniBand Fabric Problems" on page 3
- "Network Management Troubleshooting Guidelines" on page 6
- "Understanding Routing Through the Switch" on page 7
- "Switch GUIDs Overview" on page 9

InfiniBand Fabric Problems

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

Situation	Corrective Steps
After installation, no links are operational.	1. Verify that there is at least one Subnet Manager active on the InfiniBand fabric. See "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47. 2. If no Subnet Manager is active, start the Subnet Manager within the switch. Refer to *Switch Installation*, staring the Subnet Manager. 3. If the previous steps do not rectify the situation, restart the Subnet Manager. See "Disable the Subnet Manager" on page 59 and "Enable the Subnet Manager" on page 58.
After installation, not all links are operational.	1. Determine which links are nonoperational. See "Display Link Status" on page 33. 2. For links that are "`Down`", disable and re-enable the respective ports. See "Disable a Switch Chip Port" on page 36 and "Enable a Switch Chip Port" on page 37. 3. If the previous steps do not rectify the situation, disable the respective port. See "Disable a Switch Chip Port" on page 36.
There was a power outage during a firmware update.	1. If you are able to access the management controller, restart the management controller. See "Restart the Management Controller" on page 26. 2. If you are unable to access the management controller, power cycle the switch. Refer to *Switch Service*, removing the switch from the rack. Refer to *Switch Installation*, installing the switch into the rack. 3. Reperform the firmware upgrade. Refer to *Switch Remote Management*, upgrading the switch firmware.

Situation	Corrective Steps
Performance of the InfiniBand fabric seems diminished.	1. Determine if there are errors or problems with the InfiniBand fabric. See: "Perform Comprehensive Diagnostics for the Entire Fabric" on page 94 "Find 1x, SDR, or DDR Links in the Fabric" on page 96 "Determine Which Links Are Experiencing Significant Errors" on page 96 2. Locate the affected nodes by the GUID provided in the output of the `ibdiagnet` command. See "Locate a Switch Chip or Connector From the GUID and Port" on page 31. 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. Refer to *Switch Service*, servicing InfiniBand cables. 4. If the problem still remains at the cable connection, disable and re-enable the respective port and repeat Step 1. See "Disable or Enable an External Port" on page 100. Temporary solution: • If the problem still remains, disable the affected port. See "Disable or Enable an External Port" on page 100. Permanent solution: • If the problem still remains, replace the affected component or the switch. Refer to *Switch Service*, servicing InfiniBand cables. Refer to remote port's documentation for replacement procedures. Refer to *Switch Service*, removing the switch from the rack. Refer to *Switch Installation*, installing the switch into the rack.
An InfiniBand Link LED is blinking.	1. Disconnect and properly reconnect both ends of the respective InfiniBand cable. See *Switch Service*, servicing the InfiniBand cables. 2. If the LED is still blinking, determine the significance of the errors through use of the `ibdiagnet` command. See "Determine Which Links Are Experiencing Significant Errors" on page 96. 3. Determine which connectors map to the affected link by deconstructing the node's GUID and port. See "Locate a Switch Chip or Connector From the GUID and Port" on page 31. 4. If some of the links are running at 1x or SDR, use that situation elsewhere in this table to rectify the problem. 5. Disable and re-enable the respective ports. See "Disable or Enable an External Port" on page 100. 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 Step 2. 7. Depending upon what does or does not rectify the problem, replace that component. Refer to *Switch Service*, servicing the InfiniBand cables. Refer to remote port's documentation for replacement procedures.

Situation	Corrective Steps

Some InfiniBand links are running at 1x or SDR.	For a temporary solution:
	1. Identify the suspect links using the `ibdiagnet` command.
	See "Find 1x, SDR, or DDR Links in the Fabric" on page 96. Look for text like this:
	`-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`
	2. Determine which connectors map to the affected link by deconstructing the node's GUID and port.
	See "Locate a Switch Chip or Connector From the GUID and Port" on page 31.
	3. Verify the cable connection at both ends.
	Refer to *Switch Service*, servicing the InfiniBand cables.
	4. Disable and re-enable the respective ports.
	See "Disable or Enable an External Port" on page 100.
	5. If the previous steps do not rectify the problem, disable the port.
	See "Disable or Enable an External Port" on page 100.
	For a permanent solution:
	1. Perform the steps for a temporary solution, Steps 1 to Step 4.
	2. Swap the cable with a known good cable or reconnect the cable to a known good remote port, and repeat from Step 1.
	3. Depending upon what does or does not rectify the problem, replace that component or the switch.
	Refer to *Switch Service*, servicing the InfiniBand cables.
	Refer to the remote port's documentation for replacement procedures.
	Refer to *Switch Service*, removing the switch from the rack.
	Refer to *Switch Installation*, installing the switch into the rack.

There are errors on some InfiniBand links.	1. Clear the error counters.
	See "Clear Data and Error Counters" on page 97.
	2. Start a fabric stress test.
	3. Identify the suspect links using the `ibdiagnet` command.
	See "Determine Which Links Are Experiencing Significant Errors" on page 96. Look for text like this:
	`-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)`
	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 commands provides only GUID and port, not switch chip or QSFP connectors.	1. You can find the location of a node in the switch by deconstructing the node's GUID and port. See "Locate a Switch Chip or Connector From the GUID and Port" on page 31. 2. Use the `dcsport` command to provide port-to-connector and connector-to-port mapping. See "Display the Switch Chip Port to QSFP Connector Mapping" on page 30.

Related Information

- "Switch Hardware Problems" on page 1
- "Network Management Troubleshooting Guidelines" on page 6
- "Understanding Routing Through the Switch" on page 7
- "Switch GUIDs Overview" on page 9

Network Management Troubleshooting Guidelines

Follow these guidelines when experiencing difficulty connecting to the switch through its NET MGT port.

- Ensure that the `state` property of the `/SP/network` Oracle ILOM target is set to `enabled`. For example:

```
-> set /SP/network state=enabled
```

- Verify the integrity of the physical connection between the NET MGT port of the server, the network, and the management device.
- Ensure that at least one of the appropriate network services are enabled in Oracle ILOM:

Service	Command
HTTP	`set /SP/services/http secureredirect=disabled servicestate=enabled`
HTTPS	`set /SP/services/https servicestate=enabled`

- Use an industry-standard network diagnostic tool, like IPv4 Ping or Traceroute to test the network connections between the server and the management device.

Related Information

- “Switch Hardware Problems” on page 1
- “InfiniBand Fabric Problems” on page 3
- “Understanding Routing Through the Switch” on page 7
- “Switch GUIDs Overview” on page 9

Understanding Routing Through the Switch

The tables in these topics describe the routing through the switch. The first table maps the switch chip port to a QSFP connector. The second table provides a reverse map. When command output provides a switch chip port, you can use these tables to determine the route that link is following.

- “Switch Chip Port to QSFP Connectors and Link LED Routes” on page 7
- “QSFP Connectors and Link LEDs to Switch Chip Port Routes” on page 8
- “Signal Route Through the Switch” on page 9

Related Information

- “Switch Hardware Problems” on page 1
- “InfiniBand Fabric Problems” on page 3
- “Network Management Troubleshooting Guidelines” on page 6
- “Switch GUIDs Overview” on page 9

Switch Chip Port to QSFP Connectors and Link LED Routes

Port	Connector	Port	Connector	Port	Connector	Port	Connector
1	17A	10	13B	19	0B	28	4A
2	17B	11	12A	20	0A	29	5B
3	16A	12	12B	21	1B	30	5A
4	16B	13	9B	22	1A	31	8A

Port	Connector	Port	Connector	Port	Connector	Port	Connector
5	15A	14	9A	23	2B	32	8B
6	15B	15	10B	24	2A	33	7A
7	14A	16	10A	25	3B	34	7B
8	14B	17	11B	26	3A	35	6A
9	13A	18	11A	27	4B	36	6B

Related Information

- "Display the Switch Chip Port to QSFP Connector Mapping" on page 30
- "Display Link Status" on page 33
- "Display Switch Chip Port Status" on page 34
- "QSFP Connectors and Link LEDs to Switch Chip Port Routes" on page 8
- "Signal Route Through the Switch" on page 9

QSFP Connectors and Link LEDs to Switch Chip Port Routes

Connector Group	Connector A	Connector B	Connector Group	Connector A	Connector B
0	20	19	9	14	13
1	22	21	10	16	15
2	24	23	11	18	17
3	26	25	12	11	12
4	28	27	13	9	10
5	30	29	14	7	8
6	35	36	15	5	6
7	33	34	16	3	4
8	31	32	17	1	2

Related Information

- "Display the Switch Chip Port to QSFP Connector Mapping" on page 30
- "Display Link Status" on page 33

- "Display Switch Chip Port Status" on page 34
- "Switch Chip Port to QSFP Connectors and Link LED Routes" on page 7
- "Signal Route Through the Switch" on page 9

Signal Route Through the Switch

By combining the information from the tables in "Understanding Routing Through the Switch" on page 7, 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 connector 2A. The LED blinks.
2. Using "QSFP Connectors and Link LEDs to Switch Chip Port Routes" on page 8, it is determined that the link routes to the switch chip through port 24.
3. The Subnet Manager instructs the switch chip to use port 4 to forward the link.
4. Using "Switch Chip Port to QSFP Connectors and Link LED Routes" on page 7, it is determined that the link routes to connector 16B. The LED blinks.

Related Information

- "Switch Chip Port to QSFP Connectors and Link LED Routes" on page 7
- "QSFP Connectors and Link LEDs to Switch Chip Port Routes" on page 8
- "Display the Switch Chip Port to QSFP Connector Mapping" on page 30
- "Display a Route Through the Fabric" on page 83

Switch GUIDs Overview

Global unique identifiers (GUIDs) are unique 64-bit strings that identify nodes such as switches and channel adapters. For the Sun Datacenter InfiniBand switches and this switch, GUIDs are modified to identify the node's role and location. This table describes the GUID's structure.

63 16	15 12	11 8	7 4	3 0
MAC (48 bits)	Type	Pos	Device	Num

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 modify the GUID. In most occurrences, the value is 0x2.

This 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)
Gateway	0xC	0x0	0xA (I4 chip 0), 0x0 (BX chip 0), 0x4 (BX chip 1)
36-Port	0xA	0x0	0xA (I4 chip 0)
72-Port	0xB	0x0	0xA (I4 chip 0)–0xF (I4 chip 5)

For example, given this output from the `ibswitches` command:

```
[16] -> switch port {0x0021283a831da0a2}[21] lid 214-214 "Sun DCS 36 QDR FC
switch 2.0"
```

The GUID is 0x0021283a831da0a2 or 0021283A831D A 0 A 2. Using the information provided in this topic:

- The MAC address is 0x0021283A831D or 00:21:28:3A:83:1D.
- The type is 0xA, or a 36-port board.
- The position is 0x0.
- The device is 0xA, or a solitary I4 chip.
- The number is 2.

Related Information

- *Switch Reference*, `ibnetdiscover` command
- *Switch Reference*, `ibnodes` command
- *Switch Reference*, `ibswitches` command
- *Switch Reference*, `ibhosts` command

- "Identify All Switches in the Fabric" on page 80
- "Switch Hardware Problems" on page 1
- "InfiniBand Fabric Problems" on page 3
- "Understanding Routing Through the Switch" on page 7

Understanding Administrative Commands

These topics provide an overview of administrative tasks and the command sets to perform those tasks. Administering the switch requires accessing the management controller.

- "Configuration Overview" on page 13
- "Oracle ILOM Command Overview" on page 15
- "Hardware Command Overview" on page 16
- "InfiniBand Command Overview" on page 16

Related Information

- "Troubleshooting the Switch" on page 1
- "Administering the Chassis" on page 17
- "Administering the I4 Switch Chip" on page 29
- "Administering the InfiniBand Fabric" on page 79
- "Administering the Subnet Manager" on page 45

Configuration Overview

After you have installed the switch, you can configure its functionality. This table lists functions of the switch to configure, in the order of dependency. If you do not want to configure a particular functionality, skip that step.

Step	Functionality	Description	Links
1.	Configure the Subnet Manager.	Configure the priority.	"Set the Subnet Manager Priority" on page 52
		Configure the prefix.	"Set the Subnet Manager Prefix" on page 53
		Enable controlled handover.	"Enable Subnet Manager Controlled Handover" on page 54
		Configure the M_Key.	"Set the Subnet Manager Readable Management Key" on page 55
		Set the routing algorithm, if not Fat Tree.	"Set the Subnet Manager Routing Algorithm" on page 57
2.	Configure the InfiniBand fabric.	Create a fabric configuration.	"Create a Fabric Configuration" on page 122
3.	Secure fabric management.	Create the `smnodes` list.	"Create the smnodes List" on page 128
		(optional) Configure IPOIB functionality.	"Create the IPoIB Interface" on page 130
		Create the fabric element configuration.	"Create a Fabric Element Configuration" on page 115
		Set the replication password.	"Set the Replication Password" on page 72
		Enable M_Key persistence.	"Enable or Disable Secret M_Key Persistence" on page 72
		Configure secret M_Keys.	"Configuring Secret M_Keys" on page 63
		Enable secret M_Keys.	"Enable or Disable Secret M_Key Functionality" on page 73
4.	Partition the InfiniBand fabric.	Collect GUIDs.	"Identify All Switches in the Fabric" on page 80 "Identify All CAs in the Fabric" on page 81
		Update the `smnodes` file.	"Update the smnodes List" on page 129
		Create the partition information file.	"Determine the Partitions and P_Keys" on page 104
		Configure partitions and P_Keys.	"Create a User Partition" on page 105
		Add GUIDs to the partitions.	"Add or Remove a Port From a Partition" on page 107
9.	Save the configuration.	Back up the configuration for restoration at a later time, should it become necessary.	*Switch Remote Management*, backing up the configuration.

Related Information

- "Oracle ILOM Command Overview" on page 15
- "Hardware Command Overview" on page 16
- "InfiniBand Command Overview" on page 16

Oracle ILOM Command Overview

The Oracle ILOM CLI, web, SNMP, and IPMI interfaces enable additional administration features and capabilities. The Oracle ILOM CLI, SNMP, and IPMI interfaces use remote clients for command-line administration of many Oracle ILOM targets. Alternatively, the Oracle ILOM web interface permits point-and-click administration of the Oracle ILOM components and services.

After you log in to the `ilom-admin` account, the Oracle ILOM prompt (`->`) appears, and you can enter Oracle ILOM commands in this format:

`->` *command* `[`*option*`][`*target*`][`*property=value*`]...`

Information about Oracle ILOM support of the switch is available in the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Supplement for the Sun Datacenter InfiniBand Switch 36*, available online at:

`http://www.oracle.com/pls/topic/lookup?ctx=E36265-01`

Overall information about Oracle ILOM 3.0 is available online at:

`http://www.oracle.com/pls/topic/lookup?ctx=E19860-01`

Related Information

- *Switch Remote Management*, understanding Oracle ILOM commands
- "Hardware Command Overview" on page 16
- "InfiniBand Command Overview" on page 16

Hardware Command Overview

The management controller uses a simplified Linux OS and file system. From the `FabMan@`*switch_name*`->` prompt on the management controller, you can type hardware commands to perform administrative and management tasks and configure partitions, the InfiniBand fabric, and the Subnet Manager.

After you log in to the `ilom-admin` account, use the `show /SYS/Fabric_Mgmt` command to obtain the restricted Linux shell prompt `FabMan@`*switch_name*`->`. You can now enter hardware commands in this format:

`FabMan@`*switch_name*`->`*command* `[`*arguments*`][`*arguments*`]...`

Related Information

- *Switch Reference*, understanding hardware commands
- "Oracle ILOM Command Overview" on page 15
- "InfiniBand Command Overview" on page 16

InfiniBand Command Overview

The InfiniBand commands are a means of monitoring and controlling aspects of the InfiniBand fabric. These commands are also installed on and run from the management controller, which is also the host of the Subnet Manager. Use of these commands requires thorough knowledge of InfiniBand architecture and technology.

After you log in to the `ilom-admin` account, use the `show /SYS/Fabric_Mgmt` command to obtain the restricted Linux shell prompt `FabMan@`*switch_name*`->`. You can now enter InfiniBand commands in this format:

`FabMan@`*switch_name*`->`*command* `[`*option*`][`*option*`] ...`

Related Information

- *Switch Reference*, understanding InfiniBand commands
- "Oracle ILOM Command Overview" on page 15
- "Hardware Command Overview" on page 16

Administering the Chassis

These topics describe the administration of the switch.

- "Monitoring the Chassis" on page 17
- "Controlling the Chassis" on page 26

Related Information

- "Troubleshooting the Switch" on page 1
- "Understanding Administrative Commands" on page 13
- "Administering the I4 Switch Chip" on page 29
- "Administering the InfiniBand Fabric" on page 79
- "Administering the Subnet Manager" on page 45

Monitoring the Chassis

These topics enable you to display and check the operation and status of the switch.

- "Display Switch General Health" on page 18
- "Display Power Supply Status" on page 18
- "Check Board-Level Voltages" on page 19
- "Display Internal Temperatures" on page 20
- "Display Fan Status" on page 20
- "Display Switch Environmental and Operational Data" on page 21
- "Display Flash Drive Information" on page 22
- "Display Free Memory and File System Space" on page 23
- "Display Chassis FRU ID" on page 23
- "Display Power Supply FRU ID" on page 24
- "Display Switch Firmware Versions" on page 25
- "Oracle Auto Service Request" on page 25

Related Information

- "Monitoring the I4 Switch Chip" on page 29
- "Monitoring the InfiniBand Fabric" on page 79
- "Monitoring the Subnet Manager" on page 45
- "Controlling the Chassis" on page 26

▼ Display Switch General Health

An easy way to perform a quick check of the switch's operation is with the `showunhealthy` command.

- **On the management controller, type.**

```
FabMan@switch_name->showunhealthy
OK - No unhealthy sensors
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `showunhealthy` command
- "Display Switch Environmental and Operational Data" on page 21

▼ Display Power Supply Status

The `checkpower` command performs a simple pass-fail test on the power supplies.

1. **On the management controller, type.**

```
FabMan@switch_name->checkpower
PSU 0 present status: OK
PSU 1 present status: Alert
FabMan@switch_name->
```

2. **If you see the words `Alert` or `Not Present` in the command output, perform these steps for the affected power supply.**

   a. **Remove the power cord and affected power supply from the switch chassis.**

      Refer to *Switch Service*, powering off the power supply and *Switch Service*, removing the power supply.

   b. **Wait one minute.**

c. **Reinstall the power supply.**

   Refer to *Switch Service*, installing the power supply.

d. **Reattach the power cord.**

   Refer to *Switch Service*, powering on the power supply.

e. **If the power supply's Attention LED lights or the `checkpower` command still reports `Alert` or `Not Present` for the power supply, replace the power supply.**

   Refer to *Switch Service*, servicing the power supplies.

Related Information

- *Switch Reference*, `checkpower` command
- "Check Board-Level Voltages" on page 19
- "Display Switch Environmental and Operational Data" on page 21

▼ Check Board-Level Voltages

When you want to know if the voltages used by the various components within the switch are within nominal values, use the `checkvoltages` command. If a voltage deviates more than 10% of its nominal value, there is a problem. This check also reports if a battery has failed.

- **On the management controller, type.**

```
FabMan@switch_name->checkvoltages
Voltage ECB OK
Measured 3.3V Main = 3.28 V
Measured 3.3V Standby = 3.40 V
Measured 12V = 11.90 V
Measured 5V = 4.99 V
Measured VBAT = 3.01 V
Measured 2.5V = 2.49 V
Measured 1.8V = 1.78 V
Measured I4 1.2V = 1.22 V
All voltages OK
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `checkvoltages` command
- "Display Power Supply Status" on page 18
- "Display Switch Environmental and Operational Data" on page 21

▼ Display Internal Temperatures

If you are concerned with the internal temperature of the switch, you can use the `showtemps` command to display nominal and measured temperatures. If there is a great deviation, there might not be enough cooling.

