Sun Datacenter InfiniBand Switch 36

Administration Guide for Firmware Version 2.0

Part No.: E26430-01
November 2011, Revision A

Contents

Using This Documentation

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

Product Notes

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

http://www.oracle.com/pls/topic/lookup?ctx=E26698-01

Related Documentation

Documentation	Links
All Oracle products	http://www.oracle.com/documentation
Sun Datacenter InfiniBand Switch 36	http://www.oracle.com/pls/topic/lookup?ctx=E26698-01
Oracle Integrated Lights Out Manager (ILOM) 3.0	http://www.oracle.com/pls/topic/lookup?ctx=E19860-01

Feedback

Provide feedback on this documentation at:

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

Support and Accessibility

Description	Links
Access electronic support through My Oracle Support	http://support.oracle.com For hearing impaired: http://www.oracle.com/accessibility/support.html
Learn about Oracle's commitment to accessibility	http://www.oracle.com/us/corporate/accessibility/index.html

Troubleshooting the Switch

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

Related Information

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 22. 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. See *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. See "Check Chassis Status LEDs" on page 8 . 2. Verify that there is a LinkUp condition at the network management port. See "Check Network Management Port Status LEDs" on page 9. 3. Verfiy 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. See *Switch Installation*, accessing the management controller and "Restart the Management Controller" on page 28. 5. If you are unable to access the management controller through the USB connector, power cycle the switch. See *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. See *Switch Service*, removing the switch from the rack. See *Switch Installation*, installaing 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. See *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 28. 3. Check overall switch health. See "Display Switch General Health" on page 22 4. Verify that the switch is within operating temperatures and voltages. See "Display Switch Environmental and Operational Data" on page 25. 5. If the previous steps do not rectify the situation, replace the switch. See *Switch Service*, removing the switch from the rack. See *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 24. 2. If the previous step does not rectify the situation, replace the fan. See *Switch Service*, servicing fans. 3. If the previous step does not rectify the situation, install the fan into another available slot. See *Switch Service*, servicing fans. 4. If no other slots are available, replace the switch. See *Switch Service*, removing the switch from the rack. See *Switch Installation*, installing the switch into the rack.

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, and Management Key" on page 41. 2. If no Subnet Manager is active, start the Subnet Manager within the switch. See *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 49 and "Enable the Subnet Manager" on page 48.
After installation, not all links are operational.	1. Determine which links are nonoperational. See "Display Link Status" on page 34. 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 28. 2. If you are unable to access the management controller, power cycle the switch. See *Switch Service*, removing the switch from the rack. See *Switch Installation*, installing the switch into the rack. 3. Reperform the firmware upgrade. See *Switch Remote Administration*, upgrading the switch firmware.

Related Information

- "InfiniBand Fabric Problems" on page 3
- "Identifying LEDs" on page 6
- "Understanding Routing Through the Switch" on page 12
- "Switch GUIDs Overview" on page 14

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
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 62 "Find 1x, SDR, or DDR Links in the Fabric" on page 63 "Determine Which Links Are Experiencing Significant Errors" on page 64 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" on page 33. 3. If the problem is at a cable connection, swap the suspect cable with a known good cable or reconnect the cable to a known good remote port and repeat Step 1. See *Switch Service*, servicing data cables. 4. If the problem still remains at the cable connection, disable and re-enable the respective port and repeat Step 1. See "Disable a Port" on page 67 and "Enable a Port" on page 68. Temporary solution: • If the problem still remains, disable the affected port. See "Disable a Port" on page 67. Permanent solution: • If the problem still remains, replace the affected component or the switch. See *Switch Service*, servicing data cables. See remote port's documentation for replacement procedures. See *Switch Service*, removing the switch from the rack. See *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 data 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 64. 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" on page 33. 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 a Port" on page 67 and "Enable a Port" on page 68. 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. See *Switch Service*, servicing the data cables. See 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 63. 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" on page 33.
	3. Verify the cable connection at both ends.
	See *Switch Service*, servicing the data cables.
	4. Disable and re-enable the respective ports.
	See "Disable a Port" on page 67 and "Enable a Port" on page 68.
	5. If the previous steps do not rectify the problem, disable the port.
	See "Disable a Port" on page 67.
	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.
	See *Switch Service*, servicing the data cables.
	See the remote port's documentation for replacement procedures.
	See *Switch Service*, removin gthe switch from the rack.
	See *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 65.
	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 64. 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" on page 33. 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 32.

Related Information

- "Switch Hardware Problems" on page 1
- "Identifying LEDs" on page 6
- "Understanding Routing Through the Switch" on page 12
- "Switch GUIDs Overview" on page 14

Identifying LEDs

Status LEDs are used on many components of the switch chassis as a means of indicating the component's state.

- "Front Status LEDs" on page 7
- "Rear Status LEDs" on page 8
- "Check Chassis Status LEDs" on page 8
- "Check Network Management Port Status LEDs" on page 9
- "Check Link Status LEDs" on page 10
- "Check Power Supply Status LEDs" on page 10
- "Check Fan Status LEDs" on page 11

Related Information

- "Switch Hardware Problems" on page 1
- "InfiniBand Fabric Problems" on page 3
- "Understanding Routing Through the Switch" on page 12
- "Switch GUIDs Overview" on page 14

Front Status LEDs

FIGURE: Front Status LEDs

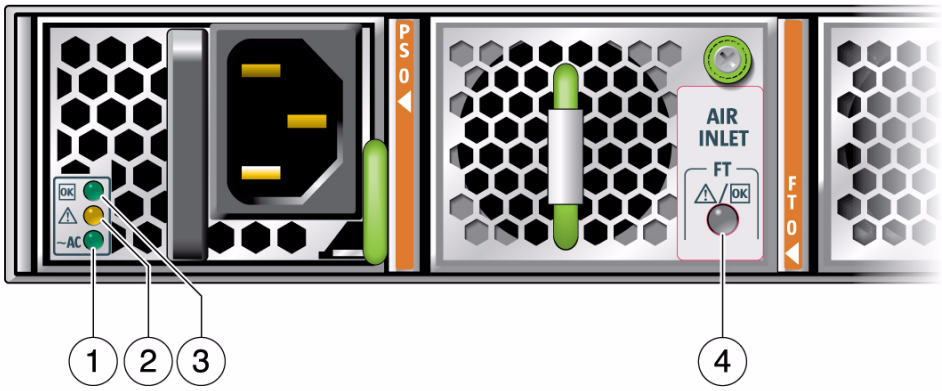

Figure Legend

1	Power supply AC LED
2	Power supply Attention LED
3	Power supply OK LED
4	Fan Attention LED

Related Information

- "Check Power Supply Status LEDs" on page 10
- "Check Fan Status LEDs" on page 11

Rear Status LEDs

FIGURE: Rear Status LEDs

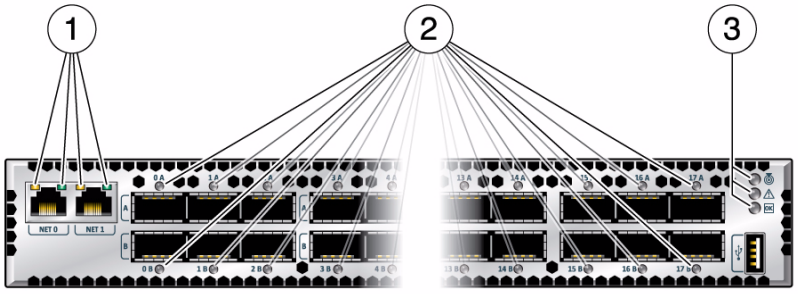

Figure Legend

1	NET status LEDs
2	Link status LEDs
3	Chassis status LEDs

Related Information

- "Check Chassis Status LEDs" on page 8
- "Check Network Management Port Status LEDs" on page 9
- "Check Link Status LEDs" on page 10

▼ Check Chassis Status LEDs

The chassis status LEDs are located on the left side of the rear panel. See "Rear Status LEDs" on page 8.

1. **Visually inspect the chassis status LEDs.**
2. **Compare what you see to this table.**

Glyph	Location	Name	Color	State and Meaning
	Top	Locator	White	On – No function. Off – Disabled. Flashing – The switch is identifying itself.
	Middle	Attention	Amber	On – Normal fault detected. Off – No faults detected. Flashing – No function.
	Bottom	OK	Green	On – Switch is functional without fault. Off – Switch is off or initializing. Flashing – No function.