- **On the management controller, type.**

```
FabMan@switch_name->showtemps
Back temperature 29
Front temperature 29
SP temperature 41
Switch temperature 32, maxtemperature 32
All temperatures OK
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `showtemps` command
- "Display Fan Status" on page 20
- "Display Switch Environmental and Operational Data" on page 21

▼ Display Fan Status

The output of the `getfanspeed` command can help you determine fan speed inconsistencies, which might indicate a future failure.

- **On the management controller, type.**

```
FabMan@switch_name->getfanspeed
Fan 0 not present
Fan 1 running at rpm 12099
Fan 2 running at rpm 11772
Fan 3 running at rpm 11772
Fan 4 not present
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `getfanspeed` command
- "Display Switch Environmental and Operational Data" on page 21

▼ Display Switch Environmental and Operational Data

The `env_test` command enables you to perform several investigative tasks, helping you determine the internal environment and operational status of the switch.

- **On the management controller, type.**

```
FabMan@switch_name->env_test
Environment test started:
Starting Environment Daemon test:
Environment daemon running
Environment Daemon test returned OK
Starting Voltage test:
Voltage ECB OK
Measured 3.3V Main = 3.28 V
Measured 3.3V Standby = 3.37 V
Measured 12V = 12.06 V
Measured 5V = 5.03 V
Measured VBAT = 3.25 V
Measured 2.5V = 2.49 V
Measured 1.8V = 1.78 V
Measured I4 1.2V = 1.22 V
Voltage test returned OK
Starting PSU test:
PSU 0 present
PSU 1 present
PSU test returned OK
Starting Temperature test:
Back temperature 28
Front temperature 28
SP temperature 42
Switch temperature 32, maxtemperature 32
Temperature test returned OK
Starting FAN test:
Fan 0 not present
Fan 1 running at rpm 11881
Fan 2 running at rpm 11663
Fan 3 running at rpm 11881
Fan 4 not present
FAN test returned OK
Starting Connector test:
Connector test returned OK
Starting Onboard ibdevice test:
Switch OK
All Internal ibdevices OK
Onboard ibdevice test returned OK
```

```
Starting SSD test:
SSD test returned OK
Environment test PASSED
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `env_test` command

▼ Display Flash Drive Information

Included in the output of the `showdisk` command is the estimated remaining life of the SSD drive, which can help in troubleshooting.

- **On the management controller, type.**

```
FabMan@gateway_name->showdisk
 Device Model    :  512MB NANDrive
 Serial Number   :  0000000000A12o3ANCtx
 Firmware Version:  ND B642J

 SMART supported :  yes
 SMART enabled   :  yes

 Average age (4 groups):
  group 0:    221
  group 1:    190
  group 2:   3098
  group 3:      9

 Block erase count (4 groups):
  group 0:    216900
  group 1:    190567
  group 2:   3104711
  group 3:      9097
  Total  :   3521275

 Total physical blocks: 4096

 Estimated Life Usage    :  3.098 %

 Estimated Remaining Life: 96.902 %

FabMan@gateway_name->
```

Related Information

- *Switch Reference*, `showdisk` command
- "Display Switch Environmental and Operational Data" on page 21

▼ Display Free Memory and File System Space

The `showfree` command gives output similar to the Linux `free` and `df` commands and can help with drive housekeeping chores.

- **On the management controller, type.**

```
FabMan@switch_name->showfree -md
Memory:                  463.7 MB   93.1%

Filesystems:
 /dev/hda2    /           197.3 MB   41.9%
 tmpfs        /dev/shm    249.1 MB  100.0%
 tmpfs        /tmp        248.2 MB   99.6%
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `showfree` command
- "Display Switch Environmental and Operational Data" on page 21

▼ Display Chassis FRU ID

The `showfruinfo` command displays switch chassis FRU ID information.

- **On the management controller, type.**

```
FabMan@switch_name->showfruinfo
Sun_Man1R:
 UNIX_Timestamp32              :  Fri Dec 14 01:24:28 2012
 Sun_Fru_Description           :  ASSY,NM2-36P
 Vendor_ID_Code                :  01 2C
 Vendor_ID_Code_Source         :  01
 Vendor_Name_And_Site_Location :  Celestica San Jose
 Sun_Part_Number               :  5111232
 Sun_Serial_Number             :  0110SJC-09183P0020
 Serial_Number_Format          :  4V3F1-2Y2W2X4S
 Initial_HW_Dash_Level         :  07
 Initial_HW_Rev_Level          :  01
```

```
 Sun_Fru_Shortname              :  NM2, 36 ports
 Sun_Hazard_Class_Code          :  Y
 Sun_SpecPartNo                 :  885-1507-08
Sun_FRU_LabelR:
 Sun_Serial_Number              :  0924AK2285
 FRU_Part_Dash_Number           :  541-3495-03
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `showfruinfo` command
- "Display Power Supply FRU ID" on page 24

▼ Display Power Supply FRU ID

The `showpsufru` command displays power supply FRU ID information.

- **On the management controller, type.**

```
FabMan@switch_name->showpsufru slot
```

where *slot* is the power supply slot (`0` or `1`). For example:

```
FabMan@switch_name->showpsufru 0
 Sun_SpecPartNo           :  885-1165-02
 UNIX_Timestamp32         :  Sun Jan  3 15:35:39 2010
 Sun_Part_Number          :  3002143
 Sun_Serial_No            :  BF15WA
 Vendor_ID_Code           :  02a2
 Initial_HW_Dash_Level    :  02
 PSU_Voltage 1            :  0x04b0 (1200)
 PSU_Current 1            :  0x186a (6250)
 PSU_Voltage 2            :  0x014a (330)
 PSU_Current 2            :  0x012c (300)
 Sun_Hazard_Class_Code    :  Y
 IPMI_Board_Manufacturer  :  EMERSON
 IPMI_Board_Product_Name  :  A237
 IPMI_Board_Serial_Number:  1357ZHO-0952BF15WA
 IPMI_Board_Part_Number   :  300-2143-02
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `showpsufru` command
- "Display Chassis FRU ID" on page 23

▼ Display Switch Firmware Versions

In time, newer or updated switch firmware might become available. The `version` command displays the versions of the firmware within the switch chassis.

- **On the management controller, type.**

```
FabMan@switch_name->version
SUN DCS 36p version: 2.1.2-1
Build time: Dec  7 2012 09:33:54
SP board info:
Manufacturing Date: 2009.02.20
Serial Number: "NCD2T0307"
Hardware Revision: 0x0100
Firmware Revision: 0x0102
BIOS version: NOW1R112
BIOS date: 04/24/2009
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `version` command
- *Switch Remote Management*, `version` command

Oracle Auto Service Request

Oracle Auto Service Request is a secure, customer-installable support feature that automatically generates a service request for specific hardware faults. The software, available as a feature of the Oracle hardware warranty and Oracle Premier Support for Systems, can improve system availability through expedited diagnostics and priority service request handling.

Please contact Oracle support for details on setting up ASR on your switch.

Related Information

- "Troubleshooting the Switch" on page 1
- *Switch Service*, detecting and managing faults

Controlling the Chassis

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

Note – To use the commands described in these topics, you must access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Restart the Management Controller" on page 26
- "Change the root User Password" on page 27

Related Information

- "Controlling the I4 Switch Chip" on page 36
- "Controlling the InfiniBand Fabric" on page 93
- "Controlling the Subnet Manager" on page 52
- "Monitoring the Chassis" on page 17

▼ Restart the Management Controller

Should the management controller enter an indeterminate state, you can restart it. The Subnet Manager and other services are also restarted. Additionally, the links to the I4 switch chip are disabled and then re-enabled when the management controller reboots.

Note – This procedure severs any management console link to the management controller. You must reaccess the management controller to regain administrative control.

- **At the Oracle ILOM prompt, type.**

```
-> reset /SP
Are you sure you want to reset /SP (y/n)? y
Performing reset on /SP
Broadcast message from root (Wed Sep  5 03:15:55 2012):
The system is going down for reboot NOW!
-> Connection to switch_name closed by remote host.
Connection to switch_name closed.
```

The management controller power cycles and reinitializes.

Related Information

- "Change the root User Password" on page 27

▼ Change the `root` User Password

The default password for the `root` user is `changeme`. You can use the `passwd` command to make the password unique and more secure.

Note – For instructions to change the Oracle ILOM user passwords, refer to *Switch Remote Management*, changing Oracle ILOM user passwords.

---

1. **On the management controller, log in as the `root` user.**
2. **Type.**

```
# passwd root
Changing password for user root.
New UNIX password: new_password
Retype new UNIX password: new_password
passwd: all authentication tokens updated successfully.
#
```

Related Information

- *Switch Installation*, accessing the management controller
- "Restart the Management Controller" on page 26

Administering the I4 Switch Chip

These topics describe the administration of the I4 switch chip.

- "Monitoring the I4 Switch Chip" on page 29
- "Controlling the I4 Switch Chip" on page 36

Related Information

- "Troubleshooting the Switch" on page 1
- "Understanding Administrative Commands" on page 13
- "Administering the Chassis" on page 17
- "Administering the InfiniBand Fabric" on page 79
- "Administering the Subnet Manager" on page 45

Monitoring the I4 Switch Chip

These topics enable you to display and check the operation and status of the I4 switch chip.

Note – To use all the commands described in these topics, you must access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Display the Switch Chip Port to QSFP Connector Mapping" on page 30
- "Locate an InfiniBand Device and Connector From the GUID and Port" on page 31
- "Locate a Switch Chip or Connector From the GUID and Port" on page 31
- "Display Switch Chip Boot Status" on page 32
- "Display Link Status" on page 33
- "Display Connectors Configured With Autodisable Functionality" on page 33
- "Display Switch Chip Port Status" on page 34
- "Display Switch Chip Port Counters" on page 35

Related Information

- "Monitoring the Chassis" on page 17
- "Monitoring the InfiniBand Fabric" on page 79
- "Monitoring the Subnet Manager" on page 45
- "Controlling the I4 Switch Chip" on page 36

▼ Display the Switch Chip Port to QSFP Connector Mapping

You can use the `dcsport` command to display the mapping of a switch chip port to its respective QSFP connector and to display the mapping of a QSFP connector to its switch chip port.

- **On the management controller, type.**

```
FabMan@switch_name->dcsport -port port
```

where *port* is the number of the port (`1` to `36`).

Alternatively, type.

```
FabMan@switch_name->dcsport -connector connector
```

where *connector* is the number of the QSFP connector (`0A` to `17B`).

For example:

```
FabMan@switch_name->dcsport -port 22
DCS-36P Switch port 22 maps to connector 1A
FabMan@switch_name->dcsport -connector 1A
Connector 1A maps to Switch port 22
FabMan@switch_name->
```

Note – For a complete mapping of all ports to all connectors, use the `dcsport -printconnnectors` command.

Related Information

- *Switch Reference*, `dcsport` command
- "Understanding Routing Through the Switch" on page 7

▼ Locate an InfiniBand Device and Connector From the GUID and Port

The `dcsport` command can also identify the connector of an InfiniBand device that maps to a GUID and port.

- **On the management controller, type.**

```
FabMan@switch_name->dcsport -guid guid -port port
```

where:

- *guid* is the global unique identifier.
- *port* is the number of the port (`1` to `36`).

For example:.

```
FabMan@switch_name->dcsport -guid 0x0021283a8980a0a0 -port 36
DCS-36P Switch port 36 maps to connector 6B
FabMan@switch_name->
```

Related Information

- "Switch GUIDs Overview" on page 9
- *Switch Reference*, `dcsport` command

▼ Locate a Switch Chip or Connector From the GUID and Port

The output of some InfiniBand commands identify a node by its GUID. You can determine the switch chip or InfiniBand cable connection associated with that GUID using this procedure.

1. **In the output of the InfiniBand command, identify the GUID and port in question.**

   For example, given this output from the `ibdiagnet` command:

```
-W- lid=0x000f guid=0x002128312345c0a0 dev=48438 Port=21
```

   In this example, the GUID is `0x002128312345c0a0` and the port is 21.

2. **Use the `ibswitches` command to verify the GUID is that of a switch.**

   For example:

```
FabMan@switch_name->ibswitches
Switch  : 0x03ba01234567a0a0 ports 36 "SUN DCS 36P QDR mnm-36p-2 10.172.144.66"
enhanced port 0 lid 1 lmc 0
Switch  : 0x002128312345c0a0 ports 36 "SUN IB QDR GW switch mnm-gw-1
10.172.144.69" enhanced port 0 lid 15 lmc 0
FabMan@switch_name->
```

3. **If the GUID is that of a switch, truncate all digits of the GUID except the four digits on the right.**

   For example, GUID `0x002128312345c0a0` is truncated to `c0a0`.

   From "Switch GUIDs Overview" on page 9, the `c0` means that the node is inside of a Sun Network QDR InfiniBand Gateway Switch. The `a` means that the node is located in I4 switch chip A, the only switch chip. The port is 21.

4. **Use the tables in "Understanding Routing Through the Switch" on page 7 to cross-reference the routing between ports and to connections.**

   For example, using "Switch Chip Port to QSFP Connectors and Link LED Routes" on page 7, I4 switch chip, port 21 routes to connector 1B.

Related Information

- "Switch GUIDs Overview" on page 9
- *Switch Reference*, `dcsport` command

▼ Display Switch Chip Boot Status

If you believe there is a problem with the switch chip, the `checkboot` command provides a quick pass-fail check of the switch chip. Should the switch chip fail the test, you can use the `ibdevreset` command in an attempt to reboot the switch chip.

- **On the management controller, type.**

```
FabMan@switch_name->checkboot
Switch OK
All Internal ibdevices OK
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `checkboot` command
- "Display Switch Environmental and Operational Data" on page 21

- *Switch Reference*, `ibdevreset` command

▼ Display Link Status

In some situations, you might need to know the status of each route through the switch. Additionally, the `listlinkup` command displays where InfiniBand cables are connected to the switch.

- **On the management controller, type.**

```
FabMan@switch_name->listlinkup connected
Connector  1A Present <-> Switch Port 22 is up (Enabled)
Connector  7A Present <-> Switch Port 33 is up (Enabled)
Connector  9A Present <-> Switch Port 14 is up (Enabled)
Connector  5B Present <-> Switch Port 29 is up (Enabled)
Connector  7B Present <-> Switch Port 34 is up (Enabled)
Connector 14B Present <-> Switch Port 08 is up (Enabled)
FabMan@switch_name->
```

If connector and its link have been autodisabled, the `listlinkup` command will append the connector status with `(AutomaticHighErrorRate)`, `(AutomaticBadSpeedOrWidth)`, or both. For example:

```
Connector  0A Present<->  Switch Port 20 down (AutomaticHighErrorRate)
```

In this situation, check the condition of the connector and cable. Refer to *Switch Service*, servicing cables.

Related Information

- *Switch Reference*, `listlinkup` command
- "Display the Link Status of a Node" on page 84

▼ Display Connectors Configured With Autodisable Functionality

You can display the connectors that are configured to automatically disable in the presence of high error rates or suboptimal link speed or width.

- **On the management controller, type.**

```
FabMan@switch_name->autodisable list
Connectors which will be disabled on high error rate:
 0A 0B 1A 1B 2A 2B 3A 3B 4A 4B 5A 5B 6A 6B 7A 7B 8A 8B 9A 9B 10A 10B 11A 11B 12A
12B 13A 13B 14A 14B 15A 15B 16A 16B 17A 17B
Connectors which will be disabled on suboptimal link speed or width:
 0A 0B 1A 1B 2A 2B 3A 3B 4A 4B 5A 5B 6A 6B 7A 7B 8A 8B 9A 9B 10A 10B 11A 11B 12A
12B 13A 13B 14A 14B 15A 15B 16A 16B 17A 17B
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `autodisable` command
- "Autodisabling Ports and Connectors" on page 39

▼ Display Switch Chip Port Status

If by using other procedures, it seems that a particular port of the switch chip is problematic, you can use the `getportstatus` command to provide the port state, width, and speed.

- **On the management controller, type.**

```
FabMan@switch_name->getportstatus connector | Switch port
```

where:

- *connector* is the name of the connector (`0A` to `17A` and `0B` to `17B`).
- *port* is the number of the port (`1` to `36`).

For example:

```
FabMan@switch_name->getportstatus Switch 22
Port status for connector 14A Switch Port 22
Adminstate:.....Enabled
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
LinkSpeedActive:.....10.0 Gbps
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
```

```
NeighborMTU:.....................4096
OperVLs:.........................VL0-7
FabMan@switch_name->
```

Note – The parameters with the string `Active` indicate the current conditions.

Related Information

- *Switch Reference*, `getportstatus` command
- "Display the Link Status of a Node" on page 84

▼ Display Switch Chip Port Counters

You can check the errors and throughput of a switch chip port using the `getportcounters` command.

- **On the management controller, type.**

```
FabMan@switch_name->getportcounters port|connector
```

where:

- *port* is the number of the port (`1` to `36`).
- *connector* is the name of the connector (`0A` to `17A` and `0B` to `17B`).

For example, for port 22:

```
FabMan@switch_name->getportcounters 22
Port counters for connector 1A Switch port 22
SymbolErrors.....................0
LinkRecovers.....................0
LinkDowned.......................0
RcvErrors........................0
RcvRemotePhysErrors..............0
RcvSwRelayErrors.................0
XmtDiscards......................0
XmtConstraintErrors..............0
RcvConstraintErrors..............0
LinkIntegrityErrors..............0
ExcBufOverrunErrors..............0
VL15Dropped......................0
XmtData..........................909556
RcvData..........................737760
XmtPkts..........................12830
```

```
RcvPkts.........................10368
XmtWait.........................0
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `getportcounters` command
- "Display Counters for a Node" on page 85

Controlling the I4 Switch Chip

These topics describe how you can manage the I4 switch chip.

Note – To use the commands described in these topics, you must access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Disable a Switch Chip Port" on page 36
- "Enable a Switch Chip Port" on page 37
- "Autodisabling Ports and Connectors" on page 39

Related Information

- "Controlling the Chassis" on page 26
- "Controlling the InfiniBand Fabric" on page 93
- "Controlling the Subnet Manager" on page 52
- "Monitoring the I4 Switch Chip" on page 29

▼ Disable a Switch Chip Port

On occasion, you might need to turn off a port. For example, a cable might become damaged and cause symbol errors that affect the switch chip's performance. Use the `disableswitchport` command to disable a switch chip port and its respective cable connection.

Note – The `disableswitchport` command is persistent and survives reboots.

- **On the management controller, type.**

```
FabMan@switch_name->disableswitchport [--reason=reason] Switch port
```

where:

- *reason* is the reason for disabling the port, `Blacklist` or `Partition`.
- *port* is the number of the port (`1` to `36`).

Alternatively, type.

```
FabMan@switch_name->disableswitchport [--reason=reason] connector
```

where *connector* is the number of the QSFP connector (`0A` to `17B`).

For example:

```
FabMan@switch_name->disableswitchport Switch 22
Disable connector 1A Switch port 22
Adminstate:.....................Disabled
LinkWidthEnabled:...............1X or 4X
LinkWidthSupported:.............1X or 4X
LinkWidthActive:................4X
LinkSpeedSupported:.............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkState:......................Down
PhysLinkState:..................Disabled
LinkSpeedActive:................2.5 Gbps
LinkSpeedEnabled:...............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
NeighborMTU:....................4096
OperVLs:........................VL0-7
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `disableswitchport` command
- "Disable or Enable an External Port" on page 100
- "Enable a Switch Chip Port" on page 37
- "Autodisabling Ports and Connectors" on page 39

▼ Enable a Switch Chip Port

You can enable a disabled switch chip port with the `enableswitchport` command.

- **On the management controller, type.**

```
FabMan@switch_name->enableswitchport [--reason=reason] Switch port
```

where:

- *reason* is the reason for disabling the port, `Blacklist` or `Partition`.
- *port* is the number of the port (`1` to `36`).

Note – If the port was disabled with the `--reason` option, (as seen with the `listlinkup` command) it can only be enabled with the same reason option.

Alternatively, type.

```
FabMan@switch_name->enableswitchport [--reason=reason] connector
```

where *connector* is the name of the connector (`0A` to `17A` and `0B` to `17B`).

For example:

```
FabMan@switch_name->enableswitchport Switch 22
Enable connector 1A Switch port 22
Adminstate:.....................Enabled
LinkWidthEnabled:...............1X or 4X
LinkWidthSupported:.............1X or 4X
LinkWidthActive:................4X
LinkSpeedSupported:.............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkState:......................Down
PhysLinkState:..................PortConfigurationTraining
LinkSpeedActive:................2.5 Gbps
LinkSpeedEnabled:...............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
NeighborMTU:....................4096
OperVLs:........................VL0-7
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `enableswitchport` command
- "Disable or Enable an External Port" on page 100
- "Disable a Switch Chip Port" on page 36
- "Autodisabling Ports and Connectors" on page 39

Autodisabling Ports and Connectors

These topics instruct how to manage autodisabled ports and connectors.

- "Autodisable Functionality" on page 39
- "Add Connectors to the Autodisable List" on page 40
- "Reenable Autodisabled Connectors" on page 41
- "Delete Connectors From the Autodisable List" on page 42

Related Information

- "Disable a Switch Chip Port" on page 36
- "Enable a Switch Chip Port" on page 37

Autodisable Functionality

Switch chip ports and their connectors can be configured to automatically disable should their links exhibit high error rates or sub-optimal link speed or width. You use the `autodisable` command to add the connectors to the autodisable list, which has two parts; one for connectors whose links fail from high error rates, and another for connectors whose links fail from suboptimal link speed or width. A connector can be configured for both parts.

The autodisable feature monitors the following to determine if a connector and its respective link are experiencing high error rates:

- SNMP traps
- Oracle ILOM event log
- Syslog
- Email alerts

The autodisable feature also monitors the link speed and width, and if any of the following combinations are discovered, the link is considered suboptimal:

- 1x SDR
- 1x DDR
- 1x QDR
- 4x SDR
- 4x DDR

Note – If you intentionally configure a link for less than 4x QDR, do not add that connector to the link speed and width part of the autodisable list. It is automatically disabled.

Autodisable functionality is not enabled by default, you must actively configure it by adding connectors to the autodisable list. Additionally, you must actively monitor its operation and take corrective action if necessary. Failure to respond to autodisabled connectors and their links results in a degraded InfiniBand fabric.

Note – Because of its silent nature, have redundancy for autodisable configured links.

You can monitor autodisable configured connectors and their links using the `listlinkup` command or the Fabric Monitor. For both interfaces, a connector which has been autodisabled is identified with the text `AutomaticHighErrorRate`, `AutomaticBadSpeedOrWidth`, or both.

Removing a connector from the autodisable list stops autodisable functionality for that connector and its link. However, it does not re-enable an autodisabled connector or its link. You must investigate the condition of the disabled cable and connector and replace it if necessary. You can then re-enable an autodisabled connector with the `--automatic` option of the `enableswitchport` command.

Related Information

- *Switch Reference*, `autodisable` command
- *Switch Reference*, `enableswitchport` command
- "Add Connectors to the Autodisable List" on page 40
- "Reenable Autodisabled Connectors" on page 41
- "Delete Connectors From the Autodisable List" on page 42
- *Switch Service*, servicing cables

▼ Add Connectors to the Autodisable List

You add connectors to either or both parts of the autodisable list with the `autodisable` command.

- **On the management controller, type.**

```
FabMan@switch_name->autodisable add connector cause
```

where:

- *connector* is the name of the connector (`0A` to `17A` and `0B` to `17B`).
- *cause* is the reason to disable the connector:
  - `H` – The error rate is higher than what the InfiniBand specification permits.
  - `L` – The link speed or width is less than optimal (less than 4x QDR).

Alternatively, to add all connectors to the autodisable list with both causes, type.

```
FabMan@switch_name->autodisable all
```

For example:, to add connector `0A` to the autodisable list for high error rates:

```
FabMan@switch_name->autodisable add 0A H
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `autodisable` command
- "Autodisable Functionality" on page 39
- "Reenable Autodisabled Connectors" on page 41
- "Delete Connectors From the Autodisable List" on page 42

▼ Reenable Autodisabled Connectors

1. **Investigate why the connector was autodisabled.**

   Inspect the condition of the cable and connectors at both ends, and replace if necessary. Refer to *Switch Service*, servicing cables.

2. **On the management controller, type.**

```
FabMan@switch_name->enableswitchport --automatic port|connector
```

where:

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

- *connector* is the name of the connector (`0A` to `17A` and `0B` to `17B`).

For example:

```
FabMan@switch_name->enableswitchport --automatic 0A
```

3. **If you did not replace the cable in** Step 1**, you must retrain the links for the port.**

   a. **Disable the port.**

      See "Disable a Switch Chip Port" on page 36.

   b. **Reenable the port.**

      Do not use the `--automatic` option when you reenable the port. See "Enable a Switch Chip Port" on page 37.

Related Information

- *Switch Reference*, `enableswitchport` command
- "Autodisable Functionality" on page 39
- "Add Connectors to the Autodisable List" on page 40
- "Delete Connectors From the Autodisable List" on page 42

▼ Delete Connectors From the Autodisable List

- **On the management controller, type.**

```
FabMan@switch_name->autodisable del connector cause
```

where:

- *connector* is the name of the connector (`0A` to `17A` and `0B` to `17B`).
- *cause* is the reason to disable the connector:
  - `H` – The error rate is higher than what the InfiniBand specification permits.

- `L` – The link speed or width is less than optimal (less than 4x QDR).

Alternatively, to delete all connectors, type.

```
FabMan@switch_name->autodisable none
```

In this situation, all connectors are deleted from the autodisable list regardless of cause.

For example, to delete connector `0A` from the autodisable list for suboptimal link width:

```
FabMan@switch_name->autodisable del 0A L
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `autodisable` command
- "Autodisable Functionality" on page 39
- "Add Connectors to the Autodisable List" on page 40
- "Reenable Autodisabled Connectors" on page 41

Administering the Subnet Manager

These topics describe the administration of the Subnet Manager.

- "Monitoring the Subnet Manager" on page 45
- "Controlling the Subnet Manager" on page 52
- "Configuring Secure Fabric Management" on page 60

Related Information

- "Troubleshooting the Switch" on page 1
- "Understanding Administrative Commands" on page 13
- "Administering the Chassis" on page 17
- "Administering the I4 Switch Chip" on page 29
- "Administering the InfiniBand Fabric" on page 79

Monitoring the Subnet Manager

These topics describe how to monitor the Subnet Manager.

Note – To use all the commands described in these topics, you must access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Display Subnet Manager Status" on page 46
- "Display Subnet Manager Location History" on page 46
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Display the Secret M_Key Configuration" on page 48
- "Verify M_Key Persistence" on page 49
- "Verify the Subnet Manager Configuration Integrity" on page 50
- "Display the Subnet Manager Log" on page 51

Related Information

- "Monitoring the Chassis" on page 17
- "Monitoring the I4 Switch Chip" on page 29
- "Monitoring the InfiniBand Fabric" on page 79
- "Controlling the Subnet Manager" on page 52
- "Configuring Secure Fabric Management" on page 60

▼ Display Subnet Manager Status

If you want to quickly determine the master Subnet Manager's state, the `getmaster` command provides the LID and GUID of the InfiniBand device hosting the master Subnet Manager.

- **On the management controller, type.**

```
FabMan@switch_name->getmaster
Local SM enabled and running, state MASTER
20120831 09:45:17 Master SubnetManager on sm lid 1 sm guid 0x21283a8980a0a0 :
SUN DCS 36P QDR mnm2 10.132.124.33
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `getmaster` command
- "Display Subnet Manager Location History" on page 46
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Verify the Subnet Manager Configuration Integrity" on page 50
- "Display the Subnet Manager Log" on page 51

▼ Display Subnet Manager Location History

- **On the management controller, type.**

```
FabMan@switch_name->getmaster -l
Local SM enabled and running, state MASTER
SM monitoring history:
20121209 23:15:39 whereismaster started
20121209 23:15:40 No Master SubnetManager seen in the system
20121209 23:15:50 Master SubnetManager on sm lid 1 sm guid 0x21283a8980a0a0 :
SUN DCS 36P QDR mnm 10.132.124.33
```

```
20121209 23:27:09 No Master SubnetManager seen in the system
.
.
.
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `getmaster` command
- "Display Subnet Manager Status" on page 46
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Verify the Subnet Manager Configuration Integrity" on page 50
- "Display the Subnet Manager Log" on page 51

▼ Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm

If you want to know your Subnet Manager's priority, controlled handover state, prefix, management key, and routing algorithm, use the `list` option of the `setsmpriority` command.

1. On the management controller, type.

```
FabMan@switch_name->setsmpriority list
Current SM settings:
smpriority 8
controlled_handover TRUE
subnet_prefix 0xfe80000000000000
M_Key 0xabbababe0001
Routing engine FatTree
FabMan@switch_name->
```

Note – The management key displayed in the output is the *readable* M_Key and not the *secret* M_Key.

---

2. If the information displayed is incorrect, modify the parameters.

See:

- "Set the Subnet Manager Priority" on page 52

- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Readable Management Key" on page 55
- "Set the Subnet Manager Routing Algorithm" on page 57

Related Information

- *Switch Reference*, `setsmpriority` command
- "Controlling the Subnet Manager" on page 52
- "Display Subnet Manager Status" on page 46
- "Display Subnet Manager Location History" on page 46
- "Verify the Subnet Manager Configuration Integrity" on page 50
- "Display the Subnet Manager Log" on page 51

▼ Display the Secret M_Key Configuration

You can display the `active` (currently in use) or `modified` (pending) secret M_Key configuration with the `smsubnetprotection list` *state* command.

Note – You must be in a configuration session to display the modified configuration.

- **On the management controller that hosts the master Subnet Manager, display the secret M_Key configuration.**

```
FabMan@switch_name->smsubnetprotection list state
```

where *state* is either `active` or `modified`. For example:

```
FabMan@switch_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=mnm
# time=15 Dec 04:03:00
# checksum=0
#! commit_number : 4
Mkey                  Untrusted Mkey          Smkey                  Attribute
------------------    ------------------      ------------------     ---------
0x00abcdefabcdef01      0x1aa45124fee612ae       0x15fc26aea300f831
0x00abcdefabcdef02      0x4ccd8230de6cd348       0x3fc7e6ad701a8a2a
0x00abcdefabcdef03      0x9baa1debcc74de5e       0x1b253003600d137b      C
FabMan@switch_name->
```

In the output, the `Attribute` heading identifies:

- `C` – The current secret M_Key
- `S` – The standby secret M_Key (if any)

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Configuring Secret M_Keys" on page 63

▼ Verify M_Key Persistence

You can check whether the M_Key survives a local power cycle, restart, or reboot.

1. On the management controller, type.

```
FabMan@gateway_name->localmkeypersistence show
Local M_Key persistence is disabled.
No local M_Key is set before a Subnet Manager is operational.
FabMan@gateway_name->
```

2. **If the information displayed is incorrect, change it.**

   See "Enable or Disable Secret M_Key Persistence" on page 72.

Related Information

- *Switch Reference*, `localmkeypersistence` command
- "Enable or Disable Secret M_Key Persistence" on page 72

▼ Verify the Subnet Manager Configuration Integrity

The `smconfigtest` command performs a check of the Subnet Manager's configuration and verifies the integrity.

Note – For the `spine` option of the `smconfigtest` command, the Subnet Manager is expected to have a priority of 5 or 8. For the `leaf` option, the expected priority is 5.

- **Verify the Subnet Manager configuration.**

```
FabMan@switch_name->smconfigtest [spine|leaf]
```

where:

- `spine` – for spine switches.
- `leaf` – for leaf switches.

For example, when the Subnet Manager has priority 6:

```
FabMan@switch_name->smconfigtest
SubnetManager configuration is valid.
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smconfigtest` command
- "Display Subnet Manager Status" on page 46
- "Display Subnet Manager Location History" on page 46
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Display the Subnet Manager Log" on page 51

▼ Display the Subnet Manager Log

1. **On the management controller, type.**

```
FabMan@switch_name->showsmlog
Jul 31 15:27:56 061000 [B7D328D0] 0x03 -> OpenSM 3.2.6_20120730 - Oracle patch
11.3
-----------------------------------------------
OpenSM 3.2.6_20120730 - Oracle patch 11.3
 Reading Cached Option File: /etc/opensm/opensm.conf
 Loading Cached Option: sm_sl = 1
 Loading Cached Option: packet_life_time = 15
.
.
.
Sep 05 04:49:30 976000 [B765AB90] 0x02 -> osm_report_notice: Reporting
Informational Notice "CapabilityMask, NodeDescription, Link [Width|Spe
ed] Enabled changed" from LID 1, GUID 0x00212856d102c0a0, new CapabilityMask
0x4250084a
Sep 05 04:49:30 993000 [B6658B90] 0x02 -> Fabric has 1 switches - topology is
not fat-tree. Falling back to default routing
Sep 05 04:49:30 993000 [B6658B90] 0x01 -> ucast_mgr_route: ftree: cannot build
lid matrices.
(END)
FabMan@switch_name->
```

The output first displays the parameters of the Subnet Manager configuration file, and then a chronological listing of Subnet Manager activity.

2. **Tap the space bar to display the next screen of the log.**
3. **Press the Q key to quit.**

Related Information

- *Switch Reference*, `showsmlog` command
- "Display Subnet Manager Status" on page 46
- "Display Subnet Manager Location History" on page 46
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Verify the Subnet Manager Configuration Integrity" on page 50

Controlling the Subnet Manager

You can enable the Subnet Manager with the `enablesm` command. When the Subnet Manager starts, it reads the configuration file for configuration information. You can disable the Subnet Manager with the `disablesm` command.

These topics describe how to control the Subnet Manager:

Note – To use the commands described in these topics, you must access them through the `/SYS/Fabric_Mgmt` Linux shell targets of the Oracle ILOM CLI.

- "Set the Subnet Manager Priority" on page 52
- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Readable Management Key" on page 55
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Enable the Subnet Manager" on page 58
- "Disable the Subnet Manager" on page 59

Related Information

- "Controlling the Chassis" on page 26
- "Controlling the I4 Switch Chip" on page 36
- "Controlling the InfiniBand Fabric" on page 93
- "Monitoring the Subnet Manager" on page 45
- "Controlling the Subnet Manager" on page 52

▼ Set the Subnet Manager Priority

By default, the Subnet Manager within the management controller is set to 0 priority. If there is more than one Subnet Manager in your InfiniBand fabric, you must set the priority of each Subnet Manager appropriately. The Subnet Manager with the highest priority is the primary (or Master) Subnet Manager.

1. **On the management controller, disable the Subnet Manager.**

   See "Disable the Subnet Manager" on page 59.

2. **Set the Subnet Manager priority.**

```
FabMan@switch_name->setsmpriority priority
```

where *priority* is `0` (lowest) to `13` (highest). For example:

```
FabMan@switch_name->setsmpriority 5
Current SM settings:
smpriority 5
controlled_handover FALSE
subnet_prefix 0xfe80000000000000
M_Key None
Routing engine FatTree
FabMan@switch_name->
```

3. **Enable the Subnet Manager.**

See "Enable the Subnet Manager" on page 58.

Related Information

- *Switch Reference*, `setsmpriority` command
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Readable Management Key" on page 55
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Enable the Subnet Manager" on page 58
- "Disable the Subnet Manager" on page 59

▼ Set the Subnet Manager Prefix

The `setsubnetprefix` command writes a prefix value to the `subnet_prefix` parameter of the configuration file.

1. **On the management controller, disable the Subnet Manager.**

See "Disable the Subnet Manager" on page 59.

2. **Set the Subnet Manager prefix.**

```
FabMan@switch_name->setsubnetprefix 0xabbababe
Current SM settings:
smpriority 5
controlled_handover FALSE
subnet_prefix 0xabbababe
M_Key None
Routing engine FatTree
FabMan@switch_name->
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 58.

Related Information

- *Switch Reference*, `setsubnetprefix` command
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Set the Subnet Manager Priority" on page 52
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Readable Management Key" on page 55
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Enable the Subnet Manager" on page 58
- "Disable the Subnet Manager" on page 59

▼ Enable Subnet Manager Controlled Handover

If your InfiniBand fabric has two or more Subnet Managers, you can force a constrained fallback protocol should the master Subnet Manager fail. Refer to *Switch Reference*, `setcontrolledhandover` command for more information.

1. **On the management controller, disable the Subnet Manager.**

   See "Disable the Subnet Manager" on page 59.

2. **Enable controlled handover.**

```
FabMan@switch_name->setcontrolledhandover TRUE
Current SM settings:
smpriority 5
controlled_handover TRUE
```

```
subnet_prefix 0xabbababe
M_Key None
Routing engine FatTree
FabMan@switch_name->
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 58.

Related Information

- *Switch Reference*, `setcontrolledhandover` command
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Set the Subnet Manager Priority" on page 52
- "Set the Subnet Manager Prefix" on page 53
- "Set the Subnet Manager Readable Management Key" on page 55
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Enable the Subnet Manager" on page 58
- "Disable the Subnet Manager" on page 59

▼ Set the Subnet Manager Readable Management Key

For added security, an easily readable management key, or readable M_Key is used by the Subnet Manager Agent to initialize and configure network nodes. Refer to *Switch Reference*, `setsmmkey` command for more information.

Note – All Subnet Managers in the InfiniBand fabric must be configured with the same readable M_Key value.

Note – For a more secure fabric, configure *secret* M_Key functionality. See "Configuring Secure Fabric Management" on page 60. Secret M_Key functionality overrides readable M_Key functionality.

---

1. **On the management controller, disable the Subnet Manager.**

   See "Disable the Subnet Manager" on page 59.

2. **Set the readable management key.**

```
FabMan@switch_name->setsmmkey m_key
```

where *m_key* is the management key (16 hexadecimal digits)

For example:

```
FabMan@switch_name->setsmmkey 0xabbababe0001
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xabbababe
M_Key 0xabbababe0001
Routing engine FatTree
FabMan@switch_name->
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 58.

4. **Repeat** Step 1 **to** Step 3 **for all Subnet Managers of the InfiniBand fabric.**

Related Information

- *Switch Reference*, `setsmmkey` command
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Set the Subnet Manager Priority" on page 52
- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Enable the Subnet Manager" on page 58
- "Disable the Subnet Manager" on page 59

▼ Clear the Subnet Manager Readable Management Key

It might be necessary to remove the restrictions of the Subnet Manager readable management key. This procedure describes how to do so.

1. **On the management controller, disable the Subnet Manager.**

   See "Disable the Subnet Manager" on page 59.