Related Information

- "Display Switch Environmental and Operational Data" on page 25

▼ Check Network Management Port Status LEDs

The network management port status LEDs are located on the network management connector of the rear panel. See "Rear Status LEDs" on page 8.

1. **Visually inspect the NET status LEDs.**
2. **Compare what you see to this table.**

Name	Location	Color	State and Meaning
Link speed	Left	Amber or green	Amber on – 100BASE-T link. Green on – 1000BASE-T link. Off – No link or link down. Flashing – No function.
Activity	Right	Green	On – No function. Off – No activity. Flashing – Packet activity.

▼ Check Link Status LEDs

The link status LEDs are located at the data cable connectors of the rear panel. See "Rear Status LEDs" on page 8.

1. **Visually inspect the link status LEDs.**
2. **Compare what you see for a particular link to this table.**

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

Related Information

- "Display Link Status" on page 34
- "Display the Link Status of a Node" on page 55
- "Display Counters for a Node" on page 57

▼ Check Power Supply Status LEDs

The power supply status LEDs are located on the power supply at the front of the chassis. See "Front Status LEDs" on page 7.

1. **Visually inspect the power supply's status LEDs.**
2. **Compare what you see on the power supply to this table.**

Glyph	Location	Name	Color	State and Meaning
OK	Top	OK	Green	On – 12 VDC is supplied. Off – No DC voltage is present. Flashing – No function.
⚠	Middle	Attention	Amber	On – Fault detected, 12 VDC shut down. Off – No faults detected. Flashing – No function.
~ AC	Bottom	AC	Green	On – AC power present and good. Off – AC power not present. Flashing – No function.

Caution – If a power supply has shut down because of a thermal or overcurrent condition, signified by the amber Attention LED lighting, remove the respective power cord from the chassis. Allow the power supply to completely cool for at least

15 minutes. A shorter cooling time might cause damage to the power supply when the power cord is reattached. If the Attention LED lights amber upon reattaching the power cord, replace the power supply.

Related Information

- "Display Power Supply Status" on page 22
- "Check Board-Level Voltages" on page 23

▼ Check Fan Status LEDs

The fan status LEDs are located in the lower right corner of the fans at the front of the switch chassis. See "Front Status LEDs" on page 7.

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

Related Information

- "Display Fan Status" on page 24

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 12
- "QSFP Connectors and Link LEDs to Switch Chip Port Routes" on page 13
- "Signal Route Through the Switch" on page 14

Related Information

- "Switch Hardware Problems" on page 1
- "InfiniBand Fabric Problems" on page 3
- "Identifying LEDs" on page 6

- "Switch GUIDs Overview" on page 14

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
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 32
- "Display Link Status" on page 34
- "Display Switch Chip Port Status" on page 35
- "QSFP Connectors and Link LEDs to Switch Chip Port Routes" on page 13
- "Signal Route Through the Switch" on page 14

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

Connector Group	Connector A	Connector B	Connector Group	Connector A	Connector B
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 32
- "Display Link Status" on page 34
- "Display Switch Chip Port Status" on page 35
- "Switch Chip Port to QSFP Connectors and Link LED Routes" on page 12
- "Signal Route Through the Switch" on page 14

Signal Route Through the Switch

By combining the information from the tables in "Understanding Routing Through the Switch" on page 12, 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 13, 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 12, 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 12
- "QSFP Connectors and Link LEDs to Switch Chip Port Routes" on page 13
- "Display the Switch Chip Port to QSFP Connector Mapping" on page 32
- "Display a Route Through the Fabric" on page 55

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.

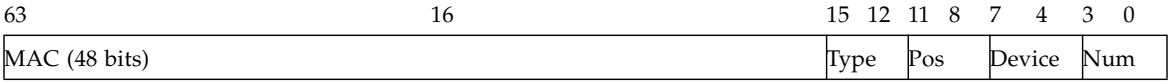

The five fields of the GUID are described as follows:

- **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 52
- "Switch Hardware Problems" on page 1
- "InfiniBand Fabric Problems" on page 3
- "Identifying LEDs" on page 6
- "Understanding Routing Through the Switch" on page 12

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.

Related Information

- "Troubleshooting the Switch" on page 1
- "Administering the Chassis" on page 21
- "Administering the I4 Switch Chip" on page 31
- "Administering the InfiniBand Fabric" on page 51
- "Administering the Subnet Manager" on page 39

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 44
		Configure the prefix.	"Set the Subnet Manager Prefix" on page 45
		Enable controlled handover.	"Enable Subnet Manager Controlled Handover" on page 46
		Configure the M_Key.	"Set the Subnet Manager Management Key" on page 46
2.	Configure the InfiniBand fabric.	Create a fabric configuration.	"Create a Fabric Configuration" on page 72
3.	Partition the InfiniBand fabric.	Collect GUIDs.	"Identify All Switches in the Fabric" on page 52 "Identify All HCAs in the Fabric" on page 53
		Create the `smnodes` file.	"Use the `smnodes` Command" on page 83
		Create the partition information file.	"Determine Partitions and P_Keys" on page 84
		Configure partitions and P_Keys.	"Create a User Partition" on page 85
		Add GUIDs to the partitions.	"Add or Remove a Port From a Partition" on page 87
4.	Save the configuration.	Back up the configuration for restoration at a later time, should it become necessary.	*Switch Remote Administration*, backing up the configuration.

Related Information

- "Hardware Command Overview" on page 19
- "InfiniBand Command Overview" on page 19
- "ILOM Command Overview" on page 20

Hardware Command Overview

The management controller uses a simplified Linux OS and file system. From the # prompt on the management controller, you can type hardware commands to perform some administrative and management tasks. Alternatively, you can use the Oracle ILOM restricted Linux shells, which are the prefferred method of issuing hardware commands. See *Switch Remote Administration*, Linux restricted shells.

Hardware commands are user-friendly and can perform some testing upon the switch chip, enabling greater control of the switch and its operation.

After you log in to the `root` account, the shell prompt (`#`) appears, and you can enter shell commands. Enter the hardware commands in this format:

`#` *command* [*arguments*] [*arguments*] . . .

Related Information

- *Switch Reference*, understanding hardware commands
- "Configuration Overview" on page 17
- "InfiniBand Command Overview" on page 19
- "ILOM Command Overview" on page 20

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 `root` account, the shell prompt (`#`) appears, and you can enter shell commands. Alternatively, you can use the Oracle ILOM restricted Linux shells, which are the prefferred method of issuing InfiniBand commands. See *Switch Remote Administration*, Linux restricted shells.

Enter the InfiniBand commands in this format:

`#` *command* [*option*] [*option*] . . .

Related Information

- *Switch Reference*, understanding InfiniBand commands
- "Configuration Overview" on page 17
- "Hardware Command Overview" on page 19
- "ILOM Command Overview" on page 20

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. Enter the Oracle ILOM commands in this format:

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

Information about 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=E19197-01`

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 Administration*, understanding Oracle ILOM commands
- "Configuration Overview" on page 17
- "Hardware Command Overview" on page 19
- "InfiniBand Command Overview" on page 19

Administering the Chassis

These topics describe the administration of the switch.

Related Information

Monitoring the Chassis

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

Related Information

- "Monitoring the InfiniBand Fabric" on page 51
- "Monitoring the Subnet Manager" on page 39
- "Controlling the Chassis" on page 28

▼ 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.**

```
# showunhealthy
OK - No unhealthy sensors
#
```

Related Information

- *Switch Reference*, `showunhealthy` command
- "Display Switch Environmental and Operational Data" on page 25

▼ Display Power Supply Status

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

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

```
# checkpower
PSU 0 present OK
PSU 1 present Alert
#
```

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.**

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

   b. **Wait one minute.**

   c. **Reinstall the power supply.**

      See *Switch Service*, installing the power supply.

d. Reattach the power cord.

See *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.

See *Switch Service*, servicing the power supplies.

Related Information

- *Switch Reference*, `checkpower` command
- "Check Board-Level Voltages" on page 23
- "Display Switch Environmental and Operational Data" on page 25

▼ 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.**

```
# checkvoltages
Voltage ECB OK
Measured 3.3V Main = 3.27 V
Measured 3.3V Standby = 3.35 V
Measured 12V = 11.97 V
Measured 5V = 5.02 V
Measured VBAT = 3.24 V
Measured 2.5V = 2.52 V
Measured 1.8V = 1.78 V
Measured I4 1.2V = 1.22 V
All voltages OK
#
```

Related Information

- *Switch Reference*, `checkvoltages` command
- "Display Power Supply Status" on page 22
- "Display Switch Environmental and Operational Data" on page 25

▼ 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.**

```
# showtemps
Back temperature 32
Front temperature 32
SP temperature 44
Switch temperature 45, maxtemperature 46
All temperatures OK
#
```

Related Information

- *Switch Reference*, `showtemps` command
- "Display Fan Status" on page 24
- "Display Switch Environmental and Operational Data" on page 25

▼ 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.**

```
# getfanspeed
Fan 0 not present
Fan 1 running at rpm 11212
Fan 2 running at rpm 11313
Fan 3 running at rpm 11521
Fan 4 not present
#
```

Related Information

- *Switch Reference*, `getfanspeed` command
- "Display Switch Environmental and Operational Data" on page 25

▼ 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.**

```
# 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 = 3.27 V
Measured 3.3V Standby = 3.35 V
Measured 12V = 11.97 V
Measured 5V = 5.02 V
Measured VBAT = 3.24 V
Measured 2.5V = 2.52 V
Measured 1.8V = 1.78 V
Measured I4 1.2V = 1.22 V
Voltage test returned OK
Starting PSU test:
PSU 0 present OK
PSU 1 present OK
PSU test returned OK
Starting Temperature test:
Back temperature 32
Front temperature 32
SP temperature 44
Switch temperature 44, maxtemperature 46
Temperature test returned OK
Starting FAN test:
Fan 0 not present
Fan 1 running at rpm 12075
Fan 2 running at rpm 11960
Fan 3 running at rpm 12075
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
Environment test PASSED
#
```

Related Information

- *Switch Reference*, `env_test` command

▼ Display Chassis FRU ID

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

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

```
# showfruinfo
 UNIX_Timestamp32              :  Tue Mar 23 17:02:15 2010
 Sun_Fru_Description           :  ASSY,NM2-36P
 Vendor_ID_Code                :  11 E1
 Vendor_ID_Code_Source         :  01
 Vendor_Name_And_Site_Location :  4577 CELESTICA CORP. SAN JOSE CA US
 Sun_Part_Number               :  5413495
 Sun_Serial_Number             :  0110SJC-1010NG0045
 Serial_Number_Format          :  4V3F1-2Y2W2X4S
 Initial_HW_Dash_Level         :  01
 Initial_HW_Rev_Level          :  50
 Sun_Fru_Shortname             :  NM2, 36 ports
 Sun_Hazard_Class_Code         :  YL
 Sun_SpecPartNo                :  885-1507-01
#
```

Related Information

- *Switch Reference*, `showfruinfo` command
- "Display Power Supply FRU ID" on page 26

▼ Display Power Supply FRU ID

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

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

```
# showpsufru slot
```

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

```
# showpsufru 0
 Sun_SpecPartNo           :  885-1390-01
 UNIX_Timestamp32         :  Sun Jan  3 15:35:39 2010
 Sun_Part_Number          :  3002234
 Sun_Serial_No            :  BF15WA
 Vendor_ID_Code           :  03ad
 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  :  Delta Energy Systems
 IPMI_Board_Product_Name  :  A236
 IPMI_Board_Serial_Number:  1357ZHO-0952BF15WA
 IPMI_Board_Part_Number   :  300-2234-01
#
```

Related Information

- *Switch Reference*, `showpsufru` command
- "Display Chassis FRU ID" on page 26

▼ 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.**

```
# version
SUN DCS 36p version: 2.0.5-1
Build time: Nov 25 2011 12:20:22
SP board info:
Manufacturing Date: 2009.02.19
Serial Number: "NCD2S0133"
Hardware Revision: 0x0100
Firmware Revision: 0x0102
```

```
BIOS version: NOW1R112
BIOS date: 04/24/2009
#
```

Related Information

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

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 either be the `root` user of the management controller or access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Restart the Management Controller" on page 28
- "Change the Administrator Password" on page 29

Related Information

- "Controlling the I4 Switch Chip" on page 36
- "Controlling the InfiniBand Fabric" on page 61
- "Controlling the Subnet Manager" on page 43
- "Monitoring the Chassis" on page 21

▼ Restart the Management Controller

Should the management controller enter an indeterminate state, you can reboot 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 – The `reboot` command severs any management console link to the management controller. You must reaccess the management controller to regain administrative control.

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

```
# reboot
Broadcast message from root (pts/0) (Tue Oct 18 23:24:50 2011):
The system is going down for reboot NOW!
# Connection to 123.45.67.89 closed by remote host.
Connection to 123.45.67.89 closed.
#
```

The management controller power cycles and reinitializes.

Related Information

- "Change the Administrator Password" on page 29

▼ Change the Administrator Password

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

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

```
# passwd user
```

where *user* is name of the user who's password is to be changed. For example:

```
# 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 28

Administering the I4 Switch Chip

These topics describe the administration of the I4 switch chip.

Related Information

- "Troubleshooting the Switch" on page 1
- "Understanding Administrative Commands" on page 17
- "Administering the Chassis" on page 21
- "Administering the InfiniBand Fabric" on page 51
- "Administering the Subnet Manager" on page 39

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 either be the `root` user of the management controller or access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

Related Information

- "Monitoring the Chassis" on page 21
- "Monitoring the InfiniBand Fabric" on page 51
- "Monitoring the Subnet Manager" on page 39
- "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.**

```
# dcsport -port port
```

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

Alternatively, type.

```
# dcsport -connector connector
```

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

For example:

```
# dcsport -port 7
DCS-36P Switch port 7 maps to connector 14A
# dcsport -connector 14A
Connector 14A maps to Switch port 7
#
```

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 12

▼ Locate a Switch Chip or Connector From the GUID

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=0x0001 guid=0x03ba01234567a0a0 dev=48438 Port=21
```

   In this example, the GUID is `0x03ba01234567a0a0` and the port is 21.

2. **Truncate all digits of the GUID except the four digits on the right.**

   For example, GUID `0x03ba01234567a0a0` is truncated to `a0a0`.

   The `a0` means that the node is inside of a Sun Datacenter InfiniBand Switch 36. The `a` means that the node is located in I4 switch chip A, the only switch chip. The port is 21.

3. **Use the `ibswitches` command to verify the information.**

```
# 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
#
```

4. **Use the tables in** "Understanding Routing Through the Switch" on page 12 **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 12, I4 switch chip, port 21 routes to connector 1B.

5. **Alternatively, you can use the `dcsport` command to display routing between ports and connections.**

   See "Display the Switch Chip Port to QSFP Connector Mapping" on page 32.

Related Information

- "Switch GUIDs Overview" on page 14
- *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.**

```
# checkboot
Switch OK
#
```

Related Information

- *Switch Reference*, `checkboot` command
- "Display Switch Environmental and Operational Data" on page 25
- *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.**

```
# listlinkup
Connector  0A Present <-> Switch Port 20 up (Enabled)
Connector  1A Present <-> Switch Port 22 up (Enabled)
Connector  2A Present <-> Switch Port 24 up (Enabled)
.
.
.
Connector 17A Not present
Connector  0B Present <-> Switch Port 19 up (Enabled)
Connector  1B Present <-> Switch Port 21 up (Enabled)
.
.
.
Connector 17B Not present
#
```

Related Information

- *Switch Reference*, `listlinkup` command

- "Display the Link Status of a Node" on page 55

▼ 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.**

```
# getportstatus connector | Switch port
```

where:

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

For example:

```
# getportstatus Switch 7
Port status for connector 14A Switch Port 7
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
#
```

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 55

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 either be the `root` user of the management controller or 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

Related Information

- "Controlling the Chassis" on page 28
- "Controlling the InfiniBand Fabric" on page 61
- "Controlling the Subnet Manager" on page 43
- "Monitoring the I4 Switch Chip" on page 31

▼ 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.

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

```
# disableswitchport [--reason=reason] Switch port
```

where:

- *reason* is the reason for disabling the port, `Blacklist` or `Partition`.

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

Alternatively, type.

```
# disableswitchport [--reason=reason] connector
```

where *connector* is the number of the QSFP connector (0A–15B).

For example:

```
# disableswitchport Switch 14
Disable connector 9A Switch port 14
MKEY 8 0000000000000000
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
#
```

Related Information

- *Switch Reference*, `disableswitchport` command
- "Disable a Port" on page 67
- "Enable a Switch Chip Port" on page 37

▼ Enable a Switch Chip Port

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

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

```
# enableswitchport [--reason=reason] Switch port
```

where:

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

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.

```
# enableswitchport [--reason=reason] connector
```

where *connector* is the number of the QSFP connector (`0A–15B`).