2. **Clear the readable management key.**

```
FabMan@switch_name->setsmmkey none
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xabbababe
M_Key None
Routing engine FatTree
FabMan@switch_name->
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 58.

4. **Repeat Step 1 to Step 3 for all Subnet Managers requiring the change.**

Related Information

- *Switch Reference*, `setsmmkey` command
- "Set the Subnet Manager Readable Management Key" on page 55
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, Management Key, and Routing Algorithm" on page 47
- "Set the Subnet Manager Priority" on page 52
- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Enable the Subnet Manager" on page 58
- "Disable the Subnet Manager" on page 59

▼ Set the Subnet Manager Routing Algorithm

The default routing algorithm is Fat Tree. To change to the Min Hop routing algorithm, use the `setsmrouting` command.

1. **On the management controller, disable the Subnet Manager.**

   See "Disable the Subnet Manager" on page 59.

2. **Set the routing algorithm.**

```
FabMan@switch_name->setsmrouting fattree|minhop
```

For example:

```
FabMan@switch_name->setsmrouting minhop
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xabbababe
M_Key 0xabbababe0001
Routing engine minhop
FabMan@switch_name->
```

3. **Enable the Subnet Manager.**

See "Enable the Subnet Manager" on page 58.

Related Information

- *Switch Reference*, `setsmrouting` command
- "Set the Subnet Manager Priority" on page 52
- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Readable Management Key" on page 55
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Enable the Subnet Manager" on page 58
- "Disable the Subnet Manager" on page 59

▼ Enable the Subnet Manager

The Subnet Manager within the management controller is not enabled by default. Use the `enablesm` command to enable and start the Subnet Manager.

- **On the management controller, type.**

```
FabMan@switch_name->enablesm
Starting IB Subnet Manager.                                [  OK  ]
Starting partitiond daemon.                                [  OK  ]
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `enablesm` command
- "Disable the Subnet Manager" on page 59
- "Set the Subnet Manager Priority" on page 52
- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Readable Management Key" on page 55
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Disable the Subnet Manager" on page 59

▼ Disable the Subnet Manager

If your InfiniBand fabric has too many Subnet Managers, you can disable and stop the Subnet Manager within the management controller by using the `disablesm` command.

- **On the management controller, type.**

```
FabMan@switch_name->disablesm
Stopping partitiond daemon.                                    [  OK  ]
Stopping IB Subnet Manager..                                   [  OK  ]
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `disablesm` command
- "Enable the Subnet Manager" on page 58
- "Set the Subnet Manager Priority" on page 52
- "Set the Subnet Manager Prefix" on page 53
- "Enable Subnet Manager Controlled Handover" on page 54
- "Set the Subnet Manager Readable Management Key" on page 55
- "Clear the Subnet Manager Readable Management Key" on page 56
- "Set the Subnet Manager Routing Algorithm" on page 57
- "Enable the Subnet Manager" on page 58

Configuring Secure Fabric Management

These topics enable you to increase security of your InfiniBand fabric.

- "Secret M_Key Overview" on page 60
- "smsubnetprotection Command" on page 61
- "Prepare for Secret M_Key Functionality" on page 62
- "Configuring Secret M_Keys" on page 63
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73
- "Considering Situations In Regard to Secret M_Key" on page 74

Related Information

- "Monitoring the Subnet Manager" on page 45
- "Controlling the Subnet Manager" on page 52

Secret M_Key Overview

The secret M-Key feature of the version 2.1 firmware adds an extra layer of security to the InfiniBand fabric through use of a hidden passphrase, or secret M_Key.

The secret M_Key enforces a policy that only a Subnet Manager authorized to manage the InfiniBand fabric is able to do so, and that write requests from unauthorized sources are rejected. Conveniently, information can be read from the InfiniBand devices without restriction, except for the secret M_Key itself.

In the process of initiating the secret M_Key functionality, you create the secret M_Key on the master Subnet Manager using the `smsubnetprotection` command and its subcommands in a configuration session. During the session, one or more secret M_Keys are added to a configuration list. There is a maximum of 10 secret M_Keys. Upon committing to the configuration, the list and active secret M_Key is automatically propagated to the standby Subnet Managers in the InfiniBand fabric. This action authorizes and enables the standby Subnet Managers to control the fabric in the event of a failover or handover. The master Subnet Manager also sets the value of the secret M_Key to the InfiniBand nodes of the subnet.

If a switch or gateway having a different secret M_Key is brought into the fabric, the fabrics current secret M_Key must be assigned manually.

Once implemented, the secret M_Key functionality is transparent to the user. Attempts to hijack the fabric by a rogue Subnet Manager are refused because the secret M_Key is missing, or there is a secret M_Key mismatch. This condition prevents the intruder from configuring the InfiniBand fabric.

There are prerequisites to implementing the secret M_Key functionality. You must:

- Correctly define and distribute the `smnodes` list to all `smnodes`.
- For redundancy and up-time, configure and enable IPoIB across the fabric.
- Verify and propagate the fabric element configuration.

Configuring secret M_Key for the first time requires:

- Setting the replication (and encryption) password
- Enabling M_Key persistence
- Enabling secret M_Key functionality

These requirements are implemented before configuring secret M_Key functionality so that the security of the fabric is maintained during the configuration.

Related Information

- "smsubnetprotection Command" on page 61
- "Prepare for Secret M_Key Functionality" on page 62
- "Configuring Secret M_Keys" on page 63
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73
- "Considering Situations In Regard to Secret M_Key" on page 74

`smsubnetprotection` Command

You can enable secret M_Key functionality using the `smsubnetprotection` command and its subcommands. The `smsubnetprotection` command is primarily issued from the switch or gateway that has the master Subnet Manager, and acts upon the secret M_Key configuration. There are two configurations, the *active* configuration is the one currently in use, and the *modified* configuration is the one which you can affect.

Configuring the secret M_Key consists of starting a configuration session adding or deleting M_Keys from the list, setting the current M_Key, and committing the list to active implementation. During the configuration session, you can also view the configuration list, undo your last action, or abort the configuration session altogether.

If you will only add or only delete secret M_Keys from the list during the configuration session, you must initiate the session with the `smsubnetprotection start -addonly` or `smsubnetprotection start -deleteonly` command.

Note – You cannot both add and delete secret M_Keys within a single configuration session. You must perform these actions in separate configuration sessions.

The `smsubnetprotection` command can also enable and disable secret M_Key functionality, display the active M_Keys, set and clear the local secret M_Key for I4 switch chips without a corresponding Subnet Manager, and set the replication password.

Related Information

- "Secret M_Key Overview" on page 60
- "Prepare for Secret M_Key Functionality" on page 62
- "Configuring Secret M_Keys" on page 63
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73
- "Considering Situations In Regard to Secret M_Key" on page 74

▼ Prepare for Secret M_Key Functionality

1. **Update the `smnodes` list for all `smnodes` in the InfiniBand fabric.**
   See "Update the smnodes List" on page 129.
2. **(Option) Configure IPoIB functionality.**
   See "Create the IPoIB Interface" on page 130.
3. **Create and propagate the fabric element configuration.**
   See "Create a Fabric Element Configuration" on page 115.
4. **Set the replication password.**
   See "Set the Replication Password" on page 72.
5. **Enable M_Key persistence.**
   See "Enable or Disable Secret M_Key Persistence" on page 72.
6. **Enable secret M_Key functionality.**
   See "Enable or Disable Secret M_Key Functionality" on page 73.

7. **Configure the secret M_Keys.**

   See "Configuring Secret M_Keys" on page 63.

Related Information

- "Secret M_Key Overview" on page 60
- "smsubnetprotection Command" on page 61
- "Configuring Secret M_Keys" on page 63
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73
- "Considering Situations In Regard to Secret M_Key" on page 74

Configuring Secret M_Keys

Use these topics to create and administer the secret M_Key configuration.

- "Create the Secret M_Key Configuration" on page 63
- "Add Secret M_Keys to the Configuration" on page 65
- "Set the Secret M_Key to Use" on page 68
- "Delete Secret M_Keys From the Configuration" on page 69
- "Configure a Local Secret M_Key" on page 71

Related Information

- "Secret M_Key Overview" on page 60
- "smsubnetprotection Command" on page 61
- "Prepare for Secret M_Key Functionality" on page 62
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73
- "Considering Situations In Regard to Secret M_Key" on page 74

▼ Create the Secret M_Key Configuration

1. **Prepare for secret M_Key functionality.**

   See "Prepare for Secret M_Key Functionality" on page 62.

2. **Initiate a secret M_Key configuration session on the management controller that hosts the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection start -addonly
FabMan@switch_name->
```

3. **Display the active secret M_Key configuration.**

   See "Display the Secret M_Key Configuration" on page 48.

4. **Add an M_Key to the configuration.**

```
FabMan@switch_name->smsubnetprotection add m_key
```

   where *m_key* is the management key (16 hexadecimal digits).

Note – There is a maximum of 10 secret M_Keys for the configuration.

   For example:

```
FabMan@switch_name->smsubnetprotection add 0x00abcdefabcdef01
Mkey added
FabMan@switch_name->smsubnetprotection add 0x00abcdefabcdef02
Mkey added
FabMan@switch_name->
```

5. **Select the M_Key to use.**

```
FabMan@switch_name->smsubnetprotection set-current m_key
```

   where *m_key* is the management key (16 hexadecimal digits).

   For example:

```
FabMan@switch_name->smsubnetprotection set-current
0x00abcdefabcdef01
Mkey 0x00abcdefabcdef01 updated
FabMan@switch_name->
```

Note – You can only set an M_Key that has been added to the configuration.

6. **Display the pending M_Keys.**

```
FabMan@switch_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=mnm
# time=15 Dec 03:38:33
# checksum=0
#! commit_number : 2
Mkey                  Untrusted Mkey         Smkey                  Attribute
------------------    ------------------     ------------------     ---------
0x00abcdefabcdef01      0x1aa45124fee612ae      0x15fc26aea300f831
0x00abcdefabcdef02      0x4ccd8230de6cd348      0x3fc7e6ad701a8a2a      S
FabMan@switch_name->
```

Note – The `S` in the `Attribute` column means that M_Key is in standby state, and is to become the current secret M_Key.

7. **End the session and commit the secret M_Key configuration to the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection commit
FabMan@switch_name->
```

The configuration is distributed to the trusted standby Subnet Managers and the current secret M_Key to the InfiniBand Fabric.

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Add Secret M_Keys to the Configuration" on page 65
- "Set the Secret M_Key to Use" on page 68
- "Delete Secret M_Keys From the Configuration" on page 69
- "Configure a Local Secret M_Key" on page 71

▼ Add Secret M_Keys to the Configuration

Note – There is a maximum of 10 secret M_Keys for the configuration.

1. **Initiate a secret M_Key configuration session on the management controller that hosts the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection start -addonly
FabMan@switch_name->
```

2. **Display the active secret M_Key configuration.**

   See "Display the Secret M_Key Configuration" on page 48.

3. **Add an M_Key to the configuration.**

```
FabMan@switch_name->smsubnetprotection add m_key
```

   where *m_key* is the management key (16 hexadecimal digits).

   For example:

```
FabMan@switch_name->smsubnetprotection add 0x00abcdefabcdef03
Mkey added
FabMan@switch_name->
```

4. **Repeat Step 3 for any additional M_Keys you want to add to the configuration.**

5. **(Optional) If you will use one of the just added secret M_Keys, set the M_Key to use.**

```
FabMan@switch_name->smsubnetprotection set-current m_key
```

   where *m_key* is the management key (16 hexadecimal digits).

   For example:

```
FabMan@switch_name->smsubnetprotection set-current
0x00abcdefabcdef03
Mkey 0x00abcdefabcdef03 updated
FabMan@switch_name->
```

Note – You can only set an M_Key which has been added to the configuration.

6. **Display the pending M_Keys.**

```
FabMan@switch_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=mnm
# time=15 Dec 03:47:12
# checksum=0
#! commit_number : 3
Mkey                  Untrusted Mkey         Smkey                  Attribute
------------------    ------------------     ------------------     ---------
0x00abcdefabcdef01      0x1aa45124fee612ae      0x15fc26aea300f831
0x00abcdefabcdef02      0x4ccd8230de6cd348      0x3fc7e6ad701a8a2a      C
0x00abcdefabcdef03      0x9baa1debcc74de5e      0x1b253003600d137b      S
FabMan@switch_name->
```

In the output, the `Attribute` heading identifies:

- `C` – The current secret M_Key
- `S` – The standby secret M_Key

7. **End the session and commit the secret M_Key configuration to the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection commit
FabMan@switch_name->
```

The configuration is distributed to the trusted standby Subnet Managers and the current secret M_Key to the InfiniBand Fabric.

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Create the Secret M_Key Configuration" on page 63
- "Set the Secret M_Key to Use" on page 68
- "Delete Secret M_Keys From the Configuration" on page 69
- "Configure a Local Secret M_Key" on page 71

▼ Set the Secret M_Key to Use

1. **Initiate a secret M_Key configuration session on the management controller that hosts the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection start
FabMan@switch_name->
```

2. **Display the active secret M_Key configuration.**

   See "Display the Secret M_Key Configuration" on page 48.

3. **Determine which M_Key you will set as the new current.**

4. **Set the current M_Key.**

```
FabMan@switch_name->smsubnetprotection set-current m_key
```

   where *m_key* is the management key (16 hexadecimal digits). For example:

```
FabMan@switch_name->smsubnetprotection set-current
0x00abcdefabcdef02
Mkey 0x00abcdefabcdef02 updated
FabMan@switch_name->
```

Note – You can only set an M_Key which is in the configuration.

5. **Display the pending M_Keys.**

```
FabMan@switch_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=mnm
# time=15 Dec 03:38:33
# checksum=0
#! commit_number : 2
Mkey                 Untrusted Mkey          Smkey                  Attribute
------------------    ------------------     ------------------     ---------
0x00abcdefabcdef01      0x1aa45124fee612ae      0x15fc26aea300f831
0x00abcdefabcdef02      0x4ccd8230de6cd348      0x3fc7e6ad701a8a2a      S
0x00abcdefabcdef03      0x9baa1debcc74de5e      0x1b253003600d137b      C
FabMan@switch_name->
```

Note – The `S` in the `Attribute` column means that M_Key is in standby state, and is to become the current secret M_Key.

6. **End the session and commit the secret M_Key configuration to the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection commit
FabMan@switch_name->
```

The configuration is distributed to the trusted standby Subnet Managers and the current secret M_Key to the InfiniBand Fabric.

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Create the Secret M_Key Configuration" on page 63
- "Add Secret M_Keys to the Configuration" on page 65
- "Delete Secret M_Keys From the Configuration" on page 69
- "Configure a Local Secret M_Key" on page 71

▼ Delete Secret M_Keys From the Configuration

Note – If you delete the secret M_Key also used by a node temporarily unavailable to the fabric, when that node becomes available, there is a secret M_Key mismatch, and the node is isolated. See "Configure a Local Secret M_Key" on page 71.

1. **(Optional) If you want to delete the current secret M_Key in use, set the new current M_Key to one you will not delete.**

   See "Set the Secret M_Key to Use" on page 68.

2. **Display the active secret M_Key configuration.**

   See "Display the Secret M_Key Configuration" on page 48.

3. **Determine which M_Key you will delete.**

4. **Initiate a secret M_Key configuration session on the management controller that hosts the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection start -deleteonly
FabMan@switch_name->
```

5. **Delete the M_Key from the configuration.**

```
FabMan@switch_name->smsubnetprotection delete m_key
```

where *m_key* is the management key (16 hexadecimal digits).

For example:

```
FabMan@switch_name->smsubnetprotection delete 0x00abcdefabcdef03
Element deleted
FabMan@switch_name->
```

6. **Repeat Step 5 for any additional M_Keys you want to delete from the configuration.**

7. **Display the pending M_Keys.**

```
FabMan@switch_name->smsubnetprotection list modified
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=mnm
# time=15 Dec 04:18:20
# checksum=0
#! commit_number : 5
Mkey                  Untrusted Mkey          Smkey                   Attribute
------------------    ------------------      ------------------      ---------
0x00abcdefabcdef01      0x1aa45124fee612ae       0x15fc26aea300f831
0x00abcdefabcdef02      0x4ccd8230de6cd348       0x3fc7e6ad701a8a2a      C
FabMan@switch_name->
```

8. **End the session and commit the secret M_Key configuration to the master Subnet Manager.**

```
FabMan@switch_name->smsubnetprotection commit
FabMan@switch_name->
```

The configuration is distributed to the trusted standby Subnet Managers and the current secret M_Key to the InfiniBand Fabric.

9. **Display the active secret M_Key configuration.**

See "Display the Secret M_Key Configuration" on page 48.

Related Information

- *Switch Reference*, `smsubnetprotection` command

- "Create the Secret M_Key Configuration" on page 63
- "Add Secret M_Keys to the Configuration" on page 65
- "Set the Secret M_Key to Use" on page 68
- "Configure a Local Secret M_Key" on page 71

▼ Configure a Local Secret M_Key

When a switch (or gateway) from an InfiniBand subnet using secret M_Key functionality is brought into a fabric using a different secret M_Key configuration, the switch is isolated from traffic because the secret M_Keys mismatch. In this situation, you must provide a recognized secret M_Key to the added switch, so that the switch is merged into the fabric and secured with the current secret M_key configuration.

Conversely, by manually changing the current secret M_Key of a switch (or gateway) to a persistent value not in the secret M_Key configuration, you isolate the switch (or gateway) from the fabric.

1. **Log in to the management controller of the switch (or gateway) to receive the local secret M_Key.**
2. **Set the local secret M_Key.**

```
FabMan@switch_name->smsubnetprotection setlocalsecretmkey m_key
```

where *m_key* is the management key (16 hexadecimal digits). For example:

```
FabMan@switch_name->smsubnetprotection setlocalsecretmkey 0x00abcdefabcdef04
Local mkey: 0xabcdefabcdef04
Attribute: secret
FabMan@switch_name->
```

3. **(Optional) If you want to isolate the switch (or gateway), set the secret M_Key to persist and restart the management controller.**

   See:

   - "Enable or Disable Secret M_Key Persistence" on page 72
   - "Restart the Management Controller" on page 26

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Create the Secret M_Key Configuration" on page 63
- "Add Secret M_Keys to the Configuration" on page 65

- "Set the Secret M_Key to Use" on page 68
- "Delete Secret M_Keys From the Configuration" on page 69

▼ Set the Replication Password

For further security, a replication password is used to authenticate secret M_Key configurations distributed by the master Subnet Manager.

- **On the management controller that hosts the master Subnet Manager, type.**

```
FabMan@switch_name->smsubnetprotection setreplicationpassword password
```

where *password* is eight alphanumeric characters. For example:

```
FabMan@switch_name->smsubnetprotection setreplicationpassword reindeer
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Secret M_Key Overview" on page 60
- "smsubnetprotection Command" on page 61
- "Prepare for Secret M_Key Functionality" on page 62
- "Configuring Secret M_Keys" on page 63
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73
- "Considering Situations In Regard to Secret M_Key" on page 74

▼ Enable or Disable Secret M_Key Persistence

When secret M_Key persistence is enabled, the InfiniBand nodes retain the active secret M_Key should a power failure, reset, or reboot occur.

- **On the management controller, type.**

```
FabMan@switch_name->localmkeypersistence state
```

where *state* is either `enable` or `disable`. For example:

```
FabMan@switch_name->localmkeypersistence enable
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `localmkeypersistence` command
- "Verify M_Key Persistence" on page 49
- "Secret M_Key Overview" on page 60
- "smsubnetprotection Command" on page 61
- "Prepare for Secret M_Key Functionality" on page 62
- "Configuring Secret M_Keys" on page 63
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73
- "Considering Situations In Regard to Secret M_Key" on page 74

▼ Enable or Disable Secret M_Key Functionality

Secret M_Key functionality is enabled or disabled independent of its configuration.