For example:

```
# enableswitchport Switch 14
Enable connector 9A Switch port 14
MKEY 8 0000000000000000
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:..................Polling
LinkSpeedActive:................2.5 Gbps
LinkSpeedEnabled:...............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
#
```

Related Information

- *Switch Reference*, `enableswitchport` command
- "Enable a Port" on page 68
- "Disable a Switch Chip Port" on page 36

Administering the Subnet Manager

These topics describe the administration of the Subnet Manager.

Related Information

- "Troubleshooting the Switch" on page 1
- "Understanding Administrative Commands" on page 17
- "Administering the Chassis" on page 21
- "Administering the I4 Switch Chip" on page 31
- "Administering the InfiniBand Fabric" on page 51

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 either be the `root` user of the management controller or access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

Related Information

- "Monitoring the Chassis" on page 21
- "Monitoring the I4 Switch Chip" on page 31
- "Monitoring the InfiniBand Fabric" on page 51
- "Controlling the Subnet Manager" on page 43

▼ Display Subnet Manager Status

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

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

```
# getmaster
Local SM enabled and running
20111021 13:46:23 Master SubnetManager on sm lid 1 sm guid 0x3ba01234567a0a0 :
SUN DCS 36P QDR o4nm2-36p-2 10.172.144.66
#
```

Related Information

- *Switch Reference*, `getmaster` command
- "Display Recent Subnet Manager Activity" on page 40
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Verify the Subnet Manager Configuration Integrity" on page 41
- "Display the Subnet Manager Log" on page 42

▼ Display Recent Subnet Manager Activity

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

```
# getmaster -l
Local SM enabled and running
Last ring buffer history listed:
20111018 19:16:12 whereismaster started
20111018 19:16:12 Master SubnetManager on sm lid 0 sm guid 0x212856cfe2c0a0 :
20111018 19:16:29 Master SubnetManager on sm lid 6 sm guid 0x212856cfe2c0a0 :
SUN IB QDR 36p switch mnm34-98
#
```

Related Information

- *Switch Reference*, `getmaster` command
- "Display Subnet Manager Status" on page 40
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Verify the Subnet Manager Configuration Integrity" on page 41
- "Display the Subnet Manager Log" on page 42

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

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

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

```
# setsmpriority list
Current SM settings:
smpriority 8
controlled_handover TRUE
subnet_prefix 0xfe80000000000000
M_Key None
#
```

Related Information

- *Switch Reference*, `setsmpriority` command
- "Controlling the Subnet Manager" on page 43
- "Display Subnet Manager Status" on page 40
- "Display Recent Subnet Manager Activity" on page 40
- "Verify the Subnet Manager Configuration Integrity" on page 41
- "Display the Subnet Manager Log" on page 42

▼ 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.**

```
# smconfigtest [spine|leaf]
```

where:

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

For example:

```
# smconfigtest leaf
#
```

Note – Output is only displayed if errors are found.

Related Information

- *Switch Reference*, `smconfigtest` command
- "Display Subnet Manager Status" on page 40
- "Display Recent Subnet Manager Activity" on page 40
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Display the Subnet Manager Log" on page 42

▼ Display the Subnet Manager Log

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

```
# showsmlog
Jun 30 14:44:11 008981 [B6694B90] 0x02 -> osm_ucast_mgr_process: minhop tables
configured on all switches
Jun 30 14:44:11 011981 [B6694B90] 0x02 -> SUBNET UP
Jun 30 14:44:11 025981 [B6694B90] 0x02 -> Fabric has 1 switches - topology is
not fat-tree. Falling back to default routing
Jun 30 14:44:11 025981 [B6694B90] 0x01 -> ucast_mgr_route: ftree: cannot build
lid matrices.
```

```
Jun 30 14:44:11 025981 [B6694B90] 0x02 -> osm_ucast_mgr_process: minhop tables
configured on all switches
Jun 30 14:44:11 027981 [B6694B90] 0x02 -> SUBNET UP
.
.
.
```

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 40
- "Display Recent Subnet Manager Activity" on page 40
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Verify the Subnet Manager Configuration Integrity" on page 41

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 either be the `root` user of the management controller or access them through the `/SYS/Fabric_Mgmt` Linux shell targets of the Oracle ILOM CLI.

- "Set the Subnet Manager Priority" on page 44
- "Set the Subnet Manager Prefix" on page 45
- "Enable Subnet Manager Controlled Handover" on page 46
- "Set the Subnet Manager Management Key" on page 46
- "Clear the Subnet Manager Management Key" on page 48
- "Enable the Subnet Manager" on page 48
- "Disable the Subnet Manager" on page 49

Related Information

- "Controlling the Chassis" on page 28
- "Controlling the I4 Switch Chip" on page 36
- "Controlling the InfiniBand Fabric" on page 61
- "Monitoring the Subnet Manager" on page 39

▼ 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 managment controller, disable the Subnet Manager.**

   See "Disable the Subnet Manager" on page 49.

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

```
# setsmpriority priority
```

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

```
# setsmpriority 5
Current SM settings:
smpriority 5
controlled_handover FALSE
subnet_prefix 0xfe80000000000000
M_Key None
#
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 48.

Related Information

- *Switch Reference*, `setsmpriority` command
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Set the Subnet Manager Prefix" on page 45
- "Enable Subnet Manager Controlled Handover" on page 46
- "Set the Subnet Manager Management Key" on page 46

- "Clear the Subnet Manager Management Key" on page 48
- "Enable the Subnet Manager" on page 48
- "Disable the Subnet Manager" on page 49

▼ Set the Subnet Manager Prefix

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

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

   See "Disable the Subnet Manager" on page 49.

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

```
# setsubnetprefix 0xabbababe
Current SM settings:
smpriority 5
controlled_handover FALSE
subnet_prefix 0xabbababe
M_Key None
#
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 48.

Related Information

- *Switch Reference*, `setsubnetprefix` command
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Set the Subnet Manager Priority" on page 44
- "Enable Subnet Manager Controlled Handover" on page 46
- "Set the Subnet Manager Management Key" on page 46
- "Clear the Subnet Manager Management Key" on page 48
- "Enable the Subnet Manager" on page 48
- "Disable the Subnet Manager" on page 49

▼ 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. See *Switch Reference*, `setcontrolledhandover` command for more information.

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

   See "Disable the Subnet Manager" on page 49.

2. **Enable controlled handover.**

```
# setcontrolledhandover TRUE
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xabbababe
M_Key None
#
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 48.

Related Information

- *Switch Reference*, `setcontrolledhandover` command
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Set the Subnet Manager Priority" on page 44
- "Set the Subnet Manager Prefix" on page 45
- "Set the Subnet Manager Management Key" on page 46
- "Clear the Subnet Manager Management Key" on page 48
- "Enable the Subnet Manager" on page 48
- "Disable the Subnet Manager" on page 49

▼ Set the Subnet Manager Management Key

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

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

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

   See "Disable the Subnet Manager" on page 49.

2. **Set the management key.**

```
# setsmmkey m_key
```

   where *m_key* is a 12-digit hexadecimal number

   For example:

```
# setsmmkey 0xabbababe0001
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xabbababe
M_Key 0xabbababe0001
#
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 48.

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, and Management Key" on page 41
- "Set the Subnet Manager Priority" on page 44
- "Set the Subnet Manager Prefix" on page 45
- "Enable Subnet Manager Controlled Handover" on page 46
- "Clear the Subnet Manager Management Key" on page 48
- "Enable the Subnet Manager" on page 48
- "Disable the Subnet Manager" on page 49

▼ Clear the Subnet Manager Management Key

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

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

   See "Disable the Subnet Manager" on page 49.

2. **Clear the management key.**

```
# setsmmkey none
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xabbababe
M_Key None
#
```

3. **Enable the Subnet Manager.**

   See "Enable the Subnet Manager" on page 48.

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

Related Information

- *Switch Reference*, `setsmmkey` command
- "Set the Subnet Manager Management Key" on page 46
- "Display Subnet Manager Priority, Controlled Handover State, Prefix, and Management Key" on page 41
- "Set the Subnet Manager Priority" on page 44
- "Set the Subnet Manager Prefix" on page 45
- "Enable Subnet Manager Controlled Handover" on page 46
- "Enable the Subnet Manager" on page 48
- "Disable the Subnet Manager" on page 49

▼ 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.**

```
# enablesm
Starting IB Subnet Manager.                                [  OK  ]
Starting partitiond daemon.                                [  OK  ]
#
```

Related Information

- *Switch Reference*, `enablesm` command
- "Disable the Subnet Manager" on page 49
- "Set the Subnet Manager Priority" on page 44
- "Set the Subnet Manager Prefix" on page 45
- "Enable Subnet Manager Controlled Handover" on page 46
- "Set the Subnet Manager Management Key" on page 46
- "Clear the Subnet Manager Management Key" on page 48

▼ 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.**

```
# disablesm
Stopping partitiond daemon.                                [  OK  ]
Stopping IB Subnet Manager..                               [  OK  ]
#
```

Related Information

- *Switch Reference*, `disablesm` command
- "Enable the Subnet Manager" on page 48
- "Set the Subnet Manager Priority" on page 44
- "Set the Subnet Manager Prefix" on page 45
- "Enable Subnet Manager Controlled Handover" on page 46
- "Set the Subnet Manager Management Key" on page 46
- "Clear the Subnet Manager Management Key" on page 48

Administering the InfiniBand Fabric

These topics describe the administration of the InfiniBand fabric.

Related Information

- "Troubleshooting the Switch" on page 1
- "Understanding Administrative Commands" on page 17
- "Administering the Chassis" on page 21
- "Administering the I4 Switch Chip" on page 31
- "Administering the Subnet Manager" on page 39

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 either be the `root` user of the management controller or access them through the `/SYS/Fabric_Mgmt` Linux shell target of the Oracle ILOM CLI.

- "Display the Link Status of a Node" on page 55
- "Display Counters for a Node" on page 57
- "Display Low-Level Detailed Information About a Node" on page 58
- "Display Low-Level Detailed Information About a Port" on page 58
- "Display the InfiniBand Fabric Configuration" on page 60
- "Display the InfiniBand Fabric Partition Configuration" on page 60

Related Information

- "Monitoring the Chassis" on page 21
- "Monitoring the I4 Switch Chip" on page 31
- "Monitoring the Subnet Manager" on page 39
- "Controlling the InfiniBand Fabric" on page 61
- "Configuring the Fabric Director Node List" on page 69
- "Partitioning the InfiniBand Fabric" on page 80

▼ 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.**

```
# 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
#
```

Related Information

- *Switch Reference*, `ibswitches` command
- *Switch Reference*, `ibhosts` command
- "Identify All HCAs in the Fabric" on page 53

▼ Identify All HCAs in the Fabric

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

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

---

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.

```
# 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"
.
.
.
#
```

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 52

▼ 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.**

```
# showtopology
```

Related Information

- *Switch Reference*, `showtopology` command
- "Determine Changes to the InfiniBand Fabric Topology" on page 62
- "Perform Comprehensive Diagnostics for the Entire Fabric" on page 62

▼ 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 must be the `root` user to run the `ibnetdiscover` command.

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

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

```
# ibnetdiscover
```

Related Information

- *Switch Reference*, `ibnetdiscover` command
- "Perform Comprehensive Diagnostics for the Entire Fabric" on page 62

▼ 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.**

```
# ibtracert slid dlid
```

where:

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

For example:

```
# ibtracert 15 14
From switch {0x0021283a8389a0a0} portnum 0 lid 15-15 "Sun DCS 36 QDR switch
localhost"
[23] -> ca port {0x0003ba000100e38a}[2] lid 14-14 "mnm-43 HCA-1"
To ca {0x0003ba000100e388} portnum 2 lid 14-14 "mnm-43 HCA-1"
#
```

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

For this example:

- The route starts at switch with GUID `0x0021283a8389a0a0` and is using port `0`. The switch is LID `15` and in the description, the switch host's name is `Sun DCS 36 QDR switch localhost`.
- The route enters at port `23` of the CA with GUID `0x0003ba000100e38a` and exits at port `2`. The CA is LID `14`.

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.**

```
# ibportstate lid port
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example:

```
# ibportstate 15 28
PortInfo:
# Port info: Lid 15 port 28
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 15 DR path slid 1; dlid 65535; 0,28 port 28
LinkState:.....................Active
PhysLinkState:.................LinkUp
LinkWidthSupported:............1X or 4X
LinkWidthEnabled:..............1X or 4X
LinkWidthActive:...............4X
LinkSpeedSupported:............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedEnabled:..............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkSpeedActive:...............10.0 Gbps
#
```

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

Related Information

- *Switch Reference*, `ibportstate` command
- "Display Switch Chip Port Status" on page 35
- "Display Link Status" on page 34
- "Find 1x, SDR, or DDR Links in the Fabric" on page 63
- "Set Port Speed" on page 66

▼ 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.**

```
# 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.

```
# perfquery 15 28
# Port counters: Lid 15 port 28
PortSelect:......................28
CounterSelect:...................0x1b01
SymbolErrors:....................0
.
.
.
VL15Dropped:.....................0
XmtData:.........................33791184
RcvData:.........................33789168
XmtPkts:.........................469322
RcvPkts:.........................469294
#
```

Related Information

- *Switch Reference*, `perfquery` command
- "Clear Data and Error Counters" on page 65

▼ 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.**

```
# smpquery switchinfo lid
```

where *lid* is the LID of the node.

For example, to see detailed information about a switch with LID 15, type.

```
# smpquery switchinfo 15
# Switch info: Lid 15
LinearFdbCap:.....................49152
RandomFdbCap:.....................0
McastFdbCap:......................4096
LinearFdbTop:.....................15
DefPort:..........................0
DefMcastPrimPort:.................255
DefMcastNotPrimPort:..............255
LifeTime:.........................18
StateChange:......................0
LidsPerPort:......................0
PartEnforceCap:...................32
InboundPartEnf:...................1
OutboundPartEnf:..................1
FilterRawInbound:.................1
FilterRawOutbound:................1
EnhancedPort0:....................1
#
```

Related Information

- *Switch Reference*, `smpquery` command
- "Display Low-Level Detailed Information About a Port" on page 58

▼ Display Low-Level Detailed Information About a Port

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

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

```
# smpquery portinfo lid port
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example, to see detailed information about port 23 on the switch with LID 15, type.

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

```
# smpquery portinfo 15 28
# Port info: Lid 15 port 28
Mkey:.............................0x0000000000000000
GidPrefix:........................0x0000000000000000
Lid:..............................0x0000
SMLid:............................0x0000
CapMask:..........................0x0
DiagCode:.........................0x0000
MkeyLeasePeriod:..................0
LocalPort:........................28
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
#
```

Related Information

- *Switch Reference*, `smpquery` command
- "Display Low-Level Detailed Information About a Node" on page 58

▼ 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.

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

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

In this example, the element named `primary` has IP address 123.45.67.89 and is a Sun Datacenter InfiniBand Switch 36. It 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 Network QDR InfiniBand Gateway Switch.

Related Information

- *Switch Reference*, `fdconfig` command
- "Configuring the Fabric Director Node List" on page 69

▼ Display the InfiniBand Fabric Partition Configuration

If you have partitioned the InfiniBand fabric, you can display the active partition configuration or the pending modified partition configuration with the `smpartition` command.

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

```
# smpartition list partition
```

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

For example:

```
# smpartition list active
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 12
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
#
```

In this example, this is the 12th configuration file to be created. The partition key is set to a default of `7fff`. IP protocol over InfiniBand is configured. All hosts (`CAS`), all switches, and the management controller itself are granted `full` membership priviledges.

Related Information

- *Switch Command Reference*, `smpartition` command
- "Partitioning the InfiniBand Fabric" on page 80
- "Display the InfiniBand Fabric Topology (Detailed)" on page 54

Controlling the InfiniBand Fabric

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

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

- "Reset a Port" on page 65
- "Set Port Speed" on page 66
- "Disable a Port" on page 67
- "Enable a Port" on page 68

Related Information

- "Controlling the Chassis" on page 28
- "Controlling the I4 Switch Chip" on page 36
- "Controlling the Subnet Manager" on page 43
- "Monitoring the InfiniBand Fabric" on page 51
- "Configuring the Fabric Director Node List" on page 69
- "Partitioning the InfiniBand Fabric" on page 80

▼ 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.**

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

Related Information

- *Switch Reference*, `ibdiagnet` command
- "Display the InfiniBand Fabric Topology (Detailed)" on page 54

▼ Determine Changes to the InfiniBand Fabric Topology

If your fabric has a number of nodes that are suspect, 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.**