- **On the management controller that hosts the master Subnet Manager, type.**

```
FabMan@switch_name->smsubnetprotection state
```

where *state* is either `enablesecretmkey` or `disable`secretmkey. For example:

```
FabMan@switch_name->smsubnetprotection enablesecretmkey
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Secret M_Key Overview" on page 60
- "smsubnetprotection Command" on page 61
- "Prepare for Secret M_Key Functionality" on page 62

- "Configuring Secret M_Keys" on page 63
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Considering Situations In Regard to Secret M_Key" on page 74

Considering Situations In Regard to Secret M_Key

When secret M_Key functionality is enabled, certain situations require additional considerations:

- "Add a New Subnet Manager to the Secret M_Key Fabric" on page 74
- "Merge Subnets Into One Secret M_Key Fabric" on page 75
- "Secret M_Key and Firmware Upgrade and Downgrade Guidelines" on page 76

Related Information

- "Secret M_Key Overview" on page 60
- "smsubnetprotection Command" on page 61
- "Prepare for Secret M_Key Functionality" on page 62
- "Configuring Secret M_Keys" on page 63
- "Set the Replication Password" on page 72
- "Enable or Disable Secret M_Key Persistence" on page 72
- "Enable or Disable Secret M_Key Functionality" on page 73

▼ Add a New Subnet Manager to the Secret M_Key Fabric

When you add a switch or gateway to your existing fabric, you might also enable a Subnet Manager on that switch. Alternatively, you might enable a previously disabled Subnet Manager. In this situation, follow these instructions.

1. **If not installed, install the switch or gateway hosting the Subnet Manager.**

   Refer to *Switch Installation* or *Gateway Installation*.

Note – Do not cable signal or InfiniBand cables to the switch or gateway at this time.

---

2. **If the firmware is less than version 2.1, upgrade it.**

   Refer to *Switch Remote Management* or *Gateway Remote Management*, upgrading the firmware.

3. **Verify that the Subnet Manager is disabled.**

   See "Disable the Subnet Manager" on page 59.

4. **Set the Subnet Manager Priority to a low value, so that it is a standby Subnet Manager.**

   See "Set the Subnet Manager Priority" on page 52.

5. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 58.

6. **If you installed a new switch or gateway, attach the signal and InfiniBand cables now.**

   Refer to *Switch Installation* or *Gateway Installation*.

7. **Prepare the new Subnet Manager for M_Key functionality.**

   See "Prepare for Secret M_Key Functionality" on page 62.

8. **Optional, change the priority of the new Subnet Manager.**

   See "Set the Subnet Manager Priority" on page 52.

Note – If you are setting the new Subnet Manager as the master Subnet Manager, you must set the older master Subnet Manager to a lower priority.

Related Information

- "Merge Subnets Into One Secret M_Key Fabric" on page 75
- "Secret M_Key and Firmware Upgrade and Downgrade Guidelines" on page 76

▼ Merge Subnets Into One Secret M_Key Fabric

You might want to combine two or more subnets into a larger fabric with secret M_Key functionality.

1. **If the firmware on the switches and gateways are less than version 2.1, upgrade them.**

   Refer to *Switch Remote Management* or *Gateway Remote Management*, upgrading the firmware.

2. **Configure identical secret M_Key lists on each subnet.**

   See "Prepare for Secret M_Key Functionality" on page 62 and "Create the Secret M_Key Configuration" on page 63.

Note – The `smnodes` list must have all `smnodes` of the future fabric, and not just those of the individual subnets. Additionally, the subnets cannot communicate with each other because they have not been merged. For these reasons, you must use the `-force` option for the `smsubnetprotection` command, otherwise the partition daemons will prevent configuration because the nodes are unavailable.

3. **Determine which Subnet Manager is to be the master Subnet Manager and which are to be the standby Subnet Managers.**
4. **Set the priority for each Subnet Manager.**

   See "Set the Subnet Manager Priority" on page 52.
5. **Physically merge the subnets into one fabric.**

Related Information

- "Add a New Subnet Manager to the Secret M_Key Fabric" on page 74
- "Secret M_Key and Firmware Upgrade and Downgrade Guidelines" on page 76

Secret M_Key and Firmware Upgrade and Downgrade Guidelines

Follow these guidelines for firmware upgrades and downgrades with respect to secret M_Key functionality:

- For ease of use, upgrade all firmware to version 2.1 or later before implementing secret M_Key.
- If you downgrade to firmware version 2.0 or earlier, you lose secret M_Key functionality. Disable secret M_Key functionality before the downgrade, otherwise the remaining secret M_Key portions of the fabric become undiscoverable by the downgraded switch or gateway.
- This table describes the firmware type of M_Key functionality and its support by the firmware versions.

Functionality	Firmware 1.3 and Earlier	Firmware 2.0	Firmware 2.1 and Later
Readable M_Key	No	Supported	Supported
Secret M_Key	No	No	Supported

Note – Readable M_Keys created in firmware versions 2.0 or 2.1 are functional under either firmware versions.

Related Information

- "Add a New Subnet Manager to the Secret M_Key Fabric" on page 74
- "Merge Subnets Into One Secret M_Key Fabric" on page 75

Administering the InfiniBand Fabric

These topics describe the administration of the InfiniBand fabric for the switch.

- "Monitoring the InfiniBand Fabric" on page 79
- "Controlling the InfiniBand Fabric" on page 93

Related Information

- "Troubleshooting the Switch" on page 1
- "Understanding Administrative Commands" on page 13
- "Administering the Chassis" on page 17
- "Administering the I4 Switch Chip" on page 29
- "Administering the Subnet Manager" on page 45

Monitoring the InfiniBand Fabric

These topics enable you to display and check the operation and status of the InfiniBand fabric and components.

Note – To use all the commands described in these topics, you must access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Displaying Information About Nodes" on page 80
- "Displaying Fabric Functionality Configuration Information" on page 89

Related Information

- "Monitoring the Chassis" on page 17
- "Monitoring the I4 Switch Chip" on page 29
- "Monitoring the Subnet Manager" on page 45
- "Controlling the InfiniBand Fabric" on page 93

Displaying Information About Nodes

These tasks display information about the nodes in the InfiniBand fabric.

- "Identify All Switches in the Fabric" on page 80
- "Identify All CAs in the Fabric" on page 81
- "Display the InfiniBand Fabric Topology (Simple)" on page 82
- "Display the InfiniBand Fabric Topology (Detailed)" on page 82
- "Display a Route Through the Fabric" on page 83
- "Display the Link Status of a Node" on page 84
- "Display Counters for a Node" on page 85
- "Display Low-Level Detailed Information About a Node" on page 86
- "Display Low-Level Detailed Information About a Port" on page 87

Related Information

- "Investigating Nodes" on page 94
- "Controlling Nodes" on page 97
- "Displaying Fabric Functionality Configuration Information" on page 89

▼ 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.

Note – The output for your InfiniBand fabric will differ greatly from the output in the example.

- **On the management controller, type.**

```
FabMan@switch_name->ibswitches
Switch  : 0x03ba01234567a0a0 ports 36 "SUN DCS 36P QDR mnm-36p-2 10.172.144.66"
enhanced port 0 lid 1 lmc 0
Switch  : 0x002128312345c0a0 ports 36 "SUN IB QDR GW switch mnm-gw-1
10.172.144.69" enhanced port 0 lid 15 lmc 0
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `ibswitches` command

- *Switch Reference*, `ibhosts` command
- "Identify All CAs in the Fabric" on page 81

▼ Identify All CAs in the Fabric

Similar to the `ibswitches` command, the `ibhosts` command displays identity information about the channel adapters (CAs), inclusive of the host channel adapters (HCAs), in the InfiniBand fabric. The output contains the GUID and name for each CA.

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

---

1. On the management controller, type.

Note – The output in the example is just a portion of the full output and varies for each InfiniBand topology.

```
FabMan@switch_name->ibhosts
Ca      : 0x00212800013e9312 ports 2 "mnm56 HCA-1"
Ca      : 0x00212800013e93f6 ports 2 "mnm57 HCA-1"
Ca     : 0x002128312345c000 ports 2 "SUN IB QDR GW switch mnm-gw-1 10.172.144.69
Bridge 0"
Ca     : 0x002128312345c040 ports 2 "SUN IB QDR GW switch mnm-gw-1 10.172.144.69
Bridge 1"
.
.
.
Ca      : 0x00212800013e9312 ports 2 "mnm56 HCA-1"
FabMan@switch_name->
```

2. (Optional) Save the output of `ibhosts` command as a text file named `hostGUIDs.txt`.

Related Information

- *Switch Reference*, `ibhosts` command
- *Switch Reference*, `ibswitches` command
- "Identify All Switches in the Fabric" on page 80

▼ Display the InfiniBand Fabric Topology (Simple)

To understand the routing that happens within your InfiniBand fabric, the `showtopology` command displays the node-to-node connectivity. The output of the command depends on the size of your fabric.

- **On the management controller, type.**

```
FabMan@switch_name->showtopology
SUNIBQDRGW mnm34-98 Bridge 0
  BX1/P1 -4x-10G-> I4_GENERIC mnm34-98 P4
      -4x-10G-> I4_GENERIC mnm34-98 P3
  BX3/P1 -4x-10G-> I4_GENERIC mnm34-98 P2
      -4x-10G-> I4_GENERIC mnm34-98 P1

DEV26428_02P mnm34-55_HCA-1
  P1 -4x-10G-> SUNIBQDRGW mnm34-97 C-7A
  P2 -4x-10G-> I4_GENERIC mnm34-98 P33

SUNIBQDRGW mnm34-97
  A-SW/P2 -4x-10G-> SUNIBQDRGW mnm34-97 BX3/P1
  A-SW/P4 -4x-10G-> SUNIBQDRGW mnm34-97 BX1/P1
  C-0B -4x-10G-> I4_GENERIC mnm34-98 P21
.
.
.
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `showtopology` command
- "Determine Changes to the InfiniBand Fabric Topology" on page 95
- "Perform Comprehensive Diagnostics for the Entire Fabric" on page 94

▼ Display the InfiniBand Fabric Topology (Detailed)

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

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

- **On the management controller, type.**

```
FabMan@switch_name->ibnetdiscover
#
# Topology file: generated on Sat Apr 13 22:28:55 2002
#
# Max of 1 hops discovered
# Initiated from node 0021283a8389a0a0 port 0021283a8389a0a0
vendid=0x2c9
devid=0xbd36
sysimgguid=0x21283a8389a0a3
switchguid=0x21283a8389a0a0(21283a8389a0a0)
Switch  36 "S-0021283a8389a0a0" # "Sun DCS 36 QDR switch localhost" enhanced port
0 lid 15 lmc 0
[23]    "H-0003ba000100e388"[2](3ba000100e38a) # "mnm33-43 HCA-1" lid 14 4xQDR
.
.
.
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `ibnetdiscover` command
- "Perform Comprehensive Diagnostics for the Entire Fabric" on page 94

▼ Display a Route Through the Fabric

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

- **On the management controller, type.**

```
FabMan@switch_name->ibtracert slid dlid
```

where:

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

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

For example:

```
FabMan@switch_name->ibtracert 15 14
From switch {0x00212856cd22c0a0} portnum 0 lid 15-15 "SUN IB QDR GW switch
mnm34-97"
[1] -> ca port {0x00212856cd22c042}[2] lid 14-14 "SUN IB QDR GW switch mnm34-97
Bridge 1"
To ca {0x00212856cd22c040} portnum 2 lid 14-14 "SUN IB QDR GW switch mnm34-97
Bridge 1"
FabMan@switch_name->
```

Note – The output for your InfiniBand fabric will differ from that in the example.

For this example:

- The route starts at the switch with GUID `0x00212856cd22c0a0` and is using port `0`. The switch is LID `15` and in the description, the switch host's name is mnm34-97.
- The route enters at port 1 of the CA with GUID `0x00212856cd22c042` and exits at port `2`. The CA is LID `14` and is within bridge chip 1.

Related Information

- *Switch Reference*, `ibtracert` command

▼ 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 management controller, type.**

```
FabMan@switch_name->ibportstate lid port
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example:

```
FabMan@switch_name->ibportstate 1 22
PortInfo:
# Port info: Lid 1 port 22
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
Peer PortInfo:
# Port info: Lid 1 DR path slid 65535; dlid 65535; 0,22 port 8
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
FabMan@switch_name->
```

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

Related Information

- *Switch Reference*, `ibportstate` command
- "Display Switch Chip Port Status" on page 34
- "Display Link Status" on page 33
- "Find 1x, SDR, or DDR Links in the Fabric" on page 96
- "Set Port Speed" on page 99

▼ Display Counters for a Node

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

- **On the management controller, type.**

```
FabMan@switch_name->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:

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

```
FabMan@switch_name->perfquery 1 22
# Port counters: Lid 1 port 22
PortSelect:......................22
CounterSelect:..................0x1b01
SymbolErrors:...................0
.
.
.
VL15Dropped:....................0
XmtData:........................20232
RcvData:........................20232
XmtPkts:........................281
RcvPkts:........................281
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `perfquery` command
- "Display Switch Chip Port Counters" on page 35
- "Clear Data and Error Counters" on page 97

▼ 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.

- **On the management controller, type.**

```
FabMan@switch_name->smpquery switchinfo lid
```

where *lid* is the LID of the node.

For example, to see detailed information about a switch with LID 1, type.

```
FabMan@switch_name->smpquery switchinfo 1
# Switch info: Lid 1
LinearFdbCap:.....................49152
RandomFdbCap:.....................0
McastFdbCap:......................4096
LinearFdbTop:.....................1
DefPort:..........................0
DefMcastPrimPort:.................255
DefMcastNotPrimPort:..............255
LifeTime:.........................15
StateChange:......................0
LidsPerPort:......................0
PartEnforceCap:...................32
InboundPartEnf:...................1
OutboundPartEnf:..................1
FilterRawInbound:.................1
FilterRawOutbound:................1
EnhancedPort0:....................1
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smpquery` command
- "Display Low-Level Detailed Information About a Port" on page 87

▼ Display Low-Level Detailed Information About a Port

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

- **On the management controller, type.**

```
FabMan@switch_name->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 22 on the switch with LID 1, type.

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

```
FabMan@switch_name->smpquery portinfo 15 23
# Port info: Lid 1 port 22
Mkey:.............................0x0000000000000000
GidPrefix:........................0x0000000000000000
Lid:..............................0x0000
SMLid:............................0x0000
CapMask:..........................0x0
DiagCode:.........................0x0000
MkeyLeasePeriod:..................0
LocalPort:........................0
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
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smpquery` command
- "Display Low-Level Detailed Information About a Node" on page 86

Displaying Fabric Functionality Configuration Information

These tasks display the configurations used for added functionality of the InfiniBand fabric.

- "Display the InfiniBand Fabric Partition Configuration" on page 89
- "Display the InfiniBand Fabric Configuration" on page 90
- "Display the smnodes List" on page 91
- "Display the Fabric Element Configuration" on page 92
- "Display the IPoIB Interface" on page 92
- "Display Configured Secret M_Keys" on page 93

Related Information

- "Displaying Information About Nodes" on page 80
- "Partitioning the InfiniBand Fabric" on page 101
- "Correlating Fabric Elements" on page 114
- "Configuring the Fabric Director Node List" on page 120
- "Performing Supportive Tasks" on page 127

▼ Display the InfiniBand Fabric Partition Configuration

If you have partitioned the InfiniBand fabric, you can display the active (in use) partition configuration or the modified (in standby) partition configuration with the `smpartition` command.

1. On the management controller, type.

```
FabMan@gateway_name->smpartition list configuration
```

where *configuration* is the partition configuration name (`active` or `modified`).

For example:

```
FabMan@gateway_name->smpartition list active
# Sun DCS IB partition config file
#! version_number : 1
#! version_number : 12
Default=0x7fff, ipoib :
ALL_CAS=full,
ALL_SWITCHES=full,
```

```
SELF=full;
SUN_DCS=0x0001, ipoib :
ALL_SWITCHES=full;
 = 0x0004,ipoib:
0x0021280001cf3787=full,
0x0021280001cf205b=full;
FabMan@gateway_name->
```

Note – The output is displayed one page at a time.

In this short example, in addition to the default partition (`Default=0x7fff`), there is a new partition with P_Key `0x0004`. This partition is configured for IPoIB and has two member ports that are assigned full membership of the partition.

2. **To advance to the next page of output, press the spacebar.**

   Repeat as necessary.

3. **If the information displayed is incorrect, modify the partition configuration.**

   See "Modify a Partition or Port" on page 109.

Related Information

- *Switch Reference*, `smpartition` command
- "Partitioning the InfiniBand Fabric" on page 101
- "Display the InfiniBand Fabric Topology (Detailed)" on page 82

▼ Display the InfiniBand Fabric Configuration

If you have configured the InfiniBand fabric for SNMP operations, you can display the elements configured with the `fdconfig` command.

1. **On the management controller, type.**

```
FabMan@switch_name->fdconfig list-current-fabric-config
Name             IP Addr           Type    Role(s)          isMaster  Fabricname
---------------  ----------------  ------  ---------------  --------  ----------
primary          123.45.67.89      sw36gw  gateway                    test
secondary        123.45.67.90      sw36    leafSwitch       no        test
FabMan@switch_name->
```

In this example, the element named `primary` has IP address `123.45.67.89` and is a Sun Network QDR InfiniBand Gateway Switch. It is being used as a gateway and is hosting the master Fabric Director daemon. The other element named `secondary` has an IP address of `123.45.67.90` and is a Sun Datacenter InfiniBand Switch 36. It is being used as a leaf switch.

2. **If the information displayed is incorrect, modify the fabric configuration.**

   See:

   - "Add or Remove an Element From the Fabric Configuration" on page 124.
   - "Modify an Element of the Fabric Configuration" on page 126

Related Information

- *Switch Reference*, `fdconfig` command
- "Configuring the Fabric Director Node List" on page 120

▼ Display the `smnodes` List

1. **On the management controller, type.**

```
FabMan@switch_name->smnodes list
123.45.67.89
123.45.67.90
FabMan@switch_name->
```

2. **If the information displayed is incorrect, modify the `smnodes` configuration list.**

   See "Update the smnodes List" on page 129.

Related Information

- *Switch Reference*, `smnodes` command
- "Create the smnodes List" on page 128
- "Update the smnodes List" on page 129

▼ Display the Fabric Element Configuration

1. **On the management controller, type.**

```
FabMan@switch_name->createfabric list active
Hostname                              ILOM IP Addr            Redundant IP Addr
--------------------------------      --------------------    -----------------
mnm12-36p-1                             123.45.67.89            -
mnm13-36p-1                             123.45.67.91            -
FabMan@switch_name->
```

2. **If the information displayed is incorrect, modify the fabric element configuration.**

   See "Modify an Element of the Configuration" on page 118.

Related Information

- *Switch Reference*, `createfabric` command
- "Create a Fabric Element Configuration" on page 115

▼ Display the IPoIB Interface

1. **On the management controller, type.**

```
FabMan@switch_name->create_ipoib -s
Name        PKey   IP-address      Subnetmask      Interface  Interface-type
-------------------------------------------------------------------------------
None        0x7fff 123.45.67.89    255.255.255.0   ib0        external
FabMan@switch_name->
```

2. **If the information displayed is incorrect, modify the IPoIB interfaces.**

   See:

   - "Delete the IPoIB Interface" on page 131
   - "Create the IPoIB Interface" on page 130.

Related Information

- *Switch Reference*, `create_ipoib` command
- "Delete the IPoIB Interface" on page 131
- "Create the IPoIB Interface" on page 130

▼ Display Configured Secret M_Keys

1. **On the management controller, type.**

```
FabMan@switch_name->smsubnetprotection list active
# File_format_version_number 1
# Sun DCS IB mkey config file
# This file is generated, do not edit
# secretmkey=enabled
# nodeid=mnm
# time=15 Dec 03:54:46
# checksum=378d9b09744e1d8b8ba6ae868c99d0c9
#! commit_number : 3
Mkey                  Untrusted Mkey          Smkey                   Attribute
------------------     ------------------     ------------------      ---------
0x00abcdefabcdef01      0x1aa45124fee612ae      0x15fc26aea300f831
0x00abcdefabcdef02      0x4ccd8230de6cd348      0x3fc7e6ad701a8a2a
0x00abcdefabcdef03      0x9baa1debcc74de5e      0x1b253003600d137b      C
FabMan@switch_name->
```

2. **If the information displayed is incorrect, reconfigure the secret M_Keys.**

   See "Configuring Secret M_Keys" on page 63.

Related Information

- *Switch Reference*, `smsubnetprotection` command
- "Set the Secret M_Key to Use" on page 68
- "Add Secret M_Keys to the Configuration" on page 65
- "Delete Secret M_Keys From the Configuration" on page 69

Controlling the InfiniBand Fabric

You can perform these tasks to manage the InfiniBand fabric and its components.

Note – To use all the commands described in these topics, you must access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Investigating Nodes" on page 94
- "Controlling Nodes" on page 97
- "Partitioning the InfiniBand Fabric" on page 101

- "Correlating Fabric Elements" on page 114
- "Configuring the Fabric Director Node List" on page 120
- "Performing Supportive Tasks" on page 127

Related Information

- "Controlling the Chassis" on page 26
- "Controlling the I4 Switch Chip" on page 36
- "Controlling the Subnet Manager" on page 52
- "Monitoring the InfiniBand Fabric" on page 79

Investigating Nodes

These tasks help you determine the condition of the InfiniBand fabric.

- "Perform Comprehensive Diagnostics for the Entire Fabric" on page 94
- "Determine Changes to the InfiniBand Fabric Topology" on page 95
- "Find 1x, SDR, or DDR Links in the Fabric" on page 96
- "Determine Which Links Are Experiencing Significant Errors" on page 96

Related Information

- "Controlling Nodes" on page 97
- "Partitioning the InfiniBand Fabric" on page 101
- "Correlating Fabric Elements" on page 114
- "Configuring the Fabric Director Node List" on page 120
- "Performing Supportive Tasks" on page 127

▼ Perform Comprehensive Diagnostics for the Entire Fabric

If 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 management controller, type.**

```
FabMan@switch_name->ibdiagnet -v -r
```

Related Information

- *Switch Reference*, `ibdiagnet` command
- *Switch Reference*, `ibdiagpath` command
- "Display the InfiniBand Fabric Topology (Detailed)" on page 82

▼ Determine Changes to the InfiniBand Fabric Topology

The `generatetopology` command enables you to take a snapshot of your fabric. At a later time, use the `matchtopology` command to compare the topology file to the present conditions.

1. Take a snapshot of the fabric topology.

```
FabMan@switch_name->generatetopology date.topo
```

where *date* is the date of the snapshot. For example:

```
FabMan@switch_name->generatetopology Sept8.topo
It will take some time to generate a topology file. Please wait!
topo.conf exist! will move it to topo.conf.old
will create new topo.conf
Topo.conf file is created. Will now start generating the topo file
Wrote Topology file:Sept8.topo
FabMan@switch_name->
```

2. After an event or as part of scheduled network management, compare the present topology to that saved in the topology file.

```
FabMan@switch_name->matchtopology date.topo
```

where *date* is the date of the snapshot. For example:

```
FabMan@switch_name->matchtopology Sept8.topo
Topology matching will take some time. Please wait!
-I-------------------------------------------------
-I- Topology matching results
-I-------------------------------------------------
-I- The topology defined in ib_topology.topo perfectly matches the
discovered fabric.
----------------------------------------------------------------
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `generatetopology` command

- *Switch Reference*, `matchtopology` command
- "Display the InfiniBand Fabric Topology (Simple)" on page 82

▼ Find 1x, 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 management controller, type.**

```
FabMan@switch_name->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 links that are not 4x QDR.

Related Information

- *Switch Reference*, `ibdiagnet` command
- "Display the Link Status of a Node" on page 84
- "Determine Which Links Are Experiencing Significant Errors" on page 96

▼ 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 management controller, type.**

```
FabMan@switch_name->ibdiagnet -c 500 -P all=1
```

In this instance of the `ibdiagnet` command, 500 test packets are injected into each link. The `-P all=1` option returns all Performance Monitor counters that increment during the test, respective to the GUID and port of the InfiniBand device.

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. If there are relatively few minor errors during the diagnostic, they can be monitored.

Note – According to 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. The respective links must be investigated for the cause of the rapid symbol error propagation.

Related Information

- *Switch Reference*, `ibdiagnet` command
- "Find 1x, SDR, or DDR Links in the Fabric" on page 96
- "Display Counters for a Node" on page 85

Controlling Nodes

These tasks enable you to manage ports in the InfiniBand fabric.

- "Clear Data and Error Counters" on page 97
- "Reset a Port" on page 98
- "Set Port Speed" on page 99
- "Disable or Enable an External Port" on page 100

Related Information

- "Investigating Nodes" on page 94
- "Partitioning the InfiniBand Fabric" on page 101
- "Correlating Fabric Elements" on page 114
- "Configuring the Fabric Director Node List" on page 120
- "Performing Supportive Tasks" on page 127

▼ Clear Data and Error 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. Alternatively, if you are troubleshooting a port, the `perfquery` command

provides counters of errors occurring at that port. To determine the throughput as a function of time, or if the problem at the port has been resolved, the `ibdiagnet -pc` command enables you to reset the data counters and error counters to 0.

- **On the management controller, type.**

```
FabMan@switch_name->ibdiagnet -pc
```

Related Information

- *Switch Reference*, `ibdiagnet` command
- *Switch Reference*, `perfquery` command
- "Display Counters for a Node" on page 85

▼ Reset a Port

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

Note – If secret M_Key functionality is enabled, use the `-M` *m_key* option to the `ibportstate` command for external ports.

- **On the management controller, type.**

```
FabMan@switch_name->ibportstate [-M m_key] lid port reset
```

where:

- *m_key* is the management key (16 hexadecimal digits)(optional)
- *lid* is the LID of the node.
- *port* is the port of the node.

For example:

```
FabMan@switch_name->ibportstate 15 23 reset
Initial PortInfo:
# Port info: Lid 1 port 22
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 1 port 22
LinkState:.....................Down
PhysLinkState:.................Disabled

After PortInfo set:
# Port info: Lid 1 port 22
LinkState:.....................Down
PhysLinkState:.................Polling
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `ibportstate` command
- "Disable or Enable an External Port" on page 100

▼ Set Port Speed

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 Gbit/sec.

Note – If secret M_Key functionality is enabled, use the `-M` *m_key* option to the `ibportstate` command for external ports.

- **On the management controller, type.**

```
FabMan@switch_name->ibportstate [-M m_key] lid port speed speed
```

where:

- *m_key* is the management key (16 hexadecimal digits)(optional)
- *lid* is the LID of the node.
- *port* is the port of the node.
- *speed* is the signaling rate per lane of the port (there are 4 lanes per port): `1` for 2.5 Gbit/sec, `2` for 5.0 Gbit/sec, and `4` for 10.0 Gbit/sec.

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

For example:

```
FabMan@switch_name->ibportstate 1 22 speed 1
Initial PortInfo:
# Port info: Lid 1 port 22
LinkSpeedEnabled:.................2.5 Gbps or 5.0 Gbps or 10.0 Gbps

After PortInfo set:
# Port info: Lid 1 port 22
LinkSpeedEnabled:.................2.5 Gbps
FabMan@switch_name->ibportstate 1 22 speed 7
Initial PortInfo:
# Port info: Lid 1 port 22
LinkSpeedEnabled:.................2.5 Gbps
After PortInfo set:
# Port info: Lid 1 port 22
LinkSpeedEnabled:.................2.5 Gbps or 5.0 Gbps or 10.0 Gbps
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `ibportstate` command
- "Display the Link Status of a Node" on page 84

▼ Disable or Enable an External Port

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

Note – If secret M_Key functionality is enabled, use the `-M` *m_key* option to the `ibportstate` command for external ports. Use the `disableswitchport` and `enableswitchport` commands for internal ports when secret M_Key functionality is enabled. See "Disable a Switch Chip Port" on page 36 and "Enable a Switch Chip Port" on page 37.

- **On the management controller, type.**

```
FabMan@switch_name->ibportstate [-M m_key] lid port state
```

where:

- *m_key* is the management key (16 hexadecimal digits)(optional)

- *lid* is the LID of the node.
- *port* is the port of the node.
- *state* is either `disable` or `enable`.

For example:

```
FabMan@switch_name->ibportstate 1 22 disable
Initial PortInfo:
# Port info: Lid 1 port 22
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 1 port 22
LinkState:.....................Down
PhysLinkState:.................Disabled
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `ibportstate` command
- "Disable a Switch Chip Port" on page 36
- "Disable or Enable an External Port" on page 100
- "Reset a Port" on page 98

Partitioning the InfiniBand Fabric

These topics enable you to partition the InfiniBand fabric and modify the partition's configuration.

- "smpartition Command" on page 102
- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109

- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

Related Information

- "Investigating Nodes" on page 94
- "Controlling Nodes" on page 97
- "Correlating Fabric Elements" on page 114
- "Configuring the Fabric Director Node List" on page 120
- "Performing Supportive Tasks" on page 127

`smpartition` Command

You can partition your InfiniBand fabric using the `smpartition` command and its subcommands. The `smpartition` command is issued from the switch or gateway that has the master Subnet Manager and acts upon the partition configuration. There are two configurations, the *active* configuration is the one currently in use, and the *modified* configuration is the one that you can affect.

You begin a configuration session with the `smpartition start` command. You then make changes to the modified configuration with the `smpartition create`, `smpartition add`, `smpartition remove`, `smpartition modify`, and `smpartition delete` commands. You finally make the modified configuration into the active configuration with the `smpartition commit` command.

Refer to *Switch Reference*, `smpartition` command, for more information.

Related Information

- *Switch Reference*, `smpartition` command
- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109
- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

smpartition Command Guidelines

When partitioning the InfiniBand fabric with the smpartition command, consider these guidelines during the configuration session:

- Valid P_Keys for partitioning are 0x1 through 0x7fff. However, both 0x1 and 0x7fff are predefined P_Keys and cannot be created or deleted with the smpartition command.
- Default entries (ports) in the predefined partitions cannot be removed. However, when you add ports into the predefined partitions using the smpartition add command, the added ports take precedence over the default ALL_SWITCHES and ALL_CAS port groupings.
- The SELF and ALL port groupings are not supported by the smpartition command. To specify all ports for a partition use the ALL_SWITCHES port grouping in conjunction with the ALL_CAS port grouping in the smpartition add command.
- The ALL_SWITCHES and ALL_CAS port groupings can have limited membership in the default partition (P_Key 0x7fff).
- To remove the ipoib flag from a partition, use the smpartition modify -flag command without any additional arguments.

Related Information

- "smpartition Command" on page 102
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109
- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

Partitions and P_Keys

The switch supports several partitions in the InfiniBand fabric. You assign a P_Key to each partition as a means of identifying the partition to the Subnet Manager. P_Keys are 15-bit integers and have a value of 0x1 to 0x7fff. The P_Key value is used by the master Subnet Manager to configure the InfiniBand fabric. A P_Key value of 0x7fff represents the default partition.

An additional bit, the membership bit, can identify the membership of the partition:

- Full – The membership bit is `1`. Full membership permits communication to all members within a partition.
- Limited – The membership bit is `0`. Limited membership permits communication only with a full member.

Combined together, P_Key and the membership bit comprise a 16-bit integer, and the most significant bit is the membership bit. In a full membership, the membership bit is set high. When this happens, the P_Key value is effectively increased by `0x8000`. Similarly, if you were to define a P_Key with a value greater than `0x8000`, the membership bit is automatically set to `1` and consequently is given full membership.

Related Information

- "smpartition Command" on page 102
- "smpartition Command Guidelines" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109
- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

▼ Determine the Partitions and P_Keys

1. **Use the output of the `ibswitches` and `ibhosts` commands to identify your switch and HCA node GUIDs.**

   See:

   - "Identify All Switches in the Fabric" on page 80
   - "Identify All CAs in the Fabric" on page 81

2. **Determine the partitions you will have, their names, and their respective P_Keys.**

3. **Collate the GUIDs, partition names, and P_Keys into a partition information text file.**

Related Information

- *Switch Reference*, `ibswitches` command
- *Switch Reference*, `ibhosts` command
- *Switch Reference*, `showgwports` command
- "smpartition Command" on page 102

- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109
- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

▼ Create a User Partition

Note – After creating user partitions, consider modifying the default partition to limited membership for CAs and disabling IPoIB. The reason being that nodes with full membership in different user partitions can communicate with each other through the default partition. See "Modify a Partition or Port" on page 109 for an example.

---

1. **If you have not already done so, create the Subnet Manager nodes list.**

   See "Create the smnodes List" on page 128.

2. **If the fabric element configuration already exists, update the configuration to reflect the `smnodes` list.**

   See "Add or Delete an Element From the Configuration" on page 117.

3. **Determine your partition P_Keys.**

   See "Determine the Partitions and P_Keys" on page 104.

4. **Initiate a partition configuration session on the management controller.**

```
FabMan@switch_name->smpartition start
FabMan@switch_name->
```

5. **Create the user partition.**

```
FabMan@switch_name->smpartition create [-n partition_name] -pkey p_key
[-flag ipoib [mtu mtu, rate rate, sl sl, scope scope]][-m defmember]
```

where:

- *partition_name* is an alphanumeric tag to the InfiniBand partition (optional).
- *p_key* is the partition key (`1` to `7fff` or `default`).
- *mtu* is the number that maps to the actual MTU (1 to 5).

mtu Number	1	2	3	4	5
MTU Value	256	512	1024	2048	4096

- *rate* is the number that maps to the actual throughput of a link (link width + link speed) (2 to 10).

rate Number	2	3	4	5	6	7	8	9	10
Rate Value in Gbps	2.5	10	30	5	20	40	60	80	120

- *sl* is the service level (0 to 15).

Note – Use service level 1 (`sl 1`) only for low-latency, high-priority, small-message, low-bandwidth traffic. Use other service levels for regular, high-bandwidth traffic.

- *scope* is the multicast address scope value (`1` to 14).

Note – The `mtu`, `rate`, `sl`, and `scope` parameters are for the multicast group created when `ipoib` (IP over InfiniBand) is configured for the partition. Typically, these values are not specified as the defaults are sufficient for the fabric configuration.

- *defmember* is the default membership type (`full`, `limited`, or `both`) for the partition.

Note – If ports are added to the partition without specifying the membership type, the default membership type is applied to the port

For example:

```
FabMan@switch_name->smpartition create -n testpartition -pkey 5 -m
full
FabMan@switch_name->
```

6. Verify the partition configuration.

```
FabMan@switch_name->smpartition list modified
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 13
```

```
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=
full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
testpartition = 0x0005,defmember=full:
FabMan@switch_name->
```

7. **Consider your next steps.**
   - If you want to add GUIDs, see "Add or Remove a Port From a Partition" on page 107.
   - If you want to modify the configuration, see "Modify a Partition or Port" on page 109.
   - If you want to delete the configuration altogether, see "Delete a Partition" on page 112.
   - Otherwise, go to Step 8.
8. **End the configuration session and commit the configuration to the active partition configuration.**

```
FabMan@switch_name->smpartition commit
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smpartition` command
- "Display the InfiniBand Fabric Partition Configuration" on page 89
- "smpartition Command" on page 102
- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109
- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

▼ Add or Remove a Port From a Partition

You can add or remove one or more ports from the partition at one time.

1. **Initiate a partition configuration session on the management controller.**

```
FabMan@switch_name->smpartition start
FabMan@switch_name->
```

2. **Add ports to the partition.**

```
FabMan@switch_name->smpartition add -n partition_name|-pkey p_key -port
port|ALL_CAS|ALL_SWITCHES|ALL_ROUTERS [-m member]
```

where:

- *partition_name* is an alphanumeric tag to the InfiniBand partition (optional).
- *p_key* is the partition key (`1` to `7fff`).
- *port* is the GUID of the port, or the special parameter, to add:
  - `ALL_CAS` – All CAs in the InfiniBand fabric.
  - `ALL_SWITCHES` – All switches.
  - `ALL_ROUTERS` – All routers.
- *member* is the membership type (`full`, `limited`, or `both`) for the port.

For example:

```
FabMan@switch_name->smpartition add -n testpartition -port 00212800013e9313
00212800013e9314 00212800013e93f7
FabMan@switch_name->
```

3. **Or, remove ports from the partition.**

```
FabMan@switch_name->smpartition remove -n partition_name|-pkey p_key -port
port|ALL_CAS|ALL_SWITCHES|ALL_ROUTERS
```

See Step 2 for variable definitions.

For example:

```
FabMan@switch_name->smpartition remove -n testpartition -port 00212800013e9314
FabMan@switch_name->
```

4. **Verify the partition configuration.**

```
FabMan@switch_name->smpartition list modified
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 13
```

```
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=
full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
testpartition = 0x0005,defmember=full:
0x00212800013e9313,
0x00212800013e93f7;
FabMan@switch_name->
```

5. **End the configuration session and commit the configuration to the active partition configuration.**

```
FabMan@switch_name->smpartition commit
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smpartition` command
- "smpartition Command" on page 102
- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Modify a Partition or Port" on page 109
- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

▼ Modify a Partition or Port

You can modify the configuration of the partition or the membership of the ports with the `smpartition modify` command.

1. **Initiate a partition configuration session on the management controller.**

```
FabMan@switch_name->smpartition start
FabMan@switch_name->
```

2. Modify the partition or ports.

```
FabMan@switch_name->smpartition modify -n partition_name|-pkey p_key [-flag [ipoib,
mtu mtu, rate rate, sl sl, scope scope]]|[-port port
|ALL_CAS|ALL_SWITCHES|ALL_ROUTERS [-m member]
```

where:

- *partition_name* is an alphanumeric tag to the InfiniBand partition (optional).
- *p_key* is the partition key (`1` to `7fff`).
- *mtu* is the number that maps to the actual MTU (1 to 5).

mtu **Number**	**1**	**2**	**3**	**4**	**5**
MTU Value	256	512	1024	2048	4096

- *rate* is the number that maps to the actual throughput of a link (link width + link speed) (2 to 10).

rate **Number**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Rate Value in Gbps	2.5	10	30	5	20	40	60	80	120

- *sl* is the service level (0 to 15).

Note – Use service level 1 (`sl 1`) only for low-latency, high-priority, small-message, low-bandwidth traffic. Use other service levels for regular, high-bandwidth traffic.

- *scope* is the multicast address scope value (`1` to 14).

Note – The `mtu`, `rate`, `sl`, and `scope` parameters are for the multicast group created when `ipoib` (IP over InfiniBand) is configured for the partition. Typically, these values are not specified as the defaults are sufficient for the fabric configuration.

- *port* is the GUID of the port, or the special parameter, to modify:
  - `ALL_CAS` – All CAs in the InfiniBand fabric.
  - `ALL_SWITCHES` – All switches.
  - `ALL_ROUTERS` – All routers.

- *member* is the membership type (`full`, `limited`, or `both`) for the port.

For example, to configure the default partition for limited membership for CAs and no IP over InfiniBand support:

```
FabMan@switch_name->smpartition modify -pkey 0x7fff -port ALL_CAS
-flag -m limited
FabMan@switch_name->
```

3. **Verify the partition configuration.**

```
FabMan@switch_name->smpartition list modified
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 13
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=
full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
testpartition = 0x0005,defmember=full,ipoib:
0x00212800013e9313,
0x00212800013e93f7;
FabMan@switch_name->
```

4. **End the configuration session and commit the configuration to the active partition configuration.**

```
FabMan@switch_name->smpartition commit
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smpartition` command
- "smpartition Command" on page 102
- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Delete a Partition" on page 112
- "Remove User Partitions for Firmware Downgrade" on page 113

▼ Delete a Partition

When you delete a partition, you effectively commit a blank default configuration.