```
# generatetopology date.topo
```

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

```
# 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
#
```

2. **After an event, compare the present topology to that saved in the topology file.**

```
# matchtopology date.topo
```

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

```
# 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.
----------------------------------------------------------------
#
```

Related Information

- *Switch Reference*, `generatetopology` command
- *Switch Reference*, `matchtopology` command
- "Display the InfiniBand Fabric Topology (Simple)" on page 54

▼ 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.**

```
# 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 55
- "Determine Which Links Are Experiencing Significant Errors" on page 64

▼ 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.**

```
# ibdiagnet -c 1000 -P all=1
```

In this instance of the `ibdiagnet` command, 1000 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 63
- "Display Counters for a Node" on page 57

▼ 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.**

```
# ibdiagnet -pc
```

Related Information

- *Switch Reference*, `ibdiagnet` command
- *Switch Reference*, `perfquery` command
- "Display Counters for a Node" on page 57

▼ Reset a Port

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

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

```
# ibportstate lid port reset
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example:

```
# ibportstate 1 28 reset
Initial PortInfo:
# Port info: Lid 1 port 28
LinkState:.....................Down
PhysLinkState:.................Polling
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 28
LinkState:.....................Down
PhysLinkState:.................Disabled

After PortInfo set:
# Port info: Lid 1 port 28
LinkState:.....................Down
PhysLinkState:.................Polling
#
```

Related Information

- *Switch Reference*, `ibportstate` command
- "Disable a Port" on page 67
- "Enable a Port" on page 68

▼ 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 Gbyte/sec.

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

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

where:

- *lid* is the LID of the node.

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

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

For example:

```
# ibportstate 15 28 speed 1
Initial PortInfo:
# Port info: Lid 15 port 28
LinkSpeedEnabled:.............2.5 Gbps or 5.0 Gbps or 10.0 Gbps

After PortInfo set:
# Port info: Lid 15 port 28
LinkSpeedEnabled:.............2.5 Gbps
# ibportstate 15 28 speed 7
Initial PortInfo:
# Port info: Lid 15 port 28
LinkSpeedEnabled:.............2.5 Gbps
After PortInfo set:
# Port info: Lid 15 port 28
LinkSpeedEnabled:.............2.5 Gbps or 5.0 Gbps or 10.0 Gbps
#
```

Related Information

- *Switch Reference*, `ibportstate` command
- "Display the Link Status of a Node" on page 55

▼ Disable a Port

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

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

```
# ibportstate lid port disable
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example:

```
# ibportstate 15 28 disable
Initial PortInfo:
# Port info: Lid 15 port 28
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 15 port 28
LinkState:.....................Down
PhysLinkState:.................Disabled
#
```

Related Information

- *Switch Reference*, `ibportstate` command
- "Disable a Switch Chip Port" on page 36
- "Enable a Port" on page 68
- "Reset a Port" on page 65

▼ Enable a Port

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

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

```
# ibportstate lid port enable
```

where:

- *lid* is the LID of the node.

- *port* is the port of the node.

For example:

```
# ibportstate 15 28 enable
Initial PortInfo:
# Port info: Lid 15 port 28
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 15 port 28
LinkState:.......................Down
PhysLinkState:...................PortConfigurationTraining
#
```

Related Information

- *Switch Reference*, `ibportstate` command
- "Enable a Switch Chip Port" on page 37
- "Disable a Port" on page 67
- "Reset a Port" on page 65

Configuring the Fabric Director Node List

These topics enable you to configure the Fabric Director node list.

Related Information

- "Monitoring the InfiniBand Fabric" on page 51
- "Controlling the InfiniBand Fabric" on page 61
- "Partitioning the InfiniBand Fabric" on page 80

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 eachother 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`.

One Fabric Director 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 Directors instances on those fabric elements. This master Fabric Director creates a model of the fabric which is also in the form of an SNMP MIB, as defined by the `SUN-FABRIC-MIB` specification.

By default, Fabric Directors run in non-master mode. Only when a Fabric Director reads the fabric node list and discovers that its respective fabric element is to host the master Fabric Director, does it become the master Fabric Director.

Fabric Directors generate `fabricMIB` tables such as:

- `fabricElemPortTable`
- `fabricElemConnectorTable`

and also generate `fabricMIB` scalars such as:

- `fabricElemMgrIpAddress`
- `fabricElemName`
- `fabricElemNumPorts`
- `fabricElemOperStatus`
- `fabricElemNumConnectors`

Master Fabric Directors additionally generate `fabricMIB` tables such as:

- `fabricMgmtElemTable`
- `fabricMgmtExtElemTable`
- `fabricMgmtLinkTable`

and generate `fabricMIB` scalars such as:

- `fabricMgmtFabricDescr`

- `fabricMgmtFabricType`
- `fabricMgmtFabricTopology`
- `fabricMgmtFabricOperStatus`
- `fabricMgmtFabricName`

Related Information

- *Switch Reference*, `SUN-FABRIC-MIB` mib
- "`fdconfig` Command" on page 71
- "Create a Fabric Configuration" on page 72
- "Add or Remove an Element From the Fabric Configuration" on page 75
- "Modify an Element of the Fabric Configuration" on page 77
- "Exchange Master Fabric Director Status" on page 79

`fdconfig` Command

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 name of the fabric element. Can be the host name if desired.
- 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.

Note – At this time, the `fdconfig` command is only available to the `root` user.

Note – See *Switch Reference*, `fdconfig` command, for more information.

Related Information

- *Switch Reference*, `fdconfig` command
- "Fabric Director and Fabric Elements" on page 70
- "Create a Fabric Configuration" on page 72
- "Add or Remove an Element From the Fabric Configuration" on page 75
- "Modify an Element of the Fabric Configuration" on page 77
- "Exchange Master Fabric Director Status" on page 79

▼ Create a Fabric Configuration

Before you can use the fabric MIBs, you must first create a fabric configuration. You must create identical configurations on all elements of the fabric.

Note – At this time, the `fdconfig` command is only available to the `root` user.

1. On the management controller, initiate a fabric configuration seesion.

```
# fdconfig start-fabric-config empty
#
```

2. Define the fabric's name.

```
# fdconfig define-fabric-name name
```

where *name* is the identifier of the fabric. For example:

```
# fdconfig define-fabric-name test
Fabric name updated
#
```

3. **Define an element of the fabric.**

```
# fdconfig define-element -name name -ip IP_address -type type -role
role1 [-role role2]
```

where:

- *name* is the identifier of the element (can be host name).
- *IP_address* is the IP address of an element.
- *type* is the identifer of the type of element:
  - `sw36` – Sun Datacenter InfiniBand Switch 36
  - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch
  - `unknown` – The element's 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's role is unknown.

For example:

```
# fdconfig define-element -name primary -ip 123.45.67.89 -type sw36 -role
leafSwitch
Element added
#
```

4. **Verify the fabric configuration.**

```
# fdconfig list-in-progress-fabric-config
Name                             IP Addr             Type    Role(s)
isMaster        Fabricname
---------------------------------         -----------------        ------
--------------------    --------        ----------
primary                         123.45.67.89         sw36    leafSwitch
yes             test
#
```

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.

```
# fdconfig abort
In progress config aborted
#
```

- Otherwise, go to Step 6.

6. **End the configuration session and commit to the new configuration.**

```
# fdconfig complete-fabric-config
#
```

7. **Perform Step 1 through Step 6 on the management controllers of all elements of the fabric.**

8. **Reinitiate the fabric configuration on the element which is to host the master Fabric Director.**

```
# fdconfig start-fabric-config from-current
#
```

9. **Configure the element to host the master Fabric Director.**

Note – There must be only one element hosting the master Fabric Director within the fabric.

```
# fdconfig set-master name
```

where *name* is the identifier assigned to the element to host the master Fabric Director. For example:

```
# fdconfig set-master primary
Element primary updated
#
```

10. **End the configuration session and commit to the new configuration.**

```
# fdconfig complete-fabric-config
#
```

11. **Perform Step 8 through Step 10 on the management controllers of all elements of the fabric.**

    The configuration files and fabric MIBs are created.

Related Information

- *Switch Reference*, `fdconfig` command
- "Fabric Director and Fabric Elements" on page 70
- "`fdconfig` Command" on page 71
- "Add or Remove an Element From the Fabric Configuration" on page 75
- "Modify an Element of the Fabric Configuration" on page 77
- "Exchange Master Fabric Director Status" on page 79

▼ Add or Remove an Element From the Fabric Configuration

Note – At this time, the `fdconfig` command is only available to the `root` user.