1. Initiate a partition configuration session on the management controller.

```
FabMan@switch_name->smpartition start
FabMan@switch_name->
```

2. Delete the partition.

```
FabMan@switch_name->smpartition delete -n partition_name|-pkey p_key
```

where:

- *partition_name* is an alphanumeric tag to the InfiniBand partition (optional).
- *p_key* is the partition key (`2` to `7ffe`).

Note – You cannot delete the pre-defined partitions with P_Keys `1` and `7fff`.

For example:

```
FabMan@switch_name->smpartition delete -n testpartition
FabMan@switch_name->
```

3. End the configuration session and commit the configuration to the active partition configuration.

```
FabMan@switch_name->smpartition commit
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `smpartition` command
- "smpartition Command" on page 102
- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109
- "Remove User Partitions for Firmware Downgrade" on page 113

▼ Remove User Partitions for Firmware Downgrade

You must perform these steps before downgrading the firmware to a version before 2.0.

Note – Only perform this procedure if you are downgrading the firmware.

Note – By removing user partitions, you might lose connectivity that was available only through those partitions. Additionally, you might gain undesired connectivity because all hosts become full members of the default partition.

---

1. **(Optional) On the management controller of the gateways where VNICs using user partitions were created, remove all VNICs that are using user partitions.**
2. **(Optional) Remove all VLANs mapped to user partitions.**
3. **Remove all user-defined partitions.**

   See "Delete a Partition" on page 112.
4. **Restore the default partition to default settings.**

   See "Modify a Partition or Port" on page 109, however use this command line to modify the partition.

```
FabMan@switch_name->smpartition modify -pkey 0x7fff -port ALL_CAS
-flag ipoib -m full
```

5. **Consider your next step:**
   - If your InfiniBand fabric has just one Subnet Manager, downgrade the firmware.

     Refer to *Switch Remote Management*, upgrading the firmware.
   - If your InfiniBand fabric has multiple Subnet Managers, go to Step 6.
6. **Remove all entries from the Subnet Manager node list.**

```
FabMan@switch_name->smnodes delete IP_address [IP_address ...]
```

where *IP_address* are the IP addresses of the Subnet Manager nodes. For example:

```
FabMan@switch_name->smnodes delete 123.45.67.89 123.45.67.90
FabMan@switch_name->
```

7. **Repeat** Step 6 **on the management controllers of all Subnet Managers in the InfiniBand fabric.**

8. **Downgrade the firmware.**

   Refer to *Switch Remote Management*, upgrading the firmware.

Related Information

- *Switch Reference*, `smpartition` command
- *Switch Reference*, `smnodes` command
- "smpartition Command" on page 102
- "smpartition Command Guidelines" on page 103
- "Partitions and P_Keys" on page 103
- "Determine the Partitions and P_Keys" on page 104
- "Create a User Partition" on page 105
- "Add or Remove a Port From a Partition" on page 107
- "Modify a Partition or Port" on page 109
- "Delete a Partition" on page 112

Correlating Fabric Elements

These topics help you to create a fabric element configuration.

- "createfabric Command" on page 115
- "Create a Fabric Element Configuration" on page 115
- "Add or Delete an Element From the Configuration" on page 117
- "Modify an Element of the Configuration" on page 118

Related Information

- "Investigating Nodes" on page 94
- "Controlling Nodes" on page 97
- "Partitioning the InfiniBand Fabric" on page 101
- "Configuring the Fabric Director Node List" on page 120
- "Performing Supportive Tasks" on page 127

`createfabric` Command

The fabric element configuration file enables both the partition daemon and the Fabric Director to communicate via Ethernet or IPoIB in a redundant fashion. This is accomplished by mapping a correlation of the management controller's host name with the IP address recognized by Oracle ILOM, and the IPoIB IP address of each element.

You can use the `createfabric` command and its subcommands to create or edit the fabric configuration file, add, modify, or delete fabric elements from the file, and complete the configuration file and commit it to use. You can also display the fabric element configuration file with the `createfabric` command.

The `createfabric` command is issued on each management controller that hosts a Subnet Manager. Like the Subnet Manager nodes list, the fabric element configuration file that is created on one management controller, must be duplicated on all other management controllers.

There are two fabric element configurations, the *active* configuration is the one currently in use, and the *in-progress* configuration is the one that you can affect with the `createfabric` command.

Note – The IP addresses of the fabric element configuration file must be in synchronization with IP addresses of the `smnodes` list. If you update the `smnodes` list, you must update the fabric element configuration file.

Related Information

- *Switch Reference*, `createfabric` command
- "Display the Fabric Element Configuration" on page 92
- "Create a Fabric Element Configuration" on page 115
- "Add or Delete an Element From the Configuration" on page 117
- "Modify an Element of the Configuration" on page 118

▼ Create a Fabric Element Configuration

1. **Verify the `smnodes` list.**

   See "Display the smnodes List" on page 91.

2. **Initiate a fabric element configuration session on the management controller.**

```
FabMan@switch_name->createfabric start empty
FabMan@switch_name->
```

3. **Add a fabric element to the configuration.**

```
FabMan@switch_name->createfabric add-element -name hostname -ilomIp IP_address
[-redundantIP IP_address]
```

where:

- *hostname* is the host name of the element.
- *IP_address* is the IP address for the element recognized by Oracle ILOM or IPoIB.

For example:

```
FabMan@switch_name->createfabric add-element -name mnm12-36p-1 -ilomIp
123.45.67.89
Element added
FabMan@switch_name->
```

4. **Repeat Step 3 for any additional fabric elements you want to add to the configuration.**

5. **Display the fabric element configuration.**

```
FabMan@switch_name->createfabric list modified
Hostname                              ILOM IP Addr           Redundant IP Addr
--------------------------------      --------------------   ----------------
mnm12-36p-1                             123.45.67.89           -
FabMan@switch_name->
```

6. **End the configuration session and commit the in-progress configuration to the active fabric element configuration.**

```
FabMan@switch_name->createfabric complete
FabMan@switch_name->
```

7. **Repeat Step 2 through Step 6 on all management controllers hosting Subnet Managers in the InfiniBand fabric.**

Related Information

- *Switch Reference*, `createfabric` command
- "Display the Fabric Element Configuration" on page 92
- "createfabric Command" on page 115
- "Add or Delete an Element From the Configuration" on page 117
- "Modify an Element of the Configuration" on page 118

▼ Add or Delete an Element From the Configuration

1. **Verify the `smnodes` list.**

   See "Display the smnodes List" on page 91.

2. **Initiate a fabric element configuration session on the management controller.**

```
FabMan@switch_name->createfabric start from-current
FabMan@switch_name->
```

3. **Add a fabric element to the configuration.**

```
FabMan@switch_name->createfabric add-element -name hostname -ilomIp IP_address
[-redundantIP IP_address]
```

   where:

   - *hostname* is the host name of the element.
   - *IP_address* is the IP address for the element recognized by Oracle ILOM or IPoIB.

   For example:

```
FabMan@switch_name->createfabric add-element -name mnm13-36p-1 -ilomIp
123.45.67.90
Element added
FabMan@switch_name->
```

4. **Or, delete a fabric element from the configuration.**

```
FabMan@switch_name->createfabric delete-element -name hostname
```

   where *hostname* is the host name of the element.

   For example:

```
FabMan@switch_name->createfabric delete-element -name mnm13-36p-1
Element deleted
FabMan@switch_name->
```

5. **Display the fabric element configuration.**

```
FabMan@switch_name->createfabric list modified
Hostname                              ILOM IP Addr           Redundant IP Addr
--------------------------------      --------------------   ----------------
```

```
mnm12-36p-1                                  123.45.67.89          -
mnm13-36p-1                                  123.45.67.90          -
FabMan@switch_name->
```

6. **End the configuration session and commit the in-progress configuration to the active fabric element configuration.**

```
FabMan@switch_name->createfabric complete
FabMan@switch_name->
```

7. **Repeat** Step 2 **through** Step 6 **on all management controllers hosting Subnet Managers in the InfiniBand fabric.**

Related Information

- *Switch Reference*, `createfabric` command
- "Display the Fabric Element Configuration" on page 92
- "createfabric Command" on page 115
- "Create a Fabric Element Configuration" on page 115
- "Modify an Element of the Configuration" on page 118

▼ Modify an Element of the Configuration

1. **Verify the** `smnodes` **list.**

   See "Display the smnodes List" on page 91.

2. **Initiate a fabric element configuration session on the management controller.**

```
FabMan@switch_name->createfabric start from-current
FabMan@switch_name->
```

3. **Display the fabric element configuration.**

```
FabMan@switch_name->createfabric list modified
Hostname                                ILOM IP Addr           Redundant IP Addr
--------------------------------        --------------------   -----------------
mnm12-36p-1                               123.45.67.89           -
mnm13-36p-1                               123.45.67.90           -
FabMan@switch_name->
```

4. **Modify a fabric element of the configuration.**

```
FabMan@switch_name->createfabric modify-element -name hostname [-ilomIp IP_address]
[-redundantIP IP_address]
```

where:

- *hostname* is the host name of the element as seen in the `Hostname` column.
- *IP_address* is the IP address for the element recognized by Oracle ILOM or IPoIB.

For example:

```
FabMan@switch_name->createfabric modify-element -name mnm13-36p-1 -ilomIp
123.45.67.91
Element modified
FabMan@switch_name->
```

5. **Verify the fabric element configuration.**

```
FabMan@switch_name->createfabric list modified
Hostname                              ILOM IP Addr           Redundant IP Addr
--------------------------------      --------------------   -----------------
mnm12-36p-1                             123.45.67.89           -
mnm13-36p-1                             123.45.67.91           -
FabMan@switch_name->
```

6. **End the configuration session and commit the in-progress configuration to the active fabric element configuration.**

```
FabMan@switch_name->createfabric complete
FabMan@switch_name->
```

7. **Repeat Step 2 through Step 6 on all management controllers hosting Subnet Managers in the InfiniBand fabric.**

Related Information

- *Switch Reference*, `createfabric` command
- "Display the Fabric Element Configuration" on page 92
- "createfabric Command" on page 115
- "Create a Fabric Element Configuration" on page 115
- "Add or Delete an Element From the Configuration" on page 117

Configuring the Fabric Director Node List

These topics enable you to configure the Fabric Director node list.

- "Fabric Director and Fabric Elements" on page 120
- "fdconfig Command Overview" on page 121
- "Create a Fabric Configuration" on page 122
- "Add or Remove an Element From the Fabric Configuration" on page 124
- "Modify an Element of the Fabric Configuration" on page 126

Related Information

- "Investigating Nodes" on page 94
- "Controlling Nodes" on page 97
- "Partitioning the InfiniBand Fabric" on page 101
- "Correlating Fabric Elements" on page 114
- "Performing Supportive Tasks" on page 127

Fabric Director and Fabric Elements

The Fabric Director monitors the InfiniBand fabric by polling the nodes or *elements* of the fabric. Each fabric element has a Fabric Director instance, and the two communicate with each other in a one-to-one relationship. The Fabric Director gathers information about the fabric element through the `LDA` and `ENVD` daemons running on the fabric element. From the information gathered, the Fabric Director creates a local SNMP `fabricMIB`.

The Fabric Director of the element hosting the master Subnet Manager is assigned the role of master. The master Fabric Director gathers information about all of the fabric elements by communicating out-of-band (management network) with the Fabric Director instances on those fabric elements. This master Fabric Director creates a model of the fabric that is also in the form of an SNMP MIB, as defined by the `SUN-FABRIC-MIB` specification.

By default, Fabric Directors run in nonmaster mode. Only when a fabric element's Subnet Manager becomes the master Subnet Manager, does the fabric element's corresponding Fabric Director become the master Fabric Director.

This table describes the fabricMIB tables and fabricMIB scalars generated by Fabric Directors.

Type	Generated `fabricMIB` Tables	Generated `fabricMIB` Scalars
All Fabric Directors	• `fabricElemPortTable` • `fabricElemConnectorTable`	• `fabricElemMgrIpAddress` • `fabricElemName` • `fabricElemNumPorts` • `fabricElemOperStatus` • `fabricElemNumConnectors`
Master Fabric Directors	• `fabricMgmtElemTable` • `fabricMgmtExtElemTable` • `fabricMgmtGatewayTable` • `fabricMgmtLinkTable`	• `fabricMgmtFabricDescr` • `fabricMgmtFabricType` • `fabricMgmtFabricTopology` • `fabricMgmtFabricOperStatus` • `fabricMgmtFabricName`

Related Information

- *Switch Reference*, `SUN-FABRIC-MIB` mib
- *Switch Reference*, `fdconfig` command
- "Display the InfiniBand Fabric Configuration" on page 90
- "fdconfig Command Overview" on page 121
- "Create a Fabric Configuration" on page 122
- "Add or Remove an Element From the Fabric Configuration" on page 124
- "Modify an Element of the Fabric Configuration" on page 126

`fdconfig` Command Overview

The `fdconfig` command is used to configure a list of fabric elements expected to be in the fabric, the fabric node list. The list is used by the Fabric Directors, and provides these attributes for each fabric element:

- Name – The host name of the fabric element as configured by Oracle ILOM.
- IP address – The out-of-band management network IP address of the management controller hosting the fabric element and Fabric Director.
- Fabric element type – The name of the type of fabric element.
  - `sw36` – Sun Datacenter InfiniBand Switch 36
  - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch
  - `unknown` – The element's type is unknown.
- Role type – The purpose of the fabric element:
  - `gateway` – The element acts as a gateway.

- `spineSwitch` – The element acts as a spine switch.
- `leafSwitch` – The element acts as a leaf switch.
- `unknown` – The element's role is unknown.

- Master – Whether the Fabric Director instance on the fabric element is the master.

Related Information

- *Switch Reference*, `fdconfig` command
- "Display the InfiniBand Fabric Configuration" on page 90
- "Fabric Director and Fabric Elements" on page 120
- "Create a Fabric Configuration" on page 122
- "Add or Remove an Element From the Fabric Configuration" on page 124
- "Modify an Element of the Fabric Configuration" on page 126

▼ Create a Fabric Configuration

Before you can use the fabric MIBs, you must first create a fabric configuration. You must create identical an configuration on all elements of the fabric.

1. On the management controller, initiate a fabric configuration session.

```
FabMan@switch_name->fdconfig start-fabric-config empty
FabMan@switch_name->
```

2. Define the fabric name.

```
FabMan@switch_name->fdconfig define-fabric-name name
```

where *name* is the identifier of the fabric. For example:

```
FabMan@switch_name->fdconfig define-fabric-name test
Fabric name updated
FabMan@switch_name->
```

3. Define an element of the fabric.

```
FabMan@switch_name->fdconfig define-element -name name -ip IP_address
-type type -role role1 [-role role2]
```

where:

- *name* is the host name of the element as configured by Oracle ILOM.
- *IP_address* is the IP address of an element.
- *type* is the identifier of the type of element:
  - `sw36` – Sun Datacenter InfiniBand Switch 36
  - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch
  - `unknown` – The element type is unknown.
- *role* is the purpose of the element:
  - `gateway` – The element acts as a gateway.
  - `spineSwitch` – The element acts as a spine switch.
  - `leafSwitch` – The element acts as a leaf switch.
  - `unknown` – The element role is unknown.

For example:

```
FabMan@switch_name->fdconfig define-element -name primary -ip 123.45.67.89 -type
sw36 -role leafSwitch
Element added
FabMan@switch_name->
```

4. **Verify the fabric configuration.**

```
FabMan@switch_name->fdconfig list-in-progress-fabric-config
Name             IP Addr           Type    Role(s)          isMaster  Fabricname
--------------  ----------------  ------  ---------------  --------  ----------
primary          123.45.67.89      sw36    leafSwitch                 test
FabMan@switch_name->
```

5. **Consider your next step.**
   - If you want to add elements, perform Step 3 to Step 4 for each new element.
   - If you want to abort the configuration altogether, type.

```
FabMan@switch_name->fdconfig abort
In progress config aborted
FabMan@switch_name->
```

   - Otherwise, go to Step 6.

6. **End the configuration session and commit to the new configuration.**

```
FabMan@switch_name->fdconfig complete-fabric-config
FabMan@switch_name->
```

7. **Perform Step 1 through Step 6 on the management controllers of all elements of the fabric.**

   The configuration files and fabric MIBs are created.

Related Information

- *Switch Reference*, `fdconfig` command
- "Display the InfiniBand Fabric Configuration" on page 90
- "Fabric Director and Fabric Elements" on page 120
- "fdconfig Command Overview" on page 121
- "Add or Remove an Element From the Fabric Configuration" on page 124
- "Modify an Element of the Fabric Configuration" on page 126

▼ Add or Remove an Element From the Fabric Configuration

1. **Initiate a fabric configuration session.**

```
FabMan@switch_name->fdconfig start-fabric-config from-current
FabMan@switch_name->
```

2. **Add an element to the fabric configuration.**

```
FabMan@switch_name->fdconfig define-element -name name -ip IP_address
-type type -role role1 [-role role2]
```

   where:

   - *name* is the host name of the element as configured by Oracle ILOM.
   - *IP_address* is the IP address of an element.
   - *type* is the identifier of the type of element:
     - `sw36` – Sun Datacenter InfiniBand Switch 36
     - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch
     - `unknown` – The element type is unknown.
   - *role* is the purpose of the element:
     - `gateway` – The element acts as a gateway.
     - `spineSwitch` – The element acts as a spine switch.
     - `leafSwitch` – The element acts as a leaf switch.

- `unknown` – The element role is unknown.

For example:

```
FabMan@switch_name->fdconfig define-element -name secondary -ip 123.45.67.90 -type
sw36gw -role gateway
Element added
FabMan@switch_name->
```

3. **Or, remove an element from the fabric configuration.**

```
FabMan@switch_name->fdconfig remove-element -name name
```

where *name* is the identifier of the element. For example.

```
FabMan@switch_name->fdconfig remove-element -name secondary
Element deleted
FabMan@switch_name->
```

4. **Verify the fabric configuration.**

```
FabMan@switch_name->fdconfig list-in-progress-fabric-config
Name            IP Addr           Type    Role(s)          isMaster  Fabricname
--------------  ----------------  ------  ---------------  --------  ----------
primary         123.45.67.89      sw36    leafSwitch                 test
secondary       123.45.67.90      sw36gw  gateway          no        test
FabMan@switch_name->
```

5. **End the configuration session and commit to the new configuration.**

```
FabMan@switch_name->fdconfig complete-fabric-config
FabMan@switch_name->
```

6. **Perform Step 1 through Step 5 on the management controllers of all elements of the fabric.**

7. **If you added a new element, create a fabric configuration on the management controller of that element using the same configuration information.**

   See "Create a Fabric Configuration" on page 122.

Related Information

- *Switch Reference*, `fdconfig` command
- "Display the InfiniBand Fabric Configuration" on page 90
- "Fabric Director and Fabric Elements" on page 120

- "fdconfig Command Overview" on page 121
- "Create a Fabric Configuration" on page 122
- "Modify an Element of the Fabric Configuration" on page 126

▼ Modify an Element of the Fabric Configuration

You can modify the parameters of an element.

1. Initiate a fabric configuration session.