1. **Initiate a fabric configuration seesion.**

```
# fdconfig start-fabric-config from-current
#
```

2. **Add an element to the fabric configuration.**

```
# fdconfig define-element -name name -ip IP_address -type type -role
role1 [-role role2]
```

where:

- *name* is the identifier of the element (can be host name).
- *IP_address* is the IP address of an element.
- *type* is the identifer of the type of element:
  - `sw36` – Sun Datacenter InfiniBand Switch 36
  - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch
  - `unknown` – The element's 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's role is unknown.

For example:

```
# fdconfig define-element -name secondary -ip 123.45.67.90 -type
sw36gw -role gateway
Element added
#
```

3. Or, remove an element from the fabric configuration.

```
# fdconfig remove-element -name name
```

where *name* is the identifier of the element. For example.

```
# fdconfig remove-element -name secondary
Element deleted
#
```

4. Verify the fabric configuration.

```
# fdconfig list-in-progress-fabric-config
Name                              IP Addr              Type    Role(s)
isMaster        Fabricname
---------------------------------          -----------------         ------
--------------------    --------         ----------
primary                          123.45.67.89         sw36    leafSwitch
yes             test
secondary                         123.45.67.90         sw36gw  gateway
no              test
#
```

5. End the configuration session and commit to the new configuration.

```
# fdconfig complete-fabric-config
#
```

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 72.

Related Information

- *Switch Reference*, `fdconfig` command
- "Fabric Director and Fabric Elements" on page 70
- "`fdconfig` Command" on page 71
- "Create a Fabric Configuration" on page 72
- "Modify an Element of the Fabric Configuration" on page 77
- "Exchange Master Fabric Director Status" on page 79

▼ Modify an Element of the Fabric Configuration

You can modify the parameters of an element.

Note – To change the master Fabric Director status, see "Exchange Master Fabric Director Status" on page 79.

Note – At this time, the `fdconfig` command is only available to the `root` user.

---

1. **Initiate a fabric configuration seesion.**

```
# fdconfig start-fabric-config from-current
#
```

2. **Modify an element of the fabric configuration.**

```
# fdconfig redefine-element -name name [-ip IP_address|-type
type|-role role1 [-role role2]]
```

   where:

   - *name* is the identifier of the element (can be host name).
   - *IP_address* is the IP address of an element.
   - *type* is the identifer of the type of element:
     - `sw36` – Sun Datacenter InfiniBand Switch 36
     - `sw36gw` – Sun Network QDR InfiniBand Gateway Switch

- `unknown` – The element's 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's role is unknown.

  For example:

```
# fdconfig redefine-element -name primary -role spineSwitch
Element modified
#
```

3. **Verify the fabric configuration.**

```
# fdconfig list-in-progress-fabric-config
Name                              IP Addr              Type    Role(s)
isMaster        Fabricname
---------------------------------         -----------------         ------
--------------------    --------         ----------
primary                          123.45.67.89         sw36    spineSwitch
yes             test
secondary                         123.45.67.90          sw36gw  gateway
no              test
#
```

4. **End the configuration session and commit to the new configuration.**

```
# fdconfig complete-fabric-config
#
```

5. **Perform Step 1 through Step 4 on the management controllers of all elements of the fabric.**

Related Information

- *Switch Reference*, `fdconfig` command
- "Fabric Director and Fabric Elements" on page 70
- "`fdconfig` Command" on page 71
- "Create a Fabric Configuration" on page 72
- "Add or Remove an Element From the Fabric Configuration" on page 75
- "Exchange Master Fabric Director Status" on page 79

▼ Exchange Master Fabric Director Status

Should you need to move the master Fabric Director status from one element to another, follow this procedure.

Note – There must be one element in the fabric hosting the master Fabric Director.

Note – At this time, the `fdconfig` command is only available to the `root` user.

---

1. Initiate a fabric configuration seesion.

```
# fdconfig start-fabric-config from-current
#
```

2. Remove master Fabric Director status from the first element.

```
# fdconfig reset-master name
```

where *name* is the identifier of the element. For example:

```
# fdconfig reset-master primary
Element primary updated
#
```

3. Assign master Fabric Director status to the new element.

```
# fdconfig set-master name
```

where *name* is the identifier of the element. For example:

```
# fdconfig set-master secondary
Element secondary updated
#
```

4. Verify the fabric configuration.

```
# fdconfig list-in-progress-fabric-config
Name                            IP Addr             Type    Role(s)
isMaster        Fabricname
--------------------------------         -----------------        ------
--------------------    --------        ----------
```

```
primary                          123.45.67.89         sw36   spineSwitch
no              test
secondary                         123.45.67.90         sw36gw  gateway
yes             test
#
```

5. **End the configuration session and commit to the new configuration.**

```
# fdconfig complete-fabric-config
#
```

6. **Perform Step 1 through Step 5 on the management controllers of all elements of the fabric.**

Related Information

- *Switch Reference*, `fdconfig` command
- "Fabric Director and Fabric Elements" on page 70
- "`fdconfig` Command" on page 71
- "Create a Fabric Configuration" on page 72
- "Add or Remove an Element From the Fabric Configuration" on page 75
- "Modify an Element of the Fabric Configuration" on page 77

Partitioning the InfiniBand Fabric

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

- "`smpartition` Command" on page 81
- "Remove User Partitions for Firmware Downgrades" on page 82
- "Use the `smnodes` Command" on page 83
- "Determine Partitions and P_Keys" on page 84
- "Create a User Partition" on page 85
- "Add or Remove a Port From a Partition" on page 87
- "Modify a Partition or Port" on page 89
- "Delete a Partition" on page 91

Related Information

- "Monitoring the InfiniBand Fabric" on page 51
- "Controlling the InfiniBand Fabric" on page 61
- "Configuring the Fabric Director Node List" on page 69

`smpartition` Command

You can partition your InfiniBand fabric using the `smpartition` command and its subcommands. The `smpartition` command is issued from the switch 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 which 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.

See *Switch Reference*, `smpartition` command, for more information.

Note – There is a hardware limitation of 127 user-defined and 1 default partition for each port of the ConnectX2 chips (in HCAs) and for each InfiniBand port of the BridgeX chips (in gateways).

Related Information

- *Switch Reference*, `smpartition` command
- "Remove User Partitions for Firmware Downgrades" on page 82
- "Use the `smnodes` Command" on page 83
- "Determine Partitions and P_Keys" on page 84
- "Create a User Partition" on page 85
- "Add or Remove a Port From a Partition" on page 87
- "Modify a Partition or Port" on page 89
- "Delete a Partition" on page 91

▼ Remove User Partitions for Firmware Downgrades

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. **On the management controller of the master Subnet Manager, remove all user defined partitions.**

   See "Delete a Partition" on page 91.

2. **Restore the default partition to default settings.**

   See "Modify a Partition or Port" on page 89, however use this command line to modify the partition.

```
# smpartition modify -pkey 0x7fff -port ALL_CAS -flag ipoib -m full
```

3. **Consider your next step:**

   - If your InfiniBand fabric has just one Subnet Manager, downgrade the firmware.

     See *Switch Remote Administration*, upgrading the firmware.

   - If your InfiniBand fabric has multiple Subnet Managers, go to Step 4.

4. **Remove all entries from the Subnet Manager node list.**

```
# smnodes delete IP_address [IP_address ...]
```

   where *IP_address* are the IP addresses of the Subnet Manager nodes. For example:

```
# smnodes delete 123.45.67.89 123.45.67.90
#
```

5. **Repeat Step 4 on the management controllers of all Subnet Managers in the InfiniBand fabric.**

6. **Downgrade the firmware.**

   See *Switch Remote Administration*, upgrading the firmware.