```
FabMan@switch_name->fdconfig start-fabric-config from-current
FabMan@switch_name->
```

2. Modify an element of the fabric configuration.

```
FabMan@switch_name->fdconfig redefine-element -name name [-ip
IP_address|-type type|-role role1 [-role role2]]
```

where:

- *name* is the host name of the element as configured by Oracle ILOM.
- *IP_address* is the IP address of an element.
- *type* is the identifier of the type of element:
  - `sw36` – Sun Datacenter InfiniBand Switch 36
  - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch
  - `unknown` – The element type is unknown.
- *role* is the purpose of the element:
  - `gateway` – The element acts as a gateway.
  - `spineSwitch` – The element acts as a spine switch.
  - `leafSwitch` – The element acts as a leaf switch.
  - `unknown` – The element role is unknown.

For example:

```
FabMan@switch_name->fdconfig redefine-element -name primary -role
spineSwitch
Element modified
FabMan@switch_name->
```

3. **Verify the fabric configuration.**

```
FabMan@switch_name->fdconfig list-in-progress-fabric-config
Name            IP Addr           Type    Role(s)          isMaster  Fabricname
--------------  ----------------  ------  ---------------  --------  ----------
primary         123.45.67.89      sw36    spineSwitch                test
secondary       123.45.67.90      sw36gw  gateway          no        test
FabMan@switch_name->
```

4. **End the configuration session and commit to the new configuration.**

```
FabMan@switch_name->fdconfig complete-fabric-config
FabMan@switch_name->
```

5. **Perform Step 1 through Step 4 on the management controllers of all elements of the fabric.**

Related Information

- *Switch Reference*, `fdconfig` command
- "Display the InfiniBand Fabric Configuration" on page 90
- "Fabric Director and Fabric Elements" on page 120
- "fdconfig Command Overview" on page 121
- "Create a Fabric Configuration" on page 122
- "Add or Remove an Element From the Fabric Configuration" on page 124

Performing Supportive Tasks

These tasks support the extended functionality of the Sun Datacenter InfiniBand Switch 36 from Oracle.

- "Create the smnodes List" on page 128
- "Update the smnodes List" on page 129
- "Create the IPoIB Interface" on page 130
- "Delete the IPoIB Interface" on page 131

Related Information

- "Investigating Nodes" on page 94
- "Controlling Nodes" on page 97
- "Partitioning the InfiniBand Fabric" on page 101

- "Correlating Fabric Elements" on page 114
- "Configuring the Fabric Director Node List" on page 120

▼ Create the `smnodes` List

If you are partitioning your InfiniBand fabric, a list of valid Subnet Manager nodes must exist in the filesystem of every management controller running a Subnet Manager. You create this list with the `smodes` command. The list contains the IP addresses of all active management controllers running a Subnet Manager in your fabric.

Note – If you are using DHCP for your management network, you must configure the DHCP server to assign a fixed and unique IP address to the Subnet Manager node, based upon the MAC address of that node. The algorithm to determine the IP address is user-determined. It is most important that the IP address is unique and consistent.

The list should have an entry for every Sun Datacenter InfiniBand Switch 36 and Sun Network QDR InfiniBand Gateway Switch that runs a Subnet Manager in your InfiniBand fabric.

Note – If the Subnet Manager nodes of your InfiniBand fabric ever change (disabled, added, and so on), you must update all copies of the Subnet Manager nodes file.

---

1. **Determine the IP addresses of all management controllers in your InfiniBand fabric.**
2. **On the management controller hosting the master Subnet Manager, propagate the Subnet Manager nodes file with the IP addresses of all Subnet Manager nodes.**

```
FabMan@switch_name->smnodes add IP_address IP_address ...
```

where *IP_address* is the IP address of each management controller hosting a Subnet Manager. For example:

```
FabMan@switch_name->smnodes add 123.45.67.89 123.45.67.90
FabMan@switch_name->
```

3. **Repeat Step 2 for all management controllers hosting Subnet Managers in the InfiniBand fabric.**

4. **Determine how you will partition your InfiniBand fabric.**

   See "Determine the Partitions and P_Keys" on page 104.

Related Information

- *Switch Reference*, `smodes` command
- "Display the smnodes List" on page 91
- "Update the smnodes List" on page 129
- "Create a User Partition" on page 105
- "Prepare for Secret M_Key Functionality" on page 62

▼ Update the `smnodes` List

1. **On the management controller of the master Subnet Manager, display all the Subnet Manager nodes.**

```
FabMan@switch_name->smnodes list
123.45.67.89
123.45.67.90
FabMan@switch_name->
```

2. **Add any missing Subnet Manager node IP addresses for the local fabric.**

```
FabMan@switch_name->smnodes add IP_address IP_address ...
```

   where *IP_address* is the IP address of each management controller hosting a Subnet Manager. For example:

```
FabMan@switch_name->smnodes add 123.45.67.91 123.45.67.92
FabMan@switch_name->
```

3. **Or, delete any unnecessary Subnet Manager node IP addresses for the local fabric.**

```
FabMan@switch_name->smnodes delete IP_address IP_address ...
```

   where *IP_address* is the IP address of each management controller hosting a Subnet Manager. For example:

```
FabMan@switch_name->smnodes delete 123.45.67.91 123.45.67.92
FabMan@switch_name->
```

4. **If you are combining fabrics, repeat Step 2, adding all the Subnet Manager node IP addresses from the remote fabrics.**

5. **Verify the updated list.**

```
FabMan@switch_name->smnodes list
123.45.67.89
123.45.67.90
123.45.67.91
123.45.67.92
123.45.66.77
123.45.68.10
FabMan@switch_name->
```

6. **Repeat Step 1 through Step 5 for all management controllers hosting Subnet Manager nodes in the local fabric.**

7. **If you are combining fabrics, repeat Step 1 through Step 5 for all management controllers hosting Subnet Manager nodes in the remote fabrics.**

Related Information

- *Switch Reference*, `smodes` command
- "Display the smnodes List" on page 91
- "Create the smnodes List" on page 128
- "Create a User Partition" on page 105
- "Prepare for Secret M_Key Functionality" on page 62

▼ Create the IPoIB Interface

This procedure creates an Internet protocol over InfiniBand interface.

Note – Only one IPoIB interface can be defined.

- **On the management controller, type.**

```
FabMan@switch_name->create_ipoib -n if_name -m mask -t if_type p_key IP_address
```

where:

- *if_name* is the identifier of the IP interface.
- *mask* is the subnet mask.
- *if_type* is the type of interface (`external`, `internal`, or `both`).

- *p_key* is the partition key (`1` to `7fff` or `default`).
- *IP_address* is the IP address of the interface.

For example:

```
FabMan@switch_name->create_ipoib -n newipoib -m 255.255.254.0 -t both 0x0001
123.45.67.89
Create interface:
Name       PKey   IP-address      Subnetmask      Interface  Interface-type
-------------------------------------------------------------------------
newipoib   0x0001 123.45.67.89    255.255.254.0   ib0        both
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `create_ipoib` command
- "Display the IPoIB Interface" on page 92
- "Delete the IPoIB Interface" on page 131
- "Prepare for Secret M_Key Functionality" on page 62

▼ Delete the IPoIB Interface

- **On the management controller, type.**

```
FabMan@switch_name->delete_ipoib if_name | p_key
```

where:

- *if_name* is the identifier of the IP interface.
- *p_key* is the partition key (`1` to `7fff` or `default`).

For example:

```
FabMan@switch_name->delete_ipoib newipoib
Delete interface:
Name       PKey   IP-address      Subnetmask      Interface  Interface-type
-------------------------------------------------------------------------
newipoib   0x0001 123.45.67.89    255.255.254.0   ib0        both
FabMan@switch_name->
```

Related Information

- *Switch Reference*, `delete_ipoib` command
- "Display the IPoIB Interface" on page 92
- "Create the IPoIB Interface" on page 130

Index

A

adding
  autodisabled connectors, 40
  element to configuration, 124
  fabric elements to configuration, 117
  ports to partition configuration, 107
  secret M_Keys, 65
  smnodes to list, 128
administering
  chassis, 17
  I4 switch chip, 29
  InfiniBand fabric, 79
  Subnet Manager, 45
administrative command overview, 13
ASR, 25
`autodisable` command, 33, 40, 42
autodisabled connectors
  adding, 40
  deleting, 42
  displaying configured, 33
  overview, 39
  reenabling, 41
autodisabling ports and connectors, 39
automatic service request, 25

C

CA identity, 81
changing password
  `root` user, 27
chassis
  administering, 17
  controlling, 26
  FRU ID, 23
  monitoring, 17
`checkboot` command, 32
checking
  board
    temperature, 20
    voltage, 19
  chassis status, 21
  environmental data, 21
  fan
    speed, 20
    status, 20
  flash drive status, 22
  link status, 33
  port counters, 35
  port status, 34
  power supply, 18
  switch chip boot status, 32
`checkpower` command, 18
`checkvoltages` command, 19
clear
  data counters, 97
  error counters, 97
  management key, 56
command
  `autodisable`, 33, 40, 42
  `checkboot`, 32
  `checkpower`, 18
  `checkvoltages`, 19
  `create_ipoib`, 92, 130
  `createfabric`, 115
    `add-element`, 117
    `delete-element`, 117
    `list active`, 92
    `modify-element`, 118
    `start`, 115
  `dcsport`, 30
  `delete_ipoib`, 131
  `disablesm`, 59
  `disableswitchport`, 36
  `enablesm`, 58
  `enableswitchport`, 37, 41

env_test, 21
fdconfig, 121
  abort, 122
  complete-fabric-config, 122
  define-element, 124
  define-fabric-name, 122
  list-current-fabric-config, 90
  redefine-element, 126
  remove-element, 124
  start-fabric-config empty, 122
generatetopology, 95
getfanspeed, 20
getmaster, 46
getportcounters, 35
getportstatus, 34
ibdiagnet, 94, 96, 97
ibhosts, 81
ibnetdiscover, 82
ibportstate, 84, 98, 99, 100
ibswitches, 80
ibtracert, 83
listlinkup, 33
localmkeypersistence, 49, 72
matchtopology, 95
perfquery, 85
reboot, 26
setsmmkey, 55, 56
setsmpriority, 47, 52, 54
setsmrouting, 57
setsubnetprefix, 53
showdisk, 22
showfree, 23
showfruinfo, 23
showpsufru, 24
showsmlog, 51
showtemps, 20
showtopology, 82
showunhealthy, 18
smconfigtest, 50
sminfo, 46
smnodes
  add, 128
  delete, 129
  list, 91
smpartition, 102
  add, 107
  create, 105
  delete, 112
  list, 89
  modify, 109
  remove, 107
  start, 105
smpquery, 86, 87
smsubnetprotection, 61, 63
  add, 65
  delete, 69
  disablesecretmkey, 73
  enablesecretmkey, 73
  list active, 93
  list modified, 48
  set-current, 68
  setlocalsecretmkey, 71
  setreplicationpassword, 72
version, 25
configuration overview, 13
configuring
  Fabric Director node list, 120
  local secret M_Key, 71
  secret M_Key, 63
  starting configuration session, 122
connector
  autodisabled
    adding, 40
    deleting, 42
    reenabling, 41
  autodisabling, 39
  displaying autodisable configured, 33
controlling
  chassis, 26
  I4 switch chip, 36
  InfiniBand fabric, 93
  nodes, 97
  Subnet Manager, 52
create_ipoib command, 92, 130
createfabric
  add-element command, 117
  delete-element command, 117
  list active command, 92
  modify-element command, 118
  start command, 115
createfabric command, 115
creating
  fabric configuration, 122
  fabric element configuration, 115
  IPoIB interface, 130
  partition configuration, 105
  secret M_Key configuration, 63

D

data counters
  clearing, 97
`dcsport` command, 30
`delete_ipoib` command, 131
deleting
  autodisabled connectors, 42
  fabric elements from configuration, 117
  IPoIB interface, 131
  partition configuration, 112
  secret M_Keys, 69
  smnodes from list, 129
`disablesm` command, 59
`disableswitchport` command, 36
disabling
  port, 100
  secret M_Key
    functionality, 73
    persistence, 72
  Subnet Manager, 59
  switch chip port, 36
displaying
  autodisable configured connectors, 33
  disabled links, 33
  fabric configuration, 90
  fabric configuration information, 89
  fabric elements, 92
  file system space, 23
  flash drive information, 22
  free memory, 23
  FRU ID
    chassis, 23
    power supply, 24
  IPoIB interface, 92
  M_Key persistence, 49
  node
    counters, 85
    information, 80, 86
    link status, 84
  partition configuration, 89
  port information, 87
  route, 83
  secret M_Key
    configuration, 48
    persistence, 49
  secret M_Keys, 93
  smnodes, 91
  Subnet Manager
    configuration integrity, 50
    controlled handover state, 47
    location history, 46
    log, 51
    management key, 47
    prefix, 47
    priority, 47
    routing algorithm, 47
    status, 46
  switch
    firmware, 25
  switch health, 18
  topology
    detailed, 82
    simple, 82

E

element
  adding, 124
  modifying, 126
  removing, 124
`enablesm` command, 58
`enableswitchport` command, 37, 41
enabling
  secret M_Key
    functionality, 73
    persistence, 72
  Subnet Manager, 58
  switch chip port, 37
`env_test` command, 21
error counters
  clearing, 97
  displaying, 85
example route through switch, 9

F

fabric
  diagnostics, 94
  displaying configuration information, 89
  element, 120
Fabric Director, 120
fabric elements, 114
  adding elements to configuration, 117
  configurations, 115
  creating configuration, 115
  deleting elements from configuration, 117
  displaying, 92

modifying elements of configuration, 118
fan
checking
speed, 20
status, 20
fdconfig
abort command, 122
complete-fabric-config command, 122
define-element command, 124
define-fabric-name command, 122
list-current-fabric-config command, 90
redefine-element command, 126
remove-element command, 124
start-fabric-config empty command, 122
fdconfig command, 121
firmware
downgrades and partitions, 113
secret M_Key guidelines, 76
switch, 25
FRU ID
chassis, 23
power supply, 24
full membership, 104

G

generatetopology command, 95
getfanspeed command, 20
getmaster command, 46
getportcounters command, 35
getportstatus command, 34
GUID
mapping
to CXP connector, 31
to switch chip, 31
nomenclature, 9
guidelines
firmware
secret M_Key, 76
smpartition, 103

H

hardware commands
overview, 16
hardware problems, 1

I

I4 switch chip
administering, 29
controlling, 36
monitoring, 29
ibdiagnet command, 94, 96, 97
ibhosts command, 81
ibnetdiscover command, 82
ibportstate command, 84, 98, 99, 100
ibswitches command, 80
ibtracert command, 83
identifying
CAs, 81
switches, 80
InfiniBand
commands, 16
creating configuration, 122
displaying configuration, 90
fabric
administering, 79
controlling, 93
monitoring, 79
problems, 3
investigating nodes, 94
IPoIB interface
creating, 130
deleting, 131
displaying, 92

L

limited membership, 104
link
status, 33
status of node, 84
listlinkup command, 33
localmkeypersistence command, 49, 72

M

management controller
restarting, 26
mapping
GUID
to CXP connector, 31
to switch chip, 31
LED to switch chip port, 8
QSFP to switch chip port, 8, 30

switch chip port
to LED, 7
to QSFP, 7, 30
`matchtopology` command, 95
membership
full, 104
limited, 104
modifying
element, 126
fabric elements of configuration, 118
partition configuration, 109
port membership, 109
monitoring
chassis, 17
I4 switch chip, 29
InfiniBand fabric, 79
Subnet Manager, 45

N

NET MGT port
troubleshooting, 6
nodes
controlling, 97
counters, 85
displaying information, 80
information, 86
investigating, 94
link status, 84

O

Oracle ILOM
command
overview, 15
overview
administrative commands, 13
autodisabled connectors, 39
hardware commands, 16
InfiniBand commands, 16
Oracle ILOM
commands, 15
secret M_Key, 60

P

P_Keys, 103, 104
partition
creating configuration, 105
displaying configuration, 89
partitioning, 103, 104
adding ports, 107
deleting configuration, 112
firmware downgrades, 113
InfiniBand fabric, 101
modifying configuration, 109
removing ports, 107
starting configuration session, 105
password, 27
performing
supportive tasks, 127
`perfquery` command, 85
poor performing links, 96
ports
autodisabling, 39
counters, 35
disabling, 100
information, 87
resetting, 98
speed, 99
status, 34
power supply
checking
status, 18
FRU ID, 24
preparing
secret M_Key, 62
problem
hardware, 1
InfiniBand fabric, 3

Q

QSFP connector
mapping
from GUID, 31
to switch chip port, 30

R

`reboot` command, 26
recovery errors, 96
removing
element from configuration, 124
ports from partition configuration, 107
resetting
port, 98
restarting
management controller, 26

route
QSFP to switch chip port, 8, 30
switch chip port to QSFP, 7, 30
through switch, 7, 9
through the fabric, 83

S

secret M_Key, 61
adding secret M_Keys, 65
conditions for situations, 74
configuring, 63
configuring local, 71
creating configuration, 63
deleting secret M_Keys, 69
disabling
functionality, 73
persistence, 72
displaying, 93
configuration, 48
persistence, 49
enabling
functionality, 73
persistence, 72
firmware guidelines, 76
merging subnets, 75
new Subnet Manager, 74
overview, 60
preparing, 62
setting replication password, 72
setting to use, 68
secure fabric management, 60
`setsmmkey` command, 55, 56
`setsmpriority` command, 47, 52, 54
`setsmrouting` command, 57
`setsubnetprefix` command, 53
setting
port speed, 99
replication password, 72
secret M_Key to use, 68
Subnet Manager
controlled handover, 54
management key, 55
prefix, 53
priority, 52
routing algorithm, 57
`showdisk` command, 22
`showfree` command, 23
`showfruinfo` command, 23
`showpsufru` command, 24
`showsmlog` command, 51
`showtemps` command, 20
`showtopology` command, 82
`showunhealthy` command, 18
`smconfigtest` command, 50
`sminfo` command, 46
smnodes
adding to list, 128
deleting from list, 129
displaying, 91
`smnodes`
`add` command, 128
`delete` command, 129
`list` command, 91
`smpartition`
`add` command, 107
`create` command, 105
`delete` command, 112
`list` command, 89
`modify` command, 109
`remove` command, 107
`start` command, 105
`smpartition` command, 102
guidelines, 103
`smpquery` command, 86, 87
`smsubnetprotection`
`add` command, 65
`delete` command, 69
`disablesecretmkey` command, 73
`enablesecretmkey` command, 73
`list active` command, 48, 93
`set-current` command, 68
`setlocalsecretmkey` command, 71
`setreplicationpassword` command, 72
`start` command, 63
`smsubnetprotection` command, 61
status
chassis, 21
fan, 20
flash drive, 22
power supply, 18
Subnet Manager, 46
Subnet Manager
added to secret M_Key fabric, 74
clearing management key, 56

configuration integrity, 50
controlled handover state, 47
controlling, 52
disabling, 59
displaying
configuration integrity, 50
location history, 46
log, 51
management key, 47
prefix, 47
priority, 47
routing algorithm, 47
status, 46
enabling, 58
management key, 47, 55, 56
monitoring, 45
setting
controlled handover, 54
management key, 55
prefix, 53
priority, 52
routing algorithm, 57
supportive tasks, 127
switch
GUIDs, 9
identity, 80
troubleshooting, 1
switch chip
checking boot status, 32
port
disabling, 36
enabling, 37
symbol errors, 96

T

topology
change, 95
displaying
detailed, 82
simple, 82
troubleshooting
network management, 6
troubleshooting the switch, 1

V

`version` command, 25