Related Information

- *Switch Reference*, `smpartition` command
- *Switch Reference*, `smnodes` command
- "`smpartition` Command" on page 81
- "Use the `smnodes` Command" on page 83
- "Determine Partitions and P_Keys" on page 84
- "Create a User Partition" on page 85
- "Add or Remove a Port From a Partition" on page 87
- "Modify a Partition or Port" on page 89
- "Delete a Partition" on page 91

▼ Use the `smnodes` Command

If you are partitioning your InfiniBand fabric, a list of valid Subnet Manager nodes must exist in the file system of every management controller running a Subnet Manager. You create this file with the `smodes` command. The file contains a list of IP addresses of all active management controllers running a Subnet Manager in your fabric. The file 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 management controller, propagate the Subnet Manager nodes file with the IP addresses of all Subnet Manager nodes.**

```
# smnodes add IP_address IP_address ...
```

   where *IP_address* is the IP address of each management controller hosting a Subnet Manager. For example:

```
# smnodes add 123.45.67.89 123.45.67.90
#
```

3. **Repeat** Step 2 **for all management controllers in the InfiniBand fabric.**

4. **Determine how you will partition your InfiniBand fabric.**

   See "Determine Partitions and P_Keys" on page 84.

Related Information

- *Switch Reference*, `smodes` command
- "`smpartition` Command" on page 81
- "Remove User Partitions for Firmware Downgrades" on page 82
- "Determine Partitions and P_Keys" on page 84
- "Create a User Partition" on page 85
- "Add or Remove a Port From a Partition" on page 87
- "Modify a Partition or Port" on page 89
- "Delete a Partition" on page 91

▼ Determine Partitions and P_Keys

The switch supports several partitions in the InfiniBand fabric. You assign a P_Key to each partition as a simplified means of identifying the partition to the Subnet Manager. P_Keys are 15-bit integers and have a value of `0x1` to `0x7fff`. 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.

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 52
   - "Identify All HCAs in the Fabric" on page 53

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.**

4. **(Optional) If you will also configure vLANs, assign a unique vLAN number to similar P_Keys in the text file.**

Related Information

- *Switch Reference*, `ibswitches` command
- *Switch Reference*, `ibhosts` command
- "`smpartition` Command" on page 81
- "Remove User Partitions for Firmware Downgrades" on page 82
- "Use the `smnodes` Command" on page 83
- "Create a User Partition" on page 85
- "Add or Remove a Port From a Partition" on page 87
- "Modify a Partition or Port" on page 89
- "Delete a Partition" on page 91

▼ 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 89 for an example.

---

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

   See "Use the `smnodes` Command" on page 83.

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

```
# smpartition start
#
```

3. **Create the user partition.**

```
# smpartition create -n partition_name -pkey p_key[-flag [ipoib, mtu
mtu, rate rate, sl sl, scope scope]][-m defmember]
```

where:

- *partition_name* is the identifier of the InfiniBand partition.
- *p_key* is the partition key (`1` – `7fff` or `default`).
- *mtu* is the value of the MTU.
- *rate* is the throughput of a link (link width + link speed) in Gbps from `2.5`, `5`, `10`, `20`, `30` to a maximum of `40`.
- *sl* is the service level.
- *scope* is the multicast address scope value (`0` – `F`).

Note – The `mtu`, `rate`, `sl`, and `scope` parameters are for the multicast group created when `ipoib` (IP over InfiniBand) is configured for the partition.

- *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:

```
# smpartition create -n testpartition -pkey 5 -m full
#
```

4. Verify the partition configuration.

```
# 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:
#
```

5. Consider your next step.

- If you want to add GUIDs, see "Add or Remove a Port From a Partition" on page 87.
- If you want to modify the configuration, see "Modify a Partition or Port" on page 89.
- If you want to delete the configuration altogether, see "Delete a Partition" on page 91.

- Otherwise, go to Step 6.

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

```
# smpartition commit
#
```

Related Information

- *Switch Reference*, `smpartition` command
- "`smpartition` Command" on page 81
- "Remove User Partitions for Firmware Downgrades" on page 82
- "Use the `smnodes` Command" on page 83
- "Determine Partitions and P_Keys" on page 84
- "Add or Remove a Port From a Partition" on page 87
- "Modify a Partition or Port" on page 89
- "Delete a Partition" on page 91

▼ 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.**

```
# smpartition start
#
```

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

 # smpartition add -n *partition_name*|**-pkey** *p_key* **-port** *port*|**ALL**|**ALL_CAS**|**ALL_SWITCHES**|**ALL_ROUTERS**|**SELF [-m** *member***]**

where:

- *partition_name* is the identifier of the InfiniBand partition.
- *p_key* is the partition key (`1` – `7fff`).
- *port* is the GUID of the port, or the special parameter, to add:
  - `ALL` – All of the CAs, switches, and routers in the InfiniBand fabric.
  - `ALL_CAS` – All CAs in the InfiniBand fabric.
  - `ALL_SWITCHES` – All switches.

- ALL_ROUTERS – All routers.
- SELF – The Master Subnet Manager.

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

For example:

```
# smpartition add -n testpartition -port 00212800013e9313 00212800013e9314
00212800013e93f7
#
```

3. Or, remove ports from the partition.

```
# smpartition remove -n partition_name|-pkey p_key -port
port|ALL|ALL_CAS|ALL_SWITCHES|ALL_ROUTERS|SELF
```

where:

- *partition_name* is the identifier of the InfiniBand partition.
- *p_key* is the partition key (1 – 7fff).
- *port* is the GUID of the port, or the special parameter, to remove:
  - ALL – All of the CAs, switches, and routers in the InfiniBand fabric.
  - ALL_CAS – All CAs in the InfiniBand fabric.
  - ALL_SWITCHES – All switches.
  - ALL_ROUTERS – All routers.
  - SELF – The Master Subnet Manager.

For example:

```
# smpartition remove -n testpartition -port 00212800013e9314
#
```

4. Verify the partition configuration.

```
# 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;
#
```

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

```
# smpartition commit
#
```

Related Information

- *Switch Reference*, `smpartition` command
- "`smpartition` Command" on page 81
- "Remove User Partitions for Firmware Downgrades" on page 82
- "Use the `smnodes` Command" on page 83
- "Determine Partitions and P_Keys" on page 84
- "Create a User Partition" on page 85
- "Modify a Partition or Port" on page 89
- "Delete a Partition" on page 91

▼ Modify a Partition or Port

You can modify the partition's configuration or the port(s) membership with the `smpartition modify` command.

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

```
# smpartition start
#
```

2. **Modify the partition or port(s).**

```
# smpartition modify -n partition_name|-pkey p_key[-flag [ipoib, mtu mtu, rate rate,
sl sl, scope scope]]|[-port port|ALL|ALL_CAS|ALL_SWITCHES|ALL_ROUTERS|SELF [-m
member]
```

where:

- *partition_name* is the identifier of the InfiniBand partition.
- *p_key* is the partition key (`1` – `7fff`).
- *mtu* is the value of the MTU.
- *rate* is the throughput of a link (link width + link speed) in Gbps from `2.5`, `5`, `10`, `20`, `30` to a maximum of `40`.
- *sl* is the service level.

- *scope* is the multicast address scope value (`0` – `F`).

Note – The `mtu`, `rate`, `sl`, and `scope` parameters are for the multicast group created when `ipoib` (IP over InfiniBand) is configured for the partition.

- *port* is the GUID of the port, or the special parameter, to modify:
  - `ALL` – All of the CAs, switches, and routers in the InfiniBand fabric.
  - `ALL_CAS` – All CAs in the InfiniBand fabric.
  - `ALL_SWITCHES` – All switches.
  - `ALL_ROUTERS` – All routers.
  - `SELF` – The Master Subnet Manager.
- *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:

```
# smpartition modify -pkey 0x7fff -port ALL_CAS -flag -m limited
#
```

3. Verify the partition configuration.

```
# 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;
#
```

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

```
# smpartition commit
#
```

Related Information

- *Switch Reference*, `smpartition` command
- "`smpartition` Command" on page 81

- "Remove User Partitions for Firmware Downgrades" on page 82
- "Use the `smnodes` Command" on page 83
- "Determine Partitions and P_Keys" on page 84
- "Create a User Partition" on page 85
- "Add or Remove a Port From a Partition" on page 87
- "Delete a Partition" on page 91

▼ 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.

```
# smpartition start
#
```

2. Delete the partition.

```
# smpartition delete -n partition_name|-pkey p_key
```

where:

- *partition_name* is the identifier of the InfiniBand partition.
- *p_key* is the partition key (`1` – `7fff`).

For example:

```
# smpartition delete -n testpartition
#
```

3. End the configuration session and commit the configuration to the active partition configuration.

```
# smpartition commit
#
```

Related Information

- *Switch Reference*, `smpartition` command
- "`smpartition` Command" on page 81
- "Remove User Partitions for Firmware Downgrades" on page 82
- "Use the `smnodes` Command" on page 83

Index

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

